

Variational Methods For Continuous-Domain Inverse Problems: the Quest for the Sparsest Solution

Présentée le 16 septembre 2022

Faculté des sciences et techniques de l'ingénieur
Laboratoire d'imagerie biomédicale
Programme doctoral en génie électrique

pour l'obtention du grade de Docteur ès Sciences

par

Thomas Jean DEBARRE

Acceptée sur proposition du jury

Prof. D. N. A. Van De Ville, président du jury
Prof. M. Unser, Dr J. R. P. Fageot, directeurs de thèse
Dr R. Gribonval, rapporteur
Dr V. Duval, rapporteur
Prof. S. van de Geer, rapporteuse

A mes grands-parents, François et Fanfette.

Abstract

The goal of this thesis is to study continuous-domain inverse problems for the reconstruction of sparse signals and to develop efficient algorithms to solve such problems computationally. The task is to recover a signal of interest as a continuous function from a finite number of measurements. This problem being severely ill-posed, we choose to favor sparse reconstructions. We achieve this by formulating an optimization problem with a regularization term involving the total-variation (TV) norm for measures. However, such problems often lead to nonunique solutions, some of which, contrary to expectations, may not be sparse. This requires particular care to assert that we reach a desired sparse solution.

Our contributions are divided into three parts. In the first part, we propose exact discretization methods for large classes of TV-based problems with generic measurement operators for one-dimensional signals. Our methods are based on representer theorems that state that our problems have spline solutions. Our approach thus consists in performing an exact discretization of the problems in spline bases, and we propose algorithms which ensure that we reach a desired sparse solution. We then extend this approach to signals that are expressed as a sum of components with different characteristics. We either consider signals whose components are sparse in different bases or signals whose first component is sparse, and the other is smooth.

In the second part, we consider more specific TV-based problems and focus on the identification of cases of uniqueness. Moreover, in cases of nonuniqueness, we provide a precise description of the solution set, and more specifically of the sparsest solutions. We then leverage this theoretical study to design efficient algorithms that reach such a solution. In this line, we consider the problem of interpolating one-dimensional data points with second-order TV regularization. We also study this same problem with an added Lipschitz constraint to favor stable solutions. Finally, we consider the problem of the recovery of periodic splines with low-frequency Fourier measurements, which we prove to always have a unique solution.

In the third and final part, we apply our sparsity-based frameworks to various real-world problems. Our first application is a method for the fitting of sparse curves to contour data. Finally, we propose an image-reconstruction method for scanning transmission X-ray microscopy.

ABSTRACT

Keywords: continuous-domain inverse problems, sparsity, total variation, representer theorems, splines, proximal algorithms, functional analysis, optimization, duality theory.

Résumé

Cette thèse se concentre sur l'étude de problèmes inverses dans le domaine continu pour la reconstruction de signaux parcimonieux, et sur le développement d'algorithmes efficaces pour résoudre ces problèmes en pratique. L'objectif d'un problème inverse est de reconstruire un signal d'intérêt sous forme de fonction continue à partir d'un nombre fini de mesures. Ce type de problème étant extrêmement mal posé, nous choisissons de favoriser les reconstructions parcimonieuses. Pour ce faire, nous formulons un problème d'optimisation avec un terme de régularisation basé sur la variation totale (VT) au sens des mesures. Cependant, de tels problèmes conduisent souvent à des solutions multiples, dont certaines, contrairement aux attentes, peuvent ne pas être parcimonieuses. Cela nécessite une attention particulière pour s'assurer que nous atteignons une solution parcimonieuse souhaitée.

Nos contributions sont divisées en trois parties. Dans la première partie, nous proposons des méthodes de discrétisation exactes pour de grandes classes de problèmes basés sur la VT avec des opérateurs de mesure génériques pour des signaux unidimensionnels. Nos méthodes sont basées sur des théorèmes du représentant qui stipulent que nos problèmes ont des solutions splines. Notre approche consiste donc à effectuer une discrétisation exacte des problèmes dans des bases de splines, et nous proposons des algorithmes qui garantissent que nous atteignons une solution parcimonieuse souhaitée. Nous étendons ensuite cette approche à des signaux exprimés comme une somme de composantes ayant des caractéristiques différentes. Nous considérons soit des signaux dont les composantes sont parcimonieuses dans différentes bases, soit des signaux dont la première composante est parcimonieuse et l'autre est lisse.

Dans la deuxième partie, nous considérons des problèmes plus spécifiques basés sur la VT et nous nous concentrons sur l'identification des cas d'unicité. De plus, dans les cas de non-unicité, nous fournissons une description précise de l'ensemble des solutions, et en particulier de la solution la plus parcimonieuse. Nous nous appuyons ensuite sur cette étude théorique pour concevoir des algorithmes efficaces permettant d'atteindre une telle solution. Dans cette ligne, nous considérons le problème de l'interpolation de données unidimensionnelles avec une régularisation basée sur la VT du second ordre. Nous étudions également ce même problème avec une contrainte de Lipschitz additionnelle

pour favoriser les solutions stables. Enfin, nous considérons le problème de reconstruction de splines périodiques avec des mesures de Fourier à basse fréquence, dont nous prouvons qu'il a toujours une solution unique.

Dans la troisième et dernière partie, nous appliquons nos cadres basés sur la parcimonie à divers problèmes concrets. Notre première application est une méthode d'ajustement de courbes parcimonieuses à des points de contour. Enfin, nous proposons une méthode de reconstruction d'image pour la microscopie à rayons X à transmission par balayage.

Mots clés : problèmes inverses continus, parcimonie, variation totale, théorèmes du représentant, splines, algorithmes proximaux, analyse fonctionnelle, optimisation, principe de dualité

Acknowledgements

These 4.5 years at the Biomedical Imaging Group (BIG) have been an incredible experience, which would not have been possible without the help of many people whom I would like to acknowledge here.

My sincere thanks Dr. Rémi Gribonval, Dr. Vincent Duval, Prof. Sara van de Geer, Prof. Dimitri Van De Ville, Dr. Julien Fageot, and Prof. Michael Unser for agreeing to be part of my thesis jury and for taking the time to review this (long) manuscript. I appreciated their kind and constructive feedback during my private thesis defense, in particular the honesty of their heart-felt reaction upon receiving this manuscript, along the lines of: “I was very worried when I received 300 pages, but it turned out to be a pleasant read”. They turned the defense not exactly into a picnic, but into a more enjoyable moment than I would have imagined.

I would also like to thank warmly the many people who have helped me proof-read parts of this manuscript: Julien, Pakshal, Matthieu, Laurène, Alexis, Sebastian, Thanh-An, and Jonathan. A special thanks to Julien, who deserves co-authorship for Chapter 2 of this thesis which has greatly contributed to shape, and to Matthieu for his scientific enthusiasm and for always being so kind to me, as demonstrated by his spontaneous offer to proof-read parts of my thesis despite not being a lab mate. I feel very fortunate to have been able to count on these amazing colleagues and friends towards the end, when I no longer had the time or mental energy to proof-read the entire manuscript. Thanks as well to Shayan and Thanh-An, my thesis buddies, for this incongruous week of thesis writing in a greenhouse near Luzern and for their support during sometimes difficult times.

Since most BIG members come and go, its true pillars are its permanent members. Many thanks to Philippe for his very personal brand of dark and often unexpected humor, and for his machine-like proof-reading rigor that gave so much color—literally—to my papers. Thanks as well to Claudia for her secretarial efficiency and for bearing with us scientists who hate reading instructions. I am also grateful for her kindness, best illustrated by her invitation to her mansion in Ticino for a few refreshing days of “télétravail”. Finally, thanks to Daniel for his enduring bubblyness and kindness, be it for professional or personal aspects.

The BIG experience was so special in large part due to the amazing group of people who have gone through the lab, whom I would all like to thank. I feel particularly grateful

towards Julien, Manu, Laurène, and Anaïs for immediately integrating me into their social circle and for making my first steps in the lab an incredibly enjoyable experience. I knew I had chosen the right lab when I was playing Pictionary on an office whiteboard on my first day. The Pictionary fad is now long gone, but it was replaced by others throughout the years—running (briefly), 6 qui prend, Expoding Kittens, and, last but not least, Curve Fever—which have all contributed to the good atmosphere in the lab. Special thanks to Manu for his heart-warming Toulouse accent and enduring kindness; to Thanh-An for quirky personality and humor, the 4 pm fresh beers, the Balavoine duos, and for his deep-rooted gentleness despite his constant bullying; to Harshit for his wacky theories and his peculiar sense of humor; to Pakshal for his decency (no joke) and his kindness that his dark humor cannot conceal; and to Quentin for being a great collaborator, for introducing me to the magical worlds of dual certificates and of Evian, and for making me feel less alone by being an avid fellow train-misser and other travelling-related misfortunes.

Right from the start of my PhD, I was bombarded as teaching assistant (TA) for the Signals and Systems bachelor course. Although I initially did not feel very qualified, I was lucky to benefit from the mentorship of the “Dream Team” Laurène, Manu, Anaïs, who quickly made a competent (hopefully) TA out of me. The team-building lunch/foosball sessions at the LTS2 with MDeff, Yann, and Lionel clearly also helped for that. This Dream Team, later joined by Quentin, Shayan, Alexis, Virginia, and Iciar, together with the excellent students we are lucky to have at EPFL, made teaching a more fulfilling experience than I could have imagined, which completely compensated for the heavy workload.

Among the Dream Team, a special thanks to Laurène, who taught me the few organizational skills that I have when I faced the daunting task of filling her giant—figuratively—shoes as head TA. For that and for so many other things, I am extremely grateful to have overlapped with your time at the BIG. You were the heart and soul of the lab, and you deserve huge credit for turning it into a group of friends with your geniality and your dedication to organizing extra-professional activities. Your departure, even if it is only to the end of the corridor, has left a hole that we can all feel. You are a source of inspiration to me on many aspects, and I am very grateful for your friendship and for all that you have done for me both professionally and personally.

Throughout my PhD, I have only known one office, which I would have exchanged with no other: the BM 4138. Many thanks to my former office mates, who have made the sometimes monotonous life of a researcher so much livelier: to Anaïs for introducing me to the love of biathlon and for her very personal brand of humor of which I am a major fan, and to Iciar for her kindness, constant bubbiness, and for introducing the eras of the “cute office” and of Britney Spears karaoke (free Britney!). Finally, I could never thank my “pote” and thesis buddy, Shayan, enough for all he has brought to me in the past few years. We started out simply as office mates, and we evolved into so many other things: scientific collaborators, math teacher (you) and student (me), fellow TAs, chess and Curve Fever rivals, movie goers, language teachers, etc., and most importantly, great friends. It has been a privilege to watch you mature into the person that you are today,

Acknowledgements

who is quite different from the Shayan I first met, but who has kept his defining qualities: an interesting mix of taarof-based kindness and unbridled bluntness, loyalty, integrity, and being such a great person to hang out with. I deeply admire your thoughtful outlook both on science and life, vast subjects on which I can always count on your insightful advice. Incidentally, you are also an immensely talented scientist, from whom I have learned a lot. I honestly think it would be hard to find a more complementary scientific match for me, and I could not have wished for a better preferential collaborator.

I was lucky enough to have not one, but two thesis directors. My eternal gratitude goes to Julien, my internship supervisor, thesis co-director, scientific mentor, and friend (among other things). You have thoroughly redefined my vision of what being a mentor can involve, and I feel incredibly lucky to have you as mine. Part of the huge impact that you had on me as a scientist can be attributed to the fact that you are a brilliant mathematician and teacher. In addition, your enthusiastic and humble approach to science, your constant willingness to proactively share your knowledge and help others, your ability to collaborate with people with diverse scientific backgrounds, to generate and clarify research questions, and to synthesize information are a huge source of inspiration to me. You have also done wonders for my self-confidence by consistently putting forward the positive aspects of my work, a highly unusual occurrence to me as a product of the French system. Moreover, we have shared many other nonscientific things that I can safely say few PhD students share with their advisors, such as a wall—we were neighbors for some time—; a lot of my furniture and kitchen utensils; countless games of chess, go, badminton, and squash; and many seminars and discussions on varying topics such as math, biology, astrophysics, sports, games, music, or history. I deeply admire your curiosity and intelligence, which go way beyond the realm of science. I am extremely grateful to have you as a friend and a mentor, both for scientific and nonscientific aspects of life.

Finally, a huge thank you to Michael, the most exceptional thesis director I could have hoped for. I truly believe you are the most brilliant scientist I have ever encountered. To this day, I am still in awe at your ability to instantly understand virtually any scientific content, and to always come up with smart and creative comments and suggestions. Nonetheless, you remain extremely humble in your approach to science; I always felt respected during our interactions, and you often sought and valued my opinions, even at times when I myself felt that they did not deserve it. Despite your world-class scientific credentials, you are extremely present in your supervision, a highly uncommon combination of luxuries. Your macro management style, rooted in a deep trust in your students, made me feel respected and pushed me to make myself worthy of your trust, all the while knowing that I could count on your hands-on supervision if I needed it. Your ability to foster collaborations and mentorships within the lab also contributes to no one being left stranded. A pleasing outcome of your legendary bluntness is that you never fail to congratulate and show your appreciation. This together with your contagious passion and enthusiasm for science—which is refreshing to see in a man who has worked for more than 40 years in his field—have been an important motivator for me. These managerial

qualities play an important role in the sustained good atmosphere in the lab, which is by no means a given at EPFL. I feel like I won the lottery by ending up at the BIG, and it was truly a privilege to be your student.

I would now also like to acknowledge the people from outside the BIG who have also played an important role in making these past 5 years in Lausanne a truly unforgettable episode in my life. My BIG-less social life in Lausanne primarily consists of my badminton friends from the incomparable Lausanne université club (LUC), who have been so important when I needed to clear my head and get my mind off of my thesis. Many thanks to Pascal, our “président d’honneur”, for welcoming me into the club with open arms, for his geniality, and for embodying and introducing me to the LUC way of life: post-badminton restaurants, drinking Spritz, my first real (*i.e.*, Swiss) fondue in Gruyères, many hikes all across Switzerland, Trivial Pursuit games, and, most importantly, the love of saunas. For the past 3 years, I had the privilege of being captain of the best interclub team in the world, the LUC 3, which has been—and still is—an incredibly fulfilling experience. Many thanks to my team members past and present for their good spirits, the magnificent post-match apéros, the fun times on and off the courts, the wins (often), and the losses (sometimes). Special thanks to “la championne” Pilar for the gossips and the mind-blowing piano concerts; to Eric for being his uproarious Epicurean self, for his oral theses, and for being a good listener; and to Timon for always being willing to share his amazingly vast culture, for the Total Direct Energie videos, and for being such a devoted friend.

I will conclude the badminton chapter with three of my dearest friends, the members of the “LUC Chardonnerets”: my flatmate Blaise and my downstairs neighbors Mathilde and Thibault. Many thanks to Blaise for taking me out of my comfort zone by offering me to join his flat share. I deeply admire your outlook on life, and I have learned a lot from you in these past months, including about myself. I really enjoy our post-badminton pasta and pesto meals, our weekend coffee grinding and drinking sessions, our movie projections, and in general hanging out with you. Your constant good spirits and your thoughtful advice have really helped me to cheer up at times when my spirits were low, for which I am extremely grateful. Aside perhaps from Thibault himself, I am probably the person who is the most thankful that Mathilde has entered Thibault’s life; this is when she truly entered mine as well. Thank you Mama for your quirky personality, for the Mexican workout sessions, for zigzagging when we go running together so that I can keep up, for your kindness and sensitivity, for being there for me when I’m feeling down, for pushing me to come out of my emotional shell, for always having good advice, and for being a wonderful person in general. Spending time with you always puts a smile on my face, for which I will always be grateful. Finally, what can I say about Thibault? We have spent so much time and shared so many things together in the past few years, it’s hard to know where to begin. It’s best not to count the number of badminton, tennis, skiing, and coffee sessions we shared together, because the number would be scary. Your devotion to your friends, to which you selflessly give so much of your time, energy, moulaga, and emotional support, never ceases to inspire me, and I am extremely grateful to be one of them. Being with you is always an adventure; I deeply admire your ability

Acknowledgements

to enliven any situation with your charisma and your hilarious sense of humor, while also being thoughtful and sensitive when needed. I honestly don't know how I would have endured the Covid confinements without you and Mama, who were virtually the only two people I saw in person for many weeks.

I would also like to acknowledge people from outside of Lausanne who are dear to me, whom I would have liked to see much more of in the past few years. They have nonetheless played an important role in helping me revitalize during and get my mind off my thesis during my vacations. Thanks to Adrien, Nathanaël, Jaybee, and Philippine for their loyal friendship and for being awesome people to hang out with. Thanks as well to all my go-playing friends, and particularly to Antoine, Nyoshi, Jitka, Benjon, Julie, Jojo, Charlotte, Jérôme, and their adorable daughter Léonie, whom I have the privilege of being the godfather. I have known you all for so long, and you are like a second family to me. I love the feeling of instantly being transported 15 years back that I get whenever I am with you guys, despite our lives having taken different paths.

To conclude, I would like to thank my real family, starting with my grand-parents, François and Fanfette. I know no one is prouder than them to see me graduate, and I am so glad to be a source of their pride. Finally, thank you to my family inner circle; my aunt Susie, my sister Alice, and my parents, Rebecca and Olivier, to whom I owe everything and who have shaped me into the person I am today. It's not necessarily easy to do a PhD coming from a family of such brilliant and high-achieving academics. However, I truly appreciate the way you have actively refrained from intruding on my personal PhD space that you sensed that I needed, despite the fact that I know you cared deeply and were dying to help—which you always did when I asked for it. I wish I had spent more time with you in these past years (the blame falls on me, I can be a “fils infâme”!), but I love you all more than I can express.

Lausanne, June 12th 2022

T. D.

Contents

Abstract (English)	i
Résumé (Français)	iii
Acknowledgements	v
Notations	xv
Abbreviations	xxi
List of figures	xxi
List of tables	xxv
1 Introduction	1
1.1 Organization of the Thesis	3
1.2 Literature Review	6
I Background on TV-Based Inverse Problems	11
<hr/>	
2 Native Spaces	13
2.1 Preliminaries on Distribution Theory	13
2.2 The Space of Radon Measures	20
2.3 Derivative Operators	21
2.4 Definition and Topology of $\mathcal{M}_L(\mathbb{K})$	24
3 The Optimality of Splines for Generalized TV-Based Problems	35
3.1 Polynomial Splines	35
3.2 Generalized TV-Based Inverse Problem Formulations	45
3.3 Representer Theorems	46

II Exact Discretization of TV-Based Inverse Problems on a Grid 51

4	Single-Component Signal Models	53
4.1	Introduction	53
4.2	Specification of the Search Space	55
4.3	Exact Discretization in the Search Space	56
4.4	Refining the Grid	64
4.5	Experimental Results	67
4.6	Conclusion	72
5	Hybrid-Spline Signal Models	77
5.1	Introduction	77
5.2	Continuous-Domain Inverse Problem	81
5.3	Exact Discretization	85
5.4	Multiresolution Strategy	91
5.5	Experimental Results	93
5.6	Conclusion	96
6	Sparse-Plus-Smooth Signal Models	99
6.1	Introduction	99
6.2	Native Spaces	103
6.3	Continuous-Domain Inverse Problem	104
6.4	Exact Discretization	105
6.5	Practical Implementation	111
6.6	Experimental Results	116
6.7	Conclusion	118

III (Non)uniqueness and the Study of the Solution Set of TV-Based Inverse Problems 121

7	Interpolation of Data with Second-Order TV Regularization	125
7.1	Introduction	125
7.2	Mathematical Preliminaries	130
7.3	The Solutions of the (g-BPC)	137
7.4	The Sparsest Solution(s) of the (g-BPC)	145
7.5	The Solutions of the (g-BLASSO)	152
7.6	Experimental Results	158
7.7	Conclusion	159
	Appendix 7.A Proof of Proposition 7.1	160

CONTENTS

8	Interpolation of Data Under Lipschitz Constraint	169
8.1	Introduction	169
8.2	Mathematical Preliminaries	173
8.3	Lipschitz-Aware Formulations for Supervised Learning	174
8.4	Finding the Sparsest CPWL Solution	183
8.5	Experimental Results	187
8.6	Conclusion	191
9	Reconstruction of Periodic Signals from Fourier-Domain Measurements	193
9.1	Introduction	193
9.2	Mathematical Preliminaries	196
9.3	Uniqueness of TV-Based Penalized Problems	200
9.4	Uniform Convergence of Grid-Based Methods	203
9.5	B-spline-Based Algorithm	205
9.6	Experimental Results	207
9.7	Conclusion	210
<hr/>		
IV	Applications	213
<hr/>		
10	Coupled Splines for Sparse Curve Fitting	215
10.1	Introduction	215
10.2	Continuous-Domain Formulation	217
10.3	Exact Discretization	224
10.4	Extension to Hybrid B-Spline Contours	225
10.5	Experimental Results	227
10.6	Conclusion	229
11	Hessian Splines for Scanning-Transmission X-Ray Microscopy	235
11.1	Introduction	235
11.2	Imaging Model	236
11.3	Inverse Problem Formulation and Algorithm	239
11.4	Experimental Results	240
11.5	Conclusion	243
<hr/>		
12	Conclusion	245
12.1	Contributions	245
12.2	Outlook	247
<hr/>		
Appendices		249
<hr/>		

Appendix A ADMM for Discrete ℓ_1-based Problems	251
A.1 Proximity Operators	251
A.2 ADMM Iterates	252
A.3 Convergence and Computational Complexity	254
Appendix B Proofs of Representer Theorems	255
B.1 Proof of Theorem 5.1	255
B.2 Proof of Theorem 6.1	259
B.3 Proof of Theorem 8.2	262
B.4 Proof of Theorem 8.3	264
B.5 Proof of Theorem 10.2	266
B.6 Proof of Theorem 10.3	270
<hr/>	
Bibliography	292
Curriculum Vitae	293

Notations

SETS

NOTATION	DESCRIPTION
\mathbb{N}	Natural integers
\mathbb{Z}	Integers
\mathbb{R}	Real numbers
\mathbb{R}^+	Nonnegative real numbers
\mathbb{C}	Complex numbers
\mathbb{T}	One-dimensional torus $[0, 2\pi]$
\mathbb{K}	Domain of functions in function spaces (\mathbb{R} or \mathbb{T})

SEQUENCE SPACES

NOTATION	DESCRIPTION	DOMAIN \rightarrow CODOMAIN
$\mathcal{S}(\mathbb{Z})$	Space of Schwartz sequences	$\mathbb{Z} \rightarrow \mathbb{R}$
$\mathcal{S}(\mathbb{Z}, \mathbb{C})$	Space of complex-valued Schwartz sequences	$\mathbb{Z} \rightarrow \mathbb{C}$
$\mathcal{S}'(\mathbb{Z})$	Space of slowly growing sequences	$\mathbb{Z} \rightarrow \mathbb{R}$
$\mathcal{S}'(\mathbb{Z}, \mathbb{C})$	Space of slowly growing complex-valued sequences	$\mathbb{Z} \rightarrow \mathbb{C}$
$\ell_p(\mathbb{Z})$	Lebesgue space of sequences ($1 \leq p \leq +\infty$)	$\mathbb{Z} \rightarrow \mathbb{R}$
$\ell_{1,L}(\mathbb{Z})$	Native space of B-spline coefficients matched to L	$\mathbb{Z} \rightarrow \mathbb{R}$

(GENERALIZED) FUNCTION SPACES

NOTATION	DESCRIPTION	DOMAIN \rightarrow CODOMAIN
\mathcal{X}	Generic space	
\mathcal{X}'	Topological dual space of \mathcal{X}	$\mathcal{X} \rightarrow \mathbb{R}$
$\mathcal{S}(\mathbb{K})$	Schwartz space	$\mathbb{K} \rightarrow \mathbb{R}$
$\mathcal{S}(\mathbb{K}, \mathbb{C})$	Complex-valued Schwartz space	$\mathbb{K} \rightarrow \mathbb{C}$
$\mathcal{S}'(\mathbb{K})$	Space of tempered distributions	$\mathcal{S}(\mathbb{K}) \rightarrow \mathbb{R}$
$\mathcal{S}'(\mathbb{K}, \mathbb{C})$	Space of complex-valued tempered distributions	$\mathcal{S}(\mathbb{K}) \rightarrow \mathbb{C}$
$\mathcal{L}_p(\mathbb{K})$	Lebesgue space ($1 \leq p \leq +\infty$)	$\mathbb{K} \rightarrow \mathbb{R}$
$\mathcal{C}_\infty(\mathbb{K})$	Space of infinitely differentiable functions	$\mathbb{K} \rightarrow \mathbb{R}$
$\mathcal{C}_0(\mathbb{R})$	Space of decaying continuous functions	$\mathbb{R} \rightarrow \mathbb{R}$
$\mathcal{C}(\mathbb{T})$	Space of periodic continuous functions	$\mathbb{T} \rightarrow \mathbb{R}$
$\mathcal{C}_0(\mathbb{T})$	Space of zero-mean periodic continuous functions	$\mathbb{T} \rightarrow \mathbb{R}$
$\mathcal{M}(\mathbb{R})$	Space of real Radon measures	$\mathcal{C}_0(\mathbb{R}) \rightarrow \mathbb{R}$
$\mathcal{M}(\mathbb{T})$	Space of periodic Radon measures	$\mathcal{C}(\mathbb{T}) \rightarrow \mathbb{R}$
$\mathcal{M}_0(\mathbb{T})$	Space of zero-mean periodic Radon measures	$\mathcal{C}_0(\mathbb{T}) \rightarrow \mathbb{R}$
$\mathcal{M}_L(\mathbb{K})$	Native space of operator L	$\mathbb{K} \rightarrow \mathbb{R}$
$\mathcal{M}_{L,\phi}(\mathbb{R})$	$\mathcal{M}_L(\mathbb{R})$ with boundary conditions ϕ	$\mathbb{R} \rightarrow \mathbb{R}$
\mathcal{N}_L	Null space of L	$\mathbb{K} \rightarrow \mathbb{R}$
$\mathcal{C}_L(\mathbb{K})$	Measurement space of operator L	$\mathbb{K} \rightarrow \mathbb{R}$
$\mathcal{M}_{L,h}(\mathbb{K})$	Space of uniform L-splines with knot spacing h	$\mathbb{K} \rightarrow \mathbb{R}$
$\mathcal{M}_{L,h,\phi}(\mathbb{R})$	$\mathcal{M}_{L,h}(\mathbb{R})$ with boundary conditions ϕ	$\mathbb{R} \rightarrow \mathbb{R}$

NORMS

NOTATION	DESCRIPTION	DOMAIN
$\ \cdot\ _p$	ℓ_p norm for vectors	\mathbb{R}^N
$\ \cdot\ _{\ell_p}$	ℓ_p norm for sequences	$\ell_p(\mathbb{Z})$
$\ \cdot\ _{L_p}$	L_p norm for functions	$L_p(\mathbb{K})$
$\ \cdot\ _{\mathcal{M}}$	Total-variation norm for measures	$\mathcal{M}_L(\mathbb{K})$
$\ \cdot\ _{\mathcal{M}_{L,\phi}}$	Banach norm of $\mathcal{M}_L(\mathbb{K})$	$\mathcal{M}_L(\mathbb{K})$
$\ \cdot\ _{\mathcal{C}_{L,\phi,\mathcal{P}}}$	Banach norm of $\mathcal{C}_L(\mathbb{K})$	$\mathcal{C}_L(\mathbb{K})$
$\ \cdot\ _*$	Nuclear norm for matrices	$\mathbb{R}^{M \times N}$

NOTATIONS

OPERATORS

NOTATION	DESCRIPTION	DOMAIN \rightarrow CODOMAIN
L	Differential operator	$\mathcal{S}'(\mathbb{K}) \rightarrow \mathcal{S}'(\mathbb{K})$
\mathcal{F}	Generalized Fourier transform	$\mathcal{S}'(\mathbb{R}, \mathbb{C}) \rightarrow \mathcal{S}'(\mathbb{R}, \mathbb{C})$
D	Derivative operator	$\mathcal{S}'(\mathbb{K}) \rightarrow \mathcal{S}'(\mathbb{K})$
I	Identity operator	$\mathcal{S}'(\mathbb{K}) \rightarrow \mathcal{S}'(\mathbb{K})$
ν	Measurement operator	$\mathcal{M}_L(\mathbb{K}) \rightarrow \mathbb{R}^M$
L_ϕ^{-1}	Right-inverse operator of $L : \mathcal{M}_L(\mathbb{R}) \rightarrow \mathcal{M}(\mathbb{R})$	$\mathcal{M}(\mathbb{R}) \rightarrow \mathcal{M}_L(\mathbb{R})$
L^\dagger	Moore-Penrose pseudoinverse of L	$\mathcal{M}(\mathbb{R}) \rightarrow \mathcal{M}_{L,h}(\mathbb{T})$
L^*	Adjoint operator of $L : \mathcal{X}' \rightarrow \mathcal{Y}'$	$\mathcal{Y} \rightarrow \mathcal{X}$

SEQUENCES

NOTATION	DESCRIPTION
$\delta[n]$	Unit sample sequence
$u[n]$	Unit step sequence
$\widehat{L}[n]$	Frequency response of an operator $L : \mathcal{S}'(\mathbb{T}) \rightarrow \mathcal{S}'(\mathbb{T})$
$\widehat{f}[n]$	Frequency response of a periodic (generalized) function $f \in \mathcal{S}'(\mathbb{T})$
$d_L[n]$	Finite-difference filter of B-spline matched to L

(GENERALIZED) FUNCTIONS

NOTATION	DESCRIPTION	SPACE
δ	Dirac distribution	$\mathcal{S}'(\mathbb{R})$
\mathbf{III}	2π -periodic Dirac comb	$\mathcal{S}'(\mathbb{T})$
u	Heaviside step function	$\mathcal{S}'(\mathbb{R})$
rect	Unit rectangle	$\mathcal{S}'(\mathbb{R})$
$(\cdot)_+$	Rectified linear unit (ReLU)	$\mathcal{S}'(\mathbb{R})$
g_ϕ	Schwartz kernel of L_ϕ^{-1}	$\mathcal{S}'(\mathbb{R} \times \mathbb{R})$
ρ_L	Green's function of L	$\mathcal{S}'(\mathbb{R})$
g_L	Periodic Green's function of L	$\mathcal{S}'(\mathbb{T})$
β_L	Causal cardinal B-spline of L	$\mathcal{S}'(\mathbb{R})$
β^n	Centered polynomial B-spline of D^{n+1}	$\mathcal{S}'(\mathbb{R})$
$\beta_{L,h}$	Scaled B-spline of L with knot spacing h	$\mathcal{S}'(\mathbb{R})$
$\beta_{L,h}^{\text{per}}$	Scaled periodic B-spline of L with knot spacing h	$\mathcal{S}'(\mathbb{T})$
\widehat{L}	Frequency response of L	$\mathcal{S}'(\mathbb{R}, \mathbb{C})$
e_k	Complex sinusoid with frequency $k \in \mathbb{Z}$	$\mathcal{S}(\mathbb{T}, \mathbb{C})$
$\mathbb{1}_V$	Indicator function of $V \subset \mathbb{R}$ (values 0 or 1)	$\mathcal{S}'(\mathbb{K})$
ν_m	Measurement functionals	$\mathcal{C}_L(\mathbb{K})$
ϕ_n	Boundary functionals for biorthogonal system of \mathcal{N}_L	$\mathcal{C}_L(\mathbb{R})$

OPTIMIZATION

NOTATION	DESCRIPTION
$\arg \min$	Argument of minimum (set or element when it is unique)
*	Superscript for solutions of optimization problems (e.g., f^* , c^* , \mathbf{c}^*)
opt	Alternative subscript for solutions of optimization problems (e.g., f_{opt} , c_{opt} , \mathbf{c}_{opt})
$E(\cdot, \cdot)$	Data-fidelity cost functional $E : \mathbb{R}^M \times \mathbb{R}^M \rightarrow \mathbb{R}^+$
$\lambda > 0$	Regularization parameter
$\mathcal{R}(\cdot)$	Regularization functional
$\iota_V(\cdot)$	Characteristic function of a set V (values 0 or $+\infty$)
\mathcal{V}	Solution set of a continuous-domain problem
$\mathcal{J}(\cdot)$	Cost functional of a continuous-domain problem
\mathcal{J}^0	Optimal cost of a continuous-domain problem
res	Subscript for restricted continuous-domain problems (e.g., \mathcal{V}_{res} , \mathcal{J}_{res})
d	Subscript for infinite-dimensional discrete problems (e.g., \mathcal{V}_{d} , \mathcal{J}_{d})
f	Subscript for finite-dimensional problems (e.g., \mathcal{V}_{f} , \mathcal{J}_{f})
$\eta(\cdot)$	Dual pre-certificate of an optimization problem ($\eta \in \mathcal{C}_0(\mathbb{R})$ or $\mathcal{C}(\mathbb{T})$)
$\text{sat}_{\pm}(\eta)$	Signed saturation set of η
$\text{supp}_{\pm}(w)$	Signed support of a Radon measure $w \in \mathcal{M}(\mathbb{K})$

LINEAR ALGEBRA

NOTATION	DESCRIPTION	SPACE
\mathbf{M} or $(M_{m,n})_{\substack{1 \leq n \leq N \\ 1 \leq m \leq M}}$ or $(M_{m,n})$	Generic matrix	$\mathbb{R}^{M \times N}$
$M_{m,n}$ or $[\mathbf{M}]_{m,n}$	(m, n) entry of \mathbf{M}	\mathbb{R}
\mathbf{M}^T	Transpose matrix of \mathbf{M}	$\mathbb{R}^{N \times M}$
\mathbf{M}^\dagger	Moore-Penrose pseudoinverse of \mathbf{M}	$\mathbb{R}^{N \times M}$
$\sigma(\mathbf{M})_n$	n th singular value of \mathbf{M}	\mathbb{R}^+
\mathbf{v} or $(v_n)_{1 \leq n \leq N}$ or (v_n)	Generic vector	\mathbb{R}^N
v_n or $[\mathbf{v}]_n$	n th entry of \mathbf{v}	\mathbb{R}
\mathbf{v}^T	Transpose vector of \mathbf{v}	\mathbb{R}^N
$\mathbf{0}$	Vector of zeros	\mathbb{R}^N
$\mathbf{1}$	Vector of ones	\mathbb{R}^N

NOTATIONS

MISCELLANEOUS

NOTATION	DESCRIPTION
\triangleq	Definition via an equality
(\cdot)	Placeholder for a variable, <i>e.g.</i> , $f(\cdot - x_0) = (x \mapsto f(x - x_0))$
f^\vee	Flipped function of f ($f^\vee : x \mapsto f(-x)$)
$\langle \cdot, \cdot \rangle_{\mathcal{X}, \mathcal{X}'}$	Duality product between \mathcal{X} and \mathcal{X}'
$\operatorname{Re}(z)$	Real part of complex number $z \in \mathbb{C}$
$\operatorname{Im}(z)$	Imaginary part of complex number $z \in \mathbb{C}$
span	Span
supp	Support
\xrightarrow{w}	Weak convergence
$\xrightarrow{w^*}$	Weak* convergence

Abbreviations

ADMM	Alternating-direction method of multipliers
BPC	Basis pursuit in the continuum
BLASSO	Beurling LASSO (see below)
CPWL	Continuous piecewise linear
CS	Compressed sensing
<i>e.g.</i>	<i>Exempli gratia</i> (Latin for “for example”)
<i>et al.</i>	<i>Et alius</i> (Latin for “and others”)
gBLASSO	Generalized Beurling LASSO (see below)
gBPC	Generalized basis pursuit in the continuum
gTV	Generalized total variation
<i>i.e.</i>	<i>Id est</i> (Latin for “that is”)
i.i.d.	Independent and identically distributed
LASSO	Least absolute shrinkage and selection operator
ReLU	Rectified linear unit
SNR	Signal-to-noise ratio
s.t.	Such that
STXM	Scanning transmission X-ray microscopy
TV	Total variation
TVS	Topological vector space

List of Figures

1.1	Roadmap of the thesis.	3
3.1	Example of a nonuniform D-spline.	36
3.2	Cardinal polynomial B-splines of operators $L = D^{N_0}$ for $N_0 = 1, \dots, 4$. . .	40
4.1	Illustration of staircasing effect before the simplex step for our single-component model.	69
4.2	Measurement types comparison for our single-component model.	70
4.3	Illustration of the effect of grid refinement for our single-component model. . .	71
4.4	Quantitative effect of grid refinement for our single-component model. . .	74
4.5	Comparison between continuous and discrete reconstruction models for our single-component model.	75
5.1	Curve fitting comparison between single-component and hybrid-spline models. . .	94
5.2	Compressed sensing example for hybrid-spline model.	96
6.1	Boundary basis function for smooth component in sparse-plus-smooth model.	112
6.2	Comparison between sparse-plus-smooth and single-component models.	119
7.1	Example of piecewise-linear regression.	126
7.2	Example of a canonical solution and dual certificate for regression with second-order TV regularization.	138
7.3	Example of nonuniqueness for regression with second-order TV regularization.	142
7.4	Localization of the graph of all solutions of a regression problem with second-order TV regularization.	145
7.5	Example of a unique sparsest solution for regression with second-order TV regularization.	150
7.6	Example of a nonunique sparsest solution for regression with second-order TV regularization.	150
7.7	Illustration of the effect of a varying regularization parameter for regression with second-order TV regularization.	160
8.1	Localization of the graph of all solutions of a regression problem with Lipschitz regularization.	179
8.2	Example of a sparsest CPWL interpolant of data points.	186

LIST OF FIGURES

8.3	Illustration of regression with Lipschitz regularization.	188
8.4	Illustration of the limitations of Lipschitz regularization alone.	190
8.5	Illustration of the robustness of Lipschitz-aware regression models to outliers.	191
9.1	Effect of the grid size for the reconstruction of periodic signals with low-pass Fourier series measurements.	208
9.2	Average uniform approximation error with respect to the grid size for the reconstruction of periodic signals with low-pass Fourier series measurements.	209
9.3	Comparison between gTV regularization and low-pass reconstruction of a periodic signal with low-pass Fourier series measurements.	210
10.1	Illustration of the rotation invariance of our curve fitting method with RI-TV regularization.	230
10.2	Illustration of the resilience to noise of our curve fitting method with RI-TV regularization.	231
10.3	Comparison between our hybrid curve fitting method with RI-TV regular- ization and single-component models.	232
10.4	Reconstructions of the Daft Punk logo using our hybrid curve fitting method with RI-TV regularization.	233
11.1	Comparisons of Hessian-spline reconstruction with standard methods for STXM on simulated data.	240
11.2	Example of a Hessian-spline reconstruction for STXM on real data.	242



List of Tables

3.1	Characteristics of N_0 th-order derivative operators.	40
6.1	Relevant filter and their supports for our sparse-plus-smooth model.	110
6.2	SNR comparison between sparse-plus-smooth and single-component reconstructions with respect to the grid size.	118

1 Introduction

In this thesis, we study *linear inverse problems* for the reconstruction of continuous-domain signals. The task in an inverse problem is to recover an unknown *ground-truth* signal s_0 of interest based on some measured data or *observations* $\mathbf{y} \in \mathbb{R}^M$, where M is the number of observations. For example, a computed tomography (CT) scanner solves an inverse problem by reconstructing an image of the density distribution of the sample based on X-ray projection data at various angles. The assumption is that the physical model that governs the data acquisition process, denoted as the operator $\nu : s \mapsto \nu(s) \in \mathbb{R}^M$ and such that $\nu(s_0) \approx \mathbf{y}$, is known. In our CT example, this entails that if the density distribution of the sample were known, then the noiseless X-ray projections could be computed. Discrepancies between $\nu(s_0)$ and \mathbf{y} , which are assumed to be small, can be due to measurement noise and/or to the inaccurate modelling of the acquisition process ν . This physical model ν , which is typically linear, is known as the *forward model*, hence the terminology “linear inverse problem” for the reverse procedure.

Our particular focus is on inverse problem where the signal s_0 is a continuous-domain function. This differs from the vast majority of the existing inverse-problem literature, which is concerned with *discrete* problems, where the signal s_0 is a vector. This is mainly for obvious practical considerations, since vectors are easier to manipulate computationally. However, the forward model ν is often an inherently continuous operation and thus requires to be discretized (*e.g.*, a line integral approximated as a finite sum in CT) in order to formulate a discrete problem. When poorly handled, this discretization can be a massive source of error which is absent from continuous-domain formulations, as uncovered by a host of different works [1–3]. This underscores the need for inverse-problem frameworks where calculations are *exact* in the continuous domain, which require particular care in the discretization process in order to solve them computationally.

The most standard approach to solving an inverse problem is via a *variational formulation*, which consists in minimizing a cost function. The latter involves a *data-fidelity term* that enforces the consistency of the reconstructed signal with the observations, *i.e.*, $\nu(s^*) \approx \mathbf{y}$.

However, minimizing solely a data-fidelity term usually leads to poor results due to noise overfitting and to the following potential flaws in the problem formulation:

- *ill-posedness*, which means that different signals may lead to the same observations \mathbf{y} . This issue is particularly conspicuous in continuous-domain formulations, since the reconstructed signal lives in an infinite-dimensional space, whereas there are a finite number M of observations \mathbf{y} .
- *ill-conditioning*, which means that small variations in the measured data \mathbf{y} may lead to drastic variations in the reconstructed signal s^* . Then, in the presence of noise in the measurements, the reconstruction process is highly unstable;

These issues are typically alleviated by adding a *regularization term* to the cost function, which is solely based on our prior knowledge on the form of the ground-truth signal. Historically, the first kind of regularization that was proposed is ℓ_2 or Tikhonov regularization [4]; the inverse problem then typically amounts to solving a linear system. In recent years, *sparsity*, *i.e.*, the assumption that a signal can be represented with few parameters, has become a prevalent choice of prior. This is due to the well-known observation that many natural signals have sparse (approximate) representations in certain bases, *e.g.*, wavelets for natural images [5]. In discrete problems, sparsity is typically enforced via the ℓ_1 norm, which leads to problems such as the basis pursuit [6] or the LASSO [7]; this strategy is foundational to the field of compressed sensing (CS) [8–11]. The continuous counterpart of the ℓ_1 norm is the *total-variation (TV) norm for measures*, in the sense that it also promotes sparse solutions [12–14].

Sparsity-promoting regularization has proved to be very effective in practice, and is considered to be superior to Tikhonov regularization [15, 16]. However, due to the nondifferentiability of the ℓ_1 norm, the optimization procedure is more involved; it is typically solved with iterative algorithms based on proximal operators [17] such as FISTA [18] or ADMM [19]. Another difficulty is that the underlying optimization problem may not have a unique solution. Moreover, contrary to expectations, not all solutions are necessarily sparse [14, 20]. Although this is also well known in the discrete inverse-problem community [21], it is surprisingly rarely discussed. In fact, standard proximal algorithms are not designed to reach a desired sparse solution, which seemingly defeats the purpose of using sparsity-promoting regularization. This “oversight” is partially due to the fact that CS frameworks typically consider problems where the (sparse) ground-truth signal is the unique solution. However, in practical application where the measurement operator cannot be chosen freely, uniqueness does not hold in general, in which case the design of algorithms that reach a sparse solution is a critical question.

In this thesis, we attempt to contribute to these questions in the context of one-dimensional continuous-domain inverse problems with *generalized TV (gTV) regularization*, *i.e.*, where the TV norm for measures is used in combination with a differential *regularization*

operator L that specifies the transform domain in which the signal is sparse. We propose discretization methods that are *exact* in the continuous domain, and design efficient algorithms that are guaranteed to reach a desired sparse solution in various generic settings. Our discretization methods are based on *representer theorems*, either existing ones [14, 22] or our own, that guarantee that our problems have *spline* (*i.e.*, piecewise-polynomial) solutions. This motivates the use of B-spline bases in our discretization methods. More generally, B-splines are ideal bridges between the analog and digital worlds due to their simple parametric form and their advantageous properties for numerical applications—most notably, their short support [23–25]. Next, we study more particular settings, *i.e.*, with specific measurement and regularization operators ν and L , in which we are able to precisely characterize the solution set of our problem. We leverage these theoretical characterizations to design efficient algorithms that reach the *sparsest* solution. Finally, we show how our continuous-domain sparsity-driven framework can be applied to practical applications.

1.1 Organization of the Thesis

This thesis is divided into four parts; the interplay between the latter is illustrated graphically in Figure 1.1. Their contents, contributions, and relevant publications are summarized below.

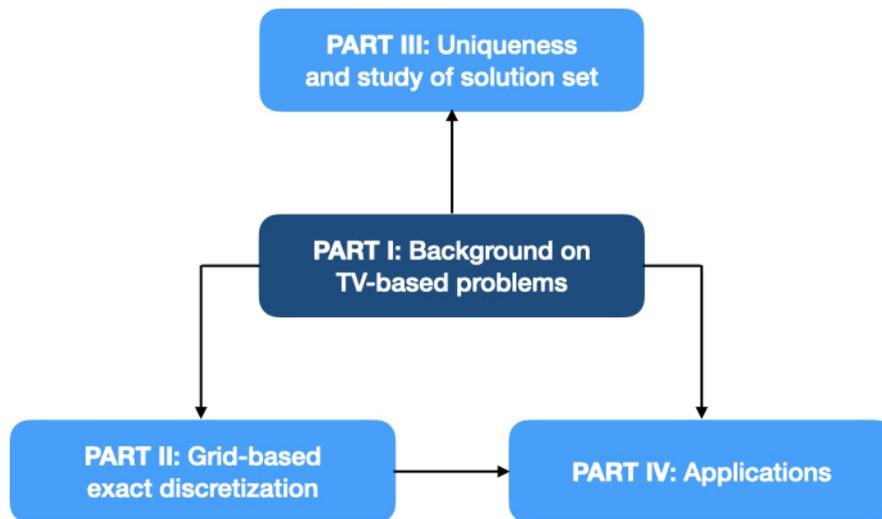


Figure 1.1 – Roadmap of the thesis.

Background on TV-based Inverse Problems (Part I)

In the first part, we provide the mathematical background for the rest of the thesis. In particular, we define the *native spaces* of our continuous-domain optimization problems

of interest in Chapter 2, which are matched to the regularization operator L . This construction relies on [14] and [22] in the nonperiodic and periodic settings, respectively. Finally, in Chapter 3, we present some background on *polynomial splines* and the representer theorems of [14, 22], on which our discretization methods are based, which state that many of our gTV-based problems of interest have spline solutions.

Exact Discretization of TV-Based Inverse Problems on a Grid (Part II)

We then present our grid-based discretization methods for various classes of gTV-based problems. In Chapter 4, we focus on piecewise-polynomial signal models with generic measurement operators ν . Inspired by the representer theorem of [14], we propose a B-spline-based discretization method which is *exact* in the continuous domain, and we propose a multiresolution algorithm that is guaranteed to reach a desired sparse solution. Next, in Chapter 5, we extend to multicomponent signal models $s = s_1 + s_2$, where s_i is assumed to be sparse in a basis specified by a regularization operator L_i for $i \in \{1, 2\}$, with $L_1 \neq L_2$. Our discretization method is again motivated by our multicomponent representer theorem which guarantees that such problems admit a solution that is a *hybrid spline*, *i.e.*, the sum of a L_1 - and a L_2 -spline. Finally, in Chapter 6, we consider another multicomponent signal model, but where the first component s_1 is assumed to be sparse in a basis specified by L_1 , and the second component s_2 is assumed to be smooth. Likewise, we prove a representer theorem for such problems which motivates our B-spline-based discretization method.

Relevant publications:

- T. Debarre, J. Fageot, H. Gupta, and M. Unser, “Solving Continuous-Domain Problems Exactly with Multiresolution B-Splines”, in *Proceedings of the Forty-Fourth IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP’19)*, Brighton, United Kingdom, May 2019, pp. 5122–5126.
- T. Debarre, J. Fageot, H. Gupta, and M. Unser, “B-Spline-Based Exact Discretization of Continuous-Domain Inverse Problems with Generalized TV Regularization”, *IEEE Transactions on Information Theory*, vol. 65, no. 7, pp. 4457–4470, Jul. 2019.
- T. Debarre, S. Aziznejad, and M. Unser, “Hybrid-Spline Dictionaries for Continuous-Domain Inverse Problems”, *IEEE Transactions on Signal Processing*, vol. 67, no. 22, pp. 5824–5836, Nov. 2019.
- T. Debarre, S. Aziznejad, and M. Unser, “Continuous-Domain Formulation of Inverse Problems for Composite Sparse-Plus-Smooth Signals”, *IEEE Open Journal of Signal Processing*, vol. 2, pp. 545–558, Sep. 2021.

(Non)uniqueness and the Study of the Solution Set of TV-Based Inverse Problems (Part III)

We then focus on more specific gTV-based problems, *i.e.*, with a specific measurement operator ν and regularization operator L . Using *duality theory* [26], we precisely describe the solution set, in particular by identifying cases of uniqueness, and we leverage this description to design efficient algorithms that are guaranteed to reach the *sparsest* solution. The first class of problems we consider in Chapter 7 are *interpolation* problems, *i.e.*, where the forward model is a sampling operator $\nu(s) = (s(x_1), \dots, s(x_M))$ (the x_m are the sampling locations), with second-order TV regularization, *i.e.*, the regularization operator L is the second derivative $L = D^2$. In Chapter 8, we also consider interpolation problems with second-order TV regularization, but with an additional constraint on the Lipschitz constant of the reconstruction. Finally, in Chapter 9, we reconsider the reconstruction problem of periodic signals based on low-pass Fourier series measurements, which we prove to always have a unique solution.

Relevant publications:

- T. Debarre, Q. Denoyelle, M. Unser, and J. Fageot, “Sparsest Piecewise-Linear Regression of One-Dimensional Data”, *Journal of Computational and Applied Mathematics*, vol. 406, p. 114044, May 2022.
- S. Aziznejad*, T. Debarre*, and M. Unser, “Sparsest Univariate Learning Models Under Lipschitz Constraint”, *IEEE Open Journal of Signal Processing*, vol. 3, pp. 140–154, Mar. 2022.
- T. Debarre, Q. Denoyelle, and J. Fageot, “On the Uniqueness of Solutions for the Basis Pursuit in the Continuum”, *arXiv preprint arXiv:2009.11855*, Feb. 2022.
- T. Debarre, Q. Denoyelle, and J. Fageot, “TV-Based Spline Reconstruction with Fourier Measurements: Uniqueness and Convergence of Grid-Based Methods”, *arXiv preprint arXiv:2202.05059*, Feb. 2022.

Applications (Part IV)

In this final part, we apply our sparsity-promoting continuous-domain inverse problem formulations and B-spline-based discretization methods to real-world applications. In Chapter 10, we consider the problem of fitting a sparse curve based on a collection of contour points. We prove a representer theorem and propose a discretization method based on Chapters 4 and 5 for single-component and hybrid curve models, respectively. Finally, in Chapter 11, we propose a 2D image-reconstruction algorithm for scanning transmission

* Equal contribution.

X-ray microscopy (STXM), where the measurement operator ν is a sampling operator on a nonuniform grid and with sparsity-promoting Hessian-Schatten regularization [27].

Relevant publications:

- I. Lloréns Jover, T. Debarre, S. Aziznejad, and M. Unser, “Coupled Splines for Sparse Curve Fitting”, *arXiv preprint arXiv:2202.01641*, Feb. 2022.
- T. Debarre, B. Watts, B. Rösner, and M. Unser, “Hessian Splines for Scanning Transmission X-Ray Microscopy”, in *Proceedings of the Seventeenth IEEE International Symposium on Biomedical Imaging (ISBI’20)*, Iowa City IA, USA, Apr. 2020, pp. 199-202.

1.2 Literature Review

We now give a more in-depth literature concerning generalized TV-based optimization problem, which are typically of the form

$$f^* \in \arg \min_f \left(\underbrace{E(\nu(f), \mathbf{y})}_{\text{Data fidelity}} + \lambda \underbrace{\|L\{f\}\|_{\mathcal{M}}}_{\text{Regularization}} \right), \text{ where:} \quad (1.1)$$

- $\mathbf{y} \in \mathbb{R}^M$ are the *observations*;
- $\nu : f \mapsto \nu(f) \in \mathbb{R}^M$ is the *forward model*;
- $E : \mathbb{R}^M \times \mathbb{R}^M \rightarrow \mathbb{R} \cup \{+\infty\}$ is the *data-fidelity cost functional* that penalizes the discrepancy between the observations \mathbf{y} and the measurements $\nu(f^*)$ of the reconstructed signal f^* ;
- $\|\cdot\|_{\mathcal{M}}$ is the sparsity-promoting *TV norm*;
- L is the differential regularization operator that specifies the transform domain in which f^* is sparse;
- $\lambda > 0$ is the regularization parameter that balances the weight between the data-fidelity and regularization terms.

1.2.1 Discrete ℓ_1 Optimization

TV-based regularization lies in the realm of sparsity-promoting regularization techniques, which were first introduced—and are, still, most commonly used—in purely discrete settings. The sparsity of a discrete signal—*i.e.*, a vector—is measured by the ℓ_0 “norm”, which counts the number of its nonzero entries. However, ℓ_0 -based optimization leads to

NP-hard problems; hence, in practice, sparsity is typically enforced via ℓ_1 norm, which can be seen as a convex relaxation of the (nonconvex) ℓ_0 “norm”. Analogously to (1.1), classical discrete problems with ℓ_1 regularization are of the form

$$\mathbf{c}^* \in \arg \min_{\mathbf{c} \in \mathbb{R}^N} \left(\underbrace{E(\mathbf{H}\mathbf{c}, \mathbf{y})}_{\text{Data fidelity}} + \underbrace{\lambda \|\mathbf{L}\mathbf{c}\|_1}_{\text{Regularization}} \right), \quad (1.2)$$

where $\mathbf{H} \in \mathbb{R}^{M \times N}$ is the sensing matrix (discrete forward model), $\mathbf{L} \in \mathbb{R}^{N \times N}$ is the regularization matrix. The choice of \mathbf{L} allows us to promote the sparsity of \mathbf{c} in a chosen transform domain, *e.g.*, a finite difference matrix (discrete TV regularization [28]) or wavelets.

The use of ℓ_1 regularization was popularized in large part by the least absolute shrinkage and selection operator (LASSO) [7]—where \mathbf{L} is the identity matrix—and the basis pursuit [6]—where E is an indicator function—, which were introduced in the late 90s. These problems are at the cornerstone of sparse statistical learning [16, 29] and sparse signal processing [30].

Such problems are also central to the field compressed sensing (CS), which surged in the mid 2000s [8, 10, 11, 20, 31]. In the CS framework, the goal is to recover a signal $\mathbf{c}_0 \in \mathbb{R}^N$ with very few observations $\mathbf{y} \in \mathbb{R}^M$ with $M \ll N$ via problems of the form (1.1). The two key conditions for the recovery to be possible are that \mathbf{c}_0 be sparse and that the sensing matrix \mathbf{H} be *incoherent* [8, 31, 32]. Infinite-dimensional extensions of the CS framework have also been proposed [3, 20, 33–36].

The benefits of ℓ_1 regularization have been extensively documented, including its sparsity-promoting effect [7, 20], CS-type recovery guarantees, or its apparent superiority over Tikhonov ℓ_2 regularization to recover many real-world signals [16, 37]. Moreover, despite its nondifferentiability, numerous efficient algorithms based on the *proximity operator* of the ℓ_1 norm have emerged to solve ℓ_1 -regularized problems [17–19, 38–41].

However, many inverse problems are inherently continuous, in which case the purely discrete approach has several important pitfalls. Firstly, discretization is typically performed in a pixel basis which is not necessarily matched to the characteristics of the underlying continuous-domain signal. Moreover, the discrete forward model \mathbf{H} is often an approximation of its continuous counterpart ν (*e.g.*, the discrete Fourier transform for the continuous Fourier transform), which introduces discretization errors [1–3]. To avoid these issues, it is crucial to formulate and to think about these problems in the continuous domain, even if they must be eventually digitalized in some way in order to be solved computationally. This continuous approach—which includes some of the infinite-dimensional CS frameworks listed above—is adopted by a growing number of works, some of which are reviewed below.

1.2.2 Discretization in Hilbert Function Spaces

Despite its esthetic appeal, discretizing continuous-domain problems with a finite number of observations \mathbf{y} leads to severely ill-posed problems. Indeed, although discrete CS problems such as (1.2) are also ill-posed when $M \ll N$, this is all the more the case when we have an infinite-dimensional search space and finitely many observations. Kernel methods based on quadratic regularization are an elegant way of removing this ill-posedness [42]; they lead to the restriction of the search space to a finite-dimensional subset of a Hilbert space [43–46]. The challenge is then to choose this Hilbert space adequately. These approaches are fruitful; however, they still ultimately revert to the finite-dimensional setting. Moreover, they do not enforce the sparsity of the reconstructed signals, which has now become the prior of choice in model-based reconstruction methods.

1.2.3 Optimization over Radon Measures

A prominent example of continuous-domain inverse-problem formulations with sparsity priors in the literature is the problem of recovering sums of Dirac masses in the continuum, which is known as *sparse spikes deconvolution* (or *super-resolution*) [12, 47–49]. This type of problem has multiple data-science applications, including radio-astronomy [50], super-resolution microscopy [51], or 3D image deconvolution [52]. This question has received considerable attention in the 21st century, including approaches that are not based on TV regularization [53], such as finite rate of innovation (FRI) techniques [54, 55] and Prony’s methods [56–58].

However, the most common approach to sparse spikes deconvolution is to formulate an optimization problem with TV regularization, *i.e.*, problems of the form (1.1) in the absence of a regularization operator L . The underlying optimization problems, either formulated in a constrained form in the noiseless case [49] or in a penalized form known as the Beurling LASSO (BLASSO) [47] in the presence of noise, are thus solved over a nonreflexive Banach space. The role of the TV norm in variational methods has a rich history [59, 60] (see [61, Section 1] for additional references). From a theoretical standpoint, many reconstruction guarantees are proved, such as exact recovery of discrete measures (sums of Dirac masses) in the noiseless case [49, 62], robustness to noise [12, 13, 63, 64], support recovery [65–68] and super-resolution for positive discrete measures [47, 69–74].

From a numerical standpoint, there exist several strategies to solve these problems. A first one is based on spatial discretization which leads back to the LASSO and algorithms such as FISTA [18]. Another approach is to use greedy algorithms [75] such as continuous-domain orthogonal matching pursuit (OMP) [76, 77]. In special setups (typically Fourier measurements), it is possible to reformulate the optimization problems as semidefinite or convex programs [49, 78–80], or to use projected gradient descent methods [81]. Recent

developments based on the Frank-Wolfe (FW) algorithm [82] stemming from [12] solve the BLASSO directly over the space of Radon measures. These FW-based methods improve on the traditional FW algorithm due to the possibility of moving the spikes in the continuous domain to further decrease the objective function [51, 52, 83–85]. Finally, we mention [86] which proposes a nonconvex gradient-descent algorithm.

1.2.4 From Sparse Measures to Splines and Beyond

In recent years, several works have extended the TV-based Dirac recovery framework to smoother continuous-domain signals by considering gTV regularization, *i.e.*, problems such as (1.1) with a nontrivial regularization operator L . Interestingly, the origins of these works predate by far the Dirac recovery literature and can be traced back to the 70's [87, 88]. They have been revived much more recently in [14], in which Unser *et al.* revealed the connection between Problem (1.1) and spline theory for general measurement functionals: the extreme-point solutions are necessarily L -splines (Theorem 3.3). This result was revisited, extended, and refined by several authors [22, 45, 61, 89–94]. On the algorithmic side, a host of recent works have proposed algorithms to solve problems with gTV regularization or closely related ones, either grid-based [95] or grid-free [90, 96, 97].

Part I Background on TV-Based Inverse Problems

In this part, we present the mathematical foundations of the remainder of the thesis. We properly introduce the functional-analysis framework of our generalized total-variation (gTV)-based inverse-problem formulations of the form (1.1). More specifically, problems related to (1.1) will be the focus of our works in Parts II and III, as well as Chapter 10. This part is organized as follows.

- In Chapter 2, we define the search spaces of our Problems (1.1), which are known as *native spaces* of the regularization operator L .
- In Chapter 3, we first provide some background on polynomial splines matched to the operator L , which are at the core of our exact discretization strategies throughout the thesis. Moreover, splines have been proved to be solutions of problems of the form (1.1) by existing representer theorems. We formulate these problems in both the nonperiodic and periodic settings, and we present their corresponding representer theorems, which are taken from [14] and [22], respectively.

2 Native Spaces

In this chapter, we introduce the function spaces over which our optimization problems are performed. The latter are Banach spaces that involve a differential operator L . The terminology *native space* refers to the domain of L . This chapter heavily relies on [14] and [22] for the nonperiodic and periodic settings, respectively. This construction of the native spaces heavily relies on *distribution theory* and dual spaces, the basics of which are presented in Section 2.1. Next, we introduce the space of Radon measures and its associated Banach norm, the total-variation (TV) norm for measures, in Section 2.2. In Section 2.3, we introduce weak derivatives, which allow us to extend the domain of our operator L via distribution theory. Finally, we introduce the native space of L in Section 2.4, and we uncover its Banach structure.

2.1 Preliminaries on Distribution Theory

The construction of our native spaces builds upon distribution theory, which was first introduced by Laurent Schwartz in [98]. The idea behind its use is that common operators are mathematically well defined over function spaces that exclude many functions of interest. By using distribution theory, the domain \mathcal{X} of an operator can be extended to the *dual space* \mathcal{X}' , which typically contains much larger—and hence more interesting—classes of (generalized) functions. For example, it is well known by engineers that the derivative operator D can be applied to piecewise-constant functions in the distributional sense, despite the fact their derivatives are undefined in the classical sense at the jump locations.

2.1.1 Topological Vector Spaces

The most general spaces we will be dealing with are *topological vector spaces* (TVS). Although we will not go into their precise definition—for which we refer to [15, Chapter 1]—, a TVS \mathcal{X} is informally the most general type of vector space that allows for the definition of limits. This is achieved by equipping \mathcal{X} with a *topology*, which specifies

the open sets of \mathcal{X} . TVSs are generalizations of normed spaces, since norms induce a topology. However, a topology is not necessarily induced by a norm or even a metric, as in the case of the Schwartz space introduced in Definitions 2.3 and 2.4.

2.1.2 Topological Dual Spaces

The definition of a topological dual space is given below.

Definition 2.1 (Topological Dual Space). *Let \mathcal{X} be a TVS. The topological dual space of \mathcal{X} , denoted by \mathcal{X}' , is the set of continuous linear functionals $w : \mathcal{X} \rightarrow \mathbb{R}$.*

Given a dual pair $(\mathcal{X}, \mathcal{X}')$, for any $\varphi \in \mathcal{X}$ and $w \in \mathcal{X}'$, the action of w on φ —or *duality product* between φ and w —is denoted by

$$w(\varphi) = \langle \varphi, w \rangle_{\mathcal{X} \times \mathcal{X}'} \in \mathbb{R} \quad (2.1)$$

In all that follows, the subscript $\mathcal{X} \times \mathcal{X}'$ that specifies the nature of the duality product will often be omitted when there is no ambiguity as to which duality product is implied, e.g., $\langle \varphi, w \rangle$. Moreover, the order in the duality product might not always be consistent, e.g., $\langle w, \varphi \rangle$ instead of $\langle \varphi, w \rangle$. Finally, a topological dual space \mathcal{X}' of \mathcal{X} will be referred to as a dual space or simply a dual throughout this thesis, and \mathcal{X} will be referred to as the *predual* of \mathcal{X}' .

2.1.3 Dual Spaces and Topologies

The existence of a dual pair \mathcal{X} and \mathcal{X}' gives rise to various classical topologies over both spaces, which we review below.

Strong Topology

When \mathcal{X} is a normed vector space over a scalar field \mathbb{K} , \mathcal{X}' inherits the normed vector-space structure of \mathcal{X} by way of the dual norm defined as

$$\|w\|_{\mathcal{X}'} \triangleq \sup_{\varphi \in \mathcal{X}: \|\varphi\|_{\mathcal{X}} \leq 1} |\langle \varphi, w \rangle_{\mathcal{X} \times \mathcal{X}'}| \quad (2.2)$$

for any $w \in \mathcal{X}'$. The topology inherited from the dual norm over \mathcal{X}' is known as the *strong topology*. The dual space \mathcal{X}' is complete with respect to the strong topology, and is thus a Banach space.

Weak* Topology

Even when \mathcal{X} is not a normed space, the dual space \mathcal{X}' inherits a topology from \mathcal{X} known as the *weak* topology*. A sequence $w_n \in \mathcal{X}'$ converges to $w \in \mathcal{X}'$ in the weak* topology if for any $\varphi \in \mathcal{X}$, we have $\lim_{n \rightarrow +\infty} \langle \varphi, w_n \rangle = \langle \varphi, w \rangle$. This convergence is written $w_n \xrightarrow{w^*} w$. Note that when \mathcal{X} is a normed space, the strong topology is—as suggested by its name—stronger than the weak* topology, in the sense that the set of weak*-open sets is included in the set of strongly-open sets. This implies that any linear functional $\nu : \mathcal{X}' \rightarrow \mathbb{R}$ that is weak*-continuous is also continuous for the strong topology.

Weak topology

Another topology deriving from distribution theory is the weak topology, which is analogous to the weak* topology, only the roles of \mathcal{X} and \mathcal{X}' are inverted. More specifically, a sequence $\varphi_n \in \mathcal{X}$ converges to $\varphi \in \mathcal{X}$ in the weak topology if for any $w \in \mathcal{X}'$, we have $\lim_{n \rightarrow +\infty} \langle \varphi_n, w \rangle = \langle \varphi, w \rangle$. This convergence is written $\varphi_n \xrightarrow{w} \varphi$.

Finally, equipping \mathcal{X}' with a topology—such as the strong or weak* topologies—enables the definition of the *bidual* (or double dual) space of \mathcal{X} .

Bidual Space

The bidual space of \mathcal{X} is defined as $\mathcal{X}'' \triangleq (\mathcal{X}')'$. The space \mathcal{X} is trivially always included in its bidual \mathcal{X}'' via the canonical linear operator $\Psi : \mathcal{X} \rightarrow \mathcal{X}''$ defined as $\Psi(\varphi) \triangleq (w \in \mathcal{X}' \mapsto w(\varphi) \in \mathbb{R})$, for both the strong and weak* topologies. This inclusion $\Psi(\mathcal{X}) \subset \mathcal{X}''$ is often slightly abusively written $\mathcal{X} \subset \mathcal{X}''$. The space \mathcal{X} is called *reflexive* when this inclusion is an equality, *i.e.*, $\mathcal{X} = \mathcal{X}''$.

2.1.4 Adjoint Operators

We now define adjoint operators in the context of distribution theory. Although this definition is well known in the context of Hilbert spaces, there are additional subtleties when dealing with *nonreflexive spaces*, *i.e.*, spaces \mathcal{X} such that $\mathcal{X}'' \neq \mathcal{X}$.

Definition 2.2 (Adjoint Operator). *Let $(\mathcal{X}, \mathcal{X}')$ and $(\mathcal{Y}, \mathcal{Y}')$ be dual pairs of topological vector spaces, and let $L : \mathcal{X} \rightarrow \mathcal{Y}$ be a linear operator. Then, the adjoint operator of L is the unique linear operator $L^* : \mathcal{Y}' \rightarrow \mathcal{X}'$ that satisfies*

$$\langle f, L^* \{g\} \rangle_{\mathcal{X} \times \mathcal{X}'} \triangleq \langle L \{f\}, g \rangle_{\mathcal{Y} \times \mathcal{Y}'} \quad \forall f, g \in \mathcal{X} \times \mathcal{Y}'. \quad (2.3)$$

Alternatively, let $L : \mathcal{X}' \rightarrow \mathcal{Y}'$ be a linear operator. Then, the adjoint of L is the operator

$L^* : \mathcal{Y} \rightarrow \mathcal{X}$ defined by

$$\langle L^*\{f\}, g \rangle_{\mathcal{Y} \times \mathcal{Y}'} \triangleq \langle f, L\{g\} \rangle_{\mathcal{X} \times \mathcal{X}'} \quad \forall f, g \in \mathcal{X}' \times \mathcal{Y}. \quad (2.4)$$

Remark 2.1. *The adjoint operator may thus involve either the dual or the predual spaces of \mathcal{X} and \mathcal{Y} , which are different spaces in the case of nonreflexive spaces, which leads to different adjoint operators L . In the thesis, we will always specify the domain of the adjoint operator, thus lifting this ambiguity. For an operator $L : \mathcal{X} \rightarrow \mathcal{Y}$, if the domain of L^* is \mathcal{Y}' , then the definition (2.3) is implied. If it is the predual of \mathcal{Y} , then the definition (2.4) is implied.*

2.1.5 The Schwartz Space

We now give a fundamental example of a dual pair $(\mathcal{X}, \mathcal{X}')$ that is central to the construction of our native spaces: the *Schwartz space* $\mathcal{S}(\mathbb{K})$ and its dual, the space of tempered distributions $\mathcal{S}'(\mathbb{K})$. Throughout Part I, we denote by \mathbb{K} the domain of the functions in our function spaces, with either $\mathbb{K} = \mathbb{R}$ or $\mathbb{K} = \mathbb{T} = [0, 2\pi]$ the one-dimensional torus when working with 2π -periodic functions. Note that virtually all the material in Part I can be extended to multivariate domains such as \mathbb{R}^d and \mathbb{T}^d with $d > 1$. However, since our work does not require it, we restrict our attention to the univariate case $d = 1$ for simplicity.

For $\mathbb{K} = \mathbb{R}$, the nonperiodic Schwartz space is defined below.

Definition 2.3 (Schwartz Space). *For any $\alpha, \beta \in \mathbb{N}$, let $\|\cdot\|_{\alpha, \beta}$ be the functional defined for any $f \in \mathcal{C}_\infty(\mathbb{K})$ as $\|f\|_{\alpha, \beta} \triangleq \sup_{x \in \mathbb{K}} (x^\beta D^\alpha \{f\}(x))$, where $\mathcal{C}_\infty(\mathbb{R})$ is the space of infinitely differentiable functions $f : \mathbb{R} \rightarrow \mathbb{R}$. The Schwartz space $\mathcal{S}(\mathbb{R})$ is the space of functions whose successive derivatives are smooth and rapidly decaying functions defined as*

$$\mathcal{S}(\mathbb{R}) = \{f : \mathbb{R} \rightarrow \mathbb{R} \in \mathcal{C}_\infty(\mathbb{R}) : \forall \alpha, \beta \in \mathbb{N}, \|f\|_{\alpha, \beta} < +\infty\}. \quad (2.5)$$

The family $\{\|\cdot\|_{\alpha, \beta}\}_{(\alpha, \beta) \in \mathbb{N}^2}$ are then seminorms over the Schwartz space, which is a topological vector space equipped with the topology induced by this family.

The definition of the periodic Schwartz space ($\mathbb{K} = \mathbb{T}$) is very similar, only without the rapid-decay aspect since the domain \mathbb{T} is compact.

Definition 2.4 (Periodic Schwartz Space). *For any $\alpha \in \mathbb{N}$, let $\|\cdot\|_\alpha$ be the functional defined for any $f \in \mathcal{C}_\infty(\mathbb{T})$ as $\|f\|_\alpha \triangleq \sup_{x \in \mathbb{T}} (D^\alpha \{f\}(x))$, where $\mathcal{C}_\infty(\mathbb{T})$ is the space of 2π -periodic infinitely differentiable functions $f : \mathbb{T} \rightarrow \mathbb{R}$. The periodic Schwartz space $\mathcal{S}(\mathbb{T})$ is the space of functions whose successive derivatives are smooth defined as*

$$\mathcal{S}(\mathbb{T}) = \{f : \mathbb{T} \rightarrow \mathbb{R} \in \mathcal{C}_\infty(\mathbb{T}) : \forall \alpha \in \mathbb{N}, \|f\|_\alpha < +\infty\}. \quad (2.6)$$

The family $\{\|\cdot\|_\alpha\}_{\alpha \in \mathbb{N}}$ are then seminorms over the periodic Schwartz space, which is a topological vector space equipped with the topology induced by this family.

For both $\mathbb{K} = \mathbb{R}$ and $\mathbb{K} = \mathbb{T}$, we define the complex-valued Schwartz space as

$$\mathcal{S}(\mathbb{K}, \mathbb{C}) \triangleq \{f_1 + jf_2 : \mathbb{K} \rightarrow \mathbb{C} : f_1, f_2 \in \mathcal{S}(\mathbb{K})\}. \quad (2.7)$$

For our purpose, the two main desirable properties of the Schwartz space are:

- differential operators applied to Schwartz functions are mathematically well defined, since $\mathcal{S}(\mathbb{K}) \subset \mathcal{C}_\infty(\mathbb{K})$;
- it is well adapted to Fourier analysis. More precisely, for $\mathbb{K} = \mathbb{R}$, the Fourier transform \mathcal{F} is well defined and is an automorphism on $\mathcal{S}(\mathbb{R}, \mathbb{C})$, i.e., $\mathcal{F} : \mathcal{S}(\mathbb{R}, \mathbb{C}) \rightarrow \mathcal{S}(\mathbb{R}, \mathbb{C})$ [98].

The topological dual of $\mathcal{S}(\mathbb{K})$ in the sense of Definition 2.1 is the *space of tempered distributions* $\mathcal{S}'(\mathbb{K})$. For $\mathbb{K} = \mathbb{R}$, the prototypical example of a tempered distribution is the Dirac distribution $\delta \in \mathcal{S}'(\mathbb{R})$, whose action on test functions $\varphi \in \mathcal{S}(\mathbb{R})$ yields $\langle \varphi, \delta \rangle \triangleq \varphi(0)$. Another example is its n -th derivative $\delta^{(n)} \in \mathcal{S}'(\mathbb{R})$, defined as $\langle \varphi, \delta^{(n)} \rangle \triangleq \varphi^{(n)}(0)$. The periodic equivalent of δ is the Dirac comb $\mathbb{III} \triangleq \sum_{k \in \mathbb{Z}} \delta(\cdot - 2k\pi) \in \mathcal{S}'(\mathbb{T})$, where $\langle \varphi, \mathbb{III} \rangle = \varphi(0)$ for any $\varphi \in \mathcal{S}(\mathbb{T})$. Similarly to the Schwartz space, we define the space of complex-valued tempered distributions as

$$\mathcal{S}'(\mathbb{K}, \mathbb{C}) \triangleq \{f_1 + jf_2 : \mathcal{S}(\mathbb{K}) \rightarrow \mathbb{C} : f_1, f_2 \in \mathcal{S}'(\mathbb{K})\}, \quad (2.8)$$

and the (complex) duality product between a pair $(\varphi, w) \in \mathcal{S}(\mathbb{K}, \mathbb{C}) \times \mathcal{S}'(\mathbb{K}, \mathbb{C})$ is defined as

$$\langle \varphi, w \rangle \triangleq \langle \operatorname{Re}(\varphi), \operatorname{Re}(w) \rangle + j \langle \operatorname{Im}(\varphi), \operatorname{Im}(w) \rangle. \quad (2.9)$$

Remark 2.2. *In standard distribution theory, duality products are directly defined with complex values instead of via real-valued distributions [98]. However, the definition in (2.9) is consistent with the complex-valued theory of Schwartz.*

Remark 2.3. *An important point to bear in mind is that some tempered distributions—and more generally, elements of the dual of function spaces—are routinely assimilated to ordinary functions, despite technically being functionals. This occurs when the integral $\int_{\mathbb{K}} f(x)\varphi(x)dx$ is well defined for any $\varphi \in \mathcal{S}(\mathbb{K})$, which is for example the case when $f \in L_p(\mathbb{K})$ for $p \geq 1$. Then, the linear functional $\varphi \mapsto \int_{\mathbb{K}} f(x)\varphi(x)dx$ is continuous and is thus an element of $\mathcal{S}'(\mathbb{K})$. A common abuse of notation, which we will adopt throughout this thesis, consists in identifying this distribution and its natural representer f , i.e., $f \in \mathcal{S}'(\mathbb{K})$. In that sense, many classical function spaces are included in $\mathcal{S}'(\mathbb{K})$, e.g., $L_p(\mathbb{K}) \subset \mathcal{S}'(\mathbb{K})$.*

However, distributions (such as the Dirac distribution δ) do not always admit such representers. We thus denote them as generalized functions, as opposed to ordinary functions $f : \mathbb{R} \rightarrow \mathbb{R}$ which have a pointwise interpretation. Nonetheless, in the engineering community, generalized functions are commonly written as if they had pointwise interpretations, e.g., $\delta(x)$, and duality product are written as integrals even when such integrals are undefined, e.g., $\langle \varphi, \delta \rangle = \int_{\mathbb{R}} \delta(x)\varphi(x)dx$ for $\varphi \in \mathcal{S}(\mathbb{R})$. These abuses of notation are less forgivable from a mathematical standpoint; we will thus try to avoid them. Note that some of our notations can be misleading in that respect, such as that of the shifted Dirac distribution $\delta(\cdot - x_0)$. Although this notation seems to imply a pointwise interpretation, $\delta(\cdot - x_0)$ is in fact a generalized function defined as $\langle \varphi, \delta(\cdot - x_0) \rangle \triangleq \varphi(x_0)$ for any $\varphi \in \mathcal{S}(\mathbb{R})$.

2.1.6 Fourier Transformations of Tempered Distributions

The Schwartz space being well-suited to the Fourier transform, the latter can be extended by duality to all tempered distributions.

The Generalized Fourier Transform

For $\mathbb{K} = \mathbb{R}$, the (ordinary) Fourier transform is the operator $\mathcal{F} : \mathcal{S}(\mathbb{R}, \mathbb{C}) \rightarrow \mathcal{S}(\mathbb{R}, \mathbb{C})$ given by

$$\mathcal{F}\{f\} = \widehat{f} \triangleq \int_{\mathbb{R}} f(x)e^{-jx(\cdot)}dx \quad (2.10)$$

for any $f \in \mathcal{S}(\mathbb{R})$, where (\cdot) is a placeholder for the variable of \widehat{f} . It is well known that \mathcal{F} is invertible, and its inverse $\mathcal{F}^{-1} : \mathcal{S}(\mathbb{R}, \mathbb{C}) \rightarrow \mathcal{S}(\mathbb{R}, \mathbb{C})$ is given by $\mathcal{F}^{-1} : f \mapsto \frac{1}{2\pi} \int_{\mathbb{R}} f(\omega)e^{j\omega(\cdot)}d\omega$. When $f \in \mathcal{S}(\mathbb{R})$ is a real-valued function, then \widehat{f} is *Hermitian-symmetric*, which means that $\widehat{f}(x) = \overline{\widehat{f}(-x)}$ for any $x \in \mathbb{R}$.

The Fourier transform can be extended by duality to $\mathcal{F} : \mathcal{S}'(\mathbb{R}, \mathbb{C}) \rightarrow \mathcal{S}'(\mathbb{R}, \mathbb{C})$, where for any $w \in \mathcal{S}'(\mathbb{R}, \mathbb{C})$, $\mathcal{F}\{w\} = \widehat{w}$ is the distribution defined as

$$\langle \varphi, \mathcal{F}\{w\} \rangle \triangleq \langle \mathcal{F}^*\{\varphi\}, w \rangle \quad \forall \varphi \in \mathcal{S}(\mathbb{R}, \mathbb{C}) \quad (2.11)$$

where \mathcal{F}^* is the adjoint operator of \mathcal{F} (Definition 2.2) given by $\mathcal{F}^* = \frac{1}{2\pi}\mathcal{F}^{-1}$. Note that the right side of (2.11) is valid since we have $\mathcal{F}^*(\mathcal{S}(\mathbb{R}, \mathbb{C})) = \mathcal{S}(\mathbb{R}, \mathbb{C})$.

Note that this somewhat abstract definition is consistent with the classical integral definition (2.10) of the Fourier transform even for functions $f \notin \mathcal{S}(\mathbb{R}, \mathbb{C})$. In fact, it generalizes the definitions of the Fourier transform over classical function spaces such as $L_p(\mathbb{R})$ spaces for $1 \leq p \leq 2$ due to their inclusion in $\mathcal{S}'(\mathbb{R})$. The prototypical example of a generalized Fourier transform is that of the Dirac distribution δ given by $\widehat{\delta} = 1 \in \mathcal{S}'(\mathbb{R})$,

which is defined as $\langle \varphi, 1 \rangle \triangleq \int_{\mathbb{R}} \varphi(x) dx$ for all $\varphi \in \mathcal{S}(\mathbb{R})$.

Generalized Fourier Series

For periodic functions, *i.e.*, $\mathbb{K} = \mathbb{T}$, the preferred Fourier-analysis tool is the Fourier series. The Fourier series $(\widehat{f}[k])_{k \in \mathbb{Z}}$ of a Schwartz function $f \in \mathcal{S}(\mathbb{T}, \mathbb{C})$ is given by

$$\widehat{f}[k] \triangleq \frac{1}{2\pi} \int_{\mathbb{T}} f(x) e_k(x) dx \quad \forall k \in \mathbb{Z}, \quad (2.12)$$

where $e_k \triangleq e^{ik} \in \mathcal{S}(\mathbb{T}, \mathbb{C})$. It is well known that we have $\widehat{f} \in \mathcal{S}(\mathbb{Z}, \mathbb{C})$, where

$$\mathcal{S}(\mathbb{Z}) \triangleq \left\{ a \in \mathbb{R}^{\mathbb{Z}} : \forall n \in \mathbb{N}, \lim_{|k| \rightarrow +\infty} k^n a[k] = 0 \right\} \quad (2.13)$$

is the space of Schwartz space of rapidly decaying sequences and $\mathcal{S}(\mathbb{Z}, \mathbb{C})$ is its complex-valued counterpart defined as in (2.7). Conversely, any Schwartz sequence $a \in \mathcal{S}(\mathbb{Z}, \mathbb{C})$ leads to a periodic Schwartz function $\sum_{k \in \mathbb{Z}} a[k] e_k \in \mathcal{S}(\mathbb{T}, \mathbb{C})$, where the sum is pointwise-convergent due to the rapid decay of a .

The dual of $\mathcal{S}(\mathbb{Z})$ is the space of slowly growing sequences

$$\mathcal{S}'(\mathbb{Z}) = \left\{ b \in \mathbb{R}^{\mathbb{Z}} : \exists n \in \mathbb{N}, \lim_{|k| \rightarrow +\infty} k^{-n} b[k] = 0 \right\}, \quad (2.14)$$

and the duality product is given by $\langle a, b \rangle_{\mathcal{S}(\mathbb{Z}) \times \mathcal{S}'(\mathbb{Z})} \triangleq \sum_{k \in \mathbb{Z}} a[k] b[k]$, the sum always being convergent due to the rapid decay of a . Following (2.8), we define $\mathcal{S}'(\mathbb{Z}, \mathbb{C})$ as the complex-valued counterpart of $\mathcal{S}'(\mathbb{Z})$, and the complex-valued duality product between $\mathcal{S}(\mathbb{Z}, \mathbb{C})$ and $\mathcal{S}'(\mathbb{Z}, \mathbb{C})$ as in (2.9). An important difference between $\mathcal{S}'(\mathbb{Z})$ and $\mathcal{S}'(\mathbb{K})$ for $\mathbb{K} = \mathbb{R}$ or \mathbb{T} is that a slowly growing sequences $b \in \mathcal{S}'(\mathbb{Z})$ always has a pointwise interpretation $b[k] \in \mathbb{R}$ for any $k \in \mathbb{Z}$, which makes them much easier to handle.

Conveniently, any distribution $w \in \mathcal{S}'(\mathbb{T}, \mathbb{C})$ is uniquely characterized by its generalized Fourier series $(\widehat{w}[k])_{k \in \mathbb{Z}} \in \mathcal{S}'(\mathbb{Z}, \mathbb{C})$ given by

$$\widehat{w}[k] \triangleq \langle e_k, w \rangle \quad \forall k \in \mathbb{Z}, \quad (2.15)$$

which satisfies $\langle w, \varphi \rangle = \sum_{k \in \mathbb{Z}} \widehat{\varphi}[k] \widehat{w}[k]$ for all $\varphi \in \mathcal{S}(\mathbb{T}, \mathbb{C})$. Analogously to the Fourier transform, when $w \in \mathcal{S}'(\mathbb{T})$ is real-valued, then \widehat{w} is Hermitian-symmetric, which means that $\widehat{w}[-k] = \overline{\widehat{w}[k]}$ for any $k \in \mathbb{Z}$. We write the generalized Fourier series decomposition of $w \in \mathcal{S}'(\mathbb{T}, \mathbb{C})$ as the sum

$$w = \sum_{k \in \mathbb{Z}} \widehat{w}[k] e_k \quad \forall w \in \mathcal{S}'(\mathbb{T}, \mathbb{C}), \quad (2.16)$$

which is a slight abuse of notation since w does not necessarily have a pointwise interpretation (in which case the sum does not converge pointwise but in the weak* sense). The prototypical example of a generalized Fourier series is that of the Dirac comb $\text{III} \triangleq \sum_{k \in \mathbb{Z}} \delta(\cdot - 2k\pi)$ given by $\widehat{\text{III}}[k] = 1$ for all $k \in \mathbb{Z}$.

2.2 The Space of Radon Measures

We now introduce the space of Radon measures $\mathcal{M}(\mathbb{K})$, which plays a crucial role in this thesis. For $\mathbb{K} = \mathbb{R}$, it is defined as the dual of the space $\mathcal{C}_0(\mathbb{R})$ of continuous functions that decay at infinity (Riesz-Markov theorem [99, Chapter 6]), *i.e.*, $\mathcal{M}(\mathbb{R}) \triangleq (\mathcal{C}_0(\mathbb{R}))'$. Here, $\mathcal{C}_0(\mathbb{R})$ is equipped with the supremum norm $\|\cdot\|_{L_\infty}$. For $\mathbb{K} = \mathbb{T}$, we similarly have $\mathcal{M}(\mathbb{T}) \triangleq (\mathcal{C}(\mathbb{T}))'$, where $\mathcal{C}(\mathbb{T})$ is simply the space of continuous functions¹, also equipped with the supremum norm $\|\cdot\|_{L_\infty}$. The inherited dual norm $\|\cdot\|_{\mathcal{M}}$ defined in (2.2) is called the *TV norm for measures*² and is given by

$$\|w\|_{\mathcal{M}} \triangleq \sup_{\varphi \in \mathcal{S}(\mathbb{K}): \|\varphi\|_{L_\infty} \leq 1} \langle \varphi, w \rangle_{\mathcal{S}(\mathbb{K}) \times \mathcal{S}'(\mathbb{K})} \quad (2.17)$$

for any $w \in \mathcal{S}'(\mathbb{K})$.

Remark 2.4. *That this definition is different from the generic definition of the dual norm (2.2), in which the supremum is taken over the predual space of $\mathcal{M}(\mathbb{K})$ ($\mathcal{C}_0(\mathbb{R})$ or $\mathcal{C}(\mathbb{T})$) instead of $\mathcal{S}(\mathbb{K})$, and the duality product is that of the dual pair $(\mathcal{C}_0(\mathbb{R}), \mathcal{M}(\mathbb{R}))$ (or $(\mathcal{C}(\mathbb{T}), \mathcal{M}(\mathbb{T}))$) instead of $(\mathcal{S}(\mathbb{K}), \mathcal{S}'(\mathbb{K}))$. However, it is well known that $\mathcal{C}_0(\mathbb{R})$ is the completion of $\mathcal{S}(\mathbb{R})$ in the supremum norm, which implies that the duality product $\langle \cdot, \cdot \rangle_{\mathcal{C}_0(\mathbb{R}) \times \mathcal{M}(\mathbb{R})}$ is compatible with the Schwartz duality product $\langle \cdot, \cdot \rangle_{\mathcal{S}(\mathbb{R}) \times \mathcal{S}'(\mathbb{R})}$ and thus that the definition (2.17) is consistent with (2.2). This stems from the fact that $\mathcal{C}_0(\mathbb{R})$ is a Schwartz-Banach space; we refer to [91, Section 2.1] for more information. The same reasoning applies to the dual pair $(\mathcal{C}(\mathbb{T}), \mathcal{M}(\mathbb{T}))$.*

*This alternative definition has the advantage of defining the optimization over the smaller and thus “nicer” space $\mathcal{S}(\mathbb{K})$, as well as properly defining the TV norm for any $w \in \mathcal{S}'(\mathbb{K})$, which is larger than $\mathcal{M}(\mathbb{K})$, *i.e.*, $\mathcal{M}(\mathbb{K}) \subset \mathcal{S}'(\mathbb{K})$. In fact, $\mathcal{M}(\mathbb{K})$ can alternatively be defined via the definition (2.17) of the TV norm as*

$$\mathcal{M}(\mathbb{K}) \triangleq \{w \in \mathcal{S}'(\mathbb{K}) : \|w\|_{\mathcal{M}} < +\infty\}. \quad (2.18)$$

For our purpose, the two main features of the Banach space $(\mathcal{M}(\mathbb{K}), \|\cdot\|_{\mathcal{M}})$ are the following:

¹Contrary to $\mathbb{K} = \mathbb{R}$, the decay at infinity is irrelevant since $\mathbb{K} = \mathbb{T}$ has compact support.

²The clarification “for measures” is to remove the ambiguity with the TV seminorm for functions, which is widely used in the signal and image processing community [28]

1. we have $L_1(\mathbb{K}) \subset \mathcal{M}(\mathbb{K})$, and the TV norm generalizes the L_1 norm in the sense that $\|f\|_{\mathcal{M}} = \|f\|_{L_1}$ for any $f \in L_1(\mathbb{K})$;
2. the $\|\cdot\|_{\mathcal{M}}$ norm of a weighted sum of Dirac impulses is $\|\sum_{k=1}^K a_k \delta(\cdot - x_k)\|_{\mathcal{M}} = \|a\|_1$, where $a_k \in \mathbb{R}$ and the $x_k \in \mathbb{R}$ are pairwise distinct.

2.3 Derivative Operators

In this thesis, we will only consider N_0 th-order derivative operators

$$L = D^{N_0}, \quad (2.19)$$

with $N_0 \geq 1$ as regularization operators. This choice is guided by a desire for simplicity and by the fact that derivatives are used in the overwhelming majority of applications for one-dimensional signals. For example, $L = D$ leads to the well-known space of functions with bounded variation $\mathcal{M}_{\text{D}}(\mathbb{K})$ —commonly denoted by $\text{BV}(\mathbb{K})$ in the literature—, which is the native space of the classical TV seminorm for functions [15].

All the results of this chapter and Theorem 3.3 can be extended to more general operators, namely *spline-admissible operators* in the sense of [91, Definition 2]. An example of a more general spline-admissible operators is the fractional derivative L^γ with $\gamma \in \mathbb{R} \setminus \mathbb{N}$ [100]. However, such a level of generality leads to significantly more technicalities. The works in Chapters 4 to Chapter 10 in this thesis can be adapted without major difficulties to ordinary differential operators of the form $L = D^N + a_{N-1}D^{N-1} + \dots + a_0I$ where $a_n \in \mathbb{R}$ [101] and even rational operators [102]. The main requirement is the existence a finitely supported B-spline matched to L (see Section 3.1.4).

2.3.1 Weak Derivatives

In order to define differential operators over nondifferentiable (in the classical sense) functions and even distributions, we introduce the notion of *weak derivatives*. This is achieved in a similar fashion as the extension of the Fourier transform to tempered distributions in Section 2.1.6³.

By definition of the Schwartz spaces (2.5) and (2.6), derivative operators $L = D^{N_0}$ are well defined in the classical sense over $\mathcal{S}(\mathbb{K})$ and clearly satisfy $L : \mathcal{S}(\mathbb{K}) \rightarrow \mathcal{S}(\mathbb{K})$. These operators then admits a unique continuous extension $L : \mathcal{S}'(\mathbb{K}) \rightarrow \mathcal{S}'(\mathbb{K})$ whose application to a distribution $w \in \mathcal{S}'(\mathbb{K})$ is defined as

$$\langle \varphi, L\{w\} \rangle \triangleq \langle L^*\{\varphi\}, w \rangle \quad \forall \varphi \in \mathcal{S}(\mathbb{K}), \quad (2.20)$$

³This technique can be used to extend any linear operator $L : \mathcal{S}(\mathbb{K}) \rightarrow \mathcal{S}(\mathbb{K})$ to $L : \mathcal{S}'(\mathbb{K}) \rightarrow \mathcal{S}'(\mathbb{K})$ by continuity.

where $L^* : \mathcal{S}(\mathbb{K}) \rightarrow \mathcal{S}(\mathbb{K})$ is the adjoint operator of L in the sense of Definition 2.2 given by $(D^{N_0})^* = (-1)^{N_0} D^{N_0}$ for any $N_0 \geq 1$. Likewise, L^* admits a continuous extension $L^* : \mathcal{S}'(\mathbb{K}) \rightarrow \mathcal{S}'(\mathbb{K})$. These extended operators $L : \mathcal{S}'(\mathbb{K}) \rightarrow \mathcal{S}'(\mathbb{K})$ are known as *weak derivatives*.

2.3.2 Nonperiodic Setting

We now focus on the nonperiodic $\mathbb{K} = \mathbb{R}$ scenario.

Frequency Response

It is well known that linear shift-invariant (LSI) operators $L : \mathcal{S}'(\mathbb{R}) \rightarrow \mathcal{S}'(\mathbb{R})$ are equivalent to convolution operators and are conveniently represented by their frequency response \widehat{L} [98, Chapter 7, §5]. The frequency response of the N_0 th-order derivative operator $L = D^{N_0}$ is given by

$$\widehat{L}(\omega) = (j\omega)^{N_0}. \quad (2.21)$$

Then, for any $w \in \mathcal{S}'(\mathbb{R})$, we have $\mathcal{F}\{L\{w\}\} = \widehat{L}\widehat{w}$, where the distributional product $\widehat{L}\widehat{w}$ is defined as

$$\langle \varphi, \widehat{L}\widehat{w} \rangle \triangleq \langle \widehat{L}\varphi, \widehat{w} \rangle \quad \forall \varphi \in \mathcal{S}(\mathbb{R}, \mathbb{C}). \quad (2.22)$$

The latter duality product in (2.22) is well defined since $\widehat{L}\varphi$ is a product of ordinary functions and $\widehat{L}\varphi \in \mathcal{S}(\mathbb{R}, \mathbb{C})$ due to the form (2.21) of \widehat{L} . Note that when \widehat{w} is an ordinary function, the definition (2.22) is consistent with the standard function product.

Null Space

We now introduce the null space \mathcal{N}_L of the operator L , which is defined as

$$\mathcal{N}_L \triangleq \{w \in \mathcal{S}'(\mathbb{R}) : L\{w\} = 0\}. \quad (2.23)$$

It is well known that the null space of $L = D^{N_0}$ is the space of polynomials of order less than N_0

$$\mathcal{N}_L = \text{span} \{x \mapsto x^n\}_{0 \leq n \leq N_0-1}, \quad (2.24)$$

which has finite dimension $N_0 \triangleq \dim(\mathcal{N}_L)$.

Green's Function

A crucial property of derivative operators are the existence of a *Green's function*, which we define below.

Definition 2.5 (Green's function). *A Green's function of a linear operator $L : \mathcal{S}'(\mathbb{R}) \rightarrow \mathcal{S}'(\mathbb{R})$ is a function $\rho_L \in \mathcal{S}'(\mathbb{R})$ such that $L\{\rho_L\} = \delta$.*

The causal Green's function of $L = D^{N_0}$ is the one-sided power function

$$\rho_L(x) = \frac{x_+^{N_0-1}}{(N_0-1)!}, \quad (2.25)$$

where $x_+ \triangleq \max(x, 0)$.

Remark 2.5. *The operator $L = D^{N_0}$ has infinitely many Green's functions. Indeed, any function $\rho_L + p$ where $p \in \mathcal{N}_L$ yields another valid Green's function. However, imposing causality makes the choice unique. Throughout this thesis, we will always choose the causal Green's function (2.25) since it has a simpler expression (and a smaller support) than the perhaps mathematically purer alternative, the canonical Green's function given by*

$$\rho_L(x) \triangleq \mathcal{F}^{-1} \left\{ \frac{1}{\widehat{L}} \right\} = \mathcal{F}^{-1} \left\{ \frac{1}{(j\cdot)^{N_0}} \right\} = \frac{\text{sign}(x)x^{N_0-1}}{2(N_0-1)!}, \quad (2.26)$$

where the distribution $\frac{1}{(j\cdot)^{N_0}}$ is understood in the Cauchy principal value sense, i.e.,

$$\left\langle \varphi, \frac{1}{(j\cdot)^{N_0}} \right\rangle \triangleq \lim_{\epsilon \rightarrow 0^+} \int_{\mathbb{R} \setminus [a-\epsilon, a+\epsilon]} \frac{\varphi(\omega)}{(j\cdot)^{N_0}} d\omega \quad \forall \varphi \in \mathcal{S}(\mathbb{R}), \quad (2.27)$$

the limit being well defined due to the smoothness of φ .

It can easily be verified that ρ_L is the N_0 th-fold convolution of the Heaviside step function $\mathbb{1}_+$ like

$$\rho_L = \underbrace{\rho_D * \dots * \rho_D}_{N_0 \text{ times}} = \underbrace{\mathbb{1}_+ * \dots * \mathbb{1}_+}_{N_0 \text{ times}}, \quad (2.28)$$

where ρ_D is the causal Green's function (2.25) of the derivative operator D .

2.3.3 Periodic Setting

We now focus on the periodic $\mathbb{K} = \mathbb{T}$ scenario, which is ideally suited to the use of Fourier series. As in the nonperiodic case, we restrict to N_d th-order derivative operators $L = D^{N_d} : \mathcal{S}'(\mathbb{T}) \rightarrow \mathcal{S}'(\mathbb{T})$, with $N_d \geq 1$. Note that we use the exponent N_d instead of N_0 because contrary to the nonperiodic case, the order of the derivative does not coincide

with the dimension of the null space \mathcal{N}_L . Again, there is no additional technical difficulty in generalizing our works to any spline-admissible operators in the sense of [22, Definition 2] (in fact, this generalization is much simpler than in the nonperiodic case).

Frequency Response

It is well known [22, Proposition 1] that a periodic LSI operator $L : \mathcal{S}'(\mathbb{T}) \rightarrow \mathcal{S}'(\mathbb{T})$ is characterized by its frequency response $\widehat{L} \in \mathcal{S}'(\mathbb{Z}, \mathbb{C})$ that satisfies

$$L\{w\} = \sum_{k \in \mathbb{Z}} \widehat{L}[k] \widehat{w}[k] e_k \quad \forall w \in \mathcal{S}'(\mathbb{T}). \quad (2.29)$$

In the case of the operator $L = D^{N_d}$, we have

$$\widehat{L}[k] = (jk)^{N_d} \quad \forall k \in \mathbb{Z}. \quad (2.30)$$

Null Space

Contrary to the nonperiodic case, the null space of $L = D^{N_d}$ does not depend on the order of differentiation N_d , since no polynomial is periodic except for the constant. Hence, the null space of L is

$$\mathcal{N}_L = \text{span}\{x \mapsto 1\} \quad (2.31)$$

for any $N_d \geq 1$, and has dimension $N_0 = 1$.

Green's function

The Green's function of $L = D^{N_d}$ is defined as

$$g_L \triangleq \sum_{k \in \mathbb{Z} \setminus \{0\}} \frac{e_k}{\widehat{L}[k]} = \sum_{k \in \mathbb{Z} \setminus \{0\}} \frac{e_k}{(jk)^{N_d}}, \quad (2.32)$$

It is the unique solution to the equation $L\{g_L\} = \text{III} - 1$. Contrary to the nonperiodic setting, there is no solution $g \in \mathcal{S}'(\mathbb{T})$ to the equation $L\{g\} = \text{III}$ since $\widehat{L\{g\}}[0] = \widehat{L}[0] \widehat{g}[0] = 0$ for any $g \in \mathcal{S}'(\mathbb{T})$, whereas $\widehat{\text{III}}[0] = 1$.

2.4 Definition and Topology of $\mathcal{M}_L(\mathbb{K})$

We now have all the required tools to properly introduce the native space of our derivative operator L , *i.e.*, the largest function space over which functional $\|L\{\cdot\}\|_{\mathcal{M}}$ takes a finite

value, which is defined as

$$\mathcal{M}_L(\mathbb{K}) \triangleq \{f \in \mathcal{S}'(\mathbb{K}) : L\{f\} \in \mathcal{M}(\mathbb{K})\} \quad \mathbb{K} = \mathbb{R} \text{ or } \mathbb{T}. \quad (2.33)$$

The native space $\mathcal{M}_L(\mathbb{K})$ inherits many crucial topological features from $\mathcal{M}(\mathbb{K})$, in particular its Banach structure and the existence of a predual space. We uncover these properties in this section, which requires the construction of some form of inverse operator of L .

2.4.1 Nonperiodic Case

We first focus on the nonperiodic $\mathbb{K} = \mathbb{R}$ scenario, which is technically more involved. This section mostly relies on the results of [14], which we present in the special case when L is the N_0 th-order derivative operator. For the most general and in-depth treatment of the topic, we refer to [91].

Biorthogonal System for \mathcal{N}_L

The main difficulty in uncovering the Banach structure of $\mathcal{M}_L(\mathbb{R})$ is the handling of the null space \mathcal{N}_L . For this, we rely on a biorthogonal system for \mathcal{N}_L , as defined in Definition 2.6.

Definition 2.6 (Biorthogonal System). *Let $\mathcal{N} \subset \mathcal{S}'(\mathbb{R})$ be a N_0 -dimensional vector space with $N_0 > 0$. Then, the pair (ϕ, \mathbf{p}) with $\mathbf{p} = (p_1, \dots, p_{N_0}) \in \mathcal{N}^{N_0}$ and $\phi = (\phi_1, \dots, \phi_{N_0}) \in \mathcal{S}(\mathbb{R})^{N_0}$ forms a biorthogonal system for \mathcal{N} if $(p_n)_{n=1}^{N_0}$ is a basis of \mathcal{N} and if any $p \in \mathcal{N}$ admits the unique expansion*

$$p = \sum_{n=1}^{N_0} \langle \phi_n, p \rangle p_n. \quad (2.34)$$

Remark 2.6. *As we will see later on in Proposition 2.2, the requirement $\phi_n \in \mathcal{S}(\mathbb{R})$ can be relaxed to $\phi_n \in \mathcal{C}_L(\mathbb{R})$, where $\mathcal{C}_L(\mathbb{R})$ is the predual of $\mathcal{M}_L(\mathbb{R})$. However, since $\mathcal{C}_L(\mathbb{R})$ is not yet properly defined—it will be in Theorem 2.3—, following [91], we initially restrict to imposing $\phi_n \in \mathcal{S}(\mathbb{R})$.*

Proposition 2.1 then ensures the existence of such biorthogonal systems.

Proposition 2.1 (Existence of a Biorthogonal System [91, Proposition 3]). *Let $\mathcal{N} \subset \mathcal{S}'(\mathbb{R})$ be a N_0 -dimensional vector space with $N_0 > 0$. Then, there always exists a biorthogonal system (ϕ, \mathbf{p}) for \mathcal{N} in the sense of Definition 2.6.*

An explicit construction of a valid biorthogonal system, based on Hermite polynomials, is given in [91, Example 3].

Stable Right-Inverse of L

Next, we introduce a right-inverse operator of L matched to a biorthogonal system, which is a crucial element to uncover the Banach structure of $\mathcal{M}_L(\mathbb{R})$.

Theorem 2.1 (Stable Right-Inverse of L [14, Theorem 3]). *Let $L = D^{N_0} : \mathcal{S}'(\mathbb{R}) \rightarrow \mathcal{S}'(\mathbb{R})$ with $N_0 \geq 1$, and let $(\mathbf{p}, \boldsymbol{\phi}) = (p_n, \phi_n)_{n=1}^{N_0}$ be a biorthogonal system for \mathcal{N}_L in the sense of Definition 2.6. Then, there exists a unique operator $L_\phi^{-1} : \mathcal{S}(\mathbb{R}) \rightarrow \mathcal{M}_L(\mathbb{R})$ that satisfies the following properties for any $\varphi \in \mathcal{S}(\mathbb{R})$:*

- *Right-inverse property: $LL_\phi^{-1}\{\varphi\} = \varphi$;*
- *Boundary conditions: $\langle \boldsymbol{\phi}, L_\phi^{-1}\{\varphi\} \rangle = \mathbf{0}$,*

where $\langle \boldsymbol{\phi}, f \rangle \triangleq (\langle \phi_n, f \rangle)_{n=1}^{N_0} \in \mathbb{R}^{N_0}$ for any $f \in \mathcal{S}'(\mathbb{R})$. This operator L_ϕ^{-1} is defined via its Schwartz kernel $g_\phi : \mathbb{R}^2 \rightarrow \mathbb{R}$ as

$$L_\phi^{-1}\{\varphi\} \triangleq \int_{\mathbb{R}} g_\phi(\cdot, y)\varphi(y)dy \quad \forall \varphi \in \mathcal{S}(\mathbb{R}), \quad (2.35)$$

where the Schwartz kernel g_ϕ is given by

$$g_\phi(x, y) \triangleq \rho_L(x - y) - \sum_{n=1}^{N_0} p_n(x)q_n(y) \quad (2.36)$$

with $q_n(y) \triangleq \langle \phi_n, \rho_L(\cdot - y) \rangle$. The operator L_ϕ^{-1} then admits a continuous extension $L_\phi^{-1} : \mathcal{M}(\mathbb{R}) \rightarrow \mathcal{M}_L(\mathbb{R})$, which still satisfies the two original properties of L_ϕ^{-1} for any $\varphi \in \mathcal{M}(\mathbb{R})$.

Although we do not rigorously prove Theorem 2.1, the following observations give an intuition of this result.

- The Schwartz kernel g_ϕ satisfies the stability condition

$$\operatorname{ess\,sup}_{x, y \in \mathbb{R}} |g_\phi(x, y)|(1 + |x|)^{-(N_0-1)} < +\infty, \quad (2.37)$$

which implies that the integral (2.35) is well defined. This condition also enables the extension of the domain of L_ϕ^{-1} from $\mathcal{S}(\mathbb{R})$ to $\mathcal{M}(\mathbb{R})$ via [14, Theorem 3].

- The right-inverse property can be proved by observing that

$$LL_\phi^{-1}\{\varphi\} = L \left\{ \int_{\mathbb{R}} g_\phi(\cdot, y)\varphi(y)dy \right\} = \int_{\mathbb{R}} L\{g_\phi(\cdot, y)\}\varphi(y)dy = \int_{\mathbb{R}} \delta(\cdot - y)\varphi(y)dy = \varphi, \quad (2.38)$$

for any $\varphi \in \mathcal{S}(\mathbb{R})$. This property also explains why the image of $\mathcal{M}(\mathbb{R})$ by L_ϕ^{-1} is $\mathcal{M}_L(\mathbb{R})$, since for any $w \in \mathcal{M}(\mathbb{R})$, we have $LL_\phi^{-1}\{w\} = w \in \mathcal{M}(\mathbb{R})$ which by definition (2.33) of the native space implies that $L_\phi^{-1}\{w\} \in \mathcal{M}_L(\mathbb{R})$.

- The boundary condition property can be proved by computing

$$\langle \phi_n, L_\phi^{-1}\{\varphi\} \rangle = \int_{\mathbb{R}} \langle \phi_n, g_\phi(\cdot, y) \rangle \varphi(y) dy \quad (2.39)$$

$$= \int_{\mathbb{R}} \left(\langle \phi_n, \rho_L(\cdot - y) \rangle - q_n(y) \right) \varphi(y) dy = 0 \quad (2.40)$$

for any $\varphi \in \mathcal{S}(\mathbb{R})$, using the biorthogonality of (\mathbf{p}, ϕ) and by definition of q_n .

Banach Structure of $\mathcal{M}_L(\mathbb{R})$

We now have all the necessary elements to uncover the direct-sum structure of the native space $\mathcal{M}_L(\mathbb{R})$, from which we can deduce its Banach structure.

Theorem 2.2 (Banach Structure of $\mathcal{M}_L(\mathbb{R})$ [14, Theorem 5]). *Let L be a derivative operator of the form (2.19), and let $(\mathbf{p}, \phi) = (p_n, \phi_n)_{n=1}^{N_0}$ be a biorthogonal system for \mathcal{N}_L in the sense of Definition 2.6. Next, let $\mathcal{M}_{L,\phi}(\mathbb{R}) \subset \mathcal{M}_L(\mathbb{R})$ be defined as*

$$\mathcal{M}_{L,\phi}(\mathbb{R}) \triangleq \{f \in \mathcal{M}_L(\mathbb{R}) : \langle \phi, f \rangle = \mathbf{0}\}. \quad (2.41)$$

Then, L_ϕ^{-1} (defined in 2.2) is a bijection from $\mathcal{M}(\mathbb{R})$ to $\mathcal{M}_{L,\phi}(\mathbb{R})$, and $\mathcal{M}_{L,\phi}(\mathbb{R})$ is a Banach space equipped with the norm $\|L\{\cdot\}\|_{\mathcal{M}}$. Moreover, the native space $\mathcal{M}_L(\mathbb{R})$ admits the direct-sum decomposition

$$\mathcal{M}_L(\mathbb{R}) = \mathcal{M}_{L,\phi}(\mathbb{R}) \oplus \mathcal{N}_L. \quad (2.42)$$

Hence, any $f \in \mathcal{M}_L(\mathbb{R})$ admits the unique decomposition

$$f = L_\phi^{-1}\{w\} + p, \quad (2.43)$$

with $w = L\{f\} \in \mathcal{M}(\mathbb{R})$ and $p = \sum_{n=1}^{N_0} \langle \phi_n, f \rangle p_n \in \mathcal{N}_L$. Finally, $\mathcal{M}_L(\mathbb{R})$ is a Banach space equipped with the norm

$$\|\cdot\|_{\mathcal{M}_L,\phi} \triangleq \|L\{\cdot\}\|_{\mathcal{M}} + \|\langle \phi, \cdot \rangle\|_2. \quad (2.44)$$

Remark 2.7. *The Banach structure of $\mathcal{M}_L(\mathbb{R})$ hangs upon on the existence of biorthogonal system (\mathbf{p}, ϕ) , which is guaranteed by Proposition 2.1. Moreover, the norm defined in (2.44) depends on the choice of this biorthogonal system. However, any admissible choice leads to the same native space $\mathcal{M}_L(\mathbb{R}) = \mathcal{M}_{L,\phi}(\mathbb{R}) \oplus \mathcal{N}_L$; this is clear in the definition of $\mathcal{M}_L(\mathbb{R})$ as a set in (2.33), which does not depend on ϕ or \mathbf{p} . Moreover, all possible norms $\|\cdot\|_{\mathcal{M}_L,\phi}$ are equivalent [91, Theorem 5].*

Remark 2.8. *The decomposition (2.43) proves that $\mathcal{M}_L(\mathbb{R})$ is a space of ordinary functions that have a pointwise interpretation. More precisely, we have $\mathcal{M}_L(\mathbb{R}) \subset L_{\infty, N_0-1}(\mathbb{R})$, where*

$$L_{\infty, n}(\mathbb{R}) \triangleq \left\{ f : \mathbb{R} \rightarrow \mathbb{R} \text{ measurable} : \operatorname{ess\,sup}_{x \in \mathbb{R}} |f(x)|(1 + |x|)^{-n} < +\infty \right\}. \quad (2.45)$$

Indeed, by [14, Theorem 3], due to the stability condition (2.37), the image of $\mathcal{M}(\mathbb{R})$ by L_{ϕ}^{-1} is included in $L_{\infty, N_0-1}(\mathbb{R})$, and we clearly also have $\mathcal{N}_L \subset L_{\infty, N_0-1}(\mathbb{R})$.

The Measurement Space of L

The final necessary element of the construction of our native spaces is the identification of the predual space of $\mathcal{M}_L(\mathbb{R})$. The latter is necessary to endow $\mathcal{M}_L(\mathbb{R})$ with the weak* topology, which is crucial to our inverse-problem formulations. We call it the *measurement space* of L, for reasons that will be explained in Remark 3.4.

Theorem 2.3 (Measurement Space of L [14, Theorem 6]). *Let $L = D^{N_0} : \mathcal{S}'(\mathbb{R}) \rightarrow \mathcal{S}'(\mathbb{R})$ with $N_0 \geq 1$, and let $(\mathbf{p}, \phi) = (p_n, \phi_n)_{n=1}^{N_0}$ be a biorthogonal system for \mathcal{N}_L in the sense of Definition 2.6 with $\phi_n \in \mathcal{S}(\mathbb{R})$. We define the space*

$$\mathcal{C}_L(\mathbb{R}) \triangleq L^*(\mathcal{C}_0(\mathbb{R})) \oplus \operatorname{span}\{\phi_n\}_{n=1}^{N_0}, \quad (2.46)$$

and the linear operator $L_{\phi}^{-1} : \mathcal{C}_L(\mathbb{R}) \rightarrow \mathcal{C}_0(\mathbb{R})$ where for any $f = L^*\{f_1\} + f_2 \in \mathcal{C}_L(\mathbb{R})$ with $f_1 \in \mathcal{C}_0(\mathbb{R})$ and $f_2 \in \operatorname{span}\{\phi_n\}_{n=1}^{N_0}$ (following the unique decomposition (2.46)), we have $L_{\phi}^{-1*}\{f\} \triangleq f_1$. Then, $\mathcal{C}_L(\mathbb{R})$ is a Banach space equipped with the norm*

$$\|\cdot\|_{\mathcal{C}_L, \phi, \mathbf{p}} \triangleq \|L_{\phi}^{-1*}\{\cdot\}\|_{L_{\infty}} + \|\langle f, \mathbf{p} \rangle\|_2 \quad (2.47)$$

where $\langle f, \mathbf{p} \rangle \triangleq (\langle f, p_n \rangle)_{n=1}^{N_0} \in \mathbb{R}^{N_0}$. Moreover, we have $(\mathcal{C}_L(\mathbb{R}))' = \mathcal{M}_L(\mathbb{R})$, and $\|\cdot\|_{\mathcal{M}_L, \phi}$ is the dual norm—in the sense of (2.2)—of $\|\cdot\|_{\mathcal{C}_L, \phi, \mathbf{p}}$.

Remark 2.9. *As suggested by the notation, the operator L_{ϕ}^{-1*} is the adjoint operator—in the sense of (2.4) in Definition 2.2 – of $L_{\phi}^{-1} : \mathcal{M}(\mathbb{R}) \rightarrow \mathcal{M}_L(\mathbb{R})$ defined in Theorem 2.2. However, one cannot define this adjoint operator before identifying the predual of the native space $\mathcal{M}_L(\mathbb{R})$. Yet the topology of $\mathcal{C}_L(\mathbb{R})$ is defined using L_{ϕ}^{-1*} ; this circularity can be broken by defining L_{ϕ}^{-1*} independently, which allows us to equip $\mathcal{C}_L(\mathbb{R})$ with a norm, and then verifying that it is indeed the adjoint operator of L_{ϕ}^{-1} . This subtlety is overlooked in [14, Theorem 6], which does not affect the validity of the construction, as demonstrated in [91].*

Now that the predual of $\mathcal{M}_L(\mathbb{R})$ has been identified, we can expand the definition of a valid biorthogonal system for \mathcal{N}_L from Definition 2.6.

Proposition 2.2 (Admissible Biorthogonal Systems [103, Proposition 5]). *Theorems 2.1, 2.2, and 2.3 remain valid for any pair $(\mathbf{p}, \phi) = (p_n, \phi_n)_{n=1}^{N_0}$ that satisfies (2.34) and such that $\phi_n \in \mathcal{C}_L(\mathbb{R})$. In this case, all the duality product involving the ϕ_n in these Theorems become duality products between the dual pair $(\mathcal{C}_L(\mathbb{R}), \mathcal{M}_L(\mathbb{R}))$ instead of $(\mathcal{S}(\mathbb{R}), \mathcal{S}'(\mathbb{R}))$.*

Remark 2.10. *The measurement space $\mathcal{C}_L(\mathbb{R})$ is the largest possible space \mathcal{X} such that $\phi_n \in \mathcal{X}$ in order for Theorems 2.1, 2.2, and 2.3 to remain valid. Indeed, the requirement is that $\mathcal{M}_L(\mathbb{R}) \subset \mathcal{X}'$, the dual space of \mathcal{X} , which holds for $\mathcal{X} = \mathcal{C}_L(\mathbb{R})$ and $\mathcal{X}' = \mathcal{M}_L(\mathbb{R})$. A larger space \mathcal{X} would lead to a smaller space \mathcal{X}' .*

The measurement space of L will play a major role in our inverse-problem formulations. In particular, an important question is the inclusion of the Dirac distribution in $\mathcal{C}_L(\mathbb{R})$, which we answer in the following proposition.

Proposition 2.3 (Sampling Admissibility for $N_0 \geq 2$). *Let $L = D^{N_0}$ with $N_0 \geq 1$. Then, we have $\delta \in \mathcal{C}_L(\mathbb{R})$ if and only if $N_0 \geq 2$.*

Proof. We first consider the $N_0 = 1$ case. Let (ϕ, p) be a biorthogonal system for \mathcal{N}_D in the sense of Definition 2.6 with $\phi \in \mathcal{S}(\mathbb{R})$ and $p \in \mathcal{N}_D$. Assume by contradiction that we have $\delta \in \mathcal{C}_D(\mathbb{R})$. By (2.46), this implies the existence of a function $f_1 \in \mathcal{C}_0(\mathbb{R})$ such that $\delta = D\{f_1\} + \alpha\phi$ with $\alpha \in \mathbb{R}$. Hence, due to the fact that $D\{\mathbb{1}_+\} = \delta$ where $\mathbb{1}_+$ is the Heaviside function, we have that $D\{\mathbb{1}_+ - f_1\} = \alpha\phi$. The function $\mathbb{1}_+ - f_1$ is discontinuous and thus cannot be differentiable, whereas its derivative $\alpha\phi \in \mathcal{S}(\mathbb{R})$ is smooth, which yields a contradiction. This proves that $\delta \notin \mathcal{C}_D(\mathbb{R})$.

It only remains to prove that $\delta \in \mathcal{C}_L(\mathbb{R})$ when $N_0 \geq 2$. The case $N_0 = 2$ is proved [103, Theorem 1, Item 2]. For the $N_0 > 2$ case, we refer to [22, Proposition 9], where this property is proved in the periodic setting and can be adapted to the nonperiodic one. \square

Proposition 2.3 allows us to introduce a convenient biorthogonal system for operators $L = D^{N_0}$ with $N_0 \geq 2$, which we call the *canonical biorthogonal system*. The latter will be used for theoretical purposes in this thesis.

Proposition 2.4 (Canonical Biorthogonal System for $N_0 \geq 2$). *Let $L = D^{N_0}$ with $N_0 \geq 2$. Then, (\mathbf{p}, ϕ) with $\mathbf{p} = \left(x \mapsto 1, \dots, x \mapsto \frac{x^{N_0-1}}{(N_0-1)!}\right)$ and $\phi = (\delta, \dots, \delta^{(N_0-2)}, \delta(\cdot - 1)^{(N_0-2)} - \delta^{(N_0-2)})$ is a valid biorthogonal system for \mathcal{N}_L in the sense of Proposition 2.2. We call it the canonical biorthogonal system.*

Proof. Let us first prove the following lemma.

Lemma 2.1. *Let $g \in \mathcal{C}_{D^{N_0}}(\mathbb{R})$. Then, $D\{g\} \in \mathcal{C}_{D^{N_0+1}}(\mathbb{R})$.*

Proof. Let $(\boldsymbol{\psi}, \mathbf{q})$ be a biorthogonal system for $\mathcal{N}_{\mathcal{D}^{N_0}}$ in the sense of Definition 2.6. Next, let $\tilde{q}_1, \dots, \tilde{q}_{N_0} \in \mathcal{N}_{\mathcal{D}^{N_0+1}}$ such that $\mathcal{D}\{\tilde{q}_n\} = -q_n$ for $1 \leq n \leq N_0$. We thus have the biorthonormality relations $\langle \mathcal{D}\{\psi_n\}, \tilde{q}_m \rangle = \langle \psi_n, q_m \rangle = \delta[n - m]$ for all $1 \leq n, m \leq N_0$. Next, we introduce $\tilde{q}_{N_0+1} \in \mathcal{N}_{\mathcal{D}^{N_0+1}}$ such that $\tilde{\mathbf{q}} \triangleq (\tilde{q}_1, \dots, \tilde{q}_{N_0+1})$ forms a basis of $\mathcal{N}_{\mathcal{D}^{N_0+1}}$ and such that $\langle \mathcal{D}\{\psi_n\}, \tilde{q}_{N_0+1} \rangle = 0$ for all $1 \leq n \leq N_0$. Such a \tilde{q}_{N_0+1} exists and is unique up to a constant since there are $N_0 + 1$ polynomial coefficients for \tilde{q}_{N_0+1} , which must satisfy N_0 linearly independent equations. Finally, let $\tilde{\psi}_{N_0+1} \in \mathcal{S}(\mathbb{R})$ such that $\langle \tilde{\psi}_{N_0+1}, \tilde{q}_n \rangle = \delta[n - (N_0 + 1)]$ for all $1 \leq n \leq N_0 + 1$, for which there are clearly infinitely many possible choices. Then, $(\tilde{\boldsymbol{\psi}}, \tilde{\mathbf{q}})$ forms a biorthogonal system for $\mathcal{N}_{\mathcal{D}^{N_0+1}}$ in the sense of Definition 2.6 with $\tilde{\boldsymbol{\psi}} \triangleq (\mathcal{D}\{\psi_1\}, \dots, \mathcal{D}\{\psi_{N_0}\}, \tilde{\psi}_{N_0+1})$.

Next, let $g \in \mathcal{C}_{\mathcal{D}^{N_0}}(\mathbb{R})$. Using the decomposition (2.46), we have $g = \mathcal{D}^{N_0}\{f_1\} + f_2$ with $f_1 \in \mathcal{C}_0(\mathbb{R})$ and $f_2 \in \text{span}\{\psi_n\}_{n=1}^{N_0}$. Hence, we have $\mathcal{D}\{g\} = \mathcal{D}^{N_0+1}\{f_1\} + \mathcal{D}\{f_2\}$. By definition of $\tilde{\boldsymbol{\psi}}$, we have $\mathcal{D}\{f_2\} \in \text{span}\{\tilde{\psi}_n\}_{n=1}^{N_0}$, and so $\mathcal{D}\{g\} \in \mathcal{D}^{N_0+1}\{\mathcal{C}_0(\mathbb{R})\} \oplus \text{span}\{\tilde{\psi}_n\}_{n=1}^{N_0+1} = \mathcal{C}_{\mathcal{D}^{N_0+1}}(\mathbb{R})$. \square

Using Proposition 2.3 and Lemma 2.1, we have that $\delta, \dots, \delta^{(N_0-2)} \in \mathcal{C}_L(\mathbb{R})$ for any $N_0 \geq 2$. We thus have $\boldsymbol{\phi} \in \mathcal{C}_L(\mathbb{R})^{N_0}$, and it can easily be verified that the biorthonormality relations $\langle \phi_n, p_m \rangle = \delta[n - m]$ —and thus (2.34)—are satisfied. Hence, by Proposition 2.2, $(\mathbf{p}, \boldsymbol{\phi})$ is an admissible biorthogonal system in the sense of Proposition 2.2. \square

Remark 2.11. *The case of the derivative operator $L = \mathcal{D}$ is excluded from Proposition 2.4. As pointed out in [91, Section 5.2], we have $L_1(\mathbb{R}) \subset \mathcal{C}_D(\mathbb{R})$, and thus any $\phi \in L_1(\mathbb{R})$ such that $\int_{\mathbb{R}} \phi(x) dx = 1$ leads to a valid biorthogonal system $(p \triangleq 1, \phi)$, e.g., $\phi = \text{rect}$.*

Finally, we introduce in Proposition 2.5 another valid biorthogonal system for operators $L = \mathcal{D}^{N_0}$, which is similar to the canonical one of Proposition 2.4 but is more convenient for implementation purposes.

Proposition 2.5 (Implementation-Friendly Biorthogonal System). *Let $L = \mathcal{D}^{N_0}$ with $N_0 \geq 1$. Then, for any $\epsilon > 0$, $(\mathbf{p}, \boldsymbol{\phi})$ with $\mathbf{p} = \left(1, (\cdot), \dots, \frac{(\cdot)^{N_0-1}}{(N_0-1)!}\right) * \delta(\cdot - \frac{\epsilon}{2})$ and*

$$\boldsymbol{\phi} = \begin{cases} \frac{1}{\epsilon} \text{rect}\left(\frac{\cdot}{\epsilon} - \frac{1}{2}\right) & \text{if } N_0 = 1 \\ \left(\delta, \dots, \delta^{(N_0-2)}, \delta^{(N_0-1)} * \frac{1}{\epsilon} \text{rect}\left(\frac{\cdot}{\epsilon}\right) * \delta\left(\cdot - \frac{\epsilon}{2}\right)\right) & \text{if } N_0 > 1, \end{cases} \quad (2.48)$$

where $\text{rect}(t) = 1$ for $-1/2 \leq t < 1/2$ and 0 elsewhere, is a valid biorthogonal for \mathcal{N}_L in the sense of Proposition 2.2.

Proof. For $N_0 = 1$, as pointed out in Remark 2.11, we have $\phi_1 \in L_1(\mathbb{R}) \subset \mathcal{C}_L(\mathbb{R})$. Next, for $N_0 > 1$, we have proved in Proposition 2.4 that $\phi_1, \dots, \phi_{N_0-1} \in \mathcal{C}_L(\mathbb{R})$, and using Lemma 2.1, we also have that $\phi_{N_0} \in \mathcal{C}_L(\mathbb{R})$ since $\text{rect} \in \mathcal{C}_D(\mathbb{R})$. One can easily verify that the biorthonormality relation $\langle \phi_n, p_m \rangle = \delta[n - m]$ are satisfied, which implies by Proposition 2.2 that $(\mathbf{p}, \boldsymbol{\phi})$ is a valid biorthogonal system. \square

2.4.2 Periodic Case

In the periodic setting $\mathbb{K} = \mathbb{T}$, uncovering the Banach structure of $\mathcal{M}_L(\mathbb{T})$ is a much less arduous task than in the nonperiodic case. In particular, the definition of an inverse operator is more straightforward: it does not require the introduction of biorthogonal systems. This increased simplicity is due to the convenience of the Fourier series representation for periodic signals. This section mostly relies on the results [22], which deals with more general operators L . Note that in contrast with the nonperiodic case, this generality does not significantly increase the technicality of this construction.

Pseudoinverse Operator of L

As in the nonperiodic setting, we require some kind of inverse operator of L . In the periodic setting, this inverse takes the form of a pseudoinverse instead of a right-inverse, as specified in Proposition 2.6.

Proposition 2.6 (Pseudoinverse Operator of L [22, Proposition 2]). *Let $L = D^{N_d}$ with $N_d \geq 1$. Then, the operator $L^\dagger : \mathcal{S}'(\mathbb{T}) \rightarrow \mathcal{S}'(\mathbb{T})$, characterized by its frequency response*

$$\widehat{L^\dagger}[k] = \begin{cases} \frac{1}{\widehat{L}[k]} = \frac{1}{(jk)^{N_d}} & \text{for } k \in \mathbb{Z} \setminus \{0\} \\ 0 & \text{for } k = 0, \end{cases} \quad (2.49)$$

is the only operator that satisfies the relations $LL^\dagger L = L$, $L^\dagger LL^\dagger = L^\dagger$, $(LL^\dagger)^ = LL^\dagger$, and $(L^\dagger L)^* = L^\dagger L$. The operator L^\dagger is called the Moore-Penrose pseudoinverse of L .*

Proof. We start by observing that we indeed have $L^\dagger : \mathcal{S}'(\mathbb{T}) \rightarrow \mathcal{S}'(\mathbb{T})$, since its frequency response satisfies $\widehat{L^\dagger} \in \mathcal{S}'(\mathbb{Z}, \mathbb{C})$ and is Hermitian-symmetric. Next, for any $w \in \mathcal{S}'(\mathbb{T})$, using the Fourier series expansion (2.16), we have

$$LL^\dagger L\{w\} = \sum_{k \in \mathbb{Z}} \widehat{L}[k] \widehat{L^\dagger}[k] \widehat{L}[k] \widehat{w}[k] e_k = \sum_{k \in \mathbb{Z}} \widehat{L}[k] \widehat{w}[k] e_k = L\{w\}, \quad (2.50)$$

which proves the first relation. The second one is proved in a similar fashion, and it is clear that $\widehat{L^\dagger}$ is the unique sequence that satisfies both relations. \square

Banach Structure of $\mathcal{M}_L(\mathbb{T})$

We now have all the necessary elements to uncover the direct-sum and Banach structure of $\mathcal{M}_L(\mathbb{T})$.

Theorem 2.4 (Banach Structure of $\mathcal{M}_L(\mathbb{T})$ [22, Theorem 1]). *Let $L = D^{N_d}$ with $N_d \geq 1$.*

Then, the native space $\mathcal{M}_L(\mathbb{T})$ admits the direct-sum decomposition

$$\mathcal{M}_L(\mathbb{T}) = L^\dagger(\mathcal{M}_0(\mathbb{T})) \oplus \mathcal{N}_L, \text{ where} \quad (2.51)$$

$$\mathcal{M}_0(\mathbb{T}) \triangleq \{w \in \mathcal{M}(\mathbb{T}) : \widehat{w}[0] = 0\} \quad (2.52)$$

is the space of zero-mean Radon measures. Hence, any $f \in \mathcal{M}_L(\mathbb{T})$ admits the unique decomposition

$$f = L^\dagger\{w\} + p, \quad (2.53)$$

where $w = L\{f\} \in \mathcal{M}_0(\mathbb{T})$ and $p = \widehat{f}[0] \in \mathcal{N}_L$. Moreover, $\mathcal{M}_L(\mathbb{T})$ is a Banach space equipped with the norm

$$\|f\|_{\mathcal{M}_L} \triangleq \|L\{f\}\|_{\mathcal{M}} + |\widehat{f}[0]|. \quad (2.54)$$

Remark 2.12. The direct sum (2.51) remains valid by replacing $\mathcal{M}_0(\mathbb{T})$ with $\mathcal{M}(\mathbb{T})$, as presented in [22, Theorem 1]. Indeed, we have $L^\dagger(\mathcal{M}_0(\mathbb{T})) = L^\dagger(\mathcal{M}(\mathbb{T}))$, since $\mathcal{N}_{L^\dagger} = \mathcal{M}(\mathbb{T})/\mathcal{M}_0(\mathbb{T}) \sim \text{span}\{x \mapsto 1\}$. However, the uniqueness of w in the decomposition (2.53) only holds true for $w \in \mathcal{M}_0(\mathbb{T})$, in contrast with the nonperiodic case (2.43) where $w \in \mathcal{M}(\mathbb{R})$, i.e., the full space of Radon measures. This is because in the periodic case, the pseudoinverse operator L^\dagger has a nonempty null space over $\mathcal{M}(\mathbb{T})$ since $1 \in \mathcal{N}_{L^\dagger} \cap \mathcal{M}(\mathbb{T})$, whereas L_ϕ^{-1} is injective over $\mathcal{M}(\mathbb{R})$ (in particular, $1 \notin \mathcal{M}(\mathbb{R})$).

Measurement Space of L

As in the nonperiodic case, the final required element for the construction of our native spaces is the identification of the predual space of $\mathcal{M}_L(\mathbb{T})$, which we call the *measurement space* of L. The latter is defined as follows:

$$\mathcal{C}_L(\mathbb{T}) \triangleq \{g \in \mathcal{S}'(\mathbb{T}) : L^{\dagger*}\{g\} \in \mathcal{C}(\mathbb{T})\}, \quad (2.55)$$

where $L^{\dagger*}$ is the adjoint operator of L^\dagger whose Fourier sequence is given by $\widehat{L^{\dagger*}\{g\}}[k] \triangleq \overline{\widehat{L^\dagger\{g\}}[k]}$. Analogously to $\mathcal{C}_L(\mathbb{R})$, $\mathcal{C}_L(\mathbb{T})$ admits a direct-sum decomposition and a Banach structure.

Theorem 2.5 (Banach Structure of $\mathcal{C}_L(\mathbb{T})$ [22, Theorems 2 and 3]). *Let $L = D^{N_d}$ with $N_d \geq 1$. Then, the space $\mathcal{C}_L(\mathbb{T})$ defined in (2.55) admits the direct-sum decomposition*

$$\mathcal{C}_L(\mathbb{T}) = L^*(\mathcal{C}(\mathbb{T})) \oplus \mathcal{N}_L, \quad (2.56)$$

and it is a Banach space equipped with the norm

$$\|g\|_{\mathcal{C}_L} \triangleq \left\| L^{\dagger*}\{g\} \right\|_{L^\infty} + |\widehat{g}[0]|. \quad (2.57)$$

Moreover, we have $(\mathcal{C}_L(\mathbb{T}))' = \mathcal{M}_L(\mathbb{T})$, and $\|\cdot\|_{\mathcal{M}_L}$ is the dual norm—in the sense of (2.2) of $\|\cdot\|_{\mathcal{C}_L}$.

The measurement space $\mathcal{C}_L(\mathbb{T})$ plays a key role in our inverse-problem formulations. A first useful observation is that we have $\mathcal{S}(\mathbb{T}) \subset \mathcal{C}_L(\mathbb{T})$ (this can easily be verified from (2.55) in the Fourier domain). Analogously to the $\delta \in \mathcal{C}_L(\mathbb{R})$ question, an important question is to determine whether we have $\mathbf{III} \in \mathcal{C}_L(\mathbb{T})$. Proposition 2.7—the periodic equivalent of Proposition 2.3—answers this question for N_d th-order derivative operators.

Proposition 2.7 (Sampling Admissibility in the Periodic Setting [22, Proposition 9]). *Let $L = D^{N_d}$ with $N_d \geq 1$. Then, we have $\mathbf{III} \in \mathcal{C}_L(\mathbb{T})$ if and only if $N_d \geq 2$.*

3 The Optimality of Splines for Generalized TV-Based Problems

In this chapter, we present our variational problems of interest with generalized total-variation (gTV) regularization, in both the nonperiodic and periodic settings. These problems, which are stated in Section 3.2, are formulated over the native spaces defined in Chapter 2. We also present existing representer theorems for these problems in Section 3.3, which prove that they have *spline* solutions. Hence, we first present some background on polynomial splines in Section 3.1.

3.1 Polynomial Splines

Polynomial splines, i.e., splines matched to N_0 -order derivative operators $L = D^{N_0}$ with $N_0 \geq 1$, are historically the first variety of splines to have been studied [23]. To this day, they are by far the most commonly used in practice, be it in this thesis or in the rest of the literature. In fact, the term “spline” refers specifically to polynomial splines in most communities. Polynomial splines are piecewise-polynomial functions of degree $N_0 - 1$ with smooth junctions at the knots, *i.e.*, their $(N_0 - 2)$ th derivative is continuous for $N_0 \geq 2$.

Splines are extremely useful tools to solve continuously-defined tasks computationally, as they serve as ideal bridge between the digital and analog world:

- they are parametric functions whose parameters can be handled computationally;
- their simple analytical expressions enable the computation of continuous operations such as differentiation, rotations, etc. [25];
- they admit computationally-friendly representations in the B-spline basis [23];
- they have the highest approximation power among any basis whose generator has a prescribed support [104];

- as we expose in Section 3.3, they are optimal solutions to large classes of continuous-domain inverse problems.

3.1.1 Definitions

3.1.2 Nonperiodic Setting

We define L-splines matched to derivative operators in Definition 3.1.

Definition 3.1 (Nonuniform L-Spline). *Let $L = D^{N_0} : \mathcal{M}_L(\mathbb{R}) \rightarrow \mathcal{M}(\mathbb{R})$ with $N_0 \geq 1$. A nonuniform L-spline is a function $s \in \mathcal{M}_L(\mathbb{R})$ that satisfies*

$$L\{s\} = \sum_{k=1}^K a_k \delta(\cdot - x_k), \quad (3.1)$$

where $K \geq 0$ is the number of knots, the $a_k \in \mathbb{R}$ are the amplitudes, and the $x_k \in \mathbb{R}$ are the pairwise-distinct knot locations. The distribution $w \triangleq \sum_{k=1}^K a_k \delta(\cdot - x_k) \in \mathcal{M}(\mathbb{R})$ is known as the innovation of the spline s .

An example of a D-spline (or *piecewise-constant spline*) with its innovation is given in Figure 3.1.

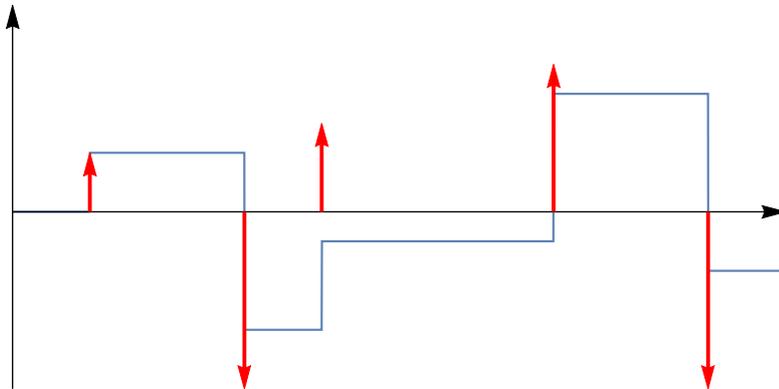


Figure 3.1 – Nonuniform D-spline (thin curve). The vertical arrows represent the innovation $D\{s\} = \sum_k a_k \delta(\cdot - x_k)$.

Any L-spline can conveniently be represented in the Green's function basis. Although this representation is well known [105], we prove it in the specific context of the native space $\mathcal{M}_L(\mathbb{R})$.

Proposition 3.1 (Green's Function Representation of Nonuniform L-splines). *Let $L = D^{N_0} : \mathcal{M}_L(\mathbb{R}) \rightarrow \mathcal{M}(\mathbb{R})$ with $N_0 \geq 1$, and $s \in \mathcal{M}_L(\mathbb{R})$ be a nonuniform L-spline in the*

sense of Definition 3.1 with $L\{s\} = \sum_{k=1}^K a_k \delta(\cdot - x_k)$. Then, s is given by

$$s = p + \sum_{k=1}^K a_k \rho_L(\cdot - x_k), \quad (3.2)$$

for some $p \in \mathcal{N}_L$ and where ρ_L is the Green's function of L in the sense of Definition 2.5.

Proof. Let $\tilde{s} \triangleq \sum_{k=1}^K a_k \rho_L(\cdot - x_k)$. By definition of the Green's function (Definition 2.5) we have $L\{s\} = L\{\tilde{s}\} = \sum_{k=1}^K a_k \delta(\cdot - x_k) \triangleq w$. Since $s, \tilde{s} \in \mathcal{M}_L(\mathbb{R})$, using the decomposition (2.43), for some stable right-inverse L_ϕ^{-1} of L given by Theorem 2.1, we have $s = L_\phi^{-1}\{w\} + q$ and $\tilde{s} = L_\phi^{-1}\{w\} + \tilde{q}$, with $q, \tilde{q} \in \mathcal{N}_L$. Hence, we have $s = \tilde{s} + q - \tilde{q}$, which proves the desired result with $p \triangleq q - \tilde{q} \in \mathcal{N}_L$. \square

We will often use the Green's function representation (3.2) rather than the decomposition (2.43) for L -splines. This is due to the simple expression of the basis function ρ_L given by (2.25), in contrast with $L_\phi^{-1}\{\delta\}$ (see Theorem 2.1).

For discretization purposes, we will rely in the works of Part II on *uniform splines*, *i.e.*, splines whose knots are on a uniform grid. When using a uniform grid, many continuous-domain operations on splines such as interpolation or differentiation can be performed using efficient digital-filtering techniques [106, 107]. Moreover, approximation theory results provide informative bounds on the approximation power of the used basis (see Theorem 3.1).

A uniform spline is a spline in the sense of Definition 3.1 whose knots x_k are on a uniform grid with step size $h > 0$, *i.e.*, $x_k \in h\mathbb{Z}$, and that can have infinitely many knots. We thus introduce the space of uniform splines with knot spacing h , which is a subspace of $\mathcal{M}_L(\mathbb{R})$.

Definition 3.2 (Space of Uniform L -Spline). *Let $L = D^{N_0} : \mathcal{M}_L(\mathbb{R}) \rightarrow \mathcal{M}(\mathbb{R})$ with $N_0 \geq 1$ and let $h > 0$. The space of uniform L -splines with knot spacing h is given by*

$$\mathcal{M}_{L,h}(\mathbb{R}) \triangleq \left\{ s = p + \sum_{k \in \mathbb{Z}} a[k] \rho_L(\cdot - kh) : a \in \ell_1(\mathbb{Z}), p \in \mathcal{N}_L \right\}, \quad (3.3)$$

where ρ_L is the Green's function of L defined in (2.25) and \mathcal{N}_L is the null space of L given by (2.24).

3.1.3 Periodic Setting

The definition of a periodic L -spline with $L = D^{N_d}$ is almost identical to the nonperiodic setting (Definition 3.1): the Dirac distribution δ is simply replaced with the Dirac comb $\text{III} \triangleq \sum_{k \in \mathbb{Z}} \delta(\cdot - 2k\pi)$.

Definition 3.3 (Periodic Nonuniform L-Spline). *Let $L = D^{N_d} : \mathcal{M}_L(\mathbb{T}) \rightarrow \mathcal{M}(\mathbb{T})$ with $N_d \geq 1$. A periodic nonuniform L-spline is a function $s \in \mathcal{M}_L(\mathbb{T})$ that satisfies*

$$L\{s\} = \sum_{k=1}^K a_k \text{III}(\cdot - x_k), \quad (3.4)$$

where $K \geq 0$ is the number of knots, the amplitudes $a_k \in \mathbb{R}$ necessarily satisfy

$$\sum_{k=1}^K a_k = 0, \quad (3.5)$$

and the $x_k \in \mathbb{T}$ are the pairwise-distinct knot locations. The distribution $w \triangleq \sum_{k=1}^K a_k \text{III}(\cdot - x_k)$ is known as the innovation of the spline.

Remark 3.1. *The constraint (3.5) does not appear in the nonperiodic case. It comes from the fact that $w = L\{s\}$ has zero mean—i.e., $\widehat{w}[0] = 0$ —since $\widehat{L}[0] = 0$, whereas the mean of the right side of (3.4) is $\sum_{k=1}^K a_k$. In particular, the Green's function $g_L = L^\dagger\{\text{III}\}$ of L is not a periodic L-spline, since we have $L\{g_L\} = \text{III} - 1$. However, $g_L - g_L(\cdot - x_0)$ is a periodic L-spline for any $x_0 \in \mathbb{T}$.*

Proposition 3.2 (Green's Function Representation of Periodic Nonuniform L-splines [22, Proposition 3]). *Let $L = D^{N_d} : \mathcal{M}_L(\mathbb{T}) \rightarrow \mathcal{M}(\mathbb{T})$ with $N_d \geq 1$, and $s \in \mathcal{M}_L(\mathbb{T})$ be a periodic nonuniform L-spline in the sense of Definition 3.3 with $L\{s\} = \sum_{k=1}^K a_k \text{III}(\cdot - x_k)$ and $\sum_{k=1}^K a_k = 0$. Then, s is given by*

$$s = \widehat{s}[0] + \sum_{k=1}^K a_k g_L(\cdot - x_k), \quad (3.6)$$

where g_L is the Green's function of L defined in (2.32).

Similarly to the nonperiodic setting, we can define the space of periodic uniform splines with knot spacing $h > 0$.

Definition 3.4 (Space of Periodic Uniform L-Spline). *Let $L = D^{N_d} : \mathcal{M}_L(\mathbb{T}) \rightarrow \mathcal{M}(\mathbb{T})$ with $N_d \geq 1$ and $h > 0$ such that $P \triangleq \frac{2\pi}{h} \in \mathbb{N}$. The space of periodic uniform L-splines with knot spacing h is given by*

$$\mathcal{M}_{L,h}(\mathbb{T}) \triangleq \left\{ s \in \mathcal{S}'(\mathbb{T}) : L\{s\} = \sum_{p=0}^{P-1} a[p] \text{III}(\cdot - ph) \right\} \subset \mathcal{M}_L(\mathbb{T}), \quad (3.7)$$

where $\mathbf{a} = (a[0], \dots, a[P-1]) \in \mathbb{R}^P$.

3.1.4 Cardinal B-Splines

The main benefit of using spline functions is arguably the availability of a basis function with short support, the B-spline. This fundamental result was first discovered by Schoenberg [23]. This short-support property is crucial for numerical conditioning in computational applications. We now introduce polynomial B-splines, *i.e.*, B-splines matched to operators $L = D^{N_0}$ with $N_0 \geq 1$. Note that B-splines can be derived for more general classes of operators, such as ordinary differential operators [101], rational operators [102], fractional derivatives [100], or Laplacian operators in higher dimensions [108]. In fact, differential and rational operators still have a compactly supported B-spline, which implies that B-spline-based works in this thesis (Part II and Chapters 9 and 10) can be extended to these classes of operators.

The causal B-spline β_D matched to the derivative $L = D$ is the causal rectangle function

$$\beta_D(x) \triangleq \mathbb{1}_{[0,1)}(x). \quad (3.8)$$

For N_0 th-order derivative operators $L = D^{N_0}$, the B-spline is defined as the convolution of first-order B-splines

$$\beta_L \triangleq \underbrace{\beta_D * \cdots * \beta_D}_{N_0 \text{ times}}. \quad (3.9)$$

The support of β_L is $\text{supp}(\beta_L) = [0, N_0]$, and it has $(N_0 + 1)$ knots $\{0, \dots, N_0\}$. It is the *cardinal* L-spline, *i.e.*, a spline in the sense of Definition 3.1 with integer knots $x_k \in \mathbb{Z}$, with the shortest support. Its innovation is given by

$$L\{\beta_L\} = \sum_{k=0}^{N_0} d_L[k] \delta(\cdot - k), \quad (3.10)$$

where d_L is the N_0 th order finite-difference sequence defined by its z transform

$$D_L(z) \triangleq (1 - z^{-1})^{N_0}, \quad (3.11)$$

which is supported in $\{0, \dots, N_0\}$. It can also be expressed as a series of discrete convolutions $d_L = \underbrace{d_D * \cdots * d_D}_{N_0 \text{ times}}$, where $d_D[k] = \delta[k] - \delta[k - 1]$ and $\delta[k]$ is the discrete

Dirac impulse. Due to the fact that it is defined as a series of convolutions in (3.9), the B-spline has a convenient expression in the Fourier domain given by

$$\widehat{\beta}_L(\omega) = \frac{D_L(e^{j\omega})}{\widehat{L}(\omega)} = \left(\frac{1 - e^{-j\omega}}{j\omega} \right)^{N_0}. \quad (3.12)$$

The analytical expression of the B-splines for small values of N_0 are provided in Table 3.1,

and their graphs are shown in Figure 3.2.

L	$\rho_L(x)$	$\beta_L(x)$	$(d_L[0], \dots, d_L[N_0])$
D	$\mathbb{1}_+(x)$	$\beta_D(x) = \begin{cases} 1, & 0 \leq x < 1 \\ 0, & \text{otherwise} \end{cases}$	$(1, -1)$
D^2	x_+	$\beta_{D^2}(x) = \begin{cases} x, & 0 \leq x < 1 \\ 2 - x, & 1 \leq x < 2 \\ 0, & \text{otherwise} \end{cases}$	$(1, -2, 1)$
D^3	$x_+^2/2$	$\beta_{D^3}(x) = \begin{cases} x^2/2, & 0 \leq x < 1 \\ -x^2 + 3x - 3/2, & 1 \leq x < 2 \\ (3 - x^2)/2, & 2 \leq x < 3 \\ 0, & \text{otherwise} \end{cases}$	$(1, -3, 3, -1)$
D^4	$x_+^3/6$	$\beta_{D^4}(x) = \begin{cases} x^3/6, & 0 \leq x < 1 \\ -x^3/2 + 2x^2 - 2x + 2/3, & 1 \leq x < 2 \\ x^3/2 - 4x^2 + 10x - 22/3, & 2 \leq x < 3 \\ (4 - x)^3/6, & 3 \leq x < 4 \\ 0, & \text{otherwise} \end{cases}$	$(1, -4, 6, -4, 1)$

Table 3.1 – Characteristics of N_0 th-order derivative operators.

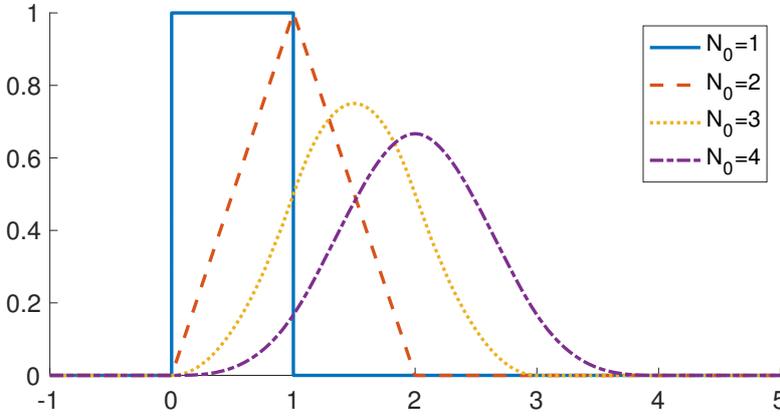


Figure 3.2 – Cardinal polynomial B-splines of operators $L = D^{N_0}$ for $N_0 = 1, \dots, 4$.

Since polynomial splines are symmetric with respect to the $x = N_0/2$ line, it is often more convenient to use centered B-splines instead of the causal ones shown in Figure 3.2. This is particularly relevant for even values of N_0 , since the centered B-spline still has integer knots, whereas they are half-integers for odd values of N_0 . Hence, in certain chapters of this thesis, we will use widely-used notation $\beta^n \triangleq \beta_{D^{N_0}}(\cdot + N_0/2)$ where $n = N_0 - 1$ is the *degree* of the spline for centered polynomial B-splines.

3.1.5 Scaled B-Splines

3.1.6 Nonperiodic Setting

The scaled B-Spline with knot spacing $h > 0$ matched to the operator $L = D^{N_0}$ is given by

$$\beta_{L,h}(x) \triangleq \beta_L\left(\frac{x}{h}\right), \quad (3.13)$$

which yields

$$\widehat{\beta_{L,h}}(\omega) = \frac{1}{h^{N_0-1}} \left(\frac{1 - e^{-j\omega h}}{j\omega} \right)^{N_0} \quad (3.14)$$

in the Fourier domain. Using (3.10), we get the innovation of the scaled B-spline

$$L\{\beta_{L,h}\} = \frac{1}{h^{N_0-1}} \sum_{k=0}^{N_0} d_L[k] \delta(\cdot - hk). \quad (3.15)$$

We now prove the fundamental property of the scaled B-spline, *i.e.*, that it is a basis function of the space of uniform splines $\mathcal{M}_{L,h}(\mathbb{R})$. Proposition 3.3 is a generalization from polynomial to exponential splines of the seminal and widely publicized result of Schoenberg [23]. Moreover, we explicitly define the resulting function space $\mathcal{M}_{L,h}(\mathbb{R})$ and sequence space for our setting, which is often overlooked.

Proposition 3.3. *The discretized native space $\mathcal{M}_{L,h}(\mathbb{R})$ defined in Definition 3.2 of the N_0 th order derivative operator $L = D^{N_0}$ with $N_0 \geq 1$ can be represented in the B-spline basis as*

$$\mathcal{M}_{L,h}(\mathbb{R}) = \left\{ s = \sum_{k \in \mathbb{Z}} c[k] \beta_{L,h}(\cdot - kh) : c \in \ell_{1,L}(\mathbb{Z}) \right\} \quad (3.16)$$

where $\beta_{L,h}$ is defined as in (3.13) and

$$\ell_{1,L}(\mathbb{Z}) \triangleq \left\{ c \in \mathcal{S}'(\mathbb{Z}) : (d_L * c) \in \ell_1(\mathbb{Z}) \right\}, \quad (3.17)$$

with $\mathcal{S}'(\mathbb{Z})$ the space of slowly growing sequences defined in (2.14).

Proof. We first prove the reverse inclusion $c \in \ell_{1,L}(\mathbb{Z}) \Rightarrow s = \sum_{k \in \mathbb{Z}} c[k] \beta_{L,h}(\cdot - kh) \in \mathcal{M}_{L,h}(\mathbb{R})$. A simple calculation using (3.15) yields

$$L \left\{ \sum_{k \in \mathbb{Z}} c[k] \beta_{L,h}(\cdot - kh) \right\} = \frac{1}{h^{N_0-1}} \sum_{k \in \mathbb{Z}} (d_L * c) \delta(\cdot - kh), \quad (3.18)$$

Chapter 3 The Optimality of Splines for Generalized TV-Based Problems

with $(d_L * c) \in \ell_1(\mathbb{Z})$. This together with the fact that $s \in \mathcal{S}'(\mathbb{R})$ (due to the fact that $c \in \mathcal{S}'(\mathbb{Z})$) implies that $s \in \mathcal{M}_{L,h}(\mathbb{R})$.

We now show the direct inclusion, *i.e.*, that any $s \triangleq p + \sum_{k \in \mathbb{Z}} a[k] \rho_L(\cdot - kh) \in \mathcal{M}_{L,h}(\mathbb{R})$ with $p \in \mathcal{N}_L$ and $a \in \ell_1(\mathbb{Z})$ is spanned by B-splines. To achieve this, we prove that the Green's function ρ_L and the null space \mathcal{N}_L are spanned by B-splines. Following standard discrete signal processing theory, there exists a unique causal sequence given by $p_L[n] \triangleq u[n]$ for $N_0 = 1$ and $p_L[n] \triangleq \frac{(n+1)\dots(n+N_0-1)}{(N_0-1)!} u[n]$ for $N_0 \geq 1$, characterized by its z transform $P_L(z) = \frac{1}{(1-z^{-1})^{N_0}}$, that is an inverse of d_L for the convolution product. It can easily be verified that we have

$$\rho_L(x) = h^{N_0-1} \sum_{k \in \mathbb{Z}} p_L[k] \beta_{L,h}(x - hk). \quad (3.19)$$

Next, it is well known that any polynomial $x \mapsto x^n$ for $0 \leq n \leq N_0 - 1$ can be reproduced using B-splines like $1 = \sum_{k \in \mathbb{Z}} \beta_{L,h}(\cdot - hk)$ for $n = 0$ and

$$x^n = h^n \sum_{k \in \mathbb{Z}} (k+1) \dots (k+n) \beta_{L,h}(x - hk) \quad 1 \leq n \leq N_0 - 1, \quad (3.20)$$

which proves that \mathcal{N}_L is spanned by B-splines. Hence, there only remains to prove that the sequence c of B-spline coefficients of $s = p + \sum_{k \in \mathbb{Z}} a[k] \rho_L(\cdot - kh) \in \mathcal{M}_{L,h}(\mathbb{R})$ satisfies $c \in \ell_{1,L}(\mathbb{Z})$. A simple calculation yields $\sum_{k \in \mathbb{Z}} a[k] \rho_L(\cdot - kh) = h^{N_0-1} \sum_{k \in \mathbb{Z}} (p_L * a) \beta_{L,h}(x - hk)$, where the convolution $(p_L * a) \in \mathcal{S}'(\mathbb{Z})$ is clearly well defined when $a \in \mathcal{S}(\mathbb{Z})$ (the space of rapidly-decreasing sequences) since p_L is slowly growing. Next, $\mathcal{S}(\mathbb{Z})$ is dense in $\ell_1(\mathbb{Z})$, which allows us to extend this definition to any $a \in \ell_1(\mathbb{Z})$ by continuity [20, Theorem 16]. We thus have $d_L * c = a \in \ell_1(\mathbb{Z})$. Finally, for the component $p \in \mathcal{N}_L$, the sequence c_p of B-spline coefficients verifies $d_L * c_p = 0 \in \ell_1(\mathbb{Z})$. This proves the direct inclusion and thus the desired result. \square

The major advantage of using the B-spline representation of $\mathcal{M}_{L,h}(\mathbb{R})$ over the Green's function representation (3.3) for computational tasks is the limited support of the basis function, the shifted B-spline $\beta_{L,h}$. In fact, $\beta_{L,h}$ is the nonzero member of $\mathcal{M}_{L,h}(\mathbb{R})$ that has minimal support. This makes the B-spline basis close to being orthogonal (it is a Riesz basis [101, Theorem 1]), which leads to efficient implementations and well-conditioned computational tasks. Conversely, the basis function in (3.3)—the Green's function ρ_L —has infinite support and is sometimes even increasing at infinity, which leads to severely ill-conditioned computational tasks.

3.1.7 Periodic Setting

Let $L = D^{N_d}$ with $N_d \geq 1$ and $h > 0$ such that $P \triangleq \frac{2\pi}{h} \in \mathbb{N}$. In a periodic setting, one can simply define the 2π -periodized B-spline as

$$\beta_{L,h}^{\text{per}} \triangleq \sum_{k \in \mathbb{Z}} \beta_{L,h}(\cdot - 2k\pi). \quad (3.21)$$

Note that when the grid is sufficiently fine, *i.e.*, $P \geq N_d$, the periodic B-spline is not aliased, since we have $\text{supp}(\beta_{L,h}) = [0, hN_d] \subset [0, 2\pi] = \mathbb{T}$: we thus have $\beta_{L,h}^{\text{per}}(x) = \beta_{L,h}(x)$ for $x \in \mathbb{T}$. One readily shows using standard Fourier analysis that the Fourier series coefficients of $\beta_{L,h}^{\text{per}}$ are given by

$$\widehat{\beta_{L,h}^{\text{per}}}[k] = \frac{1}{2\pi} \widehat{\beta_{L,h}}(k) = \frac{h}{2\pi} \left(\frac{1 - e^{-jkh}}{jkh} \right)^{N_d}. \quad (3.22)$$

Moreover, $\beta_{L,h}^{\text{per}}$ is a periodic L-spline in the sense of Definition 3.3, and using (3.15), its innovation is given by

$$L \left\{ \beta_{L,h}^{\text{per}} \right\} = \frac{1}{h^{N_d-1}} \sum_{p=0}^{P-1} d_L^{\text{per}}[p] \text{III}(\cdot - ph), \quad (3.23)$$

where d_L^{per} is the P -periodization of d_L

$$d_L^{\text{per}}[k] \triangleq \sum_{n \in \mathbb{Z}} d_L[k + nP], \quad (3.24)$$

whose discrete Fourier transform (DFT) is given by

$$D_L^{\text{per}}[k] = D_L \left(e^{jkh} \right) = \left(1 - e^{-jkh} \right)^{N_d}. \quad (3.25)$$

When $P > N_d$, since $\text{supp}(d_L) = \{0, \dots, N_d\}$, d_L^{per} does not suffer from aliasing in that $d_L^{\text{per}}[k] = d_L[k]$ for any $k \in \{0, \dots, P-1\}$. For example, for $L = D$, d_L^{per} is the P -periodized finite-difference sequence $d_L^{\text{per}}[k] = \delta_P[k] - \delta_P[k-1]$ where $\delta_P[k] = \sum_{n \in \mathbb{Z}} \delta[k + nP]$.

Periodic B-splines share the same celebrated property as their nonperiodic counterparts: they are generators of the space of uniform (periodic) splines $\mathcal{M}_{L,h}(\mathbb{T})$ introduced in (3.3), as demonstrated in Proposition 3.4. This representation leads to the same computational benefits as in the nonperiodic case.

Proposition 3.4. *Let L^{N_d} with $N_d \geq 1$ and $h > 0$ such that $P \triangleq \frac{2\pi}{h} \in \mathbb{N}$. Then, the*

periodic B-spline $\beta_{L,h}^{\text{per}}$ is a generator of the space $\mathcal{M}_{L,h}(\mathbb{T})$, i.e., we have

$$\mathcal{M}_{L,h}(\mathbb{T}) = \left\{ s = \sum_{p=0}^{P-1} c[p] \beta_{L,h}^{\text{per}}(\cdot - ph), \mathbf{c} = (c[0], \dots, c[P-1]) \in \mathbb{R}^P \right\}. \quad (3.26)$$

Proof. We first observe that the space $\mathcal{M}_{L,h}(\mathbb{T})$ is a P -dimensional vector space: there are $(P-1)$ degrees of freedom for the $a[p]$ coefficients in (3.3) (P coefficients and one linear constraint $\sum_{p=0}^{P-1} a[p] = 0$), and one for the mean $\widehat{s}[0]$ (see Proposition 3.2). Next, we prove that for any $\mathbf{c} = (c[0], \dots, c[P-1]) \in \mathbb{R}^P$, we have $s = \sum_{p=0}^{P-1} c[p] \beta_{L,h}^{\text{per}}(\cdot - ph) \in \mathcal{M}_{L,h}(\mathbb{T})$. Indeed, we have that

$$\begin{aligned} \mathbf{L}\{s\} &= \frac{1}{h^{N_d-1}} \sum_{p=0}^{P-1} \sum_{q=0}^{P-1} c[p] d_L^{\text{per}}[q] \mathbb{III}(\cdot - (p+q)h) \\ &= \frac{1}{h^{N_d-1}} \sum_{p=0}^{P-1} (\mathbf{d}_L * \mathbf{c})[p] \mathbb{III}(\cdot - ph), \end{aligned} \quad (3.27)$$

where (3.23) was used for the first line, and $(\mathbf{d}_L * \mathbf{c})$ denotes here the cyclic convolution between the vectors $\mathbf{d}_L = (d_L^{\text{per}}[0], \dots, d_L^{\text{per}}[P-1])$ and \mathbf{c} . This proves that $s \in \mathcal{M}_{L,h}(\mathbb{T})$ with coefficients $(a[0], \dots, a[P-1]) = h^{N_d-1} (\mathbf{d}_L * \mathbf{c})$, and thus that the space generated by shifts of $\beta_{L,h}^{\text{per}}$ is included in $\mathcal{M}_{L,h}(\mathbb{T})$. Yet both are P -dimensional vector spaces, which proves that they are in fact equal. \square

3.1.8 Approximation Power

In this section, we give some brief insights on uniform splines from approximation theory. More precisely, we discuss their *approximation power*, i.e., the asymptotic upper bound on the approximation error between a generic function and a uniform spline as the grid size h goes to zero.

Theorem 3.1 (Approximation Power of Polynomial Splines [109, Theorem 4.3]). *Let $L = D^{N_0}$, $h > 0$, $\mathcal{V} \triangleq \{f \in L_2(\mathbb{R}) : D^{N_0}\{f\} \in L_2(\mathbb{R})\}$,*

$$\mathcal{V}_h \triangleq \left\{ s = \sum_{k \in \mathbb{Z}} c[k] \beta_{L,h}(\cdot - kh) : c \in \ell_2(\mathbb{Z}) \right\}, \quad (3.28)$$

and let $P_h : \mathcal{V} \rightarrow \mathcal{V}_h$ be the orthogonal projector of \mathcal{V} onto \mathcal{V}_h . Then, for any $f \in \mathcal{V}$, we have

$$\|f - P_h\{f\}\|_{L_2} = \mathcal{O}(h^{N_0}). \quad (3.29)$$

Remark 3.2. *Theorem 3.1 does not directly apply to our setting, since the spaces under consideration are $L_2(\mathbb{R}) \neq \mathcal{M}_L(\mathbb{R})$ and $\mathcal{V}_h \neq \mathcal{M}_{L,h}(\mathbb{R})$. However, the important high-level*

message to remember is that the approximation power of uniform L-splines increases with the order of differentiation N_0 of L , and the approximation error is in $\mathcal{O}(h^{N_0})$, which is in keeping with the well-known Strang-Fix conditions [110].

3.2 Generalized TV-Based Inverse Problem Formulations

We now have all the necessary tools to introduce our continuous-domain inverse problems of interest. The goal of these problems is to recover a continuous-domain *ground-truth* signal $f_0 : \mathbb{R} \rightarrow \mathbb{R}$, based on an observation vector $\mathbf{y} \in \mathbb{R}^M$, which are acquired via a measurement operator $\nu : f \mapsto \nu(f) \in \mathbb{R}^M$ where $\nu(f_0) \approx \mathbf{y}$. In the case of an additive noise model, we have

$$\mathbf{y} = \nu(f_0) + \mathbf{n}, \quad (3.30)$$

where $\mathbf{n} \in \mathbb{R}^M$ models the measurement errors. To reconstruct an estimate f^* of f_0 , we solve the following optimization problem:

$$f^* \in \mathcal{V} \triangleq \arg \min_{f \in \mathcal{M}_L(\mathbb{K})} \left(\underbrace{E(\nu(f), \mathbf{y})}_{\text{Data-fidelity term}} + \underbrace{\lambda \|L\{f\}\|_{\mathcal{M}}}_{\text{Regularization term}} \right), \quad (3.31)$$

whose search space $\mathcal{M}_L(\mathbb{K})$ is defined in (2.33). The cost function \mathcal{J} can be decomposed into the following two terms.

1. A *data-fidelity* term $E(\nu(f), \mathbf{y})$, which consists of the following elements:
 - a *measurement operator* (or forward model) $\nu = (\nu_1, \dots, \nu_M) : \mathcal{M}_L(\mathbb{K}) \rightarrow \mathbb{R}^M$ whose null space is defined as $\mathcal{N}_\nu \triangleq \{f \in \mathcal{M}_L(\mathbb{K}) : \nu(f) = \mathbf{0}\}$;
 - an *observation vector* $\mathbf{y} = (y_1, \dots, y_M) \in \mathbb{R}^M$;
 - a *data-fidelity cost functional* $E : \mathbb{R}^M \times \mathbb{R}^M \rightarrow \mathbb{R}^+ \cup \{+\infty\}$.

This term enforces the consistency between the observations \mathbf{y} and measurements $\nu(f^*)$ of the reconstructed signal $f^* \in \mathcal{V}$, *i.e.*, we typically have $\nu(f^*) \approx \mathbf{y}$.

2. A *regularization term* $\lambda \|L\{f\}\|_{\mathcal{M}}$, which consists of the following elements:
 - a *regularization parameter* $\lambda > 0$;
 - a *regularization operator* $L : \mathcal{M}_L(\mathbb{K}) \rightarrow \mathcal{M}(\mathbb{K})$, where $\mathcal{M}(\mathbb{K})$ is defined in (2.18), whose null space \mathcal{N}_L defined in (2.23) has dimension N_0 ;
 - the *regularization norm* $\|\cdot\|_{\mathcal{M}}$, which is the TV norm for measures defined in (2.17).

This term enforces that the reconstructed signal $f^* \in \mathcal{V}$ be of a prescribed form based solely on our prior knowledge, *i.e.*, independent of the measured data \mathbf{y} .

Our choice of regularization norm, the TV norm for measures, promotes sparse reconstructions in a sense that will be clarified in Theorems 3.3 and 3.4. The regularization operator specifies the transform domain in which the reconstruction is sparse. Finally, the regularization parameter λ controls the balance between the data-fidelity and regularization terms, and should be chosen in accordance with the noise level and the degree of ill-posedness of the problem.

Example 3.1. *The two main examples of data-fidelity cost functional that we will use throughout this thesis are:*

- the indicator function¹ $E(\mathbf{z}, \mathbf{y}) \triangleq i_{\{\mathbf{y}\}}(\mathbf{z})$, which leads to the constrained problem

$$\mathcal{V} \triangleq \arg \min_{f \in \mathcal{M}_L(\mathbb{K}): \nu(f) = \mathbf{z}} \|\mathbb{L}\{f\}\|_{\mathcal{M}}. \quad (3.32)$$

This choice is pertinent when the observations \mathbf{y} are assumed to be noiseless; the regularization parameter λ then plays no role.

- the quadratic loss function $E(\mathbf{y}, \mathbf{z}) \triangleq \frac{1}{2} \|\mathbf{z} - \mathbf{y}\|_2^2$, which, from a Bayesian point of view, is well suited when the noise \mathbf{n} in (3.30) follows an independent and identically distributed (i.i.d.) Gaussian model (see for example [111, Section 4.3]).

Other classical data-fidelity cost functionals can be found in [92, Section 7.5].

3.3 Representer Theorems

3.3.1 Preliminaries on Convex Sets and Extreme Points

Before presenting existing representer theorems for Problem 3.31, we provide some preliminaries on the geometry of convex sets. In particular, the concept of *extreme point* will play an important role in these representer theorems.

Definition 3.5 (Extreme Point). *Let \mathcal{C} be a convex subset of a real or complex vector space. An element $c \in \mathcal{C}$ is an extreme point of \mathcal{C} if c cannot be written as $c = tc_1 + (1-t)c_2$ where $c_1, c_2 \in \mathcal{C}$ with $c_1 \neq c_2$, and with $0 < t < 1$.*

Intuitively, the extreme points are the corner points of a convex set; for example, the extreme points of a convex polygon are its vertices. However, a convex set may well have infinitely many extreme points; for example, the extreme points of a disk are *all* the points of the corresponding circle. Next, we define the *convex hull* of a set.

Definition 3.6 (Convex Hull). *Let \mathcal{S} be a subset of a real or complex vector space \mathcal{X} . The convex hull of \mathcal{S} is the set of points $x \in \mathcal{X}$ such there exist $s_1, \dots, s_N \in \mathcal{S}$ and $0 \leq t_1, \dots, t_N \leq 1$ with $\sum_{i=1}^N t_i = 1$ such that $x = \sum_{i=1}^N t_i s_i$.*

¹The indicator function of a set A is defined as $i_A(x) = 0$ if $x \in A$, and $i_A(x) = +\infty$ if $x \notin A$.

We can now recall the classical Krein-Milman theorem, which underscores the crucial role of extreme points in the geometry of convex sets.

Theorem 3.2 (Kein-Milman Theorem [15, p.75]). *Let \mathcal{C} be a compact convex subset of a Hausdorff locally convex topological vector space (Section 2.1.1). Then, \mathcal{C} is the closed convex hull of its extreme points.*

Defining Hausdorff locally convex spaces is out of the scope of this thesis. What is important for our purpose is that they include all spaces introduced in Chapter 2, in particular the native space $\mathcal{M}_L(\mathbb{K})$ defined in (2.33) equipped with the weak* topology.

3.3.2 Nonperiodic Representer Theorem

We now present existing representer theorems for Problem (3.31), which guarantee the existence of spline solutions. We first consider the nonperiodic case $\mathbb{K} = \mathbb{R}$.

Theorem 3.3 (Nonperiodic Representer Theorem [14, Theorem 1]). *Assume that the following conditions are met:*

- the regularization operator $L : \mathcal{M}_L(\mathbb{R}) \rightarrow \mathcal{M}(\mathbb{R})$ is the N_0 -th derivative operator $L = D^{N_0}$;
- the measurement operator $\nu = (\nu_1, \dots, \nu_M) : \mathcal{M}_L(\mathbb{R}) \rightarrow \mathbb{R}^M$ is a weak*-continuous linear operator;
- the observation vector is of the form $\mathbf{y} = (y_1, \dots, y_M) \in \mathbb{R}^M$;
- the data-fidelity cost functional $E : \mathbb{R}^M \times \mathbb{R}^M \rightarrow \mathbb{R}^+ \cup \{+\infty\}$ is such that $E(\cdot, \mathbf{y}) : \mathbb{R}^M \rightarrow \mathbb{R}^+ \cup \{+\infty\}$ is a proper convex function that is strictly convex over its effective domain² for any $\mathbf{y} \in \mathbb{R}^M$;
- we have $\mathcal{N}_L \cap \mathcal{N}_\nu = \{0\}$ (well-posedness assumption).

Then, the solution set $\mathcal{V} \subset \mathcal{M}_L(\mathbb{R})$ of the optimization Problem (3.31) is a nonempty weak*-compact convex set whose extreme points (Definition 3.5) s^* are nonuniform L-splines (Definition 3.1) of the form

$$s^*(x) = p + \sum_{k=1}^K a_k \rho_L(x - x_k), \quad (3.33)$$

where $p \in \mathcal{N}_L$, $a_k, x_k \in \mathbb{R}$, and whose number of knots K is bounded by $K \leq M - N_0$.

²The effective domain of a convex function $g : \mathcal{X} \rightarrow \mathbb{R}^+ \cup \{+\infty\}$ is $\{x \in \mathcal{X} : g(x) < +\infty\}$.

Theorem 3.3 has a long history in the literature, stemming from [87], where this result was proved in a compact-support setting (which differs from the periodic setting we will present in Theorem 3.4). It was revived much later in 2017 by [14], which sparked several recent similar results [45, Theorem 4], [90] and generalizations [61, 112].

Remark 3.3. *Theorem 3.3 is an adaptation of [14, Theorem 1]; this adaptation is better suited to how this theorem is used in our works. More precisely, the data-fidelity functional E is not present in [14], which instead enforces that $\nu(f) \in \mathcal{C}$, where \mathcal{C} is a convex set. It is straightforward to show that the hypotheses of Theorem 3.4 can be reformulated as such. Moreover, as stated earlier, the operator L takes a more general form in [14]: Theorem 3.3 remains valid for any spline-admissible operator in the sense of [14, Definition 1].*

Remark 3.4. *It is well known that the assumption that ν is weak* continuous is equivalent to $\nu \in \mathcal{C}_L(\mathbb{R})^M$ [113, Theorem IV.20, p. 114], where $\mathcal{C}_L(\mathbb{R})$ is the predual of $\mathcal{M}_L(\mathbb{R})$ defined in (2.46). This explains why we call $\mathcal{C}_L(\mathbb{R})$ the measurement space of L : it is the space in which the measurement functionals can lie. Here, we follow the standard practice of identifying a linear functional with a (generalized) function, i.e., $\nu_m(f) = \langle \nu_m, f \rangle_{\mathcal{C}_L(\mathbb{R}) \times \mathcal{M}_L(\mathbb{R})}$. What is less standard with this notation is that the functional ν_m is an element of the predual of $\mathcal{M}_L(\mathbb{R})$ instead of its dual; note however that we have $\mathcal{C}_L(\mathbb{R}) \subset \mathcal{M}_L(\mathbb{R})'$.*

Example 3.2. *By Proposition 2.3, an example of a weak*-continuous measurement functional when $N_0 \geq 2$ is the shifted Dirac impulse $\nu_m = \delta(\cdot - x_m)$ for any $x_m \in \mathbb{R}$, which yields $\nu_m(f) = f(x_m)$, i.e., the spatial sampling operator. These are used in the works of Part II and Chapters 7 and 8.*

Example 3.3. *Another example of a valid measurement operator is the windowed Fourier-domain sampling functional $\nu_m = \mathbb{1}_{[0,T]} \cos(\omega_m \cdot - \phi_m)$, which yields $\nu_m(f) = \int_0^T f(x) \cos(\omega_m x - \phi_m) dx$. Indeed, we have $\nu_m \in L_1(\mathbb{R}) \subset \mathcal{C}_D(\mathbb{R})$ (as pointed out in Remark 2.11), and by Lemma 2.1, we clearly have $\nu_m \in \mathcal{C}_L(\mathbb{R})$ for any D^{N_0} with $N_0 \geq 1$. These are used in the works of Part II.*

Remark 3.5. *The well-posedness assumption in Theorem 3.3 is needed to ensure that \mathcal{V} is bounded. Otherwise, for any $f^* \in \mathcal{V}$ and $p \in \mathcal{N}_L \cap \mathcal{N}_\nu$, we would have $f^* + \alpha p \in \mathcal{V}$ for any $\alpha \in \mathbb{R}$ (since we clearly have $\mathcal{J}(f^*) = \mathcal{J}(f^* + \alpha p)$). Although this would not pose any theoretical problem, having a bounded solution set is a desirable property to solve optimization problems numerically, which explains why this additional assumption is made.*

Theorem 3.3 demonstrates the sparsity-promoting effect of the TV norm for measures. Indeed, the argument of the TV norm in (3.31) for the extreme-point solutions s^* of the form (3.33) satisfies $w^* \triangleq L\{s^*\} = \sum_{k=1}^K a_k \delta(x - x_k)$. The measure w^* is sparse in the sense that it has a discrete support $\text{supp}(w^*) = \{x_1, \dots, x_K\}$, whose number of elements is “small” since it is upper-bounded by $M - N_0$. This definition of sparsity is analogous

to the more common discrete-domain definition, where a vector is considered sparse if it has a small support, *i.e.*, few nonzero elements. However, an important point to bear in mind is that this sparsity-promoting effect only applies to the *extreme-point solutions* of Problem (3.31). In cases of nonuniqueness, interior-point solutions $f^* \in \mathcal{V}$ may be highly nonsparse (*i.e.*, the support of $L\{f^*\}$ may have nonzero Lebesgue measure), as we will demonstrate for example in Chapter 7. In fact, the study of the solution set \mathcal{V} in terms of uniqueness, its sparsest solutions, and how to reach them algorithmically is a major topic of this thesis, and particularly of Part III.

Moreover, the role of L is made clear from the form of the solutions s^* given by (3.33): it specifies the transform domain in which s^* is sparse, *i.e.*, the Green's function basis $\{\rho_L(\cdot - x_0)\}_{x_0 \in \mathbb{R}}$. For example, for $L = D$, s^* is a piecewise-constant function with fewer than $M - 1$ jumps. A remarkable feature of Theorem 3.3 is that the basis in which the solution is sparse is solely determined by the regularization L . This differs for example from the L_2 —or Tikhonov—regularization scenario, where this basis also depends on the measurement operator ν [45, Theorem 3].

3.3.3 Periodic Representer Theorem

We now present the equivalent of Theorem 3.3 in the periodic setting.

Theorem 3.4 (Periodic Representer Theorem [22, Theorem 4]). *Assume that the following conditions are met:*

- *the regularization operator is the N_d -th order derivative $L = D^{N_d} : \mathcal{M}_L(\mathbb{T}) \rightarrow \mathcal{M}(\mathbb{T})$ with $N_d \geq 1$;*
- *the measurement operator (or forward model) $\nu = (\nu_1, \dots, \nu_M) : \mathcal{M}_L(\mathbb{T}) \rightarrow \mathbb{R}^M$ is a weak*-continuous linear operator;*
- *the observation vector is of the form $\mathbf{y} = (y_1, \dots, y_M) \in \mathbb{R}^M$;*
- *the cost functional $E : \mathbb{R}^M \times \mathbb{R}^M \rightarrow \mathbb{R}^+$ is such that $E(\cdot, \mathbf{y}) : \mathbb{R}^M \times \mathbb{R}^M \rightarrow \mathbb{R}^+$ is a proper convex function that is strictly convex over its effective domain for any $\mathbf{y} \in \mathbb{R}^M$.*

Then, the solution set $\mathcal{V} \subset \mathcal{M}_L(\mathbb{T})$ of Problem (3.31) is a nonempty weak-compact convex set whose extreme points (Definition 3.5) are periodic nonuniform L -splines (Definition 3.3) that satisfies*

$$s^*(x) = \widehat{s^*}[0] + \sum_{k=1}^K a_k g_L(x - x_k), \quad (3.34)$$

where $\widehat{s^}[0] \in \mathbb{R}$, $a_k, x_k \in \mathbb{R}$, and whose number of knots K is bounded by $K \leq M + 1$.*

Remark 3.6. *Once again, Theorem 3.4 is an adaptation of [22, Theorem 4], which is more general: it remains valid for any spline-admissible operator L in the sense of [22, Definition 2]. Examples of periodic spline-admissible operators are given in [22, Table 1].*

Example 3.4. *As in nonperiodic case (see Remark 3.4), the weak*-continuity of ν is equivalent to $\nu_m \in \mathcal{C}_L(\mathbb{T})$ for any $1 \leq m \leq M$. By Proposition 2.7, an example of a weak*-continuous measurement functional when $N_0 \geq 2$ is the shifted Dirac comb $\nu_m = \text{III}(\cdot - x_m)$ for any $x_m \in \mathbb{R}$, which yields $\nu_m(f) = f(x_m)$, i.e., the spatial sampling operator. These are used in Chapter 10.*

Example 3.5. *Other examples of weak*-continuous functionals are the sinusoids $\nu_m = \cos(k \cdot) \in \mathcal{S}(\mathbb{T}) \subset \mathcal{C}_L(\mathbb{T})$ and $\nu_m = \sin(k \cdot)$ with $k \in \mathbb{N}$, which lead to $\nu_m(f) = \text{Re}(\widehat{f}[k])$ and $\nu_m(f) = \text{Im}(\widehat{f}[k])$, respectively, i.e., Fourier-sampling operators. These are used in Chapter 9.*

Note that the bound $K \leq M + 1$ in Theorem 3.4 is higher than that of the nonperiodic case $K \leq M - N_0$ in Theorem 3.3. However, the constraint (3.5) on the amplitudes a_k reduces the number of degrees of freedom for s^* . More precisely, there are $2M + 2$ degrees of freedom in the periodic case (3.34) ($M + 1$ for the x_k , M for the a_k and 1 for $\widehat{s^*}[0]$), compared to $2M - N_0$ in the nonperiodic case (3.33) ($M - N_0$ for the a_k and the x_k , and N_0 for $p \in \mathcal{N}_L$). These numbers remain valid for the more general operators L presented in [14] and [22]. Despite these differences, the role of the TV norm (to promote sparsity) and of L (to specify a transform domain for this sparsity) are analogous in both settings.

Part II Exact Discretization of TV-Based Inverse Problems on a Grid

In this part, we propose exact discretization methods for continuous-domain inverse problems with generalized total-variation (gTV) regularization related to

$$\arg \min_{f \in \mathcal{M}_L(\mathbb{R})} (E(\nu(f), \mathbf{y}) + \lambda \|L\{f\}\|_{\mathcal{M}}) \quad (\text{II.1})$$

(see Part I for a definition of the notations). We discretize such problems using B-spline bases on a grid. The choice of these bases are suggested by *representer theorems*, more specifically Theorem 3.3 for Chapter 4 and our own for Chapters 5 and 6. The key feature of our approach is that contrary to purely discrete settings, our discretization is *exact* in the continuous domain, in that the forward and regularization operators ν and L in Problem (II.1) are applied in the continuous domain without any discretization error.

Our discretization methods impose the locations of the spline knots on a predefined grid. They are thus unable to locate the jumps at super-resolution, contrary to the techniques developed for the sparse spikes deconvolution problem described in Section 1.2.3. Although super-resolution is a sensible objective when the reconstructed signal consists of Dirac impulses, in our spline-based framework, finding the exact locations of the jumps is less critical since the reconstructed signals are smoother. Hence, we take the stance of using a grid and B-splines, at the expense of (arbitrarily small) localization errors on the jumps. Moreover, this choice has the following two advantages:

- the use of a grid leads to standard convex finite-dimensional problems with ℓ_1 regularization for which off-the-shelf efficient solvers can be used. By contrast, sparse spikes super-resolution methods lead to nonconvex problems which are computationally much more challenging, all the more so with an added regularization operator L [90];
- as explained in Section 3.1.4, the short support of B-spline basis functions leads to

well-conditioned optimization problems, unlike the Green's function basis [45].

Note that the reformulation of problems based on (II.1) on the grid as standard problems with ℓ_1 regularization requires a careful treatment of the boundaries in order for the discretization to remain exact.

Outline

In this part, we consider the following formulations:

- In Chapter 4, we consider problems of the form (II.1), which lead to sparse L-spline signals models as demonstrated in Theorem 3.3. We propose a multiresolution B-spline-based exact discretization method for this problem which is guaranteed to reach a desired sparse L-spline solution.
- In Chapter 5, we extend to composite signal models; more specifically, we consider *hybrid spline* models $s = s_1 + s_2$, where s_i is a sparse L_i -spline for $i \in \{1, 2\}$. We formulate a suitable continuous-domain inverse problem, for which we prove a representer theorem and apply our B-spline-based discretization method.
- In Chapter 6, we consider a different composite model, namely $s = s_1 + s_2$, where s_1 is a sparse L_1 -spline and s_2 is a smooth component. Again, we formulate the corresponding optimization problem, prove a representer theorem, and discretize using B-splines.

4 Single-Component Signal Models

This chapter is based on the following publication [114]:

T. Debarre, J. Fageot, H. Gupta, and M. Unser, “B-Spline-Based Exact Discretization of Continuous-Domain Inverse Problems with Generalized TV Regularization”, *IEEE Transactions on Information Theory*, vol. 65, no. 7, pp. 4457–4470, Jul. 2019.

4.1 Introduction

In this chapter, we consider inverse problems with generalized total-variation (gTV) regularization of the form (II.1) with a least-square data fidelity loss $E(\mathbf{y}, \boldsymbol{\nu}(f)) = \frac{1}{2} \|\boldsymbol{\nu}(f) - \mathbf{y}\|_2^2$, *i.e.*,

$$\mathcal{V} \triangleq \arg \min_{f \in \mathcal{M}_L(\mathbb{R})} \underbrace{\left(\frac{1}{2} \|\boldsymbol{\nu}(f) - \mathbf{y}\|_2^2 + \lambda \|L\{f\}\|_{\mathcal{M}} \right)}_{\mathcal{J}(f)}, \quad (4.1)$$

where $\mathcal{M}_L(\mathbb{R})$ is defined in (2.33), $\boldsymbol{\nu} : \mathcal{M}_L(\mathbb{R}) \rightarrow \mathbb{R}^M$, $\mathbf{y} \in \mathbb{R}^M$, and $\lambda > 0$. The regularization operator L is the N_0 th-order derivative¹ $L = D^{N_0}$ where $N_0 \geq 1$. For more background on Problem 4.1, we refer to Chapter 3. Theorem 3.3 states that Problem (4.1) has extreme-point solutions of the form

$$s^*(x) = \sum_{k=1}^K a_k \rho_L(x - x_k) + \sum_{n=1}^{N_0} b_n p_n(x), \quad (4.2)$$

where $K \leq M - N_0$, $a_k, x_k \in \mathbb{R}$, ρ_L is the Green’s function of L defined in (2.25) and $\{p_n\}_{n=1}^{N_0}$ form a basis of the null space \mathcal{N}_L of L defined in (2.24).

¹As in most chapters of the thesis, all the results and algorithms from this chapter can be extended to more general operators L , *e.g.*, rational operators [102].

4.1.1 Green's Function Discretization

The form of the solutions (4.2) suggests a natural basis to discretize Problem (4.1). As demonstrated in [45], by using basis functions $\{p_n\}_{n=1}^{N_0}$ and $\{\rho_L(\cdot - x_n)\}_{n=1}^N$ where the knots x_n lie on a uniform finite grid, we get the following discrete optimization problem:

$$\min_{(\mathbf{a}, \mathbf{b}) \in \mathbb{R}^{N+N_0}} \left(\frac{1}{2} \|\mathbf{H}_{\rho_L} \mathbf{a} + \mathbf{H}_p \mathbf{b} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{a}\|_1 \right) \quad (4.3)$$

with system matrices $\mathbf{H}_{\rho_L} \in \mathbb{R}^{M \times N}$ and $\mathbf{H}_p \in \mathbb{R}^{M \times N_0}$. This problem is a standard finite-dimensional problem with ℓ_1 regularization and can thus be solved using off-the-shelf convex optimization algorithms such as FISTA [18]. The major asset of this approach is that the discrete problem is *exactly* equivalent to the underlying continuous problem restricted to the search space spanned by the basis functions. By making the grid finer, this search space includes functions arbitrarily close to solutions (4.2) of the full continuous-domain problem. However, the Green's function usually has infinite support (e.g., $\rho_D = \mathbb{1}_+$), which makes the Green's function basis ill-suited for practical problems. In particular, Problem (4.3) is severely ill-conditioned, making the convergence of solvers slow and potentially numerically unstable.

4.1.2 Our Approach: B-Spline Discretization

We therefore propose to improve this discretization method by using an equivalent dictionary basis consisting of shifted B-splines, *i.e.*, $\{\beta_{L,h}(\cdot - x_n)\}_{n=1}^N$ where the knots x_n lie on a uniform finite grid with knot spacing $h > 0$, and $\beta_{L,h}$ is the scaled B-spline of L defined in Section 3.1.5. This basis leads to the following discrete optimization problem:

$$\min_{\mathbf{c} \in \mathbb{R}^N} \left(\frac{1}{2} \|\mathbf{H}_{\beta_L} \mathbf{c} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{Lc}\|_1 \right) \quad (4.4)$$

where \mathbf{L} is a finite difference-like regularization matrix. This problem is of the same form as standard discrete compressed sensing (CS)-type problems, with the advantage that the chosen basis is matched to the form of the continuous-domain solution (4.2). Moreover, Problem (4.4) shares the exact discretization property of the Green's function basis, since—setting aside boundary issues—both bases are equivalent. However, as explained in Section 3.1.5, the finite support of the B-splines makes this basis better suited for practical applications, and it induces well-conditioned problems. This leads to a rapid convergence of solvers for Problem (4.4), a prediction which will be confirmed by our experimental results.

4.1.3 Outline and Contributions

In terms of contribution, the work in this chapter extends [45] and parts of our experimental pipeline are adapted from this work. However, the use of B-splines is a critical improvement which leads to the following contributions:

- In Section 4.2, we introduce the search space consisting of L-splines with knots on a uniform grid. We have shown in Proposition 3.3 that it has an equivalent formulation in the B-spline basis;
- In Section 4.3, we show that the corresponding continuous-domain inverse problem can be recast as a finite-dimensional problem of the form (4.4) in an *exact* way. We also demonstrate that the algorithm introduced in [45] can be adapted to our framework, and that it yields sparse solutions (with lower sparsity than in [45]);
- In Section 4.4, we prove that the optimal cost of the discrete problem converges to that of the continuous problem when the grid size goes to zero. We leverage this result to devise a multiresolution algorithm which refines the grid until the desired level of accuracy is met (termination criterion);
- In Section 4.5, we demonstrate experimentally the effectiveness of our algorithm using different measurement types (ideal sampling in the spatial and Fourier domains). We also show that it compares favorably with standard purely discrete methods.

4.2 Specification of the Search Space

In order to discretize Problem (4.1), we restrict the search space to the space $\mathcal{M}_{L,h}(\mathbb{R})$ of uniform L-splines with grid size $h > 0$ introduced in (3.3). This choice is obviously guided by Theorem 3.3, which states that Problem (4.1) has L-spline solutions, although their knots are not on a uniform grid. This search space therefore contains functions which are close approximations of a solution when the grid size h is small. A more explicit mathematical justification is given in Section 4.4.1.

The key property of the search space $\mathcal{M}_{L,h}(\mathbb{R})$ is that it has an alternative representation in the B-spline basis (Proposition 3.3) given by

$$\mathcal{M}_{L,h}(\mathbb{R}) \triangleq \left\{ s = p + \sum_{k \in \mathbb{Z}} a[k] \rho_L(\cdot - kh) : a \in \ell_1(\mathbb{Z}), p \in \mathcal{N}_L \right\} \quad (4.5)$$

$$= \left\{ s = \sum_{k \in \mathbb{Z}} c[k] \beta_{L,h}(\cdot - kh) : c \in \ell_{1,L}(\mathbb{Z}) \right\} \quad (4.6)$$

where $\beta_{L,h}$ is defined as in (3.13) and

$$\ell_{1,L}(\mathbb{Z}) \triangleq \left\{ c \in \mathcal{S}'(\mathbb{Z}) : (d_L * c) \in \ell_1(\mathbb{Z}) \right\} \quad (4.7)$$

where d_L is defined in (3.11) by its z transform $D_L(z) \triangleq (1 - z^{-1})^{N_0}$; we refer to Sections 3.1.4 and 3.1.5 for more background on scaled polynomial B-splines.

The major contribution of this chapter is the use of the B-spline representation (4.6) of $\mathcal{M}_{L,h}(\mathbb{R})$ throughout the discretization process, whereas [45] uses the Green's function representation (4.5). The point of doing so is that contrary to the Green's function, B-splines have finite support, which leads to Problem (4.4) being well-conditioned (see the discussion after Proposition 3.3).

The conditioning of an inverse problem is a measure of its numerical stability: a problem is well-conditioned if a small perturbation of the signal coefficients leads to a small perturbation of its measurements. When the basis functions have limited support as in the B-spline case, it is clear that a slight disturbance of the basis coefficients does not change the signal - and thus its measurements - significantly. However, this is not the case for the Green's function basis, since not only do the basis functions have infinite support, but they are nonvanishing or even increasing as one moves away from the center, since $\rho_L(x) = x_+^{N_0-1}$ for $L = D^{N_0}$. Therefore, a small perturbation of a basis coefficient greatly affects the reconstructed signal everywhere. The measurements are thus greatly impacted and the problem is severely ill-conditioned [45, 115]. This intuition is confirmed in practice: we observe that in identical settings (*i.e.*, same regularization operator L , measurement operator ν , regularization parameter λ and grid size h), the relevant condition number is systematically greater in Problem (4.3) than in Problem (4.4). For example, in the experiment shown in Fig. 4.3, the condition number of the matrix to be inverted is $\text{cond}(\mathbf{H}_{\beta_L}^T \mathbf{H}_{\beta_L} + \lambda \mathbf{L}^T \mathbf{L}) = 5.7 \times 10^4$ using the B-spline formulation, compared to $\text{cond}(\mathbf{H}_{\rho_L}^T \mathbf{H}_{\rho_L} + \lambda \mathbf{I}) = 1.5 \times 10^{12}$ in the Green's function case. This difference of conditioning largely justifies the use of the B-spline representation of $\mathcal{M}_{L,h}(\mathbb{R})$ rather than its Green's function representation.

4.3 Exact Discretization in the Search Space

4.3.1 Discrete Problem Formulation

Let $h > 0$; in order to discretize Problem (4.1) in $\mathcal{M}_{L,h}(\mathbb{R})$, we use the B-spline representation of $s \in \mathcal{M}_{L,h}(\mathbb{R})$ given in (4.6)

$$s(x) = \sum_{k \in \mathbb{Z}} c[k] \beta_{L,h}(x - kh) \quad (4.8)$$

where $c \in \ell_{1,L}(\mathbb{Z})$. Using (3.15), the sparsity of s in the Green's function basis is given by $\|d_L * c\|_0$, where $\|\cdot\|_0$ is the ℓ_0 "norm" which counts the number of nonzero entries of a sequence or vector. When we feed (4.8) into the continuous-domain Problem (4.1), using (3.18) we get the following discretized optimization problem

$$\mathcal{V}_d \triangleq \arg \min_{c \in \ell_{1,L}(\mathbb{Z})} \left(\underbrace{\frac{1}{2} \left\| \sum_{k \in \mathbb{Z}} c[k] \boldsymbol{\nu}(\beta_{L,h}(\cdot - hk)) - \mathbf{y} \right\|_2^2}_{\mathcal{J}_d(c)} + \frac{\lambda}{h^{N_0-1}} \|d_L * c\|_1 \right). \quad (4.9)$$

By adapting Lemma 20 in [20], it can be shown that \mathcal{V}_d is a nonempty weak*-compact subset of $\ell_{1,L}(\mathbb{Z})$. Note that Problem (4.9) is *exactly* equivalent to the continuous-domain Problem (4.1) restricted to the search space $\mathcal{M}_{L,h}(\mathbb{R})$. This is the key feature of our formulation: the standard approach to discretize an inverse problem is to use an approximate discrete forward model as a surrogate for a continuous model, which leads to discretization errors. This is not the case of our method, in which the discrete forward model is equal to the continuous one; the former is simply restricted to the native space $\mathcal{M}_{L,h}(\mathbb{R})$. To the best of our knowledge, aside from [45] of which this chapter is an extension, no other work in the literature discretizes nonquadratic continuous-domain problems *exactly* by using a dictionary.

4.3.2 Finite Problem

In real-world applications, the signal of interest usually has a given finite support, which we can assume to be $I_T = [0, T]$ without loss of generality. Hence, in all that follows, we can assume that the measurement functionals ν_m , be they ordinary functions or distributions such as Dirac impulses, are supported in I_T . This assumption is completely inconsequential for signals supported in I_T , but it is necessary to express (4.9) as a finite-dimensional problem. In this case, only a finite number of B-spline coefficients affect the data fidelity term in (4.9); we denote by $I = \{i_{\min}, \dots, i_{\max}\} \subset \mathbb{Z}$ the set of their indices and $N \triangleq \#I$. Assuming that $T/h \in \mathbb{N}$, we have $i_{\min} \triangleq -N_0 + 1$, $i_{\max} \triangleq T/h - 1$ and thus $N = T/h + N_0 - 1$. To make Problem (4.9) finite, we optimize over the N B-spline coefficients in I , which are denoted by $\mathbf{c} \in \mathbb{R}^N$. By imposing natural boundary conditions for the regularization term, we get the following discrete finite-dimensional problem:

$$\mathcal{V}_f \triangleq \arg \min_{\mathbf{c} \in \mathbb{R}^N} \left(\underbrace{\frac{1}{2} \|\mathbf{H}\mathbf{c} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{L}\mathbf{c}\|_1}_{\mathcal{J}_f(\mathbf{c})} \right) \quad (4.10)$$

where the system matrix $\mathbf{H} \in \mathbb{R}^{M \times N}$ is

$$\mathbf{H} \triangleq (\mathbf{h}_{i_{\min}}, \dots, \mathbf{h}_{i_{\max}}) : \quad \mathbf{h}_k \triangleq \boldsymbol{\nu}(\beta_{L,h}(\cdot - hk)) \quad (4.11)$$

and the Toeplitz-like regularization matrix $\mathbf{L} \in \mathbb{R}^{(N-N_0) \times N}$ is a finite section of the infinite-dimensional regularization matrix in Problem (4.9), *i.e.*,

$$\mathbf{L} \triangleq \frac{1}{h^{N_0-1}} \begin{pmatrix} d_L[N_0] & \cdots & d_L[0] & 0 & \cdots & 0 \\ 0 & \ddots & & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & & \ddots & 0 \\ 0 & \cdots & 0 & d_L[N_0] & \cdots & d_L[0] \end{pmatrix}. \quad (4.12)$$

Despite the seemingly arbitrary boundary conditions on the regularization term (we are representing a convolution involving an infinite sequence as a finite matrix multiplication), remarkably, thanks to the finite support assumption, Problem (4.10) is exactly equivalent to the infinite Problem (4.9):

Proposition 4.1. *Under the assumptions of Theorem 3.3 and assuming that the functionals ν_m are supported in I_T :*

- We have $\ker \mathbf{H} \cap \ker \mathbf{L} = \{\mathbf{0}\}$, and the solution set \mathcal{V}_f of Problem (4.10) is a nonempty compact convex set;
- Problems (4.9) and (4.10) are equivalent, and there exists a natural bijection between their solution sets which maps any sequence $c^* \in \mathcal{V}_d \subset \ell_{1,L}(\mathbb{Z})$ to a vector $\mathbf{c}^* \in \mathcal{V}_f \subset \mathbb{R}^N$ such that $c^*|_I = (c^*[i_{\min}], \dots, c^*[i_{\max}]) = \mathbf{c}^*$.

Proof. The second item of Proposition 4.1 entails the existence of bijective linear map $\theta : \mathcal{V}_f \rightarrow \mathcal{V}_d$ such that $\theta(\mathbf{c}^*)|_I = \mathbf{c}^*$ for any $\mathbf{c}^* \in \mathcal{V}_f$. In order to construct this mapping, we rely on the following lemma:

Lemma 4.1. *Assume that the functionals ν_m ($1 \leq m \leq M$) are supported in I_T . Then solutions $c^* \in \mathcal{V}_d$ of Problem (4.9) are uniquely determined by their N coefficients $c^*|_I$.*

Proof. Let $c^* \in \mathcal{V}_d$ be a solution of the discrete Problem (4.9). Consider a sequence c such that $c|_I = c^*|_I$ and whose remaining coefficients are free. The latter do not affect the data fidelity term due to the finite support assumption on the ν_m . When $N > N_0$, $c[i_{\max} + 1]$ can be uniquely chosen such that $(d_L * c)[i_{\max} + 1] = \sum_{k=0}^{N_0} d_L[k]c[i_{\max} - k + 1] = 0$. Similarly, all coefficients $c[k]$ for $k > i_{\max}$ can be uniquely determined recursively to nullify $(d_L * c)[k]$ as a linear combination of the $N_0 - 1$ previous coefficients. The same can be done for coefficients $c[k]$ with $k \leq 0$, using this time the $N_0 - 1$ following coefficients of c . By construction, this sequence c yields a regularization cost smaller or equal to that

of c^* , and since both yield the same data fidelity cost, we have $\mathcal{J}_d(c) \leq \mathcal{J}_d(c^*)$. Since c^* is a solution of (4.9) and the construction of c is unique, we necessarily have $c = c^*$. \square

The proof of Lemma 4.1 details the construction of an injective linear map $\tilde{\theta} : \mathbb{R}^N \rightarrow \ell_{1,L}(\mathbb{Z})$ such that for any $\mathbf{c} \in \mathbb{R}^N$, $\tilde{\theta}(\mathbf{c})|_I = \mathbf{c}$. Let $\mathbf{c} \in \mathbb{R}^N$, and consider the corresponding sequence $\tilde{\theta}(\mathbf{c}) \in \ell_{1,L}(\mathbb{Z})$. Following the proof of Lemma 4.1, $\mathcal{J}_d(\tilde{\theta}(\mathbf{c}))$ can be computed using only the N coefficients \mathbf{c} . Indeed, all other coefficients $(\tilde{\theta}(\mathbf{c})[k])_{k \notin I}$ do not affect the data fidelity term and cancel out all the regularization terms which they affect. This implies that $\mathcal{J}_d(\tilde{\theta}(\mathbf{c})) = \mathcal{J}_f(\mathbf{c}) = \frac{1}{2} \|\mathbf{H}\mathbf{c} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{L}\mathbf{c}\|_1$, where \mathbf{H} and \mathbf{L} are defined as in (4.11) and (4.12) respectively. Since by Lemma 4.1, $\mathcal{V}_d \subset \tilde{\theta}(\mathbb{R}^N)$, problems (4.9) and (4.10) are equivalent in the sense that $\tilde{\theta}(\mathcal{V}_f) = \mathcal{V}_d$, and the restriction $\theta = \tilde{\theta}|_{\mathcal{V}_f} : \mathcal{V}_f \rightarrow \mathcal{V}_d$ is a bijective linear map.

Concerning the first item of Proposition 4.1, let $\mathbf{c} \in \ker \mathbf{H} \cap \ker \mathbf{L}$, the corresponding signal s verifies $s = \sum_{k \in \mathbb{Z}} \tilde{\theta}(\mathbf{c})[k] \beta_{L,h}(\cdot - kh) \in \mathcal{N}_{\mathcal{V}} \cap \mathcal{N}_L = \{0\}$ (well-posedness assumption in Theorem 3.3), which implies that $\mathbf{c} = 0$. Hence, $\ker \mathbf{H} \cap \ker \mathbf{L} = \{\mathbf{0}\}$, which implies that Problem (4.10) is well-posed and thus that its solution set \mathcal{V}_f is a nonempty compact set. The latter is also convex due to the convexity of the cost function \mathcal{J}_f . \square

We refer to Problem (4.10) as being *well-posed* due to the property $\ker \mathbf{H} \cap \ker \mathbf{L} = \{\mathbf{0}\}$, which implies that its solution set \mathcal{V}_f is bounded, but not necessarily unique. Proposition 4.1 demonstrates that the finite Problem (4.10) is equivalent to the continuous-domain Problem (4.1) restricted to the search space $\mathcal{M}_{L,h}(\mathbb{R})$. This is quite a remarkable outcome: we are able to solve an infinite continuous-domain problem in an exact way as a standard discrete inverse problem with ℓ_1 regularization. Once again, this is to the best of our knowledge a novelty, which is this time not present in [45].

Despite the proven equivalence between problems (4.9) and (4.10), one might wonder how to proceed in practice to reconstruct the underlying continuous-domain signal s once a solution $\mathbf{c}^* \in \mathcal{V}_f$ is reached. By Proposition 4.1, there exists a unique sequence $c^* \in \mathcal{V}_d$ such $c^*|_I = \mathbf{c}^*$. The following observations, which are direct consequences of the proof of Proposition 4.1, can be made concerning s :

- The continuous-domain reconstructed signal is $s = \sum_{k \in \mathbb{Z}} c^*[k] \beta_{L,h}(\cdot - kh)$;
- The N B-spline coefficients \mathbf{c}^* are sufficient to reconstruct s *exactly* in the interval of interest I_T ;
- The sparsity of s in the Green's function basis is given by $\|\mathbf{L}\mathbf{c}^*\|_0$.

These observations indicate that all the relevant information concerning the reconstructed signal s (*i.e.*, its expression in I_T and sparsity) is directly encoded in the vector $\mathbf{c}^* \in \mathcal{V}_f$

of Problem (4.10). Hence, computing the corresponding infinite sequence $\mathbf{c}^* \in \mathcal{V}_d$ is unnecessary. In practice, the infinite Problem (4.9) can thus be altogether forsaken in favor of the computationally feasible finite Problem (4.10).

4.3.3 Reaching a Sparse Solution

In this section, we study the so-called penalized basis pursuit (PBP) problem formulated in (4.10)

$$\mathcal{V}_f \triangleq \arg \min_{\mathbf{c} \in \mathbb{R}^N} \left(\frac{1}{2} \|\mathbf{H}\mathbf{c} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{L}\mathbf{c}\|_1 \right) \quad (4.13)$$

where the system and regularization matrices \mathbf{H} and \mathbf{L} are defined in (4.11) and (4.12) respectively. This problem is close to typical CS problems (\mathbf{L} is a finite-difference-like regularization matrix), which have been studied at length in the literature [11, 20] and are known to yield sparse solutions in a certain basis. The specificity of this problem lies in the fact that \mathbf{L} is not invertible. However, Theorem 3.3 strongly suggests that (4.10) has sparse solutions, since it is a discretized version of the continuous-domain problem. This instinct is confirmed by the following representer theorem.

Theorem 4.1 (Discrete Representer Theorem). *Let $0 \leq N_0 \leq M < N$, $\mathbf{H} \in \mathbb{R}^{M \times N}$ and $\mathbf{L} \in \mathbb{R}^{(N-N_0) \times N}$ such that $\ker \mathbf{H} \cap \ker \mathbf{L} = \{\mathbf{0}\}$ and \mathbf{L} is of full rank, i.e., $\text{ran } \mathbf{L} = N - N_0$. Then the solution set \mathcal{V}_f of the optimization problem*

$$\mathcal{V}_f \triangleq \arg \min_{\mathbf{c} \in \mathbb{R}^N} \left(\frac{1}{2} \|\mathbf{H}\mathbf{c} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{L}\mathbf{c}\|_1 \right) \quad (4.14)$$

is a compact convex set whose extreme points (Definition 3.5) \mathbf{c}^* verify $\|\mathbf{L}\mathbf{c}^*\|_0 \leq M - N_0$.

Proof. Let $\mathcal{J}_f : \mathbf{c} \mapsto \frac{1}{2} \|\mathbf{H}\mathbf{c} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{L}\mathbf{c}\|_1$. Since \mathcal{J}_f is continuous and coercive due to the well-posedness assumption $\ker \mathbf{H} \cap \ker \mathbf{L} = \{\mathbf{0}\}$, \mathcal{V}_f is a nonempty, closed compact set. Therefore, by the Krein-Milman theorem (Theorem 3.2), it is the closed convex hull of its extreme points.

Let \mathbf{c}^* be an extreme point of \mathcal{V}_f . Assume by contradiction that $\mathbf{L}\mathbf{c}^*$ has sparsity $K > M - N_0$, i.e., $\mathbf{L}\mathbf{c}^* = \sum_{k=1}^K a_{n_k} \mathbf{e}_{n_k}$ where the $n_k \in \{1, \dots, N\}$ are distinct, $a_{n_k} \neq 0$ and $\{\mathbf{e}_i\}_{i=1}^{N-N_0}$ is the canonical basis of \mathbb{R}^{N-N_0} . Consider the vector space $T = \text{ran } \mathbf{L} \cap \text{span}\{\mathbf{e}_{n_k}\}_{k=1}^K$. To find a lower bound on the dimension of T , we use the relation

$$\begin{aligned} \dim(X \cap Y) &= \dim X + \dim Y - \dim(X + Y) \\ &\geq \dim X + \dim Y - P, \end{aligned} \quad (4.15)$$

where X and Y are vector subspaces of \mathbb{R}^P . Since the rank of \mathbf{L} is $N - N_0$, for $X = \text{ran } \mathbf{L}$ and $Y = \{\mathbf{e}_{n_k}\}_{k=1}^K$, (4.15) yields $R = \dim T \geq K > M - N_0$ (with $P = N - N_0$). Let

$\{\mathbf{t}_r\}_{r=1}^R$ be a basis of T . By definition of T , there exist vectors \mathbf{g}_r and coefficients $t_k^r \in \mathbb{R}$ such that $\mathbf{t}_r = \mathbf{L}\mathbf{g}_r = \sum_{k=1}^K t_k^r \mathbf{e}_{n_k}$.

Next, we define $\mathbf{y}_r = \mathbf{H}\mathbf{g}_r \in \mathbb{R}^M$ for all $r \in \{1, \dots, R\}$, and $\mathbf{z}_n = \mathbf{H}\mathbf{p}_n \in \mathbb{R}^M$ for all $n \in \{1, \dots, N_0\}$ where $\{\mathbf{p}_n\}_{n=1}^{N_0}$ is a basis of $\ker \mathbf{L}$. The collection of vectors $\{\mathbf{y}_1, \dots, \mathbf{y}_R, \mathbf{z}_1, \dots, \mathbf{z}_{N_0}\}$ has $R + N_0 \geq K + N_0 > M$ elements, and is thus linearly dependent. Therefore, there exist coefficients $\alpha_r, \beta_n \in \mathbb{R}$ such that $\sum_{r=1}^R \alpha_r \mathbf{y}_r + \sum_{n=1}^{N_0} \beta_n \mathbf{z}_n = \mathbf{0}$ and $(\boldsymbol{\alpha}, \boldsymbol{\beta}) \neq \mathbf{0}$. We then define $\mathbf{c}_0 = \sum_{r=1}^R \alpha_r \mathbf{g}_r + \sum_{n=1}^{N_0} \beta_n \mathbf{p}_n \in \mathbb{R}^N$, which is clearly in $\ker \mathbf{H}$. Assume by contradiction $\mathbf{L}\mathbf{c}_0 = \sum_{r=1}^R \alpha_r \mathbf{t}_r = \mathbf{0}$. We thus have $\mathbf{c}_0 \in \ker \mathbf{H} \cap \ker \mathbf{L} = \{\mathbf{0}\}$. Moreover, since the \mathbf{t}_r are linearly independent, we have $\boldsymbol{\alpha} = \mathbf{0}$, and thus $\mathbf{c}_0 = \sum_{n=1}^{N_0} \beta_n \mathbf{p}_n = \mathbf{0}$. Yet the \mathbf{p}_n are also linearly independent, which means that $\boldsymbol{\beta} = \mathbf{0}$, which contradicts $(\boldsymbol{\alpha}, \boldsymbol{\beta}) \neq \mathbf{0}$. Therefore, we have $\mathbf{L}\mathbf{c}_0 \neq \mathbf{0}$, which implies that $\mathbf{c}_0 \neq \mathbf{0}$.

Finally, we pick an $\epsilon > 0$ such that

$$\epsilon < \frac{\min_k a_{n_k}}{\max_k \left| \sum_{r=1}^R \alpha_r t_k^r \right|}. \quad (4.16)$$

Note that ϵ is well defined since for all k , $a_{n_k} > 0$ and $\sum_{r=1}^R \alpha_r t_k^r = 0$ for all k would imply that $\mathbf{L}\mathbf{c}_0 = \sum_{r=1}^R \alpha_r \mathbf{t}_r = \sum_{k=1}^K \left(\sum_{r=1}^R \alpha_r t_k^r \right) \mathbf{e}_{n_k} = \mathbf{0}$, which we have proved to be false. We can then compute:

$$\begin{aligned} \|\mathbf{L}(\mathbf{c}^* \pm \epsilon \mathbf{c}_0)\|_1 &= \left\| \sum_{k=1}^K \left(a_{n_k} \pm \epsilon \sum_{r=1}^R \alpha_r t_k^r \right) \mathbf{e}_{n_k} \right\|_1 \\ &= \sum_{k=1}^K \left(a_{n_k} \pm \epsilon \sum_{r=1}^R \alpha_r t_k^r \right) \\ &= \|\mathbf{L}\mathbf{c}^*\|_1 \pm \epsilon \sum_{k=1}^K \sum_{r=1}^R \alpha_r t_k^r \end{aligned}$$

since by definition of ϵ , $a_{n_k} \pm \epsilon \left(\sum_{r=1}^R \alpha_r t_k^r \right) > 0$ for all k . Notice that both vectors $(\mathbf{c}^* \pm \epsilon \mathbf{c}_0)$ yield the same data fidelity cost as \mathbf{c}^* in Problem (4.10): indeed, $\mathbf{H}(\mathbf{c}^* \pm \epsilon \mathbf{c}_0) = \mathbf{H}\mathbf{c}^*$ since $\mathbf{c}_0 \in \ker \mathbf{H}$. Therefore, if $\sum_{k=1}^K \sum_{r=1}^R \alpha_r t_k^r \neq 0$, then either $(\mathbf{c}^* + \epsilon \mathbf{c}_0)$ or $(\mathbf{c}^* - \epsilon \mathbf{c}_0)$ yields a cost strictly smaller than that of \mathbf{c}^* in Problem (4.10), which is impossible since \mathbf{c}^* is a solution of the latter. Consequently, $\sum_{k=1}^K \sum_{r=1}^R \alpha_r t_k^r = 0$ and so $(\mathbf{c}^* \pm \epsilon \mathbf{c}_0) \in \mathcal{V}_f$. Yet $\mathbf{c}^* = \frac{1}{2}(\mathbf{c}^* + \epsilon \mathbf{c}_0) + \frac{1}{2}(\mathbf{c}^* - \epsilon \mathbf{c}_0)$, and since $\epsilon \mathbf{c}_0 \neq \mathbf{0}$, \mathbf{c}^* is not an extreme point of \mathcal{V}_f , which contradicts our initial assumption. This proves the desired result $K \leq M - N_0$. \square

Theorem 4.1 is a generalization of Theorem 6 in [20], since it allows for more general regularization matrices (\mathbf{L} must be right-invertible). It is also similar to Theorem

2.4 in [87], but with a tighter bound on the sparsity, and with an elementary proof using only standard linear algebra. This result directly applies to Problem (4.10), since $\ker \mathbf{H} \cap \ker \mathbf{L} = \{\mathbf{0}\}$ by Proposition 4.1 and \mathbf{L} in (4.12) is of full rank. Remarkably, the bound on the sparsity $M - N_0$ is the same as for the continuous-domain Problem (4.18) (Theorem 3.3), which confirms the close connection between both problems. This is not the case in the Green's function formulation of [45], where the sparsity is bounded by M .

Although Theorem 4.1 guarantees that Problem (4.10) has sparse solutions, only the extreme points of \mathcal{V}_f are known to be sparse, and in general, \mathcal{V}_f is nonunique. Therefore, while a solution of Problem (4.10) can readily be reached using standard solvers such as ADMM (see Appendix A), there is no guarantee that this solution will be sparse. In fact, we will demonstrate experimentally later on that, as observed with FISTA in [45], ADMM often converges towards nonsparse solutions, *i.e.*, vectors \mathbf{c}^* such that $\|\mathbf{L}\mathbf{c}^*\|_0 > M - N_0$. To circumvent this issue, we use the following lemma, which is well known in the absence of a regularization matrix [21, Lemma 1]. As it turns out, the latter does not make it more challenging: an elementary proof is given for the sake of completeness.

Lemma 4.2. *Let $\mathbf{H} \in \mathbb{R}^{M \times N}$, $\mathbf{L} \in \mathbb{R}^{P \times N}$, $\mathbf{y} \in \mathbb{R}^M$ and $\lambda > 0$. We assume that the problem is well posed, *i.e.*, $\ker \mathbf{H} \cap \ker \mathbf{L} = \{\mathbf{0}\}$. Then the solution set \mathcal{V}_f of the PBP Problem (4.10) is a compact convex set which has a unique measurement vector $\mathbf{y}_\lambda \in \mathbb{R}^M$ such that $\forall \mathbf{c} \in \mathcal{V}_f$, $\mathbf{H}\mathbf{c} = \mathbf{y}_\lambda$. Moreover, for any two solutions $\mathbf{c}_1, \mathbf{c}_2 \in \mathcal{V}_f$, we have*

$$(\mathbf{L}\mathbf{c}_1)_i \times (\mathbf{L}\mathbf{c}_2)_i \geq 0 \quad \forall i \in \{1, \dots, P\}. \quad (4.17)$$

Proof. Due to the well-posedness assumption, the cost function \mathcal{J}_f is coercive and since it is continuous, \mathcal{V}_f is non-empty and bounded. Let $\mathbf{c}_1, \mathbf{c}_2 \in \mathcal{V}_f$ be two (possibly identical) solutions. We have $\mathcal{J}_f(\mathbf{c}_1) = \mathcal{J}_f(\mathbf{c}_2) \triangleq \mathcal{J}_f^0$, and for any $\alpha \in [0, 1]$, we define $\mathbf{c}_\alpha = \alpha\mathbf{c}_1 + (1 - \alpha)\mathbf{c}_2$. The convexity of \mathcal{J}_f yields $\mathcal{J}_f(\mathbf{c}_\alpha) \leq \alpha\mathcal{J}_f(\mathbf{c}_1) + (1 - \alpha)\mathcal{J}_f(\mathbf{c}_2) = \mathcal{J}_f^0$. Yet, since \mathcal{J}_f^0 is the minimum of the cost function \mathcal{J}_f , we have $\mathcal{J}_f(\mathbf{c}_\alpha) = \mathcal{J}_f^0$, which implies that $\mathbf{c}_\alpha \in \mathcal{V}_f$ and thus that \mathcal{V}_f is a convex set. Another implication is that the convexity inequality is in fact an equality. For the data fidelity term, the strict convexity of the squared ℓ_2 norm implies that $\mathbf{H}\mathbf{c}_1 - \mathbf{y} = \mathbf{H}\mathbf{c}_2 - \mathbf{y} \Leftrightarrow \mathbf{H}\mathbf{c}_1 = \mathbf{H}\mathbf{c}_2 = \mathbf{y}_\lambda$.

The second property (4.17) results from the case of equality in the triangular inequality of the ℓ_1 norm: we have $\|\mathbf{L}\mathbf{c}_\alpha\|_1 = \alpha\|\mathbf{L}\mathbf{c}_1\|_1 + (1 - \alpha)\|\mathbf{L}\mathbf{c}_2\|_1$. Each coordinate can be treated separately, yielding

$$\begin{aligned} |(\mathbf{L}\mathbf{c}_\alpha)_i| &= \alpha|(\mathbf{L}\mathbf{c}_1)_i| + (1 - \alpha)|(\mathbf{L}\mathbf{c}_2)_i| \\ \Leftrightarrow (\mathbf{L}\mathbf{c}_1)_i \times (\mathbf{L}\mathbf{c}_2)_i &\geq 0 \quad \forall i \in \{1, \dots, P\}. \end{aligned}$$

□

Note that in the case of Problem (4.10), as shown in Proposition 4.1, we have $\ker \mathbf{H} \cap$

$\ker \mathbf{L} = \{\mathbf{0}\}$, which implies that Lemma 4.2 applies to our problem for $P = N - N_0$. Lemma 4.2 provides an indirect way of reaching an extreme point of \mathcal{V}_f : any solution $\mathbf{c}^* \in \mathcal{V}_f$ has a fixed measurement $\mathbf{H}\mathbf{c}^* = \mathbf{y}_\lambda$. Therefore, Problem (4.10) can be recast as a constrained optimization problem

$$\arg \min_{\mathbf{c} \in \mathbb{R}^N} \|\mathbf{L}\mathbf{c}\|_1 \quad s.t. \quad \mathbf{H}\mathbf{c} = \mathbf{y}_\lambda, \quad (4.18)$$

which clearly has the same solution set \mathcal{V}_f , since the constraint is satisfied for any $\mathbf{c} \in \mathcal{V}_f$. This constrained problem can, in turn, be recast as a linear program by introducing the slack variable $\mathbf{u} \in \mathbb{R}^{N-N_0}$:

$$\mathcal{V}_f^{\text{LP}} \triangleq \arg \min_{(\mathbf{c}, \mathbf{u}) \in \mathbb{R}^{2N-N_0}} \left(\sum_{i=1}^{N-N_0} u_i \right) \quad s.t. \quad \mathbf{u} + \mathbf{L}\mathbf{c} \geq \mathbf{0}; \quad \mathbf{u} - \mathbf{L}\mathbf{c} \geq \mathbf{0}; \quad \mathbf{H}\mathbf{c} = \mathbf{y}_\lambda, \quad (4.19)$$

where (\mathbf{c}, \mathbf{u}) is the concatenation of the vectors $\mathbf{c} \in \mathbb{R}^N$ and $\mathbf{u} \in \mathbb{R}^{N-N_0}$. The following proposition characterizes its solution set $\mathcal{V}_f^{\text{LP}}$ in terms of \mathcal{V}_f :

Proposition 4.2. $\mathcal{V}_f^{\text{LP}}$ is a compact convex set which has extreme points $(\mathbf{c}^*, \mathbf{u}^*)$ such that \mathbf{c}^* is an extreme point of \mathcal{V}_f .

Proof. We prove the following statement, which is clearly stronger than that of Proposition 4.2:

$\mathcal{V}_f^{\text{LP}}$ is a compact convex set, and \mathcal{V}_f and $\mathcal{V}_f^{\text{LP}}$ have corresponding extreme points through the one-to-one mapping $\phi : \mathcal{V}_f \rightarrow \mathcal{V}_f^{\text{LP}}$ defined by $\phi(\mathbf{c}) \triangleq (\mathbf{c}, |\mathbf{L}\mathbf{c}|)$ and its inverse $\phi^{-1} : \mathcal{V}_f^{\text{LP}} \rightarrow \mathcal{V}_f$ defined by $\phi^{-1}((\mathbf{c}, \mathbf{u})) \triangleq \mathbf{c}$.

Let us first observe that $\mathcal{V}_f^{\text{LP}}$ is of the form $\mathcal{V}_f^{\text{LP}} = \{(\mathbf{c}, |\mathbf{L}\mathbf{c}|) \in \mathbb{R}^{2N-N_0}, \mathbf{c} \in \mathbb{R}^N\}$, where $|\mathbf{x}|$ is the vector of component-wise absolute values of \mathbf{x} . This implies that $\{\mathbf{c} \in \mathbb{R}^N, (\mathbf{c}, |\mathbf{L}\mathbf{c}|) \in \mathcal{V}_f^{\text{LP}}\}$ is the solution set of the constrained optimization Problem (4.18), which is equal to \mathcal{V}_f . Therefore, we have proved that $\mathcal{V}_f = \{\mathbf{c} \in \mathbb{R}^N, (\mathbf{c}, |\mathbf{L}\mathbf{c}|) \in \mathcal{V}_f^{\text{LP}}\}$, and thus that $\mathcal{V}_f^{\text{LP}} = \{(\mathbf{c}, |\mathbf{L}\mathbf{c}|) \in \mathbb{R}^{2N-N_0}, \mathbf{c} \in \mathcal{V}_f\}$. Hence, $\mathcal{V}_f^{\text{LP}}$ is a nonempty compact set as the continuous image of the nonempty compact set \mathcal{V}_f through ϕ . Moreover, $\mathcal{V}_f^{\text{LP}}$ is convex as the solution set of a linear program.

Next, (4.17) in Lemma 4.2 implies that $\phi : \mathcal{V}_f \rightarrow \mathcal{V}_f^{\text{LP}}$ is a linear map. Moreover, ϕ is invertible and its inverse $\phi^{-1} : \mathcal{V}_f^{\text{LP}} \rightarrow \mathcal{V}_f$ is also linear. The desired result immediately follows. \square

Proposition 4.2 allows us to apply the well-known simplex or dual-simplex algorithms [116, 117] to the linear program (4.19). These algorithms are known to converge to an extreme point $(\mathbf{c}^*, \mathbf{u}^*)$ of the solution set $\mathcal{V}_f^{\text{LP}}$. Since \mathbf{c}^* is an extreme point of \mathcal{V}_f , Theorem 4.1

then ensures that it is a sparse solution we are looking for. However, to run this linear program, \mathbf{y}_λ needs to be known: hence, we must find a solution (though not necessarily an extreme point) of the PBP Problem (4.10) beforehand using ADMM or any other suitable algorithm. This solution $\mathbf{c}_{\text{ADMM}} \in \mathcal{V}_f$ is then used to compute $\mathbf{y}_\lambda = \mathbf{H}\mathbf{c}_{\text{ADMM}}$, which is used in turn to run the simplex algorithm. This procedure is adapted from [45], in which the same idea is used in the Green's function basis.

4.4 Refining the Grid

In the previous sections, we have established an experimental pipeline to solve the continuous-domain problem in $\mathcal{M}_{L,h}(\mathbb{R})$ for a fixed grid size h . We now study the behavior of the solutions when the grid size h decreases, and how they relate to solutions of the full continuous-domain Problem (4.1).

4.4.1 Convergence of the Cost Function

We place ourselves in the conditions of Theorem 3.3, which states that there exists at least one solution to the continuous-domain Problem (4.1) of the form $s(x) = p(x) + \sum_{k=1}^K a_k \rho_L(x - x_k)$ with $p \in \mathcal{N}_L$. This solution does not a priori have knots on a uniform grid, and is thus not included in $\mathcal{M}_{L,h}(\mathbb{R})$. However, by picking h sufficiently small, it can be approached arbitrarily closely by

$$s_h(x) = p(x) + \sum_{k=1}^K a_k \rho_L(x - x_k^h) \in \mathcal{M}_{L,h}(\mathbb{R}), \quad (4.20)$$

where $x_k^h \in h\mathbb{Z}$ converges to x_k .

The following lemma shows that s_h is indeed a good approximation of s in terms of cost:

Lemma 4.3. *Let all the hypotheses of Theorem 3.3 be met. Then, there exists a family of functions of the form $s_h = p + \sum_{k=1}^K a_k \rho_L(\cdot - x_k^h) \in \mathcal{M}_{L,h}(\mathbb{R})$ where $p \in \mathcal{N}_L$, $K \leq M - N_0$ and $x_k^h \in h\mathbb{Z}$ for any $h > 0$ such that*

$$\lim_{h \rightarrow 0} \mathcal{J}(s_h) = \min_{f \in \mathcal{M}_L(\mathbb{R})} \mathcal{J}(f) \triangleq \mathcal{J}^0. \quad (4.21)$$

Proof. We first recap some useful properties of $\mathcal{M}_L(\mathbb{R})$ given in Section 2.4. Let (ϕ, \mathbf{p}) be an admissible biorthogonal system for \mathcal{N}_L in the sense of Proposition 2.2. By (2.43) in Theorem 2.2, any element $f \in \mathcal{M}_L(\mathbb{R})$ has a unique decomposition as $f = L_\phi^{-1}\{w\} + q$ where $w \in \mathcal{M}(\mathbb{R})$, $q \in \mathcal{N}_L$, and the operator L_ϕ^{-1} is defined in Theorem 2.1.

By Theorem 3.3, there exists a solution s to Problem (4.1) such that $L\{s\} = \sum_{k=1}^K a_k \delta(\cdot - x_k)$ where $K \leq M - N_0$ and all x_k are pairwise distinct. Therefore, s can be represented

as $s = \mathbf{L}_\phi^{-1}\{w\} + p$ where $w = \mathbf{L}\{s\} = \sum_{k=1}^K a_k \delta(\cdot - x_k)$ and $p \in \mathcal{N}_L$. We thus have $\mathcal{J}(s) = \mathcal{J}_0 = \|\boldsymbol{\nu}(s) - \mathbf{y}\|_2^2 + \lambda \|\mathbf{a}\|_1$. For a given $h > 0$, let $x_k^h \in h\mathbb{Z}$ be the grid element closest to x_k for all $k \in \{1, \dots, K\}$, *i.e.*, $|x_k - x_k^h| \leq \frac{h}{2}$. For small enough values of h , all x_k^h are pairwise distinct; we place ourselves in this configuration. We then define $s_h = p + \sum_{k=1}^K a_k \mathbf{L}_\phi^{-1}\{\delta(\cdot - x_k^h)\}$, where $s_h \in \mathcal{M}_{L,h}(\mathbb{R})$ since it can also be written $s_h = q + \sum_{k=1}^K a_k \rho_L(\cdot - x_k^h)$ where $q \in \mathcal{N}_L$. It yields a cost $\mathcal{J}(s_h) = \frac{1}{2} \|\boldsymbol{\nu}(s_h) - \mathbf{y}\|_2^2 + \lambda \|\mathbf{a}\|_1$ since the x_k^h are pairwise distinct.

Hence, there only remains to prove that $\boldsymbol{\nu}(s_h)$ converges to $\boldsymbol{\nu}(s)$ when $h \rightarrow 0$. We now show that $s_h \xrightarrow{w^*} s$ when $h \rightarrow 0$, *i.e.*, $\langle s_h, f \rangle \rightarrow \langle s, f \rangle$ for any $f \in C_L(\mathbb{R})$ (see Section 2.1 for some background on weak* convergence). Using the decomposition of $\mathcal{C}_L(\mathbb{R})$ in (2.46) in Theorem 2.3, for any $f \in \mathcal{C}_L(\mathbb{R})$, there is a unique decomposition $f = \mathbf{L}^*\{f_1\} + f_2$ such that $f_1 \in C_0(\mathbb{R})$ and $f_2 \in \text{span}\{\phi_n\}_{n=1}^{N_0}$. We first notice that $\phi(s - s_h) = \mathbf{0}$ since $\phi(\mathbf{L}_\phi^{-1}\{w\}) = \mathbf{0}$ for any $w \in \mathcal{M}(\mathbb{R})$ by Theorem 2.1. Therefore, since $f_2 \in \text{span}\{\phi_n\}_{n=1}^{N_0}$ we have $\langle s - s_h, f_2 \rangle = 0$. Next, we have

$$\begin{aligned} \langle s - s_h, f_1 \rangle &= \langle s - s_h, \mathbf{L}^*\{f_1\} \rangle \\ &= \langle \mathbf{L}\{s - s_h\}, f_1 \rangle \\ &= \sum_{k=1}^K a_k \langle \delta(\cdot - x_k) - \delta(\cdot - x_k^h), f_1 \rangle \\ &= \sum_{k=1}^K a_k (f_1(x_k) - f_1(x_k^h)). \end{aligned}$$

Moreover, by definition of $C_0(\mathbb{R})$, f_1 is continuous, and since $\lim_{h \rightarrow 0} x_k^h = x_k$, we have $\lim_{h \rightarrow 0} \langle s - s_h, f_1 \rangle = 0$. We have thus proved that $s_h \xrightarrow{w^*} s$. Since by the hypotheses of Theorem 3.3, $\boldsymbol{\nu}$ is weak*-continuous, we have $\lim_{h \rightarrow 0} \boldsymbol{\nu}(s_h) = \boldsymbol{\nu}(s)$ and thus $\lim_{h \rightarrow 0} \mathcal{J}(s_h) = \mathcal{J}(s) = \mathcal{J}^0$. \square

Going back to polynomial B-splines, let $\mathcal{J}_d^0(h) \triangleq \min_{c \in \ell_{1,L}(\mathbb{Z})} \mathcal{J}_d(c)$ be the optimal cost of the discrete Problem (4.9). We derive the following theorem, which stems directly from Lemma 4.3 and is similar to Lemma 8 in [90].

Theorem 4.2 (Convergence of the Cost Function of the Discrete Problem). *Let all the hypotheses of Theorem 3.3 be met. Then*

$$\lim_{h \rightarrow 0} \mathcal{J}_d^0(h) = \mathcal{J}^0. \quad (4.22)$$

Proof. Firstly, we observe that for any $h > 0$, we have

$$\mathcal{J}_d^0(h) \geq \mathcal{J}^0 \quad (4.23)$$

since for any $c \in \ell_{1,L}(\mathbb{Z})$, we can define $s(x) = \sum_{k \in \mathbb{Z}} c[k] \beta_{L,h}(x - hk) \in \mathcal{M}_{L,h}(\mathbb{R})$ which verifies $\mathcal{J}_d(c) = \mathcal{J}(s) \geq \mathcal{J}^0$. Next, let s_h be a family of functions for any $h > 0$ as specified by Lemma 4.3. Since $s_h \in \mathcal{M}_{L,h}(\mathbb{R})$, by (4.6), s_h can be expressed in the B-spline basis as $s_h(x) = \sum_{k \in \mathbb{Z}} c[k] \beta_{L,h}(x - hk)$ where $c \in \ell_{1,L}(\mathbb{Z})$. Therefore, we have

$$\mathcal{J}_d^0(h) \leq \mathcal{J}_d(c) = \mathcal{J}(s_h) \xrightarrow{h \rightarrow 0} \mathcal{J}^0$$

which together with (4.23) proves the desired result. \square

Theorem 4.2 shows that the choice of $\mathcal{M}_{L,h}(\mathbb{R})$ as a search space for the continuous-domain problem is a sound one: by solving the discrete problem, we recover a solution which is arbitrarily close in terms of cost to the solution(s) of the continuous problem if h is sufficiently small. Moreover, note that there is no requirement in Theorem 4.2 that the natural gridded approximation s_h defined in (4.20) is a solution of the discrete Problem (4.9): $\mathcal{J}_d^0(h)$ might actually be smaller than $\mathcal{J}(s_h)$. We can therefore hope for a faster convergence than that of $\mathcal{J}(s_h) \rightarrow \mathcal{J}^0$.

4.4.2 Multiresolution Strategy

Although Section 4.3.3 provides an experimental pipeline to solve the continuous-domain problem in $\mathcal{M}_{L,h}(\mathbb{R})$ for a fixed grid size h , the choice of the latter is somewhat arbitrary. In practice, in order to choose the grid size, we use the convergence results of Theorem 4.2. We recursively split the grid in half by taking $h_i = T/2^i$ for increasing values of $i \in \mathbb{N}$, and we solve the corresponding finite problems. This way, the finest grid (highest value of i) contains all its coarser predecessors, which implies that the search spaces are embedded, *i.e.*, $\mathcal{M}_{L,h_i}(\mathbb{R}) \subset \mathcal{M}_{L,h_{i+1}}(\mathbb{R})$. This allows us to use the solution obtained with the previous grid as a starting point of ADMM, which leads to considerable time gains. Another consequence of this embedding is that $\mathcal{J}_d^0(h_i) \geq \mathcal{J}_d^0(h_{i+1})$, which indicates that splitting the grid in half can only improve the solution in terms of cost. Theorem 4.2 then guarantees that $\lim_{i \rightarrow +\infty} \mathcal{J}_d^0(h_i) = \mathcal{J}^0$. This gives us a natural stopping criterion: we increment i until the relative decrease of cost $(\mathcal{J}_d^0(h_{i-1}) - \mathcal{J}_d^0(h_i)) / \mathcal{J}_d^0(h_{i-1})$ is smaller than some tolerance parameter ϵ . When ϵ is sufficiently small, we consider that the cost function has converged and that there is no need to make the grid any finer. Note that the simplex step is only necessary for the final grid size. This complete procedure is detailed in Algorithm 4.1.

In Algorithm 4.1, $\text{ADMM}(\mathbf{c}_{\uparrow 2}; \mathbf{H}, \mathbf{L}, \mathbf{y}, \lambda)$ runs ADMM on Problem (4.10) with the starting point $\mathbf{c}_{\uparrow 2}$. The latter corresponds to the vector of B-spline coefficients \mathbf{c} converted to the current grid size, which is twice as fine as that of \mathbf{c} . This conversion is made possible by the embedding of the search spaces. Similarly, $\text{Simplex}(\mathbf{H}, \mathbf{L}, \mathbf{y}, \lambda, \mathbf{y}_\lambda)$ runs the simplex algorithm on the constrained Problem (4.19) (no starting point is required). The output \mathbf{c}^* of this algorithm is therefore a vector whose size is not predetermined, but which

Input: $\nu, L, T, \mathbf{y}, \lambda, i_{\min}, \epsilon$

Output: \mathbf{c}^*

$i = i_{\min}; \mathbf{c} = \mathbf{0}; cost_p = +\infty; error = \epsilon + 1;$

while $error > \epsilon$ **do**

$h = T/2^i;$	
update $\mathbf{H}, \mathbf{L};$	<i>// Depend on h, ν, L, T</i>
$\mathbf{c} = \text{ADMM}(\mathbf{c}_{\uparrow 2}; \mathbf{H}, \mathbf{L}, \mathbf{y}, \lambda);$	
$error = cost(\mathbf{c}) - cost_p /cost_p;$	
$cost_p = cost(\mathbf{c});$	
$i = i + 1;$	

end

$\mathbf{y}_\lambda = \mathbf{H}\mathbf{c};$

$\mathbf{c}^* = \text{Simplex}(\mathbf{H}, \mathbf{L}, \mathbf{y}, \lambda, \mathbf{y}_\lambda);$

Algorithm 4.1: Pseudocode of our algorithm

represents a continuous-domain signal in I_T that is sparse in the Green's function basis, and yields a cost close to \mathcal{J}_0 .

4.5 Experimental Results

We now discuss our implementation of Algorithm 4.1 and present some results. The ADMM algorithm is implemented using GlobalBioIm [118], an inverse problem library developed in our group (see Appendix A for more details on how ADMM is used to solve Problem (4.10)). The Gurobi optimizer² is used for the simplex algorithm.

4.5.1 Experimental Setting

Test signal

We attempt to reconstruct sparse signals of the form

$$s(x) = \sum_{k=1}^{K_s} a_k \rho_L(x - x_k) + \sum_{n=1}^{N_0} b_n p_n(x), \quad (4.24)$$

for which gTV regularization is an adequate prior by Theorem 3.3. The sparsity index K_s is chosen by the user and the knots x_k are drawn at random in the interval of interest I_T following a uniform distribution. The coefficients a_k and b_n are i.i.d Gaussian random variables projected on the subspace of vectors $(\mathbf{a}, \mathbf{b}) \in \mathbb{R}^{K+N_0}$ for which s is supported in I_T . This is to enforce the finite support assumption on the test signal, which is implicit in the discrete problem formulation (4.10). Therefore, aside from the approximation error on the knot locations, the test signal in (4.24) is in the span of feasible signals reconstructed

²LLC Gurobi Optimization, Gurobi optimizer reference manual, 2018.

by the discrete problem, which is obviously a desirable property.

Measurements

We implemented two types of measurement operators $\boldsymbol{\nu} = (\nu_1, \dots, \nu_M)$:

- **Ideal sampling:** This case corresponds to a measurement operator $\nu_m = \delta(\cdot - x_m)$, *i.e.*, $\boldsymbol{\nu}(f) \triangleq (f(x_1), \dots, f(x_M))$, where $x_m \in I_T$ for any $m \in \{1, \dots, M\}$. As discussed in Example 3.2, this choice is technically only admissible—in the sense that ν_m is weak* continuous and thus satisfies the assumptions of Theorem 3.3—when $N_0 \geq 2$. However, even when $N_0 = 1$, one can choose a Dirac mollifier such as (2.48), which leads numerically to taking the right limit $\nu_m(f) = f(x_m^+)$. Given the form of \mathcal{N}_L in (2.24), it can be shown that $\boldsymbol{\nu}$ satisfies the well-posedness assumption $\mathcal{N}_\boldsymbol{\nu} \cap \mathcal{N}_L = \{0\}$ as soon as $M \geq N_0$ and all sampling points x_m are pairwise distinct. We either take uniformly spaced knots or random samples following a uniform distribution in I_T for x_1, \dots, x_M .
- **Fourier-domain sampling:** We take the inner product with functions $x \mapsto \mathbb{1}_{[0,T]}(x) \cos(\omega_m x + \theta_m)$, where we have $\omega_1 = 0$ and $\omega_m \in (0, \omega_{\max}]$ for $m = \{2, \dots, M\}$. The maximum pulsation ω_{\max} is chosen so that the spectrum of s has small energy above this threshold. We take $\theta_1 = 0$ and the θ_m are drawn from a uniform distribution in $[0, \pi)$ for $m \in \{2, \dots, M\}$. We thus have

$$\boldsymbol{\nu} \triangleq \left(\mathbb{1}_{[0,T]}, \mathbb{1}_{[0,T]} \cos(\omega_2 \cdot + \theta_2), \dots, \mathbb{1}_{[0,T]} \cos(\omega_M \cdot + \theta_M) \right),$$

which leads to $\nu_m(f) = \int_0^T f(x) \cos(\omega_m x + \theta_m) dx$. As discussed in Example 3.3, this is a valid choice in the sense that ν_m is weak*-continuous.

We compare experimental results using these measurement operators, predicting that reconstructed signals should be closer to the test signals when sampling in the Fourier domain than with ideal sampling. This ensues from the theory of CS: Fourier matrices are known to have good recovery properties with few measurements [119], whereas sampling matrices clearly do not. In order to be more realistic and to verify the robustness of our algorithm, we add Gaussian noise to the measurements with standard deviation σ computed from a given signal-to-noise ratio (SNR).

Regularization Parameter

The choice of the regularization parameter λ is critical, as it greatly affects the reconstructed signal: high values of λ can lead to overly regularized solutions, whereas low values tend to suppress the effect of regularization. The value of λ should be tuned

according to the type of measurement. To this end, in our experiments, we choose a value of λ among a list of potential values such that the SNR between the reconstructed signal and the test signal is the highest for a certain value of h . The selected λ is then used for all values of h , as specified in Algorithm 4.1.

4.5.2 Experimental Results

We now present several results of our numerical implementation.

Sparsity

In our experiments, we observe that, as predicted by Theorem 4.1, the final reconstructed signal has sparsity $K \leq M - N_0$ in the Green's function basis.

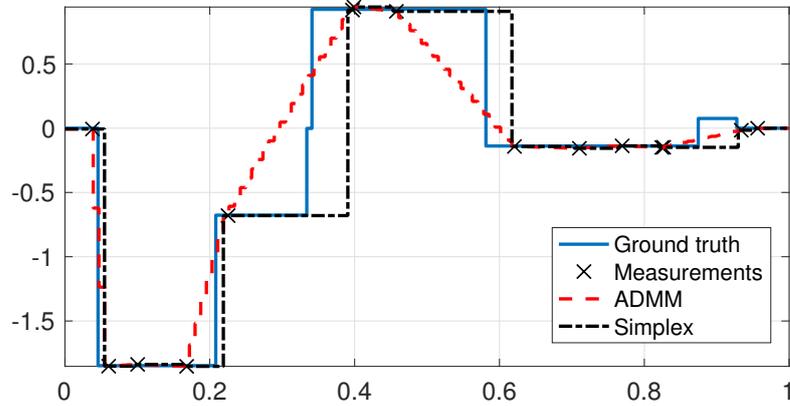


Figure 4.1 – Reconstruction results for $L = D$, $M = 15$ (ideal sampling), grid size $h = \frac{1}{128}$. Sparsity: 118 after ADMM, 13 after simplex.

However, running the simplex after ADMM is far from being superfluous: reconstructed signals after ADMM are typically not sparse at all. This is best illustrated in the case of ideal sampling, where we observe a staircase effect between measurements (Fig. 4.1). Although this phenomenon does not affect the cost function, it is clearly not optimal in terms of sparsity, and it illustrates the nonuniqueness of \mathcal{V}_f . However, after the simplex step, the sparsity improves dramatically, going from 118 to $13 \leq M - N_0$ in Fig. 4.1, as predicted by Theorem 4.1.

Measurement Types Comparison

In the experiment shown in Figure 4.2, we set $L = D^2$. We compare the reconstruction results for both types of measurement operators presented above on the same ground-truth signal with the same number $M = 10$ of measurements, and the same grid size

$h = 1/2^8$. As predicted by the theory of CS (*i.e.*, Fourier matrices are known to have good recovery properties [119]), the reconstruction using Fourier-domain samples is very accurate despite the small number of measurements and the presence of noise. Conversely, the reconstruction using ideal samples is less faithful since we have no information on the behavior of the signal between sampling points. This disparity is underlined by the difference in SNR values (12.03 dB versus 18.07 dB). Note that in both cases, the reconstructed signals have sparsity $K \leq M - N_0 = 8$ which conforms with Theorem 4.1.

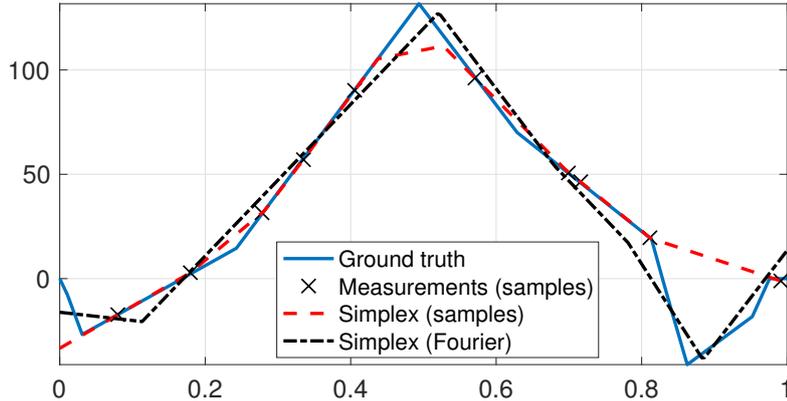


Figure 4.2 – Reconstructions using ideal samples (SNR= 12.03 dB) and Fourier-domain measurements (SNR= 18.07 dB) for $L = D^2$, $M = 10$ and $h = 1/2^8$.

Decreasing Grid Size

As the grid size decreases, the search space of our optimization problem becomes larger: we can therefore reconstruct functions in finer detail. This is illustrated in Fig. 4.3, in which we observe that very coarse grids approximate complex signals very poorly, whereas after splitting the grid in half recursively, these signals can rapidly be approximated much better (Fig. 4.3c).

To illustrate the effect of the decreasing grid size in terms of cost, we present an example run of Algorithm 4.1 with a regularization operator $L = D^3$ in Fig. 4.4. The final reconstructed signal is shown in Fig. 4.4a: notwithstanding the reasonably fine grid size ($h = 1/2^8$), the reconstruction is near-perfect. The evolution of the cost function with respect to the grid size in our example is shown in Fig. 4.4b: we observe that after an initial rapid decrease, the cost function starts plateauing, which is in line with Theorem 4.2. Given the aspect of this evolution, it is safe to assert that the cost is very close to its limit value \mathcal{J}_0 specified by Theorem 4.2. Although we could consider tightening the tolerance ϵ to get a marginally smaller cost, this is not necessarily a sensible choice. Indeed, for very fine grids (*e.g.*, $h < 1/2^{11}$), the increased scale of the problem can cause computational problems larger than the potential gain in terms of cost. We found the choice of $\epsilon = 10^{-3}$ to be a good compromise in our experimental setting: the final grid size h is typically coarser than $1/2^{10}$, even for very nonsparse test signals ($K_s \approx 100$) and

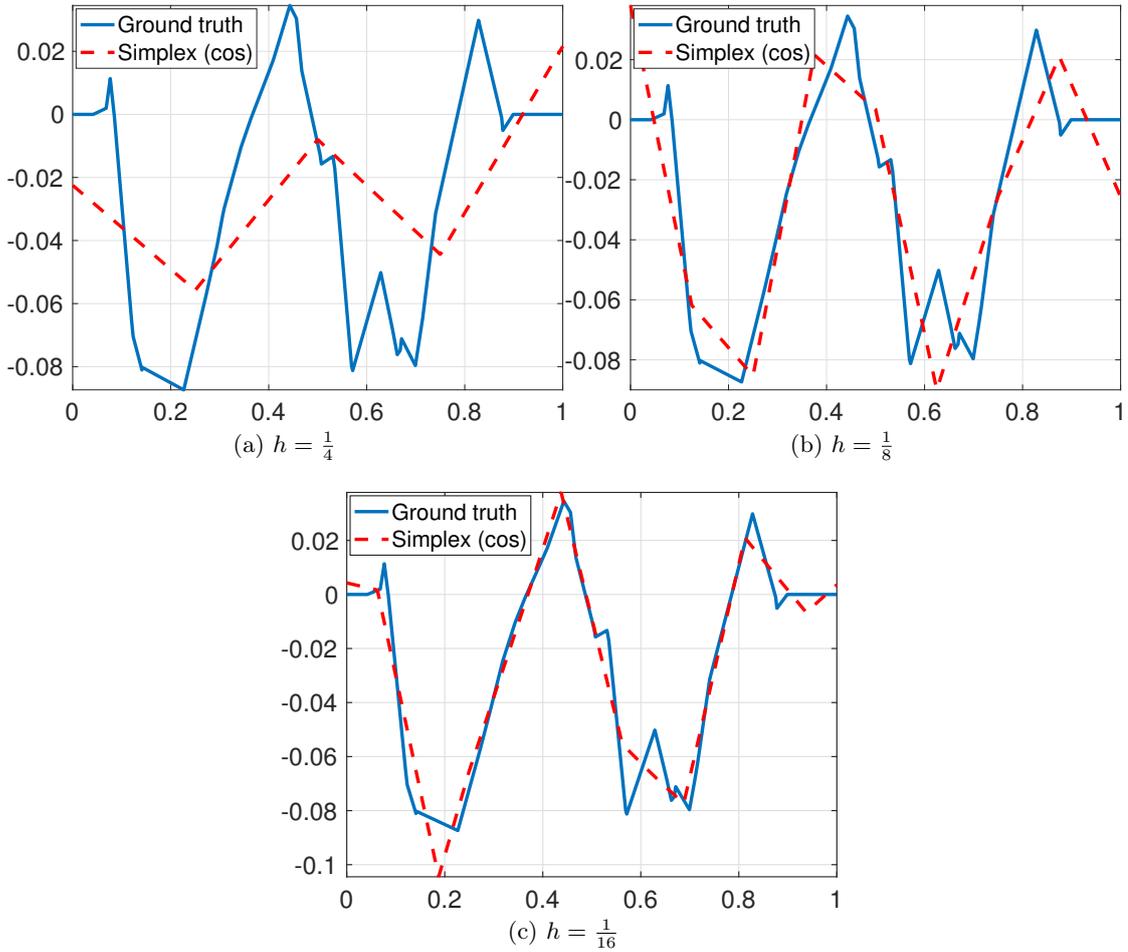


Figure 4.3 – Reconstruction results after simplex for $L = D^2$, $M = 31$ (dephased cosine sampling) with decreasing h .

with many measurements ($M \approx 100$). For such grid sizes, due to the good conditioning of the system matrix \mathbf{H} , the optimization problems are entirely feasible (ADMM typically converges in a few seconds with a properly tuned penalty parameter ρ) and computational problems are avoided.

In order to compare reconstruction results for different grid sizes, we applied the simplex step as described in Sec. 4.3.3 for every grid size h in our example. Despite the convergence of the cost function for the finer grid sizes observed in Fig. 4.4b, the variations in the sparsity (Fig. 4.4c) indicate that the reconstructed signals are not identical from one grid size to the next. However, in regard to the optimization problem, there is no reason to decrease the grid size any further or to favor one solution over another if both yield the same cost within a user-defined tolerance.

4.5.3 Comparison with Discrete Methods

In this section, we assess the pertinence of our framework by comparing it with a purely discrete method. The standard way of discretizing Problem (4.1) would be to consider uniform samples of the reconstructed function, *i.e.*, a pixel basis, and to approximate derivative operators with finite differences. Within this framework, the underlying discrete optimization problem of the form (4.10) is very similar for both methods. Indeed, in both cases, the regularization matrix \mathbf{L} is a finite difference-type matrix as in (4.12). However, since the basis functions are different, the number of coefficients N and the system matrix \mathbf{H} differ. We solve both problems using our pipeline described in 4.3.3.

We consider noiseless Fourier-domain measurements (dephased cosine sampling) with $M = 100$. As explained earlier, this bolsters the recovery properties of the reconstruction, and thus allows us to use similarity metrics between the reconstructed and test signals to compare both methods.

Such a comparison is made in Figure 4.5, with $L = D^4$ and $\lambda = 10^{-15}$. For the sake of fairness, we use a piecewise-linear test signal, since the latter does not resemble the basis functions of either method. Figure 4.5a shows the reconstruction result, using a coarse grid for visualization purposes. Our continuous-domain reconstruction is clearly a lot closer to the test signal; this observation is confirmed by the SNR values of both reconstructions in Figure 4.5b. We notice that the SNR is similar for both methods using finer grids: this is in keeping with [120], which demonstrates some form of convergence of discrete methods towards solutions of continuous-domain problems as the grid size goes to zero. However, our continuous-domain method converges much faster (*i.e.*, for coarser grids) towards a very faithful reconstruction.

Note that the observed linear regime of the blue curve in Fig. 4.5b is consistent with the approximation power of pixels, which is well known to be in $\mathcal{O}(h)$. Moreover, using finite differences instead of the derivative yields additional errors which increase with the order N_0 of the operator and when the grid gets coarser. Conversely, our method is exact in the continuous domain for *any* grid size, which explains why a grid of $h = 1/2^5$ can be sufficiently fine. Finally, note that although the discrete method leads to a slightly better-conditioned problem, the difference in speed is negligible due to the Riesz basis property of B-splines.

4.6 Conclusion

In this chapter, we have devised an efficient multiresolution algorithm to numerically compute sparse solutions of continuous-domain inverse problems with gTV regularization. Our grid-based discretization uses the B-spline dictionary basis matched to the regularization operator L . On the theoretical side, we proved that this is an exact discretization

of the underlying continuous-domain problem restricted to a search space, and that this discrete problem converges in terms of cost towards the continuous problem when the grid size decreases. On the experimental side, we implemented this discretization scheme for higher-order derivative regularization operators L , and different measurement operators. Our experimental results demonstrate that our formulation is computationally inexpensive, well-suited for practical problems and compares favorably with standard, purely discrete methods.

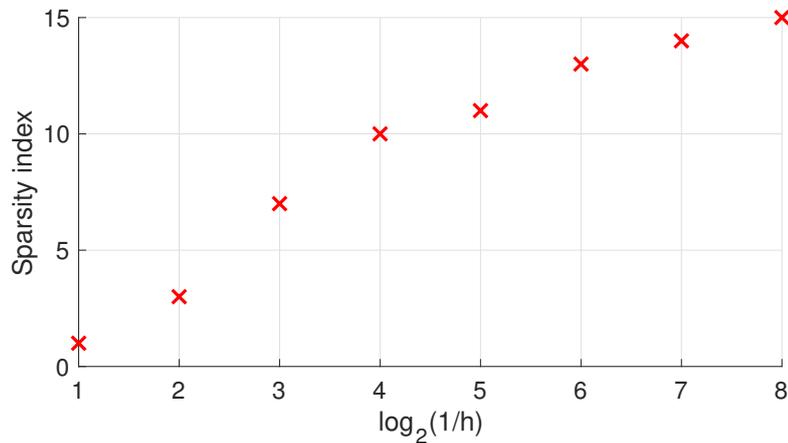
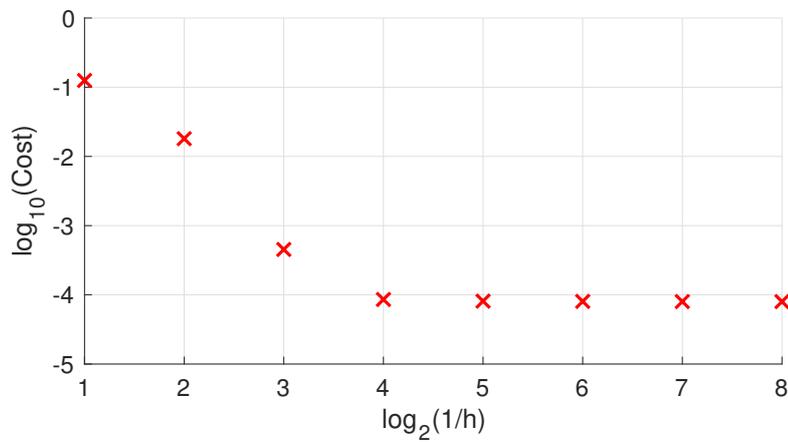
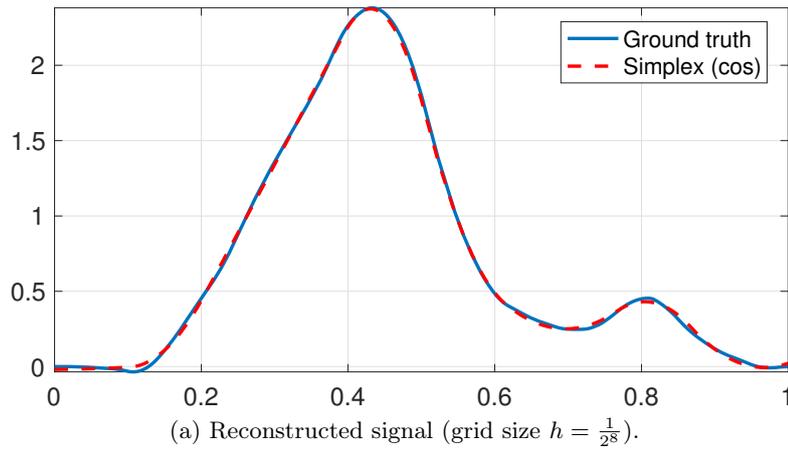


Figure 4.4 – Example run for $L = D^3$ (quadratic splines), $M = 31$ (dephased cosine sampling).

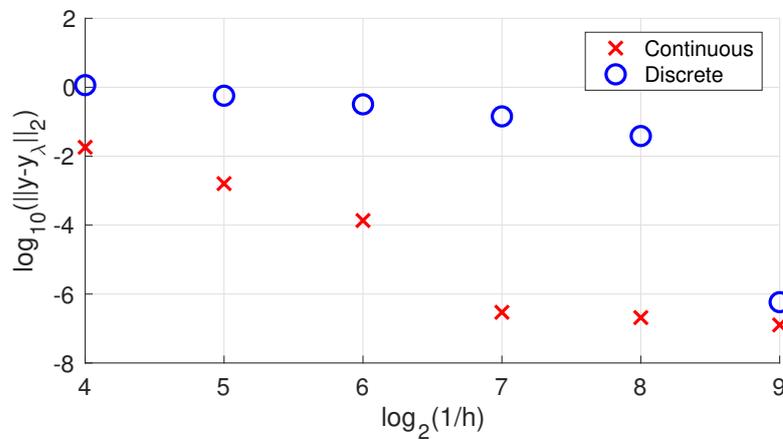
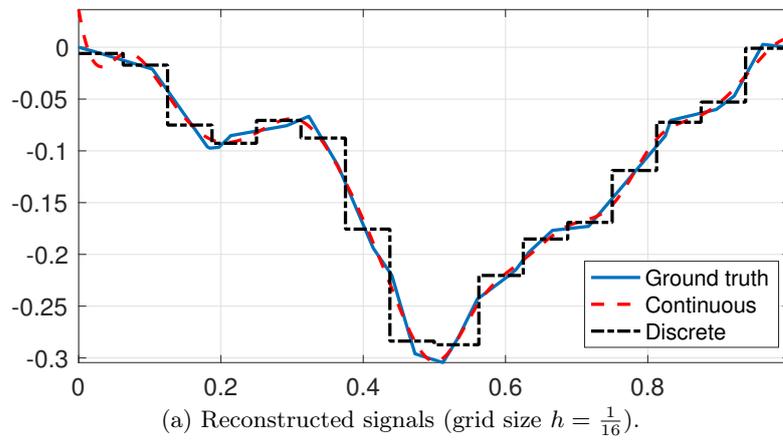


Figure 4.5 – Comparison between our continuous model and the pixel-based discrete model for $L = D^4$, $M = 100$ (dephased cosine sampling).

5 Hybrid-Spline Signal Models

This chapter is based on the following publication [121]:

T. Debarre, S. Aziznejad, and M. Unser, “Hybrid-Spline Dictionaries for Continuous-Domain Inverse Problems”, *IEEE Transactions on Signal Processing*, vol. 67, no. 22, pp. 5824–5836, Nov. 2019.

5.1 Introduction

The compressed sensing (CS) approach to find a sparse representation of a signal given its measurements is to solve an inverse problem with sparsity-promoting regularization. The aim is to reconstruct a signal $\mathbf{c} \in \mathbb{R}^N$ given measurements $\mathbf{y} = \mathbf{H}\mathbf{c} + \mathbf{n} \in \mathbb{R}^M$ (with $M \ll N$), where $\mathbf{H} : \mathbb{R}^N \rightarrow \mathbb{R}^M$ is the system (or measurement) matrix and \mathbf{n} is some additive noise which models measurement errors. The a priori assumption on the signal \mathbf{c} is that it is sparse in a certain dictionary frame $\mathbf{D} : \mathbb{R}^P \rightarrow \mathbb{R}^N$, so that $\mathbf{c} = \mathbf{D}\mathbf{a}$, where the vector of coefficients $\mathbf{a} \in \mathbb{R}^P$ is sparse. When the dictionary basis is orthonormal (*i.e.*, $P = N$ and \mathbf{D} is orthogonal), and the system matrix \mathbf{H} and the noise \mathbf{n} satisfy certain conditions, standard CS theory guarantees that the signal \mathbf{c} can be recovered exactly (see Section 1.2.1 for more information). This is achieved by solving the optimization problem

$$\hat{\mathbf{a}} \in \arg \min_{\mathbf{a} \in \mathbb{R}^P} (\|\mathbf{H}\mathbf{D}\mathbf{a} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{a}\|_1), \quad (5.1)$$

where $\lambda > 0$ is a regularization parameter, and by computing $\hat{\mathbf{c}} = \mathbf{D}\hat{\mathbf{a}}$.

In many cases, however, the sparsity of a signal is expressed in a dictionary basis that is redundant and thus not orthonormal. More precisely, in this work, we are interested in multicomponent signals $\mathbf{c} = \sum_{i=1}^Q \mathbf{c}_i$ such that each of the Q components \mathbf{c}_i is sparse in a different dictionary \mathbf{D}_i ($i \in \{1, \dots, Q\}$). For simplicity of exposition, we set $Q = 2$. A natural way of formulating the recovery problem is to solve (5.1) with a concatenated dictionary $\mathbf{D} = (\mathbf{D}_1, \mathbf{D}_2)$, where we can assume for now that \mathbf{D}_1 and \mathbf{D}_2 are orthonormal

bases. In this case, $P = 2N$ and the dictionary basis is clearly highly redundant. Although this setting introduces practical and theoretical difficulties, it is extremely useful in many applications when a single dictionary is insufficient to represent the richness of a signal.

The problem of accurately reconstructing both components \mathbf{c}_1 and \mathbf{c}_2 is known as data separation [10, Chapter 11], and has been studied extensively both theoretically and practically. In fact, some of the first theoretical works concerning sparse vector recovery using ℓ_1 -norm minimization involved a concatenated dictionary consisting of a mixture of sinusoids and spikes [122, 123]. The goal was to provide a condition under which ℓ_0 and ℓ_1 minimization yield the same solution. This sparked an abundance of research, which extended and improved these results for more general (nonorthonormal) dictionaries [124–127]. An overview is given in [128]. Later, these results were extended to images to separate point-like and curve-like structures [129]. These works mostly tackle denoising problems characterized by $\mathbf{H} = \mathbf{I}_N$ and $M = N$. In the CS field, in which we have $M \ll N$, [130] considers redundant dictionaries in general. On the practical side, data separation is intimately related to morphological component analysis (MCA), a method popularized by Starck *et al.* [131–135] with applications in inpainting removal or the separation of texture and natural parts of an image.

5.1.1 Analysis Formulation with Unions of Dictionaries

For overcomplete dictionaries \mathbf{D} in general, optimizing over the synthesis coefficients \mathbf{a} as in (5.1) seems natural, since the sparsity of the signal is precisely enforced on \mathbf{a} . This leads to a separable (if not differentiable) regularization term and, thus, to standard soft-thresholding-based algorithms. Conversely, optimizing over the analysis coefficients \mathbf{c} is less straightforward since the representation of the signal \mathbf{c} by the synthesis coefficients \mathbf{a} is not unique. Therefore, the enforcement of the sparsity of the signal would require one to find its sparsest representation in the dictionary basis, which seems to add unnecessary complications. This explains why most of the published works choose to optimize over the synthesis coefficients.

However, in the case of a concatenated dictionary $\mathbf{D} = (\mathbf{D}_1, \mathbf{D}_2)$, one can also think of optimizing over both analysis coefficients \mathbf{c}_1 and \mathbf{c}_2 . This leads to an analysis problem of the form

$$(\hat{\mathbf{c}}_1, \hat{\mathbf{c}}_2) \in \arg \min_{\mathbf{c}_1, \mathbf{c}_2 \in \mathbb{R}^N} \left(\frac{1}{2} \|\mathbf{H}(\mathbf{c}_1 + \mathbf{c}_2) - \mathbf{y}\|_2^2 + \lambda (\|\mathbf{L}_1 \mathbf{c}_1\|_1 + \|\mathbf{L}_2 \mathbf{c}_2\|_1) \right), \quad (5.2)$$

which yields the reconstructed signal $\hat{\mathbf{c}} = \hat{\mathbf{c}}_1 + \hat{\mathbf{c}}_2$. When \mathbf{D}_1 and \mathbf{D}_2 are orthonormal bases, taking $\mathbf{L}_1 = \mathbf{D}_1^T$ and $\mathbf{L}_2 = \mathbf{D}_2^T$ makes Problems (5.1) and (5.2) exactly equivalent. In other cases, the Moore-Penrose pseudoinverses of \mathbf{D}_1 and \mathbf{D}_2 can be used for \mathbf{L}_1 and \mathbf{L}_2 , in which case (5.1) and (5.2) are not equivalent. Although the analysis formulation (5.2) may seem more tedious, there are several reasons to believe that it might be appealing:

- in recent years, efficient splitting algorithms that overcome the nonseparability of the regularization terms in Problem (5.2) have been designed (*e.g.*, [136] or the well-known alternating-direction method of multipliers (ADMM) [19]);
- many successful practical CS applications favor the analysis formulation, *e.g.*, [37];
- there is an intimate connexion between regularization operators (*e.g.*, \mathbf{L}_1 and \mathbf{L}_2 in (5.2)) and splines. Indeed, it has recently been proved that continuous-domain inverse problems with generalized total-variation (gTV) regularization have sparse spline solutions (Theorem 3.3). The analysis formulation (5.2) therefore seems like the natural discrete counterpart to this continuous framework, as we clarify in Section 5.2.

The difference between the two formulations has been studied in [137] where the authors conclude that, in the general case, both problems are inherently different. Several practitioners have used the analysis formulation for data-separation problems, most notably Starck *et al.* in the context of MCA [131–134]. More recently, Problem (5.2) was applied to the task of separating cartoon and texture parts of an image in [138]. There, the authors favor the analysis formulation for general redundant dictionaries. A similar approach is used in low-rank plus sparse decomposition methods [139].

Despite these empirical works, virtually no theoretical study of Problem (5.2) has been carried out. In [140], Candès *et al.* have named it the “split-analysis” problem. They recommend solving it precisely in the case of concatenated dictionaries (without further investigation). A theoretical study was later done by Lin *et al.* in [141], where they show that the data-separation problem (*i.e.*, the recovery of both components of the original signal) can be solved via Problem (5.2). This result requires that \mathbf{H} satisfies the restricted-isometry property adapted to a dictionary (D-RIP) and that \mathbf{L}_1 and \mathbf{L}_2 satisfy a mutual coherence condition.

While the literature on the topic is scarce in the discrete setting, to the best of our knowledge, it is nonexistent in the continuous domain. Since most real-world signals are continuously defined, the reconstruction of continuous-domain solutions is a desirable objective. Moreover, although handling discrete signals is obviously appealing from a computational perspective, it introduces discretization errors in the measurements.

5.1.2 Continuous-Domain Problems for Hybrid Splines

In this chapter, we propose to use unions of dictionaries in a continuous-domain framework. Our goal is to reconstruct a multicomponent continuous-domain 1D signal $s = s_1 + s_2$, where s_1 and s_2 have different characteristics. Similarly to the discrete setting, we are given measurements $\mathbf{y} = \boldsymbol{\nu}(s) + \mathbf{n} \in \mathbb{R}^M$, where $\boldsymbol{\nu} : s \mapsto \boldsymbol{\nu}(s) \in \mathbb{R}^M$ is a (continuous-domain) linear measurement operator and $\mathbf{n} \in \mathbb{R}^M$ is some additive noise. We focus on

continuous-domain inverse problems of the form

$$s^* \in \arg \min_f \left(\frac{1}{2} \|\boldsymbol{\nu}(f) - \mathbf{y}\|_2^2 + \lambda \mathcal{R}_{\text{hyb}}(f) \right), \quad (5.3)$$

where $\lambda > 0$ is the regularization parameter. The hybrid regularization term is given by

$$\mathcal{R}_{\text{hyb}}(f) \triangleq \min_{\substack{f_1, f_2 \\ f_1 + f_2 = f}} \left((1 - \alpha) \|\mathbf{L}_1\{f_1\}\|_{\mathcal{M}} + \alpha \|\mathbf{L}_2\{f_2\}\|_{\mathcal{M}} \right), \quad (5.4)$$

where $\alpha \in (0, 1)$ controls the weighing of the two regularization terms.

Although they are formulated in slightly different forms, there is an obvious connection between Problem (5.3) and the discrete analysis Problem (5.2). Similarly to the discrete case, the sparsity of each component s_i is promoted in a different dictionary, which is determined by the choice of the regularization operator \mathbf{L}_i for $i \in \{1, 2\}$. Our main theoretical result is a representer theorem that states that, for differential operators \mathbf{L}_i , Problem (5.3) leads to spline solutions $s^* = s_1^* + s_2^*$, where each component s_i^* is an \mathbf{L}_i -spline. The reconstructed signal s^* is therefore a sum of different splines, which we coin as a *hybrid spline*. Moreover, the total sparsity of s^* in this union of spline dictionaries is no larger than the number of measurements M . This representer theorem generalizes the main result of [14].

An important characteristic of our framework is its compatibility with the class of piecewise-polynomial functions, which can accurately model a large variety of real-world signals. For example, by taking $\mathbf{L}_1 = \mathbf{D}$ (piecewise-constant splines) and $\mathbf{L}_2 = \mathbf{D}^4$ (cubic splines), the dictionary consists of piecewise-cubic polynomials that admit discontinuities.

A key feature of our continuous-domain formulation is that Problem (5.3) can be discretized in an exact way using B-splines, based on the methodology of Chapter 4.

5.1.3 Outline

This chapter is organized as follows: in Section 5.2, we present the continuous-domain framework of Problem (5.3), which is a multicomponent extension of Section 3.2. We then prove our representer theorem; it states that Problem (5.3) admits a global minimizer that is a hybrid spline. In Section 5.3, we show how we discretize Problem (5.3) in an exact way using B-splines and propose a practical solution to the corresponding discrete problem. In Section 5.4, we present our proposed multiresolution algorithm. Finally, we illustrate some applications of our algorithm in Section 5.5, including curve fitting and CS-type problems.

5.2 Continuous-Domain Inverse Problem

In this section, we present a class of continuous-domain inverse problems that involve hybrid gTV regularization. These problems extend the continuous-domain inverse-problem framework of Section 3.2.

5.2.1 Definitions and Notations

Following the notations of Section 3.2, our aim is to recover a continuous-domain signal $s : \mathbb{R} \rightarrow \mathbb{R}$ given M noisy measurements modeled as $\mathbf{y} = \boldsymbol{\nu}(s) + \mathbf{n}$, where $\mathbf{n} \in \mathbb{R}^M$ is some additive noise. The noiseless measurements $\boldsymbol{\nu}(s)$ are acquired through M linear measurement functionals $\boldsymbol{\nu} = (\nu_1, \dots, \nu_M)$, with $\boldsymbol{\nu}(s) = (\langle \nu_1, s \rangle, \dots, \langle \nu_M, s \rangle)$.

Next, we sum up all the relevant information and notations that concern the regularization operators L_i ($i \in \{1, 2\}$); for more details, we refer to Chapter 2.

1. The operator L_i is assumed to be the $N_{0,i}$ th-order derivative operator¹. It has a Green's function ρ_{L_i} defined in (2.25) which verifies $L_i\{\rho_{L_i}\} = \delta$.
2. The native space $\mathcal{M}_{L_i}(\mathbb{R})$ of L_i is defined in (2.33); we then have $L_i : \mathcal{M}_{L_i}(\mathbb{R}) \rightarrow \mathcal{M}(\mathbb{R})$, where $\mathcal{M}(\mathbb{R})$ is the space of Radon measures defined in (2.18).
3. The null space $\mathcal{N}_{L_i} \triangleq \{f \in \mathcal{S}'(\mathbb{R}) : L\{f\} = 0\}$ of L_i is given in (2.24) and has finite dimension $N_{0,i}$.
4. The intersection of the null spaces is denoted by $\mathcal{N}_0 = \mathcal{N}_{L_1} \cap \mathcal{N}_{L_2}$ and has dimension N_0 . If $N_0 > 0$, we introduce the biorthogonal system $(\boldsymbol{\phi}_0, \mathbf{p}_0) = (\phi_n, p_n)_{n=1}^{N_0}$ for \mathcal{N}_0 in the sense of Definition 2.6, where $\boldsymbol{\phi}_0 : \mathcal{M}_{L_1}(\mathbb{R}) \rightarrow \mathbb{R}^{N_0}$ and $\mathbf{p}_0 \in \mathbb{R}^{N_0}$.
5. The restricted search space for L_1 is defined as

$$\mathcal{M}_{L_1, \boldsymbol{\phi}_0}(\mathbb{R}) \triangleq \{f \in \mathcal{M}_{L_1}(\mathbb{R}) : \boldsymbol{\phi}_0(f) = \mathbf{0}\}. \quad (5.5)$$

Finally, we introduce the *hybrid regularization functional* for functions $f \in \mathcal{M}_{L_1, \boldsymbol{\phi}_0}(\mathbb{R}) + \mathcal{M}_{L_2}(\mathbb{R})$ as

$$\mathcal{R}_{\text{hyb}}(f) \triangleq \min_{\substack{f_1 \in \mathcal{M}_{L_1, \boldsymbol{\phi}_0}(\mathbb{R}) \\ f_2 \in \mathcal{M}_{L_2}(\mathbb{R}) \\ f_1 + f_2 = f}} \left((1 - \alpha) \|L_1\{f_1\}\|_{\mathcal{M}} + \alpha \|L_2\{f_2\}\|_{\mathcal{M}} \right). \quad (5.6)$$

Proposition 5.1. *The hybrid regularization functional \mathcal{R}_{hyb} defined in (5.6) has the following properties:*

¹As in most chapters of the thesis, all the results and algorithms from this chapter can be extended to more general operators L_i , e.g., rational operators [102].

1. it is well defined in the sense that, for any $f \in \mathcal{M}_{L_1, \phi_0}(\mathbb{R}) + \mathcal{M}_{L_2}(\mathbb{R})$, there exist functions $(f_1, f_2) \in \mathcal{M}_{L_1, \phi_0}(\mathbb{R}) \times \mathcal{M}_{L_2}(\mathbb{R})$ such that $f = f_1 + f_2$ and

$$\mathcal{R}_{\text{hyb}}(f) = (1 - \alpha)\|\mathbb{L}_1\{f_1\}\|_{\mathcal{M}} + \alpha\|\mathbb{L}_2\{f_2\}\|_{\mathcal{M}}; \quad (5.7)$$

2. it is convex.

Proof. Recall that (ϕ_0, \mathbf{p}_0) is a biorthogonal system in the sense of Definition 2.6 for $\mathcal{N}_0 = \mathcal{N}_{L_1} \cap \mathcal{N}_{L_2}$. The latter can be extended to two biorthogonal systems $(\tilde{\phi}_1, \tilde{\mathbf{p}}_1)$ and $(\tilde{\phi}_2, \tilde{\mathbf{p}}_2)$ for \mathcal{N}_{L_1} and \mathcal{N}_{L_2} , respectively, such that $\tilde{\phi}_i = [\phi_0 \ \phi_i]$ and $\tilde{\mathbf{p}}_i = [\mathbf{p}_0 \ \mathbf{p}_i]$ for $i \in \{1, 2\}$.

Using Theorem 2.1, there exist stable right-inverse operators $\mathbb{L}_{\tilde{\phi}_i}^{-1}$ such that any $f \in \mathcal{M}_{L_i}(\mathbb{R})$ has a unique representation as

$$f(x) = \mathbb{L}_{\tilde{\phi}_i}^{-1}\{w_i\}(x) + \mathbf{c}_0^T \mathbf{p}_0(x) + \mathbf{c}_i^T \mathbf{p}_i(x), \quad (5.8)$$

where $w_i \in \mathcal{M}(\mathbb{R})$, $\mathbf{c}_0 \in \mathbb{R}^{N_0}$, and $\mathbf{c}_i \in \mathbb{R}^{N_{0,i} - N_0}$ with $i \in \{1, 2\}$. We then proceed in two steps.

The minimum is reached in (5.6). Let us denote

$$\mathcal{C}_f \triangleq \{(f_1, f_2) \in \mathcal{M}_{L_1, \phi_0}(\mathbb{R}) \times \mathcal{M}_{L_2}(\mathbb{R}) : f = f_1 + f_2\} \quad (5.9)$$

the set of feasible pairs (f_1, f_2) and

$$\mathcal{R}_{1,2}(f_1, f_2) = (1 - \alpha)\|\mathbb{L}_1\{f_1\}\|_{\mathcal{M}} + \alpha\|\mathbb{L}_2\{f_2\}\|_{\mathcal{M}}. \quad (5.10)$$

The hybrid regularization can be rewritten as

$$\mathcal{R}_{\text{hyb}}(f) = \min_{(f_1, f_2) \in \mathcal{C}_f} \mathcal{R}_{1,2}(f_1, f_2). \quad (5.11)$$

The feasible set \mathcal{C}_f is nonempty since $f \in \mathcal{M}_{L_1, \phi_0}(\mathbb{R}) + \mathcal{M}_{L_2}(\mathbb{R})$. Therefore, there exists a sequence $(f_{1,k}, f_{2,k})_{k \in \mathbb{N}} \in \mathcal{C}_f$ that monotonically decreases to the infimum of $\mathcal{R}_{1,2}$ over \mathcal{C}_f . Using (5.8), both components can be represented as $f_{1,k} = \mathbb{L}_{\tilde{\phi}_1}^{-1}\{w_{1,k}\} + \mathbf{c}_{1,k}^T \mathbf{p}_1$ and $f_{2,k} = \mathbb{L}_{\tilde{\phi}_2}^{-1}\{w_{2,k}\} + \mathbf{c}_{0,k}^T \mathbf{p}_0 + \mathbf{c}_{2,k}^T \mathbf{p}_2$. The assumption of monotonic decrease implies that

$$\forall k \in \mathbb{N}, i \in \{1, 2\} : \quad 0 \leq \|w_{i,k}\|_{\mathcal{M}} \leq A, \quad (5.12)$$

where $A \triangleq \mathcal{R}_{1,2}(f_{1,0}, f_{2,0})$. Next, since $(f_{1,k}, f_{2,k})_{k \in \mathbb{N}} \in \mathcal{C}_f$, we have that

$$\forall k \in \mathbb{N} : f = L_{\tilde{\phi}_1}^{-1}\{w_{1,k}\} + L_{\tilde{\phi}_2}^{-1}\{w_{2,k}\} + \sum_{i=0}^2 \mathbf{c}_{i,k}^T \mathbf{p}_i. \quad (5.13)$$

Therefore, for all test functions $\varphi \in \mathcal{S}(\mathbb{R})$, we have that

$$\langle f, \varphi \rangle = \left\langle w_{1,k}, L_{\tilde{\phi}_1}^{-1*}\{\varphi\} \right\rangle + \left\langle w_{2,k}, L_{\tilde{\phi}_2}^{-1*}\{\varphi\} \right\rangle + \left\langle \sum_{i=0}^2 \mathbf{c}_{i,k}^T \mathbf{p}_i, \varphi \right\rangle, \quad (5.14)$$

where $L_{\tilde{\phi}_1}^{-1*}$ is the adjoint operator of $L_{\tilde{\phi}_1}^{-1}$. Let $B_k \triangleq \max_i(\|\mathbf{c}_{i,k}\|_\infty)$, and

$$\mathbf{b}_{i,k} \triangleq \begin{cases} \frac{\mathbf{c}_{i,k}}{B_k}, & B_k \neq 0, \\ \mathbf{1}, & \text{otherwise.} \end{cases} \quad (5.15)$$

Clearly, the vectors $\mathbf{b}_{i,k}$ are in the unit ball. Moreover, as stated in (5.12), $w_{i,k}$ is in the ball of radius A , which is weak*-compact according to the Banach-Alaoglu theorem [15, Theorem 3.15]. Therefore, there exists a subsequence $(f_{1,k_n}, f_{2,k_n})_{n \in \mathbb{N}}$ with the following properties

- the sequences $(w_{1,k_n})_{n \in \mathbb{N}}$ and $(w_{2,k_n})_{n \in \mathbb{N}}$ are converging for the weak*-topology. Their limits are denoted by $w_{1,\text{lim}}, w_{2,\text{lim}} \in \mathcal{M}(\mathbb{R})$, respectively;
- there exist vectors $\mathbf{b}_{i,\text{lim}}$ such that $\lim_{n \rightarrow +\infty} \mathbf{b}_{i,k_n} = \mathbf{b}_{i,\text{lim}}$ for $i \in \{0, 1, 2\}$, at least one of which is nonzero since, for any $k \in \mathbb{N}$, there exists an $i \in \{0, 1, 2\}$ such that $\mathbf{b}_{i,k} = \mathbf{1}$;
- we have the convergence $\lim_{n \rightarrow +\infty} B_{k_n} = B_{\text{lim}}$ where $B_{\text{lim}} \in \mathbb{R} \cup \{+\infty\}$.

Rewriting (5.14) for $k = k_n$ yields

$$\langle f, \varphi \rangle - \left\langle w_{1,k_n}, L_{\tilde{\phi}_1}^{-1*}\{\varphi\} \right\rangle - \left\langle w_{2,k_n}, L_{\tilde{\phi}_2}^{-1*}\{\varphi\} \right\rangle = B_{k_n} \left\langle \sum_{i=0}^2 \mathbf{b}_{i,k_n}^T \mathbf{p}_i, \varphi \right\rangle. \quad (5.16)$$

Assume by contradiction that $B_{\text{lim}} = +\infty$. As the left-hand side of the equality in (5.16) converges to a finite limit, so must the right-hand side, which implies that $\lim_{n \rightarrow +\infty} \langle \sum_{i=0}^2 \mathbf{b}_{i,k_n}^T \mathbf{p}_i, \varphi \rangle = 0$. Since $\lim_{n \rightarrow +\infty} \langle \sum_{i=0}^2 \mathbf{b}_{i,k_n}^T \mathbf{p}_i, \varphi \rangle = \langle q_{\text{lim}}, \varphi \rangle$ for all $\varphi \in \mathcal{S}(\mathbb{R})$, where $q_{\text{lim}} = \sum_{i=0}^2 \mathbf{b}_{i,\text{lim}}^T \mathbf{p}_i$, it follows that $q_{\text{lim}} = 0$. Yet, $\{\mathbf{p}_0, \mathbf{p}_1, \mathbf{p}_2\}$ are linearly independent while not all $\mathbf{b}_{i,\text{lim}}$ are zero, which yields a contradiction. Therefore, we must have that $B_{\text{lim}} < +\infty$.

Let $f_1 = L_{\tilde{\phi}_1}^{-1}\{w_{1,\text{lim}}\} + \mathbf{c}_{1,\text{lim}}^T \mathbf{p}_1$ and $f_{2,\text{lim}} = L_{\tilde{\phi}_2}^{-1}\{w_{2,\text{lim}}\} + \mathbf{c}_{0,\text{lim}}^T \mathbf{p}_0 + \mathbf{c}_{2,\text{lim}}^T \mathbf{p}_2$, where

$\mathbf{c}_{i,\text{lim}} = B_{\text{lim}} \mathbf{b}_{i,\text{lim}}$ for $i \in \{0, 1, 2\}$. By taking the limit in (5.16), we get that

$$\langle f, \varphi \rangle - \left\langle w_{1,\text{lim}}, L_{\tilde{\phi}_1}^{-1*} \{\varphi\} \right\rangle - \left\langle w_{2,\text{lim}}, L_{\tilde{\phi}_2}^{-1*} \{\varphi\} \right\rangle = \left\langle \sum_{i=0}^2 \mathbf{c}_{i,\text{lim}}^T \mathbf{p}_i, \varphi \right\rangle. \quad (5.17)$$

Since (5.17) is valid for all $\varphi \in \mathcal{S}(\mathbb{R})$, we have that $f = f_{1,\text{lim}} + f_{2,\text{lim}}$ and $\mathcal{R}_{\text{hyb}}(f) = \mathcal{R}_{1,2}(f_{1,\text{lim}}, f_{2,\text{lim}})$.

The regularizer \mathcal{R}_{hyb} is convex. Consider two functions $f, g \in \mathcal{M}_{L_1, \phi_0}(\mathbb{R}) + \mathcal{M}_{L_2}(\mathbb{R})$ and let $f = f_1 + f_2$ and $g = g_1 + g_2$ be their decomposition, as specified by (5.7). Denote $h_\beta = \beta f + (1 - \beta)g$ as a convex combination of f and g , where $\beta \in [0, 1]$. Since $h_\beta = (\beta f_1 + (1 - \beta)g_1) + (\beta f_2 + (1 - \beta)g_2)$, we have that

$$\begin{aligned} \mathcal{R}_{\text{hyb}}(h_\beta) &\leq (1 - \alpha) \|\beta f_1 + (1 - \beta)g_1\|_{\mathcal{M}} + \alpha \|\beta f_2 + (1 - \beta)g_2\|_{\mathcal{M}} \\ &\leq (1 - \alpha)(\beta \|f_1\|_{\mathcal{M}} + (1 - \beta) \|g_1\|_{\mathcal{M}}) + \alpha(\beta \|f_2\|_{\mathcal{M}} + (1 - \beta) \|g_2\|_{\mathcal{M}}) \\ &= \beta \mathcal{R}_{\text{hyb}}(f) + (1 - \beta) \mathcal{R}_{\text{hyb}}(g). \end{aligned} \quad (5.18)$$

□

5.2.2 Representer Theorem

We now have the necessary tools to present the main theoretical result of this chapter, on which our implementation is based.

Theorem 5.1 (Continuous-Domain Representer Theorem). *Let L_1, L_2 be derivative operators of the form 2.19, and let $\boldsymbol{\nu} = (\nu_1, \dots, \nu_M)$ be a linear measurement operator composed of the M linear functionals $\nu_m : f \mapsto \nu_m(f) \in \mathbb{R}$ which are weak*-continuous over both $\mathcal{M}_{L_1}(\mathbb{R})$ and $\mathcal{M}_{L_2}(\mathbb{R})$. Assume that $\mathcal{N}_{\boldsymbol{\nu}} \cap (\mathcal{N}_{L_1} + \mathcal{N}_{L_2}) = \{0\}$, where $\mathcal{N}_{\boldsymbol{\nu}}$ is the null space of $\boldsymbol{\nu}$ (well-posedness assumption). Then, the linear inverse problem*

$$\mathcal{V} \triangleq \arg \min_{f \in \mathcal{M}_{L_1, \phi_0}(\mathbb{R}) + \mathcal{M}_{L_2}(\mathbb{R})} \underbrace{\left(\frac{1}{2} \|\boldsymbol{\nu}(f) - \mathbf{y}\|_2^2 + \lambda \mathcal{R}_{\text{hyb}}(f) \right)}_{\mathcal{J}(f)} \quad (5.19)$$

has a solution s of the form $s = s_1 + s_2$, where s_i are nonuniform L_i -splines (in the sense of Definition 3.1) of the form

$$s_i(x) = \sum_{k=1}^{K_i} a_{i,k} \rho_{L_i}(x - x_{i,k}) + q_i(x) \quad (5.20)$$

for $q_i \in \mathcal{N}_{L_i}$, $a_{i,k}, x_{i,k} \in \mathbb{R}$. Moreover, the sparsity indices K_i verify $K_1 + K_2 \leq M$.

The proof of Theorem 5.1 is given in Appendix B.1.

Remark 5.1. *As in Theorem 3.3, we restrict to derivative operators for L_1 and L_2 . However, Theorem 5.1 and our proof technique remain valid for any spline-admissible operators in the sense of [14, Definition 1].*

Theorem 5.1 is a powerful generalization of Theorem 3.3 for hybrid regularization terms. It states that the continuous-domain Problem (5.19) has a hybrid spline solution that consists of the sum of a L_1 -spline and a L_2 -spline. The unknowns are the locations $x_{i,k}$ and amplitudes $a_{i,k}$ of the knots (which specify the innovation of the spline corresponding to the vertical arrows in Figure 3.1), as well as the null-space components q_i . The total sparsity of this solution is bounded by the number of measurements M , meaning that it is characterized by very few parameters. The following observations can be made concerning Theorem 5.1:

- We use the restricted search space $\mathcal{M}_{L_1, \phi_0}(\mathbb{R})$ defined in (5.5) instead of the complete space $\mathcal{M}_{L_1}(\mathbb{R})$ in order to ensure that Problem (5.19) is well-posed². This does not restrict the native space of the reconstructed signal s since $\mathcal{M}_{L_1, \phi_0}(\mathbb{R}) + \mathcal{M}_{L_2}(\mathbb{R}) = \mathcal{M}_{L_1}(\mathbb{R}) + \mathcal{M}_{L_2}(\mathbb{R})$.
- Theorem 5.1 can readily be extended to Q operators L_1, \dots, L_Q : however, for $Q > 2$, the handling of the pairwise null space intersections would make the general formulation more tedious. For the sake of clarity, we therefore only consider the case $Q = 2$.
- A remarkable feature of Theorem 5.1 is that the bound on the sparsity of the solutions does not increase with the number Q of operators. This is particularly appealing from a theoretical point of view since, compared to the single-operator framework of Chapter 4, we essentially enrich our dictionary at no cost in terms of sparsity.

5.3 Exact Discretization

We now detail our discretization method for Problem (5.19), which is based on the methodology of Chapter 4. For all that follows, the conditions of Assumptions 5.1 apply.

Assumptions 5.1.

- *The measurement functionals ν_m are supported in an interval $I_T = [0, T]$ where $T \in h\mathbb{Z}$.*
- *The regularization operators are $L_1 = D^{N_{0,1}}$ and $L_2 = D^{N_{0,2}}$ where $N_{0,1} < N_{0,2}$; this implies that $\mathcal{N}_0 = \mathcal{N}_{L_1}$ and thus that $N_0 = N_{0,1}$.*

²An unbounded solution set \mathcal{V} would arise if we allowed ourselves to add contributions $(p, -p)$ to a solution, with arbitrary $p \in \mathcal{N}_0$. Indeed, the extended solution would still be an element of \mathcal{V} .

- For the space $\mathcal{M}_{L_1, \phi_0}(\mathbb{R})$ defined in (5.5), we choose the implementation-friendly boundary conditions $\phi_0 : \mathcal{M}_{L_1}(\mathbb{R}) \rightarrow \mathbb{R}^{N_0}$ from Proposition 2.5 given by

$$\phi_0 = \begin{cases} \frac{1}{\epsilon} \text{rect}\left(\frac{\cdot}{\epsilon} - \frac{1}{2}\right) & \text{if } N_0 = 1 \\ (\delta, \dots, \delta^{(N_0-2)}, \delta^{(N_0-1)} * \frac{1}{\epsilon} \text{rect}\left(\frac{\cdot}{\epsilon}\right) * \delta(\cdot - \frac{\epsilon}{2})) & \text{if } N_0 > 1, \end{cases} \quad (5.21)$$

for some small $\epsilon > 0$.

The first assumption is natural and is often fulfilled in practice, for instance in imaging with a finite field of view. Since L_1 and L_2 have almost symmetrical roles, the second assumption is only significant in relation to the boundary conditions ϕ_0 . It implies that $\mathcal{N}_1 \subset \mathcal{N}_2$ and $\mathcal{N}_0 = \mathcal{N}_{L_1}$. Finally, our choice of boundary conditions ϕ_0 is guided by computational considerations, which will be detailed in Remark 5.2.

5.3.1 Specification of the Search Spaces

We now briefly recall the necessary tools to perform our exact discretization, which are taken from Section 4.2. Following Proposition 3.3, the space of uniform L_i -splines with grid size $h > 0$ for $i \in \{1, 2\}$ can be expressed in the B-spline basis as

$$\mathcal{M}_{L_i, h}(\mathbb{R}) = \left\{ \sum_{k \in \mathbb{Z}} c[k] \beta_{L_i, h}(\cdot - kh) : c \in \ell_{1, L_i}(\mathbb{Z}) \right\}, \quad (5.22)$$

where $\beta_{L_i, h}$ is the scaled B-spline of L_i defined in (3.13). The analytical expressions of polynomial B-spline for low derivative orders are given in Table 3.1.

$$\ell_{1, L_i}(\mathbb{Z}) = \left\{ c \in \mathcal{S}'(\mathbb{Z}) : (d_{L_i} * c) \in \ell_1(\mathbb{Z}) \right\}, \quad (5.23)$$

where d_{L_i} is the finite-support B-spline filter defined in (3.11). We refer to Sections 3.1.4 and 3.1.5 for more background on polynomial B-splines.

The discretized search space for L_1 with boundary conditions ϕ_0 is given by

$$\mathcal{M}_{L_1, \phi_0, h}(\mathbb{R}) = \left\{ s \in \mathcal{M}_{L_1, h}(\mathbb{R}) : \phi_0(s) = \mathbf{0} \right\} \quad (5.24)$$

$$= \left\{ \sum_{k \in \mathbb{Z}} c[k] \beta_{L_1, h}(\cdot - kh) : c \in \ell_{1, L_1}(\mathbb{Z}), \phi_{0, h}(c) = \mathbf{0} \right\}, \quad (5.25)$$

where $\phi_{0, h} : \ell_{1, L}(\mathbb{Z}) \rightarrow \mathbb{R}^{N_0}$ satisfies $\phi_{0, h}(c) = \phi_0(s)$. We thus define the matching search space for B-spline coefficients as

$$\ell_{1, L_1, \phi_0}(\mathbb{Z}) = \left\{ c \in \ell_{1, L_1}(\mathbb{Z}) : \phi_{0, h}(c) = \mathbf{0} \right\}. \quad (5.26)$$

Remark 5.2. *The reason why the boundary conditions ϕ_0 in (5.21) are implementation-friendly is that for a choice of ϵ such that $\frac{\epsilon}{h}$ is arbitrarily small, they are numerically equivalent to $\phi_0(f) = (f(0), \dots, f^{(N_0-1)}(0^+))$, where $f^{(N_0-1)}(0^+)$ is the right limit of $f^{(N_0-1)}$ at 0. It can easily be shown in this case that, for $s = \sum_{k \in \mathbb{Z}} c[k] \beta_{L_1, h}(\cdot - kh) \in \mathcal{M}_{L_1, \phi_0, h}(\mathbb{R})$, we have*

$$\phi_0(s) = \phi_{0, h}(c) = \mathbf{0} \Leftrightarrow c[-N_0 + 1] = \dots = c[0] = 0. \quad (5.27)$$

This choice simplifies the upcoming optimization task by reducing the number of coefficients $c[k]$ to be optimized, since some are set to zero. Other boundary conditions would lead to more complicated linear constraints over the coefficients $c[k]$ that would make the optimization task more difficult.

5.3.2 Discrete Problem Formulation

Next, we introduce the sets of indices $I_i = \{m_i, \dots, M_i\}$ of cardinality $N_i = \#I_i = M_i - m_i + 1$ such that $k \in I_i$ if and only if the support of the B-spline atom $\beta_{L_i}(\cdot - kh)$ and I_T have a nontrivial intersection. The indices $\{-N_{0,1} + 1, \dots, 0\}$ are excluded from I_1 due to the boundary conditions (5.27). We thus have that $m_1 = 1$, $m_2 = (-N_{0,1} + 1)$, and $M_1 = M_2 = (T/h - 1)$, which implies that $N_1 = (T/h - 1)$ and $N_2 = (T/h + N_{0,2} - 1)$. By optimizing over the B-spline coefficients in I_i , we get the following finite optimization problem

$$\mathcal{V}_f \triangleq \arg \min_{(\mathbf{c}_1, \mathbf{c}_2) \in \mathbb{R}^{N_1} \times \mathbb{R}^{N_2}} \mathcal{J}_f(\mathbf{c}_1, \mathbf{c}_2) \quad \text{where} \quad (5.28)$$

$$\mathcal{J}_f(\mathbf{c}_1, \mathbf{c}_2) \triangleq \frac{1}{2} \|\mathbf{H}_1 \mathbf{c}_1 + \mathbf{H}_2 \mathbf{c}_2 - \mathbf{y}\|_2^2 + \lambda \left((1 - \alpha) \|\mathbf{L}_1 \mathbf{c}_1\|_1 + \alpha \|\mathbf{L}_2 \mathbf{c}_2\|_1 \right). \quad (5.29)$$

The system matrices $\mathbf{H}_i \in \mathbb{R}^{M \times N_i}$ are defined as

$$\mathbf{H}_i \triangleq \begin{bmatrix} \mathbf{h}_{m_i} & \dots & \mathbf{h}_{M_i} \end{bmatrix} : \quad \mathbf{h}_k \triangleq \boldsymbol{\nu}(\beta_{L_i}(\cdot - kh)), \quad (5.30)$$

and the Toeplitz-like regularization matrices are given by $\mathbf{L}_1 \in \mathbb{R}^{N_1 \times N_1}$, with

$$\mathbf{L}_1 \triangleq \frac{1}{h^{N_{0,1}-1}} \begin{pmatrix} d_{L_1}[0] & 0 & \dots & \dots & \dots & 0 \\ \vdots & \ddots & \ddots & & & \vdots \\ d_{L_1}[N_{0,1}] & & \ddots & \ddots & & \vdots \\ 0 & \ddots & & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & & \ddots & 0 \\ 0 & \dots & 0 & d_{L_1}[N_{0,1}] & \dots & d_{L_1}[0] \end{pmatrix}, \quad (5.31)$$

and $\mathbf{L}_2 \in \mathbb{R}^{(N_2-N_{0,2}) \times N_2}$, with

$$\mathbf{L}_2 \triangleq \frac{1}{h^{N_{0,2}-1}} \begin{pmatrix} d_{L_2}[N_{0,2}] & \cdots & d_{L_2}[0] & 0 & \cdots & 0 \\ 0 & \ddots & & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & & \ddots & 0 \\ 0 & \cdots & 0 & d_{L_2}[N_{0,2}] & \cdots & d_{L_2}[0] \end{pmatrix}. \quad (5.32)$$

The cost function can thus be rewritten as

$$\mathcal{J}_f(\mathbf{c}_1, \mathbf{c}_2) = \frac{1}{2} \left\| \mathbf{H} \begin{pmatrix} \mathbf{c}_1 \\ \mathbf{c}_2 \end{pmatrix} - \mathbf{y} \right\|_2^2 + \lambda \left\| \mathbf{L} \begin{pmatrix} \mathbf{c}_1 \\ \mathbf{c}_2 \end{pmatrix} \right\|_1, \quad (5.33)$$

where the concatenated system and regularization matrices are

$$\mathbf{H} \triangleq \begin{pmatrix} \mathbf{H}_1 & \mathbf{H}_2 \end{pmatrix} \in \mathbb{R}^{M \times (N_1+N_2)} \quad \text{and} \quad (5.34)$$

$$\mathbf{L} \triangleq \begin{pmatrix} (1-\alpha)\mathbf{L}_1 & \mathbf{0} \\ \mathbf{0} & \alpha\mathbf{L}_2 \end{pmatrix} \in \mathbb{R}^{(N_1+N_2-N_{0,2}) \times (N_1+N_2)}, \quad (5.35)$$

respectively. Therefore, Problem (5.28) is a standard penalized basis pursuit (PBP) problem which can be solved using off-the-shelf algorithms such as ADMM (see Appendix A). Furthermore, it satisfies the property of Theorem 5.2.

Theorem 5.2. *Under the assumptions of Theorem 5.1 and Assumptions 5.1, the following conditions hold true.*

- We have that $\ker \mathbf{H} \cap \ker \mathbf{L} = \{\mathbf{0}\}$, and the solution set \mathcal{V}_f of Problem (5.28) is a nonempty compact convex set.
- Problem (5.28) is truly equivalent to the continuous Problem (5.19) in the discretized search spaces $\mathcal{M}_{L_1, h, \phi_0}(\mathbb{R})$ and $\mathcal{M}_{L_2, h}(\mathbb{R})$ defined in (5.22) and (5.24), respectively, with

$$\mathcal{V}_{\text{res}} \triangleq \arg \min_{f \in \mathcal{M}_{L_1, h, \phi_0}(\mathbb{R}) + \mathcal{M}_{L_2, h}(\mathbb{R})} \mathcal{J}(f), \quad (5.36)$$

in the sense that there exists a bijective linear mapping between the solution sets.

- The reconstructed signal associated to a solution $(\mathbf{c}_1, \mathbf{c}_2) \in \mathcal{V}_f$ of Problem (5.28) is a hybrid spline

$$s = s_1 + s_2 = \sum_{k \in \mathbb{Z}} \left(c_1[k] \beta_{L_1}(\cdot - kh) + c_2[k] \beta_{L_2}(\cdot - kh) \right), \quad (5.37)$$

where the sequences $(c_1, c_2) \in \ell_{1, L_1, \phi_0}(\mathbb{Z}) \times \ell_{1, L_2}(\mathbb{Z})$ satisfy $c_i|_{I_i} = (c_i[m_i], \dots, c_i[M_i]) = \mathbf{c}_i$;

- The sparsity of s_i in the Green's function basis is given by $\|\mathbf{L}_i \mathbf{c}_i\|_0$.

Proof. We first prove the following lemma.

Lemma 5.1. *Let $s^* = s_1^* + s_2^* \in \mathcal{V}_{\text{res}}$ be a solution of Problem (5.36) and let $(c_1^*, c_2^*) \in \ell_{1, L_1, \phi_0}(\mathbb{Z}) \times \ell_{1, L_2}(\mathbb{Z})$ be the corresponding sequences of B-spline coefficients. Under Assumptions 5.1, c_i^* is uniquely determined by its N_i coefficients $c^*|_{i_i}$.*

Proof. Let us first observe that Problem (5.36) can be rewritten by using the B-spline representations of $\mathcal{M}_{L_1, h, \phi_0}(\mathbb{R})$ and $\mathcal{M}_{L_2, h}(\mathbb{R})$ as in (5.22). By optimizing over the spline coefficients, we get the optimization problem

$$\mathcal{V}_d \triangleq \arg \min_{(c_1, c_2) \in \ell_{1, L_1, \phi_0}(\mathbb{Z}) \times \ell_{1, L_2}(\mathbb{Z})} \underbrace{\mathcal{J} \left(\sum_{k \in \mathbb{Z}} c_1[k] \beta_{L_1}(\cdot - kh), \sum_{k \in \mathbb{Z}} c_2[k] \beta_{L_2}(\cdot - kh) \right)}_{\mathcal{J}_d(c_1, c_2)}, \quad (5.38)$$

where the reconstructed signal for $(c_1^*, c_2^*) \in \mathcal{V}_d$ is given by $s^* = s_1^* + s_2^*$, whose components s_i^* are L_i -splines with B-spline coefficients c_i^* . Using (3.18), the regularization term associated to a signal $s_i \in \mathcal{M}_{L_i}(\mathbb{R})$ with B-spline coefficients $c_i \in \ell_{1, L_i}(\mathbb{Z})$ can be expressed as $\|\mathbf{L}_i \{s_i\}\|_{\mathcal{M}} = \frac{1}{h^{N_{0,i}-1}} \|d_{L_i} * c_i\|_1$. Next, consider a sequence c_1 such that $c_1|_{i_1} = c_1^*|_{i_1}$, $c_1[-N_{0,1} + 1] = \dots = c_1[0] = 0$ (boundary conditions (5.27)), and whose remaining coefficients are free. This freedom does not affect the data fidelity term due to the finite-support assumption on ν_m . The coefficient $c_1[M_1 + 1]$ is uniquely chosen such that $(d_{L_1} * c_1)[M_1 + 1] = \sum_{k=0}^{N_{0,1}} d_{L_1}[k] c_1[M_1 - k + 1] = 0$. Similarly, all $c[k]$ coefficients for $k > M_1$ can be uniquely determined recursively to nullify $(d_{L_1} * c_1)[k]$ as a linear combination of the $(N_{0,1} - 1)$ previous coefficients.

Next, due to the boundary conditions, we have that $(d_{L_1} * c_1)[0] = \sum_{k=0}^{N_{0,1}} d_{L_1}[k] c_1[-k] = 0 \Rightarrow c[-N_{0,1}] = 0$. Analogously, the unique way of canceling all coefficients $(d_{L_1} * c_1)[k]$ for $k \leq 0$ is to set $c_1[k] = 0$ for all $k \leq (-N_{0,1})$. By construction, this sequence c_1 yields a regularization cost no greater than that of c_1^* . Indeed, we have that $(d_{L_1} * c_1^*)[k] = (d_{L_1} * c_1)[k]$ for $1 \leq k \leq M_1$ and that $(d_{L_1} * c_1)[k] = 0$ otherwise. Since the two sequences c_1 and c_1^* yield the same measurements, we have that $\mathcal{J}_d(c_1, c_2^*) \leq \mathcal{J}_d(c_1^*, c_2^*)$. Yet, since (s_1^*, s_2^*) is a solution of (5.36) and the construction of c_1 is unique, we necessarily have that $c_1 = c_1^*$.

Finally, we construct a sequence c_2 in a similar fashion. The only difference is the absence of boundary conditions, which means that the coefficients $c_2[k]$ for $k < m_2$ are determined as a linear combination of the $(N_{0,2} - 1)$ next coefficients to cancel out $(d_{L_2} * c_2)[k]$ for all $k \leq 0$. The same argument as for c_1 can then be used to prove that $c_2 = c_2^*$. \square

The proof of Lemma 5.1 details the construction of two injective linear maps $\tilde{\theta}_1 : \mathbb{R}^{N_1} \rightarrow \ell_{1,L_1,\phi_0}(\mathbb{Z})$ and $\tilde{\theta}_2 : \mathbb{R}^{N_2} \rightarrow \ell_{1,L_2}(\mathbb{Z})$ such that, for any $\mathbf{c}_i \in \mathbb{R}^{N_i}$, $\tilde{\theta}_i(\mathbf{c}_i)|_{I_i} = \mathbf{c}_i$ and $\phi_{0,h}(\tilde{\theta}_1(\mathbf{c}_1)) = \mathbf{0}$. We then define the mapping $\tilde{\theta} : \mathbb{R}^{N_1} \times \mathbb{R}^{N_2} : (\mathbf{c}_1, \mathbf{c}_2) \mapsto (\tilde{\theta}_1(\mathbf{c}_1), \tilde{\theta}_2(\mathbf{c}_2))$. Let $(\mathbf{c}_1, \mathbf{c}_2) \in \mathbb{R}^{N_1 \times N_2}$ and consider $\tilde{\theta}(\mathbf{c}_1, \mathbf{c}_2) \in \ell_{1,L_1}(\mathbb{Z}) \times \ell_{1,L_2}(\mathbb{Z})$. Following the proof of Lemma 5.1, $\mathcal{J}_d(\tilde{\theta}(\mathbf{c}_1, \mathbf{c}_2))$ can be computed using only the $N_1 + N_2$ coefficients \mathbf{c}_1 and \mathbf{c}_2 . Indeed, all other coefficients $(\tilde{\theta}_i(\mathbf{c}_i)[k])_{k \notin I_i}$ do not affect the data fidelity term and cancel out all the regularization terms which they affect. This implies that $\mathcal{J}_d(\tilde{\theta}(\mathbf{c}_1, \mathbf{c}_2)) = \mathcal{J}_f(\mathbf{c}_1, \mathbf{c}_2) = \frac{1}{2} \left\| \mathbf{H} \begin{pmatrix} \mathbf{c}_1 \\ \mathbf{c}_2 \end{pmatrix} - \mathbf{y} \right\|_2^2 + \lambda \left\| \mathbf{L} \begin{pmatrix} \mathbf{c}_1 \\ \mathbf{c}_2 \end{pmatrix} \right\|_1$, where \mathbf{H} and \mathbf{L} are defined as in (4.11) and (4.12), respectively. Since, by Lemma 5.1, $\mathcal{V}_d \subset \tilde{\theta}(\mathbb{R}^{N_1} \times \mathbb{R}^{N_2})$, Problems (5.38) and (5.28) are equivalent in the sense that $\tilde{\theta}(\mathcal{V}_f) = \mathcal{V}_d$, and the restriction $\theta = \tilde{\theta}|_{\mathcal{V}_f} : \mathcal{V}_f \rightarrow \mathcal{V}_d$ is a bijective linear map.

Concerning the first item of Theorem 5.2, let $\begin{pmatrix} \mathbf{c}_1 \\ \mathbf{c}_2 \end{pmatrix} \in \ker \mathbf{H} \cap \ker \mathbf{L}$. Since \mathbf{L}_1 is invertible (5.31), we have that $\mathbf{c}_1 = \mathbf{0}$. Hence, the continuous-domain signal $s = \sum_{k \in \mathbb{Z}} \tilde{\theta}_2(\mathbf{c}_2)[k] \beta_{L_2}(\cdot - kh)$ verifies $s \in \mathcal{N}_\nu \cap \mathcal{N}_{L_2} = \{0\}$ (well-posedness assumption in Theorem 5.1). We thus have that $\tilde{\theta}_2(\mathbf{c}_2) = 0$, which yields $\mathbf{c}_2 = 0$. Therefore, $\ker \mathbf{H} \cap \ker \mathbf{L} = \{0\}$, which implies that Problem (5.28) is well-posed and that its solution set \mathcal{V}_f is a nonempty compact set. It is also convex due to the convexity of the cost function \mathcal{J}_f . \square

The proof of Theorem 5.2 is a multicomponent extension of Proposition 4.1. The main difference is to be found in the boundary conditions. Note that, by definition of I_i , the B-spline coefficients outside I_i do not affect the reconstructed signal outside the interval of interest I_T . Therefore, by Theorem 5.2, the expression of s over I_T and its sparsity is entirely determined by the coefficients \mathbf{c}_i . Computing the complete sequences of B-spline coefficients c_i is therefore unnecessary.

5.3.3 Reaching a Sparse Solution

Since by (5.33), Problem (5.28) can be rewritten as a standard PBP problem of the form

$$\mathcal{V}_f \triangleq \arg \min_{\mathbf{c} \in \mathbb{R}^N} \left(\frac{1}{2} \|\mathbf{H}\mathbf{c} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{L}\mathbf{c}\|_1 \right), \quad (5.39)$$

where $N = N_1 + N_2$, all the results of Section 4.3.3 are applicable. Since \mathbf{L} is clearly of full rank, its null space has dimension $N_{0,2}$, which therefore plays the role of N_0 in Chapter 4. We now sum up the properties of Problem (5.28).

- By Theorem 4.1, its solution set \mathcal{V}_f is a compact convex set whose extreme points $\mathbf{c}^* = (\mathbf{c}_1^*, \mathbf{c}_2^*)$ are sparse, in the sense that they verify $\|\mathbf{L}\mathbf{c}^*\|_0 \leq (M - N_{0,2})$.

- By Lemma 4.2, all its solutions yield the same measurements \mathbf{y}_λ , so that $\forall \mathbf{c} \in \mathcal{V}_f$, $\mathbf{H}\mathbf{c} = \mathbf{y}_\lambda$.
- By Proposition 4.2, it is equivalent to the linear program

$$\begin{aligned} \mathcal{V}_f^{\text{LP}} \triangleq & \arg \min_{(\mathbf{c}, \mathbf{u}) \in \mathbb{R}^N \times \mathbb{R}^{(N-N_0, 2)}} \left(\sum_{n=1}^{N-N_0, 2} u_n \right) \\ \text{s.t.} \quad & \mathbf{u} + \mathbf{L}\mathbf{c} \geq \mathbf{0}, \quad \mathbf{u} - \mathbf{L}\mathbf{c} \geq \mathbf{0}, \quad \mathbf{H}\mathbf{c} = \mathbf{y}_\lambda, \end{aligned} \quad (5.40)$$

where for any $\mathbf{x}, \mathbf{y} \in \mathbb{R}^P$, $\mathbf{x} \leq \mathbf{y}$ implies that $x_p \leq y_p$ for all $p \in \{1, \dots, P\}$. Moreover, an extreme point $(\mathbf{c}^*, \mathbf{u}^*)$ of $\mathcal{V}_f^{\text{LP}}$ yields an extreme point \mathbf{c}^* of \mathcal{V}_f .

The same pipeline can also be used to compute a sparse extreme point of \mathcal{V}_f . The procedure is as follows:

1. Run any iterative solver (e.g., ADMM) to find a solution $\mathbf{c}_{\text{ADMM}} \in \mathcal{V}_f$ of Problem (5.28).
2. Compute $\mathbf{y}_\lambda = \mathbf{H}\mathbf{c}_{\text{ADMM}}$.
3. Solve the linear program (5.40) using the simplex or dual-simplex algorithm [116, 117]. These algorithms are known to converge to an extreme point $(\mathbf{c}^*, \mathbf{u}^*)$ of $\mathcal{V}_f^{\text{LP}}$, where \mathbf{c}^* is an extreme point of \mathcal{V}_f .

5.4 Multiresolution Strategy

Although the algorithm introduced in Section 5.3.3 yields a sparse solution of Problem (5.28), it does so for a fixed grid size h . Yet, choosing a suitable grid size clearly depends on the problem at hand. We therefore propose a multiresolution strategy that selects the grid size automatically. We introduce the optimal costs of the continuous and discrete problems, which are defined as $\mathcal{J}^0 \triangleq \min_{f \in \mathcal{M}_{L_1, \phi_0}(\mathbb{R}) + \mathcal{M}_{L_2}(\mathbb{R})} \mathcal{J}(f)$ and $\mathcal{J}_f^0(h) \triangleq \min_{(\mathbf{c}_1, \mathbf{c}_2) \in \mathbb{R}^{N_1} \times \mathbb{R}^{N_2}} \mathcal{J}_f(\mathbf{c}_1, \mathbf{c}_2)$, respectively. Our approach is based on Theorem 5.3.

Theorem 5.3 (Convergence of the Cost Function of the Discrete Problem). *Assume that the hypotheses of Theorem 5.1 and Assumptions 5.1 are met. Then,*

$$\lim_{h \rightarrow 0} \mathcal{J}_f^0(h) = \mathcal{J}^0. \quad (5.41)$$

Proof. Let $s = s_1 + s_2 \in \mathcal{V}$, where s_1 and s_2 are of the form (5.20). Assume by contradiction that there exists a knot $x_{1, k_0} \notin I_T$. We can then construct a spline \tilde{s}_1 which has the same expression as s_1 in I_T and the same knot locations, except for the knot at x_{1, k_0} which is removed. Since the functionals ν_m are supported in I_T , we have that

$\nu(s_1) = \nu(\tilde{s}_1)$ and that $\|\mathbf{L}_1\{s_1\}\|_{\mathcal{M}} = \|\mathbf{L}_1\{\tilde{s}_1\}\|_{\mathcal{M}} + |a_{1,k_0}| > \|\mathbf{L}_1\{\tilde{s}_1\}\|_{\mathcal{M}}$. This implies that $\mathcal{J}(s) > \mathcal{J}(\tilde{s}_1 + s_2)$, which contradicts $s \in \mathcal{V}$. Therefore, we have that $x_{1,k} \geq 0$ for all k ; we can thus construct the spline

$$s_{1,h}(x) = q_1(x) + \sum_{k=1}^{K_1} a_{1,k} \rho_{L_1}(x - x_{1,k}^h), \quad (5.42)$$

where $x_{1,k}^h \in h\mathbb{Z}$ and such that $x_{1,k}^h \geq 0$ converges to $x_{1,k}$ when h goes to zero—we choose $x_{1,k}^h = 0$ if $x_{1,k} = 0$. Since this spline has the same expression as s_1 in the neighborhood of zero, we have that $\phi_0(s_1) = \mathbf{0} \Rightarrow \phi_0(s_{1,h}) = \mathbf{0}$ and, thus, that $s_{1,h} \in \mathcal{M}_{L_1,h,\phi_0}(\mathbb{R})$. Similarly, we construct $s_{2,h} \in \mathcal{M}_{L_2,h}(\mathbb{R})$. The desired result then follows from the proof of Theorem 4.2. \square

Theorem 5.3 justifies our choice of the search space (5.22) since it contains functions which yield costs arbitrarily close to the optimal cost \mathcal{J}^0 of the continuous-domain Problem (5.19). It also justifies the use of a multiresolution algorithm based on Algorithm 4.1. We provide the pseudocode of its hybrid-spline version in Algorithm 5.1.

Input: $\nu, N_{0,1}, N_{0,2}, T, \mathbf{y}, \lambda, \alpha, n_0, \epsilon$
 $n = n_0; \mathbf{c}_1 = \mathbf{0}; \mathbf{c}_2 = \mathbf{0}; cost_p = +\infty; error = \epsilon + 1$ **while** $error > \epsilon$ **do**
 $h = T/2^n$
 update \mathbf{H}, \mathbf{L} ; // Depend on $h, \nu, N_{0,1}, N_{0,2}, T$
 $(\mathbf{c}_1, \mathbf{c}_2) \leftarrow \text{ADMM}((\mathbf{c}_{1,\uparrow 2}, \mathbf{c}_{2,\uparrow 2}); \mathbf{H}, \mathbf{L}, \mathbf{y}, \lambda)$
 $error = |\mathcal{J}_f(\mathbf{c}_1, \mathbf{c}_2) - cost_p|/cost_p$
 $cost_p = \mathcal{J}_f(\mathbf{c}_1, \mathbf{c}_2)$
 $n \leftarrow n + 1$
end
 $\mathbf{y}_\lambda = \mathbf{H} \begin{pmatrix} \mathbf{c}_1 \\ \mathbf{c}_2 \end{pmatrix}$
return $(\mathbf{c}_1^*, \mathbf{c}_2^*) = \text{Simplex}(\mathbf{H}, \mathbf{L}, \mathbf{y}, \lambda, \mathbf{y}_\lambda)$

Algorithm 5.1: Pseudocode of our algorithm.

The principle of this algorithm is to split the grid in half by taking grid sizes $(h_n = T/2^n)_{n \geq n_0}$, where h_{n_0} is an initial coarse grid size. We stop the refinement as soon as the optimal cost $\mathcal{J}_f^0(h_n)$ ceases to decrease within a certain tolerance ϵ chosen by the user. Further refinement of the grid is then useless, since the reconstructed signal does not vary significantly anymore. The variable $cost_p$ denotes the cost $\mathcal{J}_f^0(h_{n-1})$ from the previous grid size. Because the search spaces are embedded like $\mathcal{M}_{L_1,h_n,\phi_0}(\mathbb{R}) \subset \mathcal{M}_{L_1,h_{n+1},\phi_0}(\mathbb{R})$ and $\mathcal{M}_{L_2,h_n}(\mathbb{R}) \subset \mathcal{M}_{L_2,h_{n+1}}(\mathbb{R})$, a decrease in the grid size can only improve the reconstruction in terms of cost, *i.e.*, $(\mathcal{J}_f^0(h_n))_{n \geq n_0}$ is decreasing. Moreover, Theorem 5.3 states that this quantity converges to \mathcal{J}^0 , which guarantees the convergence of our algorithm. This embedding also allows us to use the reconstruction from the previous grid size as a warm start for ADMM. This is done by converting the B-spline coefficients to $\mathbf{c}_{i,\uparrow 2}$ for $i \in \{1, 2\}$,

which represent the same continuous-domain signal on the finer grid. In practice, for a tolerance $\epsilon = 10^{-3}$, Algorithm 5.1 typically converges for a grid size of $h = 1/2^9$ or coarser in our experiments, which leads to reasonable computation times of a few seconds in the conditions of Section 5.5.

5.5 Experimental Results

Our algorithms are implemented using GlobalBioIm [118], a Matlab inverse-problem library developed in our group, as well as the Gurobi optimizer³ for the simplex algorithm. In order for our hybrid regularization (5.6) to be an adequate prior, the ground-truth signal in our experiments is a sparse hybrid spline $s = s_1 + s_2$. Each component s_i is a sparse L_i -spline of the form

$$s_i(x) = \sum_{k=1}^{K_{s_i}} a_{i,k} \rho_{L_i}(x - x_{i,k}) + \sum_{n=1}^{N_{0,i}} b_{i,n} p_{i,n}(x), \quad (5.43)$$

where $\{p_{i,n}\}_{n=1}^{N_{0,i}}$ form a basis of \mathcal{N}_{L_i} . The sparsity K_{s_i} of s_i is chosen by the user, the knot locations $x_{i,k} \in I_T$ are chosen at random and the coefficients $a_{i,k}$ and $b_{i,n}$ are i.i.d. Gaussian random variables. The tuning parameters λ and α are selected using grid search, and we pick a tolerance $\epsilon = 10^{-3}$.

5.5.1 Curve Fitting

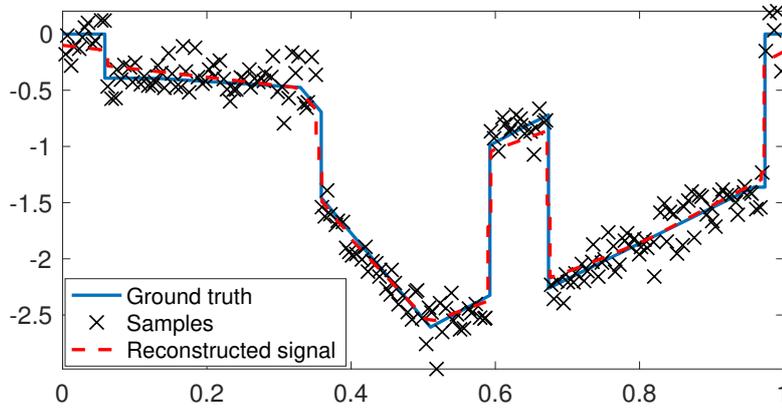
Curve fitting is particularly well-suited for smoothing problems, which consist in fitting a continuous-domain function which is sparse in a certain dictionary basis from many noisy data points. The measurement functionals are then given by $\nu_m(s) = s(x_m)$ where $x_m \in [0, T]$ are the sampling locations. As discussed in Example 3.2, this choice is technically only admissible—in the sense that ν_m is weak*-continuous and thus satisfies the assumptions of Theorem 5.1—when $N_{0,i} \geq 2$ for $i \in \{1, 2\}$. However, even when $N_{0,1} = 1$, one can choose a Dirac mollifier such as (2.48), which leads numerically to taking the right limit $\nu_m(s) = s(x_m^+)$.

Such reconstruction problems are commonly tackled in a single-operator framework, where the dictionary is associated with a single brand of splines. A limit case is polynomial regression, which is achieved by taking $L = D^{N_0}$ and picking a very high value of the regularization parameter λ .

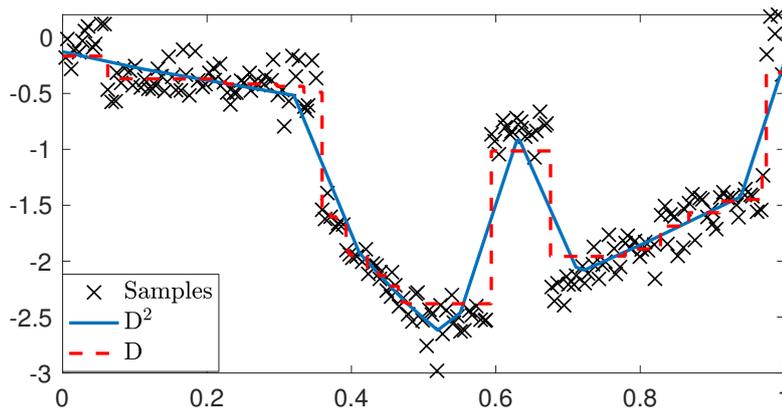
In contrast with standard single-operator frameworks, our approach allows for the joint use of several families of basis functions, and can therefore represent a richer class of signals. An example of a curve-fitting reconstruction using Algorithm 5.1 is given in

³LLC Gurobi Optimization, Gurobi optimizer reference manual, 2018.

Figure 5.1a. The chosen regularization operators are $L_1 = D$ and $L_2 = D^2$; our dictionary thus consists of both piecewise-constant and piecewise-linear splines. We compare our results to single-operator reconstructions with regularization operators $L = D$ and $L = D^2$ in Figure 5.1b.



(a) Hybrid spline reconstruction with $L_1 = D$, $L_2 = D^2$, $\lambda = 2$, $\alpha = 0.04$ and grid size $h = 1/2^9$. Sparsity of the reconstruction: $14 + 6 = 20$.



(b) Single-operator reconstructions $L = D$ ($\lambda = 3$, sparsity 26) and $L = D^2$ ($\lambda = 0.1$, sparsity 15).

Figure 5.1 – Curve fitting comparison with $M = 200$ data points.

Observe that the reconstructed signal is quite sparse ($K_1 = 14$ and $K_2 = 6$), and is satisfactory in that it is close to what a human would reconstruct, except for some shearing effects around $x = 0$ and $x = 0.6$ that are typical of regularization methods. It was not obvious *a priori* that such results could be achieved. Indeed, although we are using the sparsity-promoting ℓ_1 norm, the large number of measurements M leads us far from a CS-type framework. Theorem 4.1 states that the sparsity of the reconstructed signal is bounded by $(M - N_{0,2})$. We observe experimentally that increasing λ tends to produce sparser solutions, and that the level of sparsity can be adjusted at will by tuning λ . This is not altogether surprising, since $\lambda = +\infty$ yields a reconstructed signal in the null space of L_2 with sparsity zero.

Another promising feature of this experiment is that our algorithm is able to strike the right balance between both dictionaries (*i.e.*, D-splines and D²-splines), in that the selected dictionary elements explain the ground-truth signal well. Note that this requires careful tuning of the weight parameter α . Indeed, a badly tuned α leads to a lopsided use of one of the dictionaries, the other being too strongly penalized. We observe that a suitable balance can consistently be found when the ground truth fits the signal model as in (5.43).

By contrast, the single-operator reconstructions in Figure 5.1b do not do well in regions that call for different dictionary elements. More precisely, linear regimes in the ground truth signal lead to a staircasing effect in the piecewise-constant reconstruction (D), which results in a loss of sparsity. As for the sharp jumps in the ground truth, they lead to gradual increases in the piecewise-linear reconstruction (D²). This is because sharp increases are heavily penalized by the regularization, which is undesirable given the form of the ground truth.

5.5.2 Compressed Sensing

A second application of our framework is CS-type problems, which attempt to recover a sparse multicomponent signals given a small number of measurements⁴. We use the same type of test signals as in (5.43), namely hybrid splines with low sparsity $K_{s_1} + K_{s_2}$. The measurement functionals are assumed to take the form

$$\nu_m(s) \triangleq \int_0^T \cos(\omega_m x + \phi_m) s(x) dx, \quad (5.44)$$

where $\omega_m \in \mathbb{R}$ are the sampling frequencies and $\phi_m \in [0, 2\pi)$ some given phase offsets. This amounts to sampling in the Fourier space—a rectangular window is applied in order to satisfy Assumptions 5.1. Moreover, as discussed in Example 3.3, ν_m is weak*-continuous and thus satisfies the assumptions of Theorem 5.1. We choose Fourier measurements because they are known to have good recovery properties according to the theory of CS [119]. However, the absence of a D-RIP-type assumption prevents us from making any theoretical claims on the quality of the recovery.

In this setting, perfect recovery is illusory since the $x_{i,k}$ are not on a grid and the s_i signals are thus not included in the search spaces $\mathcal{M}_{L_1, \phi_0, h}(\mathbb{R})$ and $\mathcal{M}_{L_2, h}(\mathbb{R})$. Even in purely discrete settings, the recovery of a multicomponent discrete signal is a difficult problem, which requires conditions on the measurement matrix \mathbf{H} and some form of incoherence between dictionaries [129, 141]. Here, no such assumptions are made, which makes the recovery problem even more challenging. To avoid adding to the difficulty, we only consider noiseless measurements.

⁴Note that we are not interested in the data separation problem as in [141], but only in recovering the complete multicomponent signal.

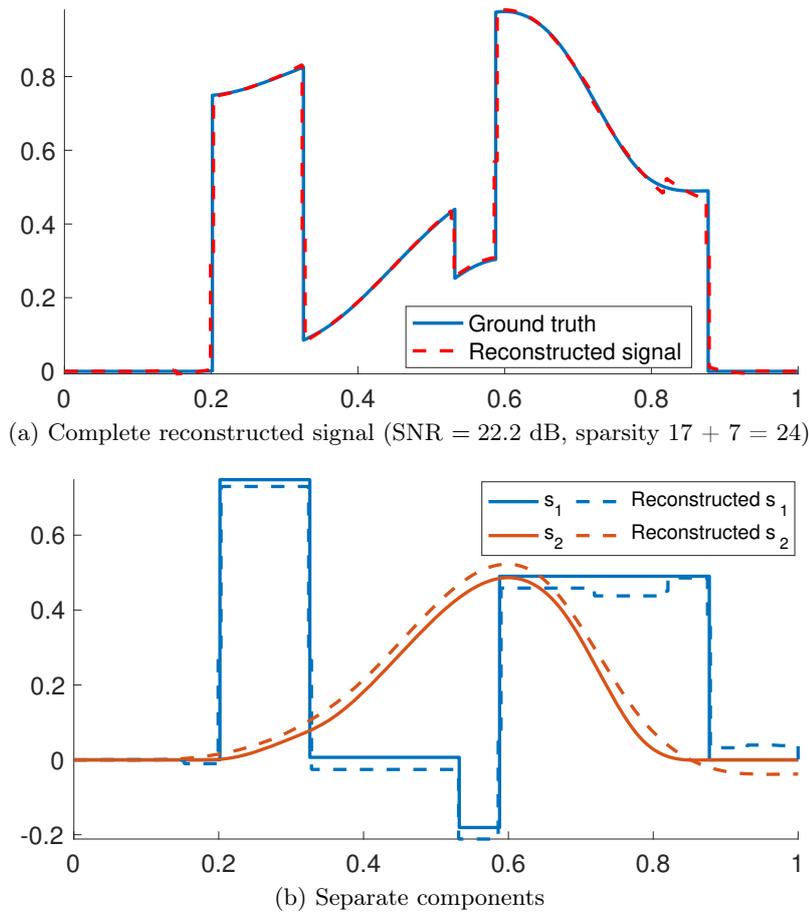


Figure 5.2 – Reconstruction result with noiseless Fourier measurements for $L_1 = D$, $L_2 = D^4$, $M = 30$, $\lambda = 10^{-15}$, $\alpha = 5 \times 10^{-5}$. Final grid size: $h = 1/2^8$.

An example run is shown in Figure 5.2 for a test signal with sparsity 15 ($K_{s_1} = 5$ and $K_{s_2} = 10$) and $M = 30$. Since the measurements are noiseless and we are interested in recovering the test signal as faithfully as possible, the data-fidelity term should be penalized much more than the regularization term. We therefore pick the regularization parameter $\lambda = 10^{-15} \ll 1$. The reconstructed signal in Figure 5.2a is remarkably close to the test signal, considering the difficulties of the problem. Notice that the final grid size selected by Algorithm 5.1 ($h = 1/2^8$) is still computationally tractable. The separate components of the reconstructed signal compared to those of the test signal are provided in Figure 5.2b. We observe that the separation is not perfect: there is a small compensation effect between the two reconstructed components.

5.6 Conclusion

We have established a representer theorem that states that hybrid splines are solutions of continuous-domain inverse problems with multiple gTV regularization. The regularization

operators L_1 and L_2 are taken to be multiple-order derivatives, which lead to piecewise-polynomial splines. This result implies that such problems can be solved exactly using a concatenated dictionary that consists of L_1 and L_2 -splines. We propose an exact B-spline-based discretization scheme and a multiresolution algorithm to solve the continuous-domain problem in a suitable search space. We then apply our algorithm to curve fitting and CS-type problems, and show that it is both computationally feasible and very successful experimentally.

To the best of our knowledge, this is the first instance of a continuous-domain inverse problem for multicomponent signals. Our algorithm can be viewed as the continuous-domain counterpart of discrete data-separation problems, such as morphological component analysis.

6 Sparse-Plus-Smooth Signal Models

This chapter is based on the following publication [142]:

T. Debarre, S. Aziznejad, and M. Unser, “Continuous-Domain Formulation of Inverse Problems for Composite Sparse-Plus-Smooth Signals”, *IEEE Open Journal of Signal Processing*, vol. 2, pp. 545–558, Sep. 2021.

6.1 Introduction

In the traditional discrete formalism of linear inverse problems, the goal is to recover a signal $\mathbf{c}_0 \in \mathbb{R}^N$ based on some observation vector $\mathbf{y} \in \mathbb{R}^M$. These observations are typically acquired via a sensing matrix $\mathbf{H} \in \mathbb{R}^{M \times N}$, so that $\mathbf{H}\mathbf{c}_0 \approx \mathbf{y}$. The recovery is often achieved by solving an optimization problem of the form

$$\mathbf{c}^* \in \arg \min_{\mathbf{c} \in \mathbb{R}^N} \left(\underbrace{E(\mathbf{H}\mathbf{c}, \mathbf{y})}_{\text{Data fidelity}} + \underbrace{\lambda \mathcal{R}(\mathbf{L}\mathbf{c})}_{\text{Regularization}} \right), \quad (6.1)$$

where $E : \mathbb{R}^M \times \mathbb{R}^M \rightarrow \mathbb{R}$ is a suitable convex data-fidelity loss functional, \mathcal{R} is the regularization functional, \mathbf{L} specifies a suitable transform domain, and $\lambda > 0$ is a tuning parameter that determines the strength of the regularization. Historically, the first instance of regularization dates back to Tikhonov [4] with a quadratic regularization functional $\mathcal{R} = \|\cdot\|_2^2$. Tikhonov regularization constrains the energy of $\mathbf{L}\mathbf{c}$ which, when \mathbf{L} is a finite-difference matrix, leads to a smooth reconstructed signal \mathbf{c}^* . Tikhonov regularization has the practical advantage of being mathematically tractable which leads to a closed-form solution.

More recently, the sparsity-promoting regularization functional $\mathcal{R} = \|\cdot\|_1$ has now surpassed the ℓ_2 norm as the regularization method of choice in most applications. However, the nondifferentiability of the ℓ_1 norm leads to more involved iterative optimization

procedures using proximal algorithms [17]. We refer to the introduction of Part II for a more in-depth discussion on discrete ℓ_1 optimization.

6.1.1 Discrete Inverse Problems for Composite Signals

Despite their success, ℓ_1 and ℓ_2 regularization methods are too simple to model many real-world signals. In this chapter, we investigate composite models of the form $s = s_1 + s_2$ where the two components have different characteristics. More precisely, s_1 is assumed to be sparse in some given domain and is treated with ℓ_1 regularization, while s_2 is assumed to be smooth and is treated with ℓ_2 regularization. In discrete settings, a natural way of reconstructing such signals is to solve the optimization problem

$$\min_{\mathbf{c}_1, \mathbf{c}_2 \in \mathbb{R}^N} \left(E(\mathbf{H}(\mathbf{c}_1 + \mathbf{c}_2), \mathbf{y}) + \lambda_1 \|\mathbf{L}_1 \mathbf{c}_1\|_1 + \lambda_2 \|\mathbf{L}_2 \mathbf{c}_2\|_2^2 \right), \quad (6.2)$$

where $\mathbf{c}_1, \mathbf{c}_2$ are the two components of the signal $\mathbf{c} = \mathbf{c}_1 + \mathbf{c}_2$, $\lambda_1, \lambda_2 > 0$ are tuning parameters, $\mathbf{L}_1 \in \mathbb{R}^{N \times N}$ is a sparsifying transform for \mathbf{c}_1 , and $\mathbf{L}_2 \in \mathbb{R}^{N \times N}$ is a low-energy-promoting transform for \mathbf{c}_2 . This modeling is considered amongst others in [143–148].

6.1.2 Continuous-Domain Formulation

As explained in the introduction of this thesis, most real-world signals are best modeled as continuous functions. Therefore, when feasible, to formulate the inverse problem in the continuous domain is a natural and desirable objective. In this chapter, we thus adapt the discrete approach of (6.2) to 1D continuous-domain composite signals by solving an optimization problem of the form

$$\min_{s_1, s_2} \left(E(\boldsymbol{\nu}(s_1 + s_2), \mathbf{y}) + \lambda_1 \|\mathbf{L}_1 \{s_1\}\|_{\mathcal{M}} + \lambda_2 \|\mathbf{L}_2 \{s_2\}\|_{L_2}^2 \right), \quad (6.3)$$

where s_1, s_2 are the two components of the signal $s = s_1 + s_2 : \mathbb{R} \rightarrow \mathbb{R}$, $\boldsymbol{\nu} = (\nu_1, \dots, \nu_M) : s \mapsto \boldsymbol{\nu}(s) \in \mathbb{R}^M$ is the linear forward model, and $\mathbf{L}_i = \mathbf{D}^{N_{0,i}}$ for $i \in \{1, 2\}$, where \mathbf{D} is the derivative operator and $N_{0,i} \geq 1$ is the order of the derivative. The regularization norm $\|\cdot\|_{\mathcal{M}}$ is the total-variation (TV) norm for measures defined in (2.17), which is the continuous sparsity-promoting counterpart of the discrete ℓ_1 norm. We refer to this term as the *generalized TV* (gTV) regularizer. Finally, $\|\cdot\|_{L_2}$ is the usual norm over the Lebesgue space $L_2(\mathbb{R})$ of signals with finite energy; we refer to the corresponding term as the *generalized Tikhonov* (gTikhonov) regularizer, which promotes smoothness in combination with the operator \mathbf{L}_2 .

6.1.3 Representer Theorems and Discretization

A classical way of discretizing a continuous-domain problem is to reformulate it as a finite-dimensional one by relying on a *representer theorem* that gives a parametric form of the solution. Prominent examples include representer theorems for problems formulated over reproducing-kernel Hilbert spaces (RKHS), which are foundational to the field of machine learning [42, 43]. As demonstrated in [45, Theorem 3], the minimization Problem (6.3) over the component s_2 (with a fixed s_1)—*i.e.*, gTikhonov regularization—falls into this category: the representer theorem states that there is a unique solution of the form

$$s_2^*(x) = p_2(x) + \sum_{m=1}^M a_{2,m} h_m(x), \quad (6.4)$$

where the additional component p_2 lies in the null space of L_2 (*i.e.*, $L_2\{p_2\} = 0$), h_m is a (typically quite smooth) kernel function that is fully determined by the choice of ν_m and L_2 , and $a_{2,m} \in \mathbb{R}$ are expansion coefficients. Therefore, to solve the continuous-domain problem, one need only optimize over the $a_{2,m}$ coefficients and the null-space component p_2 which lives in a finite-dimensional space. This leads to a standard finite-dimensional problem with Tikhonov regularization.

Concerning the minimization over the component s_1 (gTV regularization), Theorem 3.3 states that there is a L_1 -spline solution of the form

$$s_1^*(x) = p_1(x) + \sum_{k=1}^K a_{1,k} \rho_{L_1}(x - x_k), \quad (6.5)$$

where $a_{1,k}, x_k \in \mathbb{R}$, ρ_{L_1} is the Green's function of L_1 (defined in (2.25)), K is the number of atoms of s_1 which is bounded by $K \leq (M - N_{0,1})$, $N_{0,1}$ being the dimension of the null space of L_1 , and p_1 lies in the null space of L_1 . These representer theorems have paved the way for various exact discretization methods. In the gTikhonov case, one can optimize over the $a_{m,2}$ coefficients in (6.4) directly [45]. For the gTV case (6.5), grid-based techniques using a well-conditioned B-spline basis (Chapter 4) as well as grid-free techniques [90] have been proposed.

6.1.4 Our Contribution

In this chapter, we show that the representer theorems presented in the previous section can be combined into a composite one when dealing with Problem (6.3). More specifically, we prove that there exists a solution to (6.3) of the form $s_1^* = s_1^* + s_2^*$ such that s_1^* is of the form (6.5) and s_2^* is of the form (6.4): a “sparse plus smooth” solution. Building on this representation, we propose an exact discretization scheme. Both components s_i for $i \in \{1, 2\}$ are expressed in a suitable Riesz basis as $s_i = \sum_k c_i[k] \varphi_{i,k}$, where $c_i[k]$ are the coefficients to be optimized. This leads to an infinite-dimensional optimization

problem reminiscent of the infinite-dimensional compressed sensing framework of Adcock and Hansen [3] and the wavelet-based model of Daubechies *et al.* [149]. However, these frameworks differ from our original Problem (6.3) in that their native spaces admit a countable basis which leads to a more straightforward discretization process.

To solve this infinite-dimensional problem numerically, we cast it as a finite-dimensional problem under some mild assumptions. This requires a careful handling of the boundaries of our interval of interest. In our implementation, we choose basis functions $\varphi_{1,k} = \beta_{L_1}(\cdot - k)$ and $\varphi_{2,k} = \beta_{L_2^*}(\cdot - k)$, where β_L is the B-spline for the operator L . B-splines are popular choices of basis functions, in large part due to their minimal-support property (see Section 3.1). We show that optimizing over the spline coefficients leads to a discrete problem similar to (6.2) of the form

$$\min_{(\mathbf{c}_1, \mathbf{c}_2) \in \mathbb{R}^{N_1} \times \mathbb{R}^{N_2}} \left(E(\mathbf{H}_1 \mathbf{c}_1 + \mathbf{H}_2 \mathbf{c}_2, \mathbf{y}) + \lambda_1 \|\mathbf{L}_1 \mathbf{c}_1\|_1 + \lambda_2 \|\mathbf{L}_2 \mathbf{c}_2\|_2^2 \right), \quad (6.6)$$

where $\mathbf{H}_i \in \mathbb{R}^{M \times N_i}$ and $\mathbf{L}_i \in \mathbb{R}^{P_i \times N_i}$ for $i \in \{1, 2\}$. This discretization is *exact* in the sense that it is equivalent to the continuous Problem (6.3) when each component s_i lies in the space generated by the basis functions $\{\varphi_{i,k}\}_{k \in \mathbb{Z}}$. This is a consequence of our informed choice of these basis functions $\varphi_{i,k}$. Moreover, the short support of the B-splines leads to well-conditioned \mathbf{H}_i matrices and, thus, to a computationally feasible problem.

6.1.5 Related Works

The use of multiple regularization penalties is quite common in the literature. However, in most cases, each penalty is applied to the full signal instead of a component-wise [150–155]. A prominent example of such an approach is the elastic net [156], which is widely used in statistics. The spirit of these approaches is however quite different from ours: the reconstructed signal is encouraged to satisfy different priors *simultaneously*. Conversely, in (6.2), each component satisfies different priors independently of the others, which will give very different results.

Closer to our framework is the popular low-rank plus sparse matrix decomposition model [139], which imposes some form of sparsity on both components.

The model of Meyer [157] and its generalization by Vese and Osher [158, 159] follow the same idea as Problem (6.3), with the important difference that they use calculus-of-variation techniques to solve it. There is a connection as well with the Mumford-Shah functional [160], which is commonly used to segment an image in piecewise-smooth regions. The main difference lies in the fact that the optimization is not performed over the different components of the signal, but over the region boundaries. Another difference is that these models assume that one has full access to the noisy signal over a continuum, whereas (6.3) assumes that we only have access to some discrete measurements specified

by the forward model ν .

6.1.6 Outline

In Section 6.2, we defined the search spaces of our optimization problem, which are partly based on the native spaces introduced in Chapter 2. In Section 6.3, we formulate the continuous-domain problem and formulate our representer theorem, which is our main theoretical result. In Section 6.4, we detail our discretization strategy, which relies on the selection of suitable Riesz bases. Finally, in Section 6.6, we present experiments on simulated data.

6.2 Native Spaces

We now introduce the native spaces of our regularization operators L_i for $i \in \{1, 2\}$, which are $N_{0,i}$ th-order derivative operators¹ of the form (2.19). In particular, their causal Green's function ρ_{L_i} is given by (2.25), and their null space \mathcal{N}_{L_i} by (2.24); for more information, we refer to Section 3.1.2.

6.2.1 Sparse Component

The native space for s_1 in (6.3), which is defined in (2.33), is given by

$$\mathcal{M}_{L_1}(\mathbb{R}) \triangleq \{f \in \mathcal{S}'(\mathbb{R}) : L_1\{f\} \in \mathcal{M}(\mathbb{R})\}, \quad (6.7)$$

where $\mathcal{M}(\mathbb{R})$ is the space of Radon measures defined in (2.18). We refer to Chapter 2 for more details on the construction of $\mathcal{M}_{L_1}(\mathbb{R})$.

6.2.2 Smooth Component

The regularization norm $\|\cdot\|_{L_2}$ for the smooth component s_2 in (6.3) is defined over the Lebesgue space $L_2(\mathbb{R})$ of functions with finite energy. The corresponding native space of the smooth component s_2 is

$$\mathcal{H}_{L_2}(\mathbb{R}) \triangleq \{f \in \mathcal{S}'(\mathbb{R}) : L_2\{f\} \in L_2(\mathbb{R})\}, \quad (6.8)$$

which has a Hilbert-space structure [45].

¹As in most chapters of the thesis, all the results and algorithms from this chapter can be extended to more general operators L_i , *e.g.*, rational operators [102].

6.3 Continuous-Domain Inverse Problem

When the intersection $\mathcal{N}_0 = \mathcal{N}_{L_1} \cap \mathcal{N}_{L_2}$ of the null spaces of our two native spaces is nontrivial, we need to specify boundary conditions on one of the two spaces to ensure the well-posedness of our optimization problem. This is done by introducing a biorthogonal system (ϕ_0, \mathbf{p}_0) for \mathcal{N}_0 in the sense of Definition 2.6. An example of a valid choice is given in (2.48). The search space with boundary conditions ϕ_0 is then given by

$$\mathcal{M}_{L_1, \phi_0}(\mathbb{R}) \triangleq \{f \in \mathcal{M}_{L_1}(\mathbb{R}) : \phi_0(f) = \mathbf{0}\}. \quad (6.9)$$

We now present in Theorem 6.1 our problem formulation to reconstruct sparse-plus-smooth composite signals. This representer theorem gives a parametric form of a solution of our optimization problem.

Theorem 6.1 (Continuous-Domain Representer Theorem). *Let $E : \mathbb{R}^M \times \mathbb{R}^M \rightarrow \mathbb{R}$ be a nonnegative, coercive, proper, convex, and lower-semicontinuous functional. Let L_1, L_2 be derivative operators of the form (2.19), and let $\boldsymbol{\nu} = (\nu_1, \dots, \nu_M)$ be a linear measurement operator composed of the M linear functionals $\nu_m : f \mapsto \nu_m(f) \in \mathbb{R}$ that are weak*-continuous over $\mathcal{M}_{L_1}(\mathbb{R})$ and over $\mathcal{H}_{L_2}(\mathbb{R})$. We assume that $\mathcal{N}_{\boldsymbol{\nu}} \cap (\mathcal{N}_{L_1} + \mathcal{N}_{L_2}) = \{0\}$, where $\mathcal{N}_{\boldsymbol{\nu}}$ is the null space of $\boldsymbol{\nu}$ (well-posedness assumption). Then, for any $\lambda_1, \lambda_2 > 0$, the optimization problem*

$$\begin{aligned} \mathcal{V} \triangleq \arg \min_{\substack{s_1 \in \mathcal{M}_{L_1, \phi_0}(\mathbb{R}) \\ s_2 \in \mathcal{H}_{L_2}(\mathbb{R})}} \mathcal{J}(s_1, s_2) \quad \text{with} \\ \mathcal{J}(s_1, s_2) \triangleq E(\boldsymbol{\nu}(s_1 + s_2), \mathbf{y}) + \lambda_1 \|L_1\{s_1\}\|_{\mathcal{M}} + \lambda_2 \|L_2\{s_2\}\|_{L_2}^2 \end{aligned} \quad (6.10)$$

has a solution $(s_1^*, s_2^*) \in \mathcal{V}$ with the following components:

- the component s_1^* is a nonuniform L_1 -spline of the form

$$s_1^*(x) = p_1(x) + \sum_{k=1}^{K_1} a_{1,k} \rho_{L_1}(x - x_k) \quad (6.11)$$

for some $K_1 \leq (M - N_{0,1})$, where $p_1 \in \mathcal{N}_{L_1}$, and $a_{1,k}, x_k \in \mathbb{R}$;

- the component s_2^* is of the form

$$s_2^*(x) = p_2(x) + \sum_{m=1}^M a_{2,m} h_m(x), \quad (6.12)$$

where $h_m(x) \triangleq \left(\nu_m * \mathcal{F}^{-1} \left\{ \frac{1}{|\widehat{L}_2|^2} \right\} \right) (x)$, $p_2 \in \mathcal{N}_{L_2}$, $a_{2,k} \in \mathbb{R}$, and where $\sum_{m=1}^M a_{2,m} \langle q_2, \nu_m \rangle = 0$ for any $q_2 \in \mathcal{N}_{L_2}$.

Moreover, for any pair of solutions $(s_1^*, s_2^*), (\tilde{s}_1^*, \tilde{s}_2^*) \in \mathcal{V}$, s_2^* and \tilde{s}_2^* differ only up to an element of the null space \mathcal{N}_{L_2} , so that $(s_2^* - \tilde{s}_2^*) \in \mathcal{N}_{L_2}$.

The proof of Theorem 6.1 is given in Appendix B.2.

Remark 6.1. *As in Theorem 3.3, we restrict to derivative operators for L_1 and L_2 . However, Theorem 6.1 and our proof technique remain valid for any spline-admissible operators in the sense of [14, Definition 1].*

A pleasing outcome of Theorem 6.1 is that it combines Theorems 3 and 4 of [45] into one. There is, however, an added technicality due to the boundary conditions ϕ_0 . The latter are necessary to ensure the well-posedness of Problem (6.10)². Note, however, that these conditions do not restrict the search space, since $\mathcal{M}_{L_1, \phi_0}(\mathbb{R}) + \mathcal{H}_{L_2}(\mathbb{R}) = \mathcal{M}_{L_1}(\mathbb{R}) + \mathcal{H}_{L_2}(\mathbb{R})$.

6.4 Exact Discretization

In order to discretize Problem (6.10), we restrict the search spaces $\mathcal{M}_{L_1, \phi_0}(\mathbb{R})$ and $\mathcal{H}_{L_2}(\mathbb{R})$. The standard approach to achieve this is to choose a sequence of appropriate basis functions $\{\varphi_{i,k}\}_{k \in \mathbb{Z}}$ that span the reconstruction spaces

$$V_i(\mathbb{R}) \triangleq \left\{ \sum_{k \in \mathbb{Z}} c_i[k] \varphi_{i,k} : c_i \in V_i(\mathbb{Z}) \right\} \quad (6.13)$$

for $i \in \{1, 2\}$ that are subject to the constraints $V_1(\mathbb{R}) \subset \mathcal{M}_{L_1, \phi_0}(\mathbb{R})$ and $V_2(\mathbb{R}) \subset \mathcal{H}_{L_2}(\mathbb{R})$. These continuous spaces are linked to discrete spaces $V_i(\mathbb{Z})$, the choices of which will be made explicit in (6.16) and (6.27). More precisely, there is a one-to-one mapping between them using the basis functions $\varphi_{i,k}$.

6.4.1 Riesz Bases and B-Splines

For numerical purposes, a desirable property is that our basis functions satisfy the Riesz property. Riesz bases are highly important concepts in that they generalize orthonormal bases, while leaving more flexibility for other desirable properties such as short support [161].

Definition 6.1 (Riesz basis). *A sequence of functions $\{\varphi_k\}_{k \in \mathbb{Z}}$ with $\varphi_k \in L_2(\mathbb{R})$ is said to be a Riesz basis if there exist constants $0 < A \leq B$ such that, for any $c \in \ell_2(\mathbb{Z})$, we*

²Otherwise, for any $(s_1^*, s_2^*) \in \mathcal{V}$ and $p \in \mathcal{N}_0$, we would have that $(s_1^* + p, s_2^* - p) \in \mathcal{V}$ which would imply that \mathcal{V} is unbounded.

have that

$$A\|c\|_{\ell_2} \leq \left\| \sum_{k \in \mathbb{Z}} c[k] \varphi_k \right\|_{L_2} \leq B\|c\|_{\ell_2}. \quad (6.14)$$

Popular examples of Riesz bases are polynomial B-spline bases, which are introduced in (3.16). The polynomial B-spline β_L of a derivative operator L of the form (2.19) is characterized by its B-spline filter d_L defined in (3.11). This filter appears in the expression of the innovation of the B-spline, which is given by

$$L\{\beta_L\} = \sum_{k \in \mathbb{Z}} d_L[k] \delta(\cdot - k). \quad (6.15)$$

Both β_L and d_L have finite support which is determined by the order of the operator L . More precisely, the supports of β_{L_i} and d_{L_i} for $i \in \{1, 2\}$ are $[0, D_i]$ and $\{0, \dots, D_i\}$, respectively, where D_i is the order of the operator L_i . For more background on polynomial B-splines, we refer to Section (3.1.4).

6.4.2 Choice of Basis Functions

We now present and discuss our choice for the basis functions $\varphi_{1,k}$ and $\varphi_{2,k}$.

Sparse Component

For the sparse component, we choose the B-spline basis functions $\varphi_{1,k} = \beta_{L_1}(\cdot - k)$ for all $k \in \mathbb{Z}$. With this choice,

$$V_1(\mathbb{R}) \triangleq \left\{ f = \sum_{k \in \mathbb{Z}} c_1[k] \varphi_{1,k} : c_1 \in V_1(\mathbb{Z}) \right\} \subset \mathcal{M}_{L_1, \phi_0}(\mathbb{R})$$

with the digital-filter space

$$V_1(\mathbb{Z}) \triangleq \left\{ (c_1[k])_{k \in \mathbb{Z}} : (d_{L_1} * c_1) \in \ell_1(\mathbb{Z}) \text{ and } \sum_{k \in \mathbb{Z}} c_1[k] \phi_0(\varphi_{1,k}) = \mathbf{0} \right\}, \quad (6.16)$$

is the largest possible native reconstruction space, as demonstrated in the previous chapter (5.24). The choice of the basis functions $\varphi_{1,k}$ is guided by the following considerations:

- they generate the space of uniform L_1 splines. This conforms with Theorem 6.1, which states that the component s_1^* is an L_1 -spline;
- they enable exact computations in the continuous domain. In particular, by (3.18), we have that $\|L_1\{\sum_{k \in \mathbb{Z}} c_1[k] \varphi_{1,k}\}\|_{\mathcal{M}} = \|d_{L_1} * c_1\|_{\ell_1}$;

- the Riesz-basis property of B-splines leads to a well-conditioned system matrix, which is paramount in numerical applications.

B-splines are the only functions that satisfy all these properties. Based on these criteria, B-splines are thus optimal.

Smooth Component

At first glance, the most natural choice for $\varphi_{2,k}$ is to select the basis functions suggested by (6.12) in Theorem 6.1: h_m for $1 \leq m \leq M$ and a basis of \mathcal{N}_{L_2} , which yield a finite number $M + N_{0,2}$ of basis functions. However, this approach runs into the following hitches:

- the basis functions h_m are typically increasing at infinity, which contradicts the Riesz-basis requirement and leads to severely ill-conditioned optimization tasks [45];
- depending on the measurements operator ν , h_m may lack a closed-form expression.

We therefore focus on these criteria, in a spirit similar to [162]. The $\varphi_{2,k}$ are chosen to be regular shifts of a generating function φ_2 , with $\varphi_{2,k} = \varphi_2(\cdot - k)$ such that $\{\mathbb{L}_2\{\varphi_{2,k}\}\}_{k \in \mathbb{Z}}$ forms a Riesz basis in the sense of Definition 6.1. Contrary to $\varphi_{1,k}$, these requirements allow for many choices of $\varphi_{2,k}$. In order to perform exact discretization, one then only needs to compute the following autocorrelation filter.

Proposition 6.1 (Autocorrelation Filter for the Smooth Component). *Let φ_2 be a generating function such that $\varphi_{2,k} \triangleq \varphi_2(\cdot - k)$ form a Riesz basis. Then, the following two items hold:*

- the inner product $\langle \mathbb{L}_2\{\varphi_{2,k}\}, \mathbb{L}_2\{\varphi_{2,k'}\} \rangle_{L_2}$ only depends on the difference $(k - k')$. We can thus introduce the autocorrelation filter

$$\rho[k] \triangleq \langle \mathbb{L}_2\{\varphi_{2,k}\}, \mathbb{L}_2\{\varphi_{2,0}\} \rangle_{L_2} = \langle \mathbb{L}_2\{\varphi_{2,k+k'}\}, \mathbb{L}_2\{\varphi_{2,k'}\} \rangle_{L_2} \quad (6.17)$$

for any $k, k' \in \mathbb{Z}$;

- the filter ρ is positive semidefinite, with $\sum_{k,k' \in \mathbb{Z}} c[k]c[k']\rho[k - k'] \geq 0$ for any finitely supported real digital filter c .

Proof. The first item is proved with a simple change of variable in the integral that defines

the inner product. The second item is derived by observing that, for any c_2 , we have

$$\begin{aligned} \left\| \mathbb{L}_2 \left\{ \sum_{k \in \mathbb{Z}} c_2[k] \varphi_{2,k} \right\} \right\|_{L_2}^2 &= \sum_{k, k' \in \mathbb{Z}} \left(c_2[k] c_2[k'] \langle \mathbb{L}_2 \{ \varphi_{2,k} \}, \mathbb{L}_2 \{ \varphi_{2,k'} \} \rangle \right) \\ &= \sum_{k, k' \in \mathbb{Z}} c_2[k] c_2[k'] \rho[k - k'] \geq 0. \end{aligned} \quad (6.18)$$

□

For our implementation, we make a specific choice of basis function φ_2 among the many choices for which the autocorrelation filter (6.17) can be computed analytically. We choose the $L_2^* L_2$ B-spline basis $\varphi_2 = \beta_{L_2^* L_2}$ and $\varphi_{2,k} = \varphi_2(\cdot - k)$, where L_2^* denotes the adjoint operator of L_2 . This choice has the following additional advantages:

- the generator φ_2 has a simple explicit expression that does not depend on the measurement operator ν ;
- as will be shown in Proposition 6.2, the autocorrelation filter ρ can be factorized, which is convenient for implementation purposes;
- in the special case of the sampling operator $\nu_m = \delta(\cdot - x_m)$, where the x_m are the sampling locations, this choice conforms with (6.12) in Theorem 6.1 since s_2^* is then an $L_2^* L_2$ -spline. Note, however, that we do not exploit the knowledge that s_2^* has knots at the sampling locations x_m .

With this basis function $\varphi_2 = \beta_{L_2^* L_2}$, there is no straightforward choice of the digital-filter space $V_2(\mathbb{Z})$. Our practical choice is given in (6.27); it depends on our discretization method for reasons that are discussed in Remark 6.2.

Our choice of basis function $\varphi_2 = \beta_{L_2^* L_2}$ enables us to factorize the autocorrelation filter ρ via its “square root” g , as demonstrated in Proposition 6.2.

Proposition 6.2 (Factorization of the Autocorrelation Filter). *Let L_2 be a derivative operator of the form (2.19), and let $\varphi_2 = \beta_{L_2^* L_2} = \beta_{L_2} * \beta_{L_2}^\vee$. Then, the basis $\{\varphi_{2,k}\}_{k \in \mathbb{Z}}$ forms a Riesz basis as required in Section 6.4.2, and the autocorrelation filter ρ defined in Proposition 6.1 is of the form*

$$\rho = d_{L_2} * d_{L_2}^\vee * b, \quad (6.19)$$

where $b[k] \triangleq \beta_{L_2^* L_2}(k)$ is the B-spline kernel of the operator $L_2^* L_2$, which is a positive-semidefinite filter supported in $[-(D_2 - 2) \dots D_2 - 2]$, where D_2 is the order of the operators L_2 . The filter ρ can thus be factorized as $\rho = g * g^\vee$ with

$$g \triangleq d_{L_2} * b^{1/2}, \quad (6.20)$$

where the filter $b^{1/2}$ satisfies $b = b^{1/2} * (b^{1/2})^\vee$ and is of length $B \triangleq (D_2 - 1)$, and g is thus of length $G \triangleq 2B = 2(D_2 - 1)$.

Proof. We have that

$$\begin{aligned}
 \rho[k] &= \langle \mathbf{L}_2\{\varphi_{2,k}\}, \mathbf{L}_2\{\varphi_{2,0}\} \rangle_{L_2} \\
 &= \langle \mathbf{L}_2^* \mathbf{L}_2\{\varphi_{2,k}\}, \varphi_{2,0} \rangle_{\mathcal{H}'_{L_2} \times \mathcal{H}_{L_2}} \\
 &= \left\langle \sum_{k' \in \mathbb{Z}} d_{\mathbf{L}_2^* \mathbf{L}_2}[k] \delta(\cdot - (k + k')), \varphi_{2,0} \right\rangle_{\mathcal{H}'_{L_2} \times \mathcal{H}_{L_2}} \\
 &= \sum_{k' \in \mathbb{Z}} d_{\mathbf{L}_2^* \mathbf{L}_2}[k] b[k + k'] \\
 &= (d_{\mathbf{L}_2^* \mathbf{L}_2} * b^\vee)[-k] \\
 &= (d_{\mathbf{L}_2^* \mathbf{L}_2} * b)[k],
 \end{aligned}$$

where $\langle \cdot, \cdot \rangle_{\mathcal{H}'_{L_2} \times \mathcal{H}_{L_2}}$ denotes the duality product between $\mathcal{H}_{L_2}(\mathbb{R})$ and its dual $\mathcal{H}'_{L_2}(\mathbb{R})$, and the last line results from the symmetry of ρ and b .

Next, we prove that b is positive-semidefinite. Indeed, for any finitely supported filter c , we have that

$$\sum_{k, k' \in \mathbb{Z}} c[k] c[k'] b[k - k'] = \left\| \sum_{k \in \mathbb{Z}} c[k] \beta_{L_2}(\cdot - k) \right\|_{L_2}^2 \geq 0, \quad (6.21)$$

where we have used the property

$$b[k] = (\beta_{L_2} * \beta_{L_2}^\vee)(k) = \langle \beta_{L_2}, \beta_{L_2}(\cdot - k) \rangle_{L_2}. \quad (6.22)$$

Finally, to prove the existence of $b^{1/2}$, we notice that b has the finite support $[-(B - 1) \dots B - 1]$ due to the finite support $(-D_2, D_2)$ of $\beta_{L_2^* \mathbf{L}_2}$, and we have $B = (D_2 - 1)$. Since b is also symmetric, its z -transform satisfies $B(z) = B(z^{-1})$; therefore, for any zero z_k of $B(z)$, z_k^{-1} is also a zero. Moreover, it is well known that $B(\pm 1) \neq 0$, so that zeros must come in pairs $z_k \neq z_k^{-1}$. Hence, $B(z)$ can be written as $B(z) = \prod_{k=1}^B (1 - z_k z)(1 - z_k z^{-1})$. Hence, to take $b^{1/2}$ to be the inverse z -transform of $B^{1/2}(z) = \prod_{k=1}^B (1 - z_k z^{-1})$ is a valid choice (we clearly have $b = b^{1/2} * (b^{1/2})^\vee$), and (6.20) is readily obtained. \square

We summarize in Table 6.1 the different filters and their mutual relations. These filters will be useful for the definition of the regularization matrix \mathbf{L}_2 .

	d_{L_1}	d_{L_2}	$\rho = b * d_{L_2} * d_{L_2}^\vee$	$g = b^{1/2} * d_{L_2}$	$b = (\beta_{L_2^* L_2}(k))_{k \in \mathbb{Z}}$	$b^{1/2}$
Description	Finite-difference filter for L_1	Finite-difference filter for L_2	Autocorrelation filter for L_2	“Square root” of ρ ($\rho = g * g^\vee$)	Samples of basis function ($L_2^* L_2$ B-spline)	“Square root” of b ($b = b^{1/2} * (b^{1/2})^\vee$)
Introduced in	Equation (3.11)	Equation (3.11)	Proposition 6.1	Proposition 6.2	Proposition 6.2	Proposition 6.2
Support length	D_1	D_2	$2G - 1 = 4D_2 - 5$	$G = 2D_2 - 2$	$2B - 1 = 2D_2 - 3$	$B = D_2 - 1$
Support	$[0 \dots D_1 - 1]$	$[0 \dots D_2 - 1]$	$[-(G - 1) \dots G - 1]$	$[0 \dots G - 1]$	$[-(B - 1) \dots B - 1]$	$[0 \dots B - 1]$
Example $L_2 = D$		$[1, -1]$	$[-1, 2, -1]$	$[1, -1]$	$[1]$	$[1]$
Example $L_2 = D^2$		$[1, -2, 1]$	$\frac{1}{6}[1, 0, -9, 16, -9, 0, 1]$	$C[1, \sqrt{3}, -(3 + 2\sqrt{3}), 2 + \sqrt{3}]$	$\frac{1}{6}[1, 4, 1]$	$C[1, 2 + \sqrt{3}]$

Table 6.1 – Relevant filters and their supports $\left(C \triangleq \sqrt{\frac{2-\sqrt{3}}{6}}\right)$.

6.4.3 Formulation of the Discrete Problem

The autocorrelation filter introduced in Proposition 6.1 enables us to discretize Problem (6.10) in an exact way in the $V_i(\mathbb{R})$ spaces.

Proposition 6.3 (Riesz-Basis Discretization). *Let $\varphi_{i,k}$ be chosen as specified in Section 6.4.2 for $i \in \{1, 2\}$, $k \in \mathbb{Z}$, and*

$$\mathcal{V}_d \triangleq \arg \min_{(c_1, c_2) \in V_1(\mathbb{Z}) \times V_2(\mathbb{Z})} \mathcal{J}_d(c_1, c_2). \quad (6.23)$$

The cost function is given by

$$\mathcal{J}_d(c_1, c_2) \triangleq E \left(\sum_{k \in \mathbb{Z}} \left(c_1[k] \boldsymbol{\nu}(\varphi_{1,k}) + c_2[k] \boldsymbol{\nu}(\varphi_{2,k}) \right), \mathbf{y} \right) + \lambda_1 \|d_{L_1} * c_1\|_{\ell_1} + \lambda_2 \langle c_2, \rho * c_2 \rangle_{\ell_2}, \quad (6.24)$$

where d_{L_1} is the finite-difference filter given by (6.15), ρ is defined in Proposition 6.1, and $\langle \cdot, \cdot \rangle_{\ell_2}$ is the inner product over $\ell_2(\mathbb{Z})$. Then, Problem (6.23) is equivalent to

$$\mathcal{V}_{\text{res}} = \arg \min_{(s_1, s_2) \in V_1(\mathbb{R}) \times V_2(\mathbb{R})} \mathcal{J}(s_1, s_2), \quad (6.25)$$

i.e., Problem (6.10) where the search spaces $\mathcal{M}_{L_1, \phi_0}(\mathbb{R})$ and $\mathcal{H}_{L_2}(\mathbb{R})$ are restricted to the spaces $V_1(\mathbb{R})$ and $V_2(\mathbb{R})$ defined in (6.13), respectively. This equivalence is in the sense that there exists a bijective linear mapping $(c_1, c_2) \mapsto (\sum_{k \in \mathbb{Z}} c_1[k] \varphi_{1,k}, \sum_{k \in \mathbb{Z}} c_2[k] \varphi_{2,k})$ from \mathcal{V}_d to \mathcal{V}_{res} .

Proof. By plugging the expansions $s_i = \sum_{k \in \mathbb{Z}} c_i[k] \varphi_{i,k}$ into the cost function \mathcal{J} , using the linearity of $\boldsymbol{\nu}$, we get the data-fidelity term of (6.24). Using (3.18), we have that $\|L_1 \{ \sum_{k \in \mathbb{Z}} c_1[k] \varphi_{1,k} \} \|_{\mathcal{M}} = \|d_{L_1} * c_1\|_{\ell_1}$. As for the second regularization term, we observe

that

$$\langle c_2, c_2 * \rho \rangle_{\ell_2} = \sum_{k, k' \in \mathbb{Z}} c_2[k] c_2[k'] \rho[k - k'] = \left\|_{L_2} \left\{ \sum_{k \in \mathbb{Z}} c_2[k] \varphi_{2,k} \right\} \right\|_{L_2}^2, \quad (6.26)$$

using (6.18) for the last step. This proves the equivalence between Problems (6.25) and (6.23), up to the specified mapping which is indeed a bijective linear mapping due to the Riesz-basis properties of $\{\varphi_{1,k}\}_{k \in \mathbb{Z}}$ and $\{\varphi_{2,k}\}_{k \in \mathbb{Z}}$. \square

6.5 Practical Implementation

We now discuss how to solve our discretized Problem (6.23) in practice, which involves recasting it as a finite-dimensional problem.

6.5.1 Finite Domain Assumptions

To solve Problem (6.23) numerically in an exact way, similarly to Chapters 4 and 5, we must make an additional assumption that will enable us to restrict the problem to a finite interval of interest.

Assumptions 6.1. *The measurement functionals ν_m are supported in an interval $I_T = [0, T]$, where $T \in \mathbb{N}$.*

This assumption is natural and is often fulfilled in practice, for instance in imaging with a finite field of view. The support length T then roughly corresponds to the number of grid points in the interval of interest. Note that, for simplicity, we only consider integer grids. However, following Chapters 4 and 5, the fineness of the grid can be tuned at will by adjusting T and rescaling the problem over the interval of interest.

6.5.2 Formulation of the Finite-Dimensional Problem

Our choice of basis functions together with the assumptions in Section 6.5.1 enable us to restrict Problem (6.23) to the interval of interest I_T . More precisely, we introduce the indices $m_i, M_i \in \mathbb{Z}$ for $i \in \{1, 2\}$; the range $[m_i \dots M_i]$ corresponds to the set of indices k for which $\text{Supp}(\varphi_{i,k}) \cap I_T \neq \emptyset$, so that the basis function $\varphi_{i,k}$ affects the measurements. Hence, the number of active basis functions (*i.e.*, the number of spline coefficients to be optimized) is $N_i = (M_i - m_i - 1)$. It can easily be verified that we have $m_1 = (-D_1 + 2)$, $M_1 = (T - 1)$, $m_2 = (-D_2 + 2)$, and $M_2 = (T + D_2 - 2)$. See Figure 6.1 for an illustrative example.

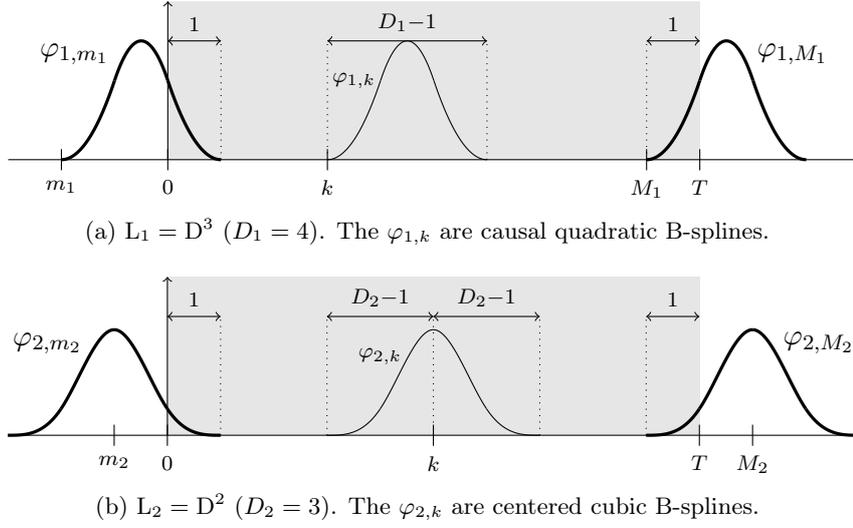


Figure 6.1 – Examples of boundary basis functions φ_{i,m_i} and φ_{i,M_i} for $i \in \{1, 2\}$.

Finally, we introduce the native digital-filter space

$$V_2(\mathbb{Z}) = \left\{ (c_2[k])_{k \in \mathbb{Z}} : \text{Supp}(d_{L_2} * c_2) \subset [1 \dots M_2] \right\}, \quad (6.27)$$

which is a valid choice because $V_2(\mathbb{R}) \subset \mathcal{H}_{L_2}(\mathbb{R})$. Indeed, we can verify that, for any $c_2 \in V_2(\mathbb{Z})$, the function $s_2 = \sum_{k \in \mathbb{Z}} c_2[k] \varphi_{2,k}$ satisfies $\|L_2\{s_2\}\|_{L_2}^2 = \|g * c_2\|_{\ell_2}^2 = \|b^{1/2} * (d_{L_2} * c_2)\|_{\ell_2}^2 < +\infty$, which proves that $s_2 \in \mathcal{H}_{L_2}(\mathbb{R})$. This is due to the finite support of both $(d_{L_2} * c_2)$ and $b^{1/2}$, where the filter $b^{1/2}$ and the decomposition $g = b^{1/2} * d_{L_2}$ are introduced in Proposition 6.2.

Remark 6.2. *Contrary to $V_1(\mathbb{Z})$ defined in (6.16), our choice of $V_2(\mathbb{Z})$ in (6.27) is not the largest valid space: there exist larger vector spaces such that $V_2(\mathbb{R}) \subset \mathcal{H}_{L_2}(\mathbb{R})$. However, the support restriction implies that for any $s_2 = \sum_{k \in \mathbb{Z}} c_2[k] \beta_{L_2^* L_2}(\cdot - k) \in V_2(\mathbb{R})$, the function $L_2\{s_2\} = \sum_{k \in \mathbb{Z}} (d_{L_2} * c_2)[k] \beta_{L_2^*}^\vee(\cdot - k)$ has a finite support. This is a desirable property both for simplicity of implementation and because it conforms with Theorem 6.1, since s_2^* in (6.12) also satisfies this property. Our specific choice of support for $(d_{L_2} * c_2)$ is guided by boundary considerations and will be justified in the proof of Proposition 6.4.*

The restriction to a finite number of active spline coefficients leads to finite-dimensional system and regularization matrices. The system matrices are of the form

$$\mathbf{H}_i \triangleq \left[\nu(\varphi_{i,m_i}) \quad \dots \quad \nu(\varphi_{i,M_i}) \right] \in \mathbb{R}^{M \times N_i}. \quad (6.28)$$

The regularization matrix for the sparse component, denoted by $\mathbf{L}_1 \in \mathbb{R}^{(N_1 - D_1 + 1) \times N_1}$, is

of the form

$$\mathbf{L}_1 \triangleq \begin{pmatrix} d_{L_1}[D_1 - 1] & \cdots & d_{L_1}[0] & 0 & \cdots & 0 \\ 0 & \ddots & & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & & \ddots & 0 \\ 0 & \cdots & 0 & d_{L_1}[D_1 - 1] & \cdots & d_{L_1}[0] \end{pmatrix}. \quad (6.29)$$

The second component requires a careful handling of the boundaries in order to achieve exact discretization, which leads to a more complicated expression for the associated regularization matrix. The regularization matrix $\mathbf{L}_2 \in \mathbb{R}^{(N_2-1) \times N_2}$ for the smooth component is given by

$$\mathbf{L}_2 \triangleq \left(\begin{array}{c|c} \mathbf{M}^- & \mathbf{0} \\ \hline & \mathbf{M} \\ \hline \mathbf{0} & \mathbf{M}^+ \end{array} \right), \quad (6.30)$$

where the central matrix $\mathbf{M} \in \mathbb{R}^{(N_2-G+1) \times N_2}$ is given by

$$\mathbf{M} \triangleq \begin{pmatrix} g[G-1] & \cdots & g[0] & 0 & \cdots & 0 \\ 0 & \ddots & & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & & \ddots & 0 \\ 0 & \cdots & 0 & g[G-1] & \cdots & g[0] \end{pmatrix}, \quad (6.31)$$

where g is defined in Proposition 6.2. The matrices $\mathbf{M}^\pm \in \mathbb{R}^{(B-1) \times (G-1)}$ are defined as $[\mathbf{M}^-]_{i,j} \triangleq g^{-(B_2-i)}[G-B+(i-1)-(j-1)]$ and $[\mathbf{M}^+]_{i,j} \triangleq g^{+i}[G+(i-1)-j]$ for $1 \leq i \leq (B_2-1)$ and $1 \leq j \leq (G-1)$, where the filter $g^{\pm k}$ are given by $g^{-k} \triangleq b^{1/2}|_{\{0, \dots, B-1-k\}} * d_{L_2}$ (supported in $[0 \dots G-1-k]$) and $g^{+k} \triangleq b^{1/2}|_{\{k, \dots, B-1\}} * d_{L_2}$ (supported in $\{k, \dots, G-1\}$). Here, the notation $a|_J$ refers to the filter a restricted to the set J of indices, with $a|_J[k] = a[k]$ if $k \in J$, and $a|_J[k] = 0$ otherwise.

As an illustration, for $L_2 = D$, we have that $B = 1$ and, hence, simply that $\mathbf{L}_2 = \mathbf{M}$. For $L_2 = D^2$, we have $B = 2$ and

$$\mathbf{L}_2 = \left(\begin{array}{cccccc} C & -2C & C & 0 & \cdots & \cdots & 0 \\ \hline & & & \mathbf{M} & & & \\ \hline 0 & \cdots & \cdots & 0 & C' & -2C' & C' \end{array} \right), \quad (6.32)$$

where $C \triangleq \sqrt{\frac{2-\sqrt{3}}{6}}$ and $C' \triangleq C(2+\sqrt{3})$.

Finally, we introduce the matrix $\mathbf{A} \in \mathbb{R}^{N_0 \times N_1}$ associated to the boundary condition functionals ϕ_0 . The constraint $\phi_0(\sum_{k \in \mathbb{Z}} c_1[k] \varphi_{1,k}) = \mathbf{0}$ leads to N_0 linear constraints on the coefficients $\mathbf{c}_1 = (c_1[m_1], \dots, c_1[M_1])$, which can be written in matrix form as

$\mathbf{A}\mathbf{c}_1 = \mathbf{0}$. Our practical choice of boundary condition functionals ϕ_0 is presented in (2.48) in the previous chapter. With this choice, these constraints simply lead to the N_0 first coefficients of \mathbf{c}_1 to be set to zero, which thus reduces the dimension of the optimization problem.

These matrices enable an exact discretization of Problem (6.23), as shown in Proposition 6.4.

Proposition 6.4 (Recasting as a Finite Problem). *Let $\varphi_{1,k} \triangleq \beta_{L_1}(\cdot - k)$, $\varphi_{2,k} \triangleq \beta_{L_2^* L_2}(\cdot - k)$, and let the assumptions in Section 6.5.1 be satisfied. Then, Problem (6.23) is equivalent to the optimization problem*

$$\mathcal{V}_f \triangleq \underset{\substack{(\mathbf{c}_1, \mathbf{c}_2) \in \mathbb{R}^{N_1} \times \mathbb{R}^{N_2} \\ \mathbf{A}\mathbf{c}_1 = \mathbf{0}}}{\arg \min} \mathcal{J}_f(\mathbf{c}_1, \mathbf{c}_2), \quad (6.33)$$

where the cost function is given by

$$\mathcal{J}_f(\mathbf{c}_1, \mathbf{c}_2) \triangleq E(\mathbf{H}_1 \mathbf{c}_1 + \mathbf{H}_2 \mathbf{c}_2, \mathbf{y}) + \lambda_1 \|\mathbf{L}_1 \mathbf{c}_1\|_1 + \lambda_2 \|\mathbf{L}_2 \mathbf{c}_2\|_2^2. \quad (6.34)$$

The matrices \mathbf{H}_i and \mathbf{L}_i for $i \in \{1, 2\}$ are defined in (6.28), (6.29), and (6.30). This equivalence holds in the sense that there exists a bijective linear mapping from \mathcal{V}_d to \mathcal{V}_f

$$(c_1, c_2) \mapsto \left((c_1[m_1], \dots, c_1[M_1]), (c_2[m_2], \dots, c_2[M_2]) \right) \quad (6.35)$$

between their solution sets.

Proof. Let $s_i = \sum_{k \in \mathbb{Z}} c_i[k] \varphi_{i,k}$ with $c_i \in V_i(\mathbb{Z})$ for $i \in \{1, 2\}$. The filters c_i are assumed to have values determined by the vector $\mathbf{c}_i = (c_i[m_i], \dots, c_i[M_i])$ at certain points. By definition of m_i and M_i , the values of c_i outside these intervals do not affect the measurements $\nu(s_i)$, and we clearly have that $\nu(s_i) = \mathbf{H}_i \mathbf{c}_i$. Therefore, these coefficients solely affect the regularization terms. We now show that, for a solution $(c_1, c_2) \in \mathcal{V}_d$ to Problem (6.23), the coefficients are uniquely determined by the vectors \mathbf{c}_i , and that the regularization terms $\|d_{L_1} * c_1\|_{\ell_1}$ and $\langle c_2, \rho * c_2 \rangle_{\ell_2}$ can thus be expressed exclusively in terms of these vectors.

Concerning the first component, this is proved in Proposition 4.1, which shows that $\|\mathbf{L}_1 \{s_1\}\|_{\mathcal{M}} = \|\mathbf{L}_1 \mathbf{c}_1\|_1$. The additional constraint $\mathbf{A}\mathbf{c}_1$ comes from $\phi_0(s_1) = \mathbf{0}$ imposed on the search space $V_1(\mathbb{Z})$ of Problem (6.23).

We now consider the regularization term for the component $\langle c_2, \rho * c_2 \rangle_{\ell_2}$. By Proposition 6.2, we have that $\rho = g * g^\vee$, and, hence, that $\langle c_2, \rho * c_2 \rangle_{\ell_2} = \langle g * c_2, g * c_2 \rangle_{\ell_2} = \|g * c_2\|_{\ell_2}^2$, where $g = b^{1/2} * d_{L_2}$. We also have $g * c_2 = b^{1/2} * a$, where $a = d_{L_2} * c_2$ is supported in $[1 \dots M_2]$ by definition of the native space $V_2(\mathbb{Z})$ given in (6.27). Since $a[n] = \sum_{k=0}^{D_2-1} d_{L_2}[k] c_2[n-k]$, $a[n]$ is entirely determined by the vector \mathbf{c}_2 for $1 \leq n \leq M_2$,

which justifies our choice of the space $V_2(\mathbb{Z})$. For values of n outside this interval, there is a unique way of setting the coefficients $c_2[k]$ in order to nullify $a[n]$ and thus obtain that $c_2 \in V_2(\mathbb{Z})$. For example, $c_2[M_2 + 1]$ can be set to nullify $a[M_2 + 1]$ based on the $(D_2 - 1)$ previous coefficients of c_2 , and, similarly, all the $c_2[n]$ for $n > M_2 + 1$ can be set recursively to nullify all the $a[k]$ for all $k > M_2 + 1$. The same argument can be made to show that there is a unique choice $c_2[n]$ for $n < m_2$ that nullifies $a[k]$ for all $k < 1$.

We now compute the values of $(g * c_2)[n]$ in different regimes for n . We have that $(g * c_2)[n] = (b^{1/2} * a)[n] = \sum_{k=0}^{B_2-1} b^{1/2}[k]a[n-k]$ where a is supported in $[1 \dots M_2]$. For $B_2 \leq n \leq M_2$, this sum is solely affected by the coefficients $\mathbf{c}_2 = (c_2[m_2], \dots, c_2[M_2])$, so that the corresponding terms can be written in matrix form as $\mathbf{M}\mathbf{c}$ (the central part of the \mathbf{L}_2 matrix defined in (6.30)). Outside this interval, for example for $n = M_2 + 1$, we have that $(g * c_2)[n] = \sum_{k \in \mathbb{Z}} b^{1/2}|_{\{1, \dots, B_2-1\}}[k]a[n-k]$, since the $k = 0$ term is anyway nullified by the fact that $a[n] = 0$. An analogous reformulation allows us to have $(g * c_2)[n]$ only depend on the \mathbf{c}_2 coefficients. The same reformulation for all the coefficients $M_2 + 1 \leq n \leq (M_2 + B_2 - 1)$ leads to the matrix \mathbf{M}^+ in (6.30), while a similar argument for coefficients $(g * c_2)[n]$ with $1 \leq n \leq (B_2 - 1)$ leads to the matrix \mathbf{M}^- .

We have thus proved that the solutions $(c_1, c_2) \in \mathcal{V}_d$ to Problem (6.23) are uniquely determined by their coefficients $\mathbf{c}_i = (c_i[m_i], \dots, c_i[M_i])$ for $i \in \{1, 2\}$, and that the regularization terms can be written $\|d_{L_1} * c_1\|_{\ell_1} = \|\mathbf{L}_1 \mathbf{c}_1\|_1$ and $\|g * c_2\|_{\ell_2}^2 = \|\mathbf{L}_2 \mathbf{c}_2\|_{L_2}^2$. This, together with the fact that $\boldsymbol{\nu}(\sum_{k \in \mathbb{Z}} c_i[k] \varphi_{i,k}) = \mathbf{H}_i \mathbf{c}_i$, proves that $\mathcal{J}_d(c_1, c_2) = \mathcal{J}_f(\mathbf{c}_1, \mathbf{c}_2)$. Conversely, for any $(\mathbf{c}_1, \mathbf{c}_2) \in \mathbb{R}^{N_1} \times \mathbb{R}^{N_2}$, there is a unique extension of these vectors to filters $c_i \in V_i(\mathbb{R})$ such that $\mathbf{c}_i = (c_i[m_i], \dots, c_i[M_i])$ and $\mathcal{J}_d(c_1, c_2) = \mathcal{J}_f(\mathbf{c}_1, \mathbf{c}_2)$. These extensions are explicated in [114, Proposition 2] for c_1 and earlier in this proof for c_2 . This proves the existence of the bijective linear mapping between the solution sets \mathcal{V}_d and \mathcal{V} specified in Proposition 6.4. \square

The combination of Propositions 6.3 and 6.4 allows us to solve the continuous-domain infinite-dimensional Problem (6.25) by finding a solution $(\mathbf{c}_1^*, \mathbf{c}_2^*) \in \mathcal{V}_f$ of the finite-dimensional Problem (6.33). We obtain the corresponding solution of (6.25) by extending these vectors to digital filters $(c_1^*, c_2^*) \in \mathcal{V}_d$ (this extension is unique as specified by Proposition 6.4), which yields the continuous-domain reconstruction $s^* = s_1^* + s_2^*$, where $s_i^* \triangleq \sum_{k \in \mathbb{Z}} c_d^*[k] \varphi_{i,k}$.

6.5.3 Sparsification Step

Although Problem (6.33) can be solved using standard solvers such as the alternating-direction method of multipliers (ADMM), there is no guarantee that such solvers will yield a solution of the desired form specified by Theorem 6.1, *i.e.*, s_1^* is an L_1 -spline with fewer than $(M - N_{0,1})$ knots, and s_2^* is a sum of M kernel functions and a null space element. This is a particularly relevant observation for the first component since, at fixed second

component s_2^* , only extreme-point solutions s_1^* of Problem (6.10) take the prescribed form [14]. This problem can be alleviated by computing a solution $(\mathbf{c}_1^*, \mathbf{c}_2^*)$ to Problem (6.33), and then finding an extreme point of the solution set $\mathbf{c}_1^{\text{extr}} \in \arg \min_{\mathbf{c}_1 \in \mathbb{R}^{N_1}} \mathcal{J}_f(\mathbf{c}_1, \mathbf{c}_2^*)$, which leads to a solution $(\mathbf{c}_1^{\text{extr}}, \mathbf{c}_2^*)$ of the prescribed form. This is achieved by recasting the problem as a linear program and using the simplex algorithm [116] to reach an extreme-point solution [45, Theorem 7].

6.6 Experimental Results

We now validate our reconstruction algorithm in a simulated setting.

6.6.1 Experimental Setting

Grid Size

Following Chapters 4 and 5, we rescale the problem by a factor T so that the interval of interest I_T is mapped into $[0, 1]$. We tune the fineness of the grid (and the dimension of the optimization task) by varying T , which amounts to varying the grid size $h = 1/T$ in the rescaled problem.

Ground Truth

We generate a ground-truth signal $s^{\text{GT}} = s_1^{\text{GT}} + s_2^{\text{GT}}$. The sparse component s_1^{GT} is chosen to be an L_1 -spline of the form (3.2) with few jumps, for which gTV is an adequate choice of regularization, as demonstrated by (6.11) in our representer theorem. For the smooth component s_2^{GT} , we generate a realization of a solution s_2 of the stochastic differential equation $L_2 s_2 = w$, where w is a Gaussian white noise with standard deviation σ_2 by following the method of [163]. The operator L_2 then acts as a whitening operator for the stochastic process s_2 . The reason for this choice is the connection between the minimum mean-square estimation of such stochastic processes and the solutions to variational problems with gTikhonov regularization $\|L_2\{s_2\}\|_{L_2}^2$ [43, 46, 164].

Forward Operator

Our forward model is the Fourier-domain cosine sampling operator of the form $\nu_1(s) = \int_0^1 s(t)dt$ (DC term) and

$$\nu_m(s) = \int_0^1 \cos(\omega_m t + \theta_m) s(t) dt \quad (6.36)$$

for $2 \leq m \leq M$, where the sampling pulsations ω_m are chosen at random within the interval $(0, \omega_{\max}]$, and the phases θ_m are chosen at random within the interval $[0, 2\pi)$. Notice that ν_m is a Fourier-domain measurement of the restriction of s to the interval of interest $[0, 1]$, in conformity with the finite-domain assumption in Section 6.5.1. As discussed in Example 3.3, ν_m is weak*-continuous and thus also conforms with the assumptions of Theorem 6.1.

For the data-fidelity term, we use the standard quadratic error $E(\mathbf{x}, \mathbf{y}) \triangleq \frac{1}{2} \|\mathbf{x} - \mathbf{y}\|_2^2$.

6.6.2 Comparison with Noncomposite Models

We now validate our new sparse-plus-smooth model against more standard noncomposite models. More precisely, for $i \in \{1, 2\}$ we solve the regularized problems

$$\arg \min_{f \in \mathcal{X}_i} \left(E(\nu(f), \mathbf{y}) + \lambda \mathcal{R}_i(f) \right) \quad (6.37)$$

with regularizers $\mathcal{R}_1(f) \triangleq \|\mathbf{L}_1\{f\}\|_{\mathcal{M}}$ (sparse model with native space $\mathcal{X}_1 \triangleq \mathcal{M}_{L_1}(\mathbb{R})$) and $\mathcal{R}_2(f) \triangleq \|\mathbf{L}_2\{f\}\|_{L_2}$ (smooth model with native space $\mathcal{X}_2 \triangleq \mathcal{H}_{L_2}(\mathbb{R})$). We discretize these problems using the reconstruction spaces $V_i(\mathbb{R})$ described in this chapter (without restricting $V_1(\mathbb{R})$ with the boundary conditions ϕ_0). The sparse model thus amounts to an ℓ_1 -regularized discrete problem which we solve using ADMM (see Appendix A), while the smooth model has a closed-form solution that can be obtained by inverting a matrix.

For this comparison, we choose regularization operators $L_1 = D$ and $L_2 = D^2$ with $M = 50$ Fourier-domain measurements (cosine sampling with $\omega_{\max} = 100$). We generate the ground-truth signal according to Section 6.6.1, with $K_1 = 5$ jumps whose i.i.d. Gaussian amplitudes have the variance $\sigma_1^2 = 1$ for s_1^{GT} . For the smooth component s_2^{GT} , we generate a realization of a Gaussian white noise w with the variance $\sigma_2^2 = 100$, such that $L_2\{s_2^{\text{GT}}\} = w$. The measurements are corrupted by some i.i.d. Gaussian white noise $\mathbf{n} \in \mathbb{R}^M$ so that $\mathbf{y} = \nu(s_{\text{GT}}) + \mathbf{n}$. We set the signal-to-noise ratio (SNR) between $\nu(s_{\text{GT}})$ and \mathbf{n} to be 50 dB. The regularization parameters are selected through a grid search with $h = 1/2^9$ to maximize the SNR of the reconstructed signal s with respect to the ground truth (defined as $\text{SNR} \triangleq 10 \log_{10} \left(\frac{\int_0^T (s^{\text{GT}}(t))^2 dt}{\int_0^T (s(t) - s^{\text{GT}}(t))^2 dt} \right)$).

The results of this comparison in terms of SNR are shown in Table 6.2 for varying grid sizes h . For all methods, the SNR values increase when the grid size decreases, which is to be expected since the grids are embedded. The only exception is the sparse-plus-smooth reconstruction for the finest grid size, which is likely due to numerical issues arising from the increased dimension ($N \approx 2^{11}$) of the optimization problem. The effect of the grid size on the quality of the reconstruction varies between the models: it is almost nonexistent for the smooth-only model, whereas it is most significant for our sparse-plus-smooth model. Over all grid sizes, due to the fact that our sparse-plus-smooth signal model

h	Sparse plus smooth	Sparse only	Smooth only
2^6	18.02	17.72	18.16
2^7	21.87	21.08	18.16
2^8	23.46	22.07	18.16
2^9	27.02	23.01	18.16
2^{10}	25.70	23.04	18.16

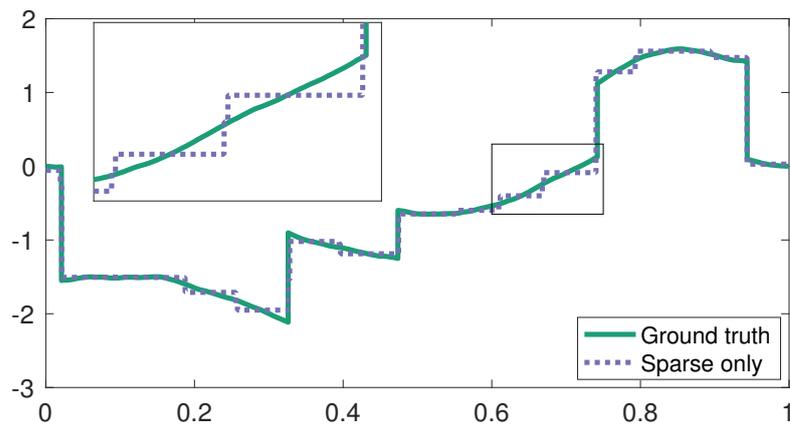
Table 6.2 – SNR values (in dB) of the reconstructed signal with respect to the ground truth with varying grid size h .

matches the ground truth, our reconstructed signal yields a higher SNR (27.02 dB) than the sparse-only (23.04 dB) and smooth-only (18.16 dB) models.

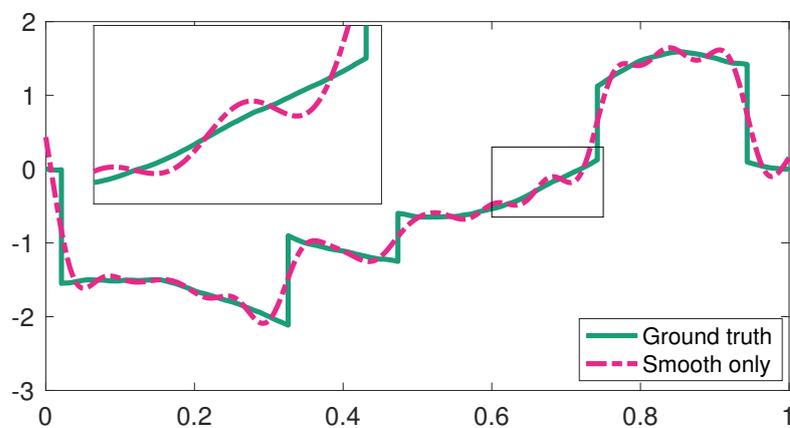
The reconstruction results for the grid size $h = 1/2^9$ are shown in Figure 6.2. Our sparse-plus-smooth reconstruction is qualitatively much more satisfactory. As can be observed in the zoomed-in section, the sparse-only model is subject to a staircasing phenomenon in the smooth regions of the ground-truth signal, a well-known shortcoming of TV regularization for functions. Conversely, our reconstruction is remarkably accurate in the smooth regions. Finally, the smooth-only model fails both visually and in terms of SNR, due to its inability to represent sharp jumps.

6.7 Conclusion

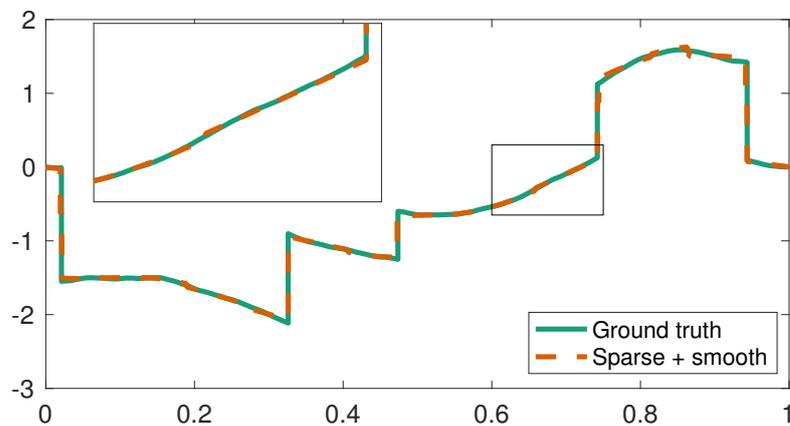
We have introduced a continuous-domain framework for the reconstruction of multicomponent signals. It assumes two additive components, the first one being sparse and the other being smooth. The reconstruction is performed by solving a regularized inverse problem, using a finite number of measurements of the signal. The form of a solution to this problem is given by our representer theorem. This form justifies the choice of the search space in which we discretize the problem. Our discretization is exact, in the sense that it amounts to solving a continuous-domain optimization problem restricted to our search space. The discretized problem is then solved using our ADMM-based algorithm, which we validate on simulated data.



(a) Sparse-only model: SNR = 23.01 dB with $\lambda = 10^{-9}$.



(b) Smooth-only model: SNR = 18.16 dB with $\lambda = 10^{-11}$.



(c) Sparse-plus-smooth model: SNR = 27.02 dB with $\lambda_1 = 8 \cdot 10^{-7}$ and $\lambda_2 = 5 \cdot 10^{-10}$.

Figure 6.2 – Comparison between our sparse-plus-smooth model and single-component models with regularization operators $L_1 = D$, $L_2 = D^2$, and $M = 50$ Fourier-domain measurements.

Part III (Non)uniqueness and the Study of the Solution Set of TV-Based Inverse Problems

Introduction

In the discrete domain, it is well known that finite-dimensional ℓ_1 regularization can lead to nonunique solutions [21, 165]. However, considering the abundant literature focusing on ℓ_1 -based optimization, there are surprisingly few works that study these cases of nonuniqueness. Conversely, in the compressed sensing (CS) literature, a host of works focus on the identification of conditions under which uniqueness of ℓ_1 -based problems does hold (see the introduction of Part II). However, the CS framework typically does not assume that the observations are generic, but that they are generated via a ground-truth signal. Under certain conditions based on the incoherence of the sensing matrix, this ground-truth signal can then be recovered as the unique solution to an ℓ_1 -based optimization problem [8, 31, 32]. Cases of nonuniqueness are thus excluded from such frameworks, which explains the scarcity of works on this topic.

Our Approach: Generic Observations $\mathbf{y} \in \mathbb{R}^M$

However, in practical applications, the sensing matrix is often determined by the acquisition system and thus cannot be chosen freely. Moreover, the assumption that the observations are generated via a (nearly) sparse ground-truth signal may not hold in practice. The same analysis holds true for our continuous-domain total-variation (TV)-based problems of interest, since the TV norm for measures is the continuous sparsity-promoting equivalent of the ℓ_1 norm. Hence, we take a different perspective from the traditional CS one by assuming that we are *agnostic* to how the observations \mathbf{y} are generated: they can be any vector $\mathbf{y} \in \mathbb{R}^M$. Our objective is then to study the solution set of our optimization problems, in particular by answering the following questions:

- identification of cases of uniqueness based on the observations \mathbf{y} ;

- whenever uniqueness does not hold, identification of the *sparsest* solution(s).

Uniqueness Studies for TV-based Optimization

We start by reviewing the existing literature that focuses on the identification of cases of uniqueness for TV-based regularization. These results are analogous to discrete CS-type results: the underlying assumption is that the observations $\mathbf{y} \in \mathbb{R}^M$ are generated via a sparse measure $w_0 \in \mathcal{M}(\mathbb{K})$, where $\mathcal{M}(\mathbb{K})$ is the space of Radon measures (defined in Section 2.2) with $\mathbb{K} = \mathbb{R}$ or \mathbb{T} . We thus have $\boldsymbol{\nu}(w_0) \approx \mathbf{y}$, where $\boldsymbol{\nu} : w \mapsto \boldsymbol{\nu}(w) \in \mathbb{R}^M$ is the forward model. The question then becomes whether w_0 can be uniquely recovered, either exactly or in a stable way, via a TV-based problem such as the *basis pursuit in the continuum* (BPC):

$$\arg \min_{w \in \mathcal{M}(\mathbb{K}): \boldsymbol{\nu}(w) = \mathbf{y}} \|w\|_{\mathcal{M}}. \quad (\text{III.1})$$

In [47], de Castro and Gamboa introduced the concept of *extrema Jordan type measure* (see [47, Definition 1]), which gives sufficient conditions on a given signed (with positive and negative weights) sparse measure to be the unique solution of a TV-based optimization problem. They also proved that when the ground-truth measure is nonnegative, k -sparse, and the number of observations is greater or equal than $2k + 1$, then it is the unique solution of (III.1) if the measurement operator $\boldsymbol{\nu}$ is defined from a T-system. Note that it has recently been proved [53, 166–168] that in the nonnegativity setting, a k -sparse nonnegative measure can be uniquely recovered from at least $2k + 1$ measurements and a nonnegativity constraint without the need for TV minimization.

Candès and Fernandez-Granda also studied the super-resolution problem of recovering a ground-truth sparse Radon measure $w_0 \in \mathcal{M}(\mathbb{T})$ from its low-frequency measurements [49, 62]. They have shown that if the minimal distance between the spikes of w_0 is large enough, then (III.1) has a unique solution, which is w_0 itself [13, Theorem 1.2]. Duval and Peyré identified the so-called *nondegenerate source condition* [65, Definition 5], under which the uniqueness of the reconstruction together with the recovery of the support of the underlying ground-truth sparse measure are shown. These results are based on the key notion of *dual certificates*, which also play an important role in our works in this part. This notion has been introduced for discrete CS problems in [169] and connected to TV-based optimization problems in [47]. However, to the best of our knowledge, ours are the first works that apply this type of approach to problems with *generalized* TV (gTV) regularization, *i.e.*, problems similar to (III.1), but with an added regularization operator L .

Outline

In this part, we study the solution set of gTV-based problems by taking our agnostic approach described above. Note that various existing representer theorems [14, 22, 61, 87, 89, 95] also take this agnostic approach, and have identified that gTV-based problems usually lead to nonunique solution. However, due to the generality of their frameworks, their description of the solution set is limited; in particular, they do not focus on identifying cases of uniqueness. Here, we study less general problems which enables a much more precise description of the solution set. More specifically, we fully characterize cases of uniqueness based on the observations \mathbf{y} , and, in cases of nonuniqueness, we identify the sparsest solution and design efficient algorithms to reach it. We consider three different formulations with specific forward models ν and operators L :

- In Chapter 7, we consider the problem of interpolating data points with second-order TV regularization, *i.e.*, with a sampling forward model $\nu(f) = (f(x_1), \dots, f(x_M))$ where $x_m \in \mathbb{R}$, and with the second-derivative regularization operator $L = D^2$. In this case, we show that uniqueness is not systematic, and we design a simple algorithm that reaches the sparsest solution in all cases.
- In Chapter 8, we consider the same problem with an added Lipschitz constraint on the reconstructed signal to favor stable solutions. Once again, we show that uniqueness is not systematic, and we design an algorithm that reaches the sparsest solution.
- In Chapter 9, we consider the reconstruction of periodic signals ($\mathbb{K} = \mathbb{T}$) where the forward model ν consists of low-frequency Fourier series coefficients, and the regularization operator is a generic N_0 th-order derivative operator $L = D^{N_0}$. We prove that this problem always has a unique solution, and we design a B-spline-based algorithm based on Chapter 4 to solve it computationally.

7 Interpolation of Data with Second-Order TV Regularization

This chapter is based on the following publication [170]:

T. Debarre, Q. Denoyelle, M. Unser, and J. Fageot, “Sparsest Piecewise-Linear Regression of One-Dimensional Data”, *Journal of Computational and Applied Mathematics*, vol. 406, p. 114044, May 2022.

7.1 Introduction

Regression problems consist in learning a function f that best approximates some data $(x_m, y_m)_{m=1}^M$, where M is the number of data points, in the sense that $f(x_m) \approx y_m$. This is typically achieved by parametrizing f with a vector of parameters $\boldsymbol{\theta}$ as the parametric function $f_{\boldsymbol{\theta}}$, and then by minimizing some objective function with respect to $\boldsymbol{\theta}$. The oldest and most basic form of regression is linear regression: f is parametrized as a linear (or affine) function. Although this model has the advantage of being simple, it is very limited due to the fact that many data distributions are poorly approximated by linear functions, as illustrated by the dotted line example in Figure 7.1. The choice of parametrization $f_{\boldsymbol{\theta}}$ is therefore crucial, as it must strike an appropriate balance between two conflicting desirable properties. Firstly, in order to be suitable for a variety of problems, the parametric model should be flexible enough to represent a large class of functions. In the field of machine learning, where regression is known as *supervised learning*, this quest for universality is for instance underlined by several universal approximation theorems for artificial neural networks [171–173]. Secondly, the model should be simple enough so that it generalizes well to input vectors \mathbf{x} that are outside the training set. Indeed, a known pitfall of machine learning algorithms is overfitting, which happens when the model is unduly complex and fits the training data too closely [174, Chapter 3]. This leads to poor generalization abilities for out-of-sample data. This pitfall is often dealt with by adding some regularization to the objective function, which tends to lead to “simpler” reconstructions. The overarching guiding principle to avoid overfitting is Occam’s razor:

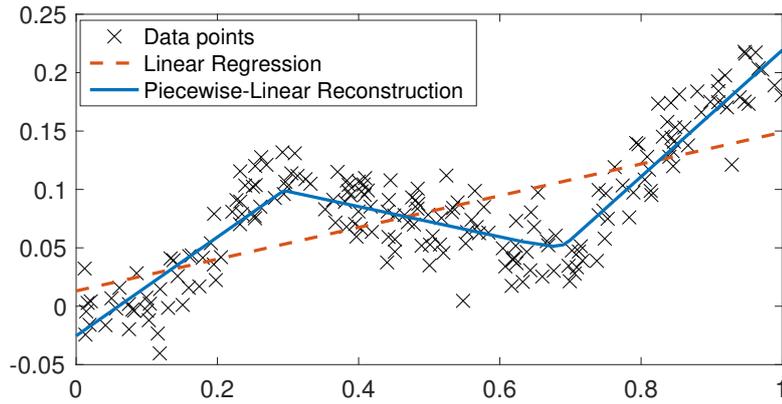


Figure 7.1 – Examples of reconstructions.

the simplest model that explains the data well will generalize better and should thus be selected.

7.1.1 Problem Formulation

In this chapter, we study the regression (or supervised learning) problem in one dimension, *i.e.*, $f : \mathbb{R} \rightarrow \mathbb{R}$ and $x_m, y_m \in \mathbb{R}$. Instead of parametrizing the reconstructed function, we formulate the learning problem as a regularized inverse problem in a continuous-domain framework. Inspired by their connection (that we discuss later on) to popular ReLU (rectified linear unit) neural networks, we focus on reconstructing piecewise-linear splines. Our metric for model simplicity is sparsity, *i.e.*, the number of spline knots. For regularization purposes, we therefore use the total-variation (TV) norm for measures $\|\cdot\|_{\mathcal{M}}$, which is defined over the space of bounded Radon measures $\mathcal{M}(\mathbb{R})$ (see Section 2.2). As demonstrated in Theorem 3.3, this norm is known to promote sparse solutions. We formulate the following optimization problem, which we refer to as the *generalized Beurling LASSO* (g-BLASSO)

$$\arg \min_{f \in \mathcal{M}_{D^2}(\mathbb{R})} \left(\sum_{m=1}^M E(f(x_m), y_m) + \lambda \|D^2 f\|_{\mathcal{M}} \right), \quad (\text{g-BLASSO})$$

where $\mathcal{M}_{D^2}(\mathbb{R})$ is the native space defined in Section 2.4 and E is a cost function that penalizes the discrepancy between $f(x_m)$ and the data $y_m \in \mathbb{R}$ (*e.g.*, a quadratic loss $E(z, y) = \frac{1}{2}(z - y)^2$). We assume that the sampling locations are ordered, *i.e.*, $x_1 < \dots < x_M$. The parameter $\lambda > 0$ balances the contribution of the data fidelity and the regularization, and D^2 is the second-derivative operator. The “generalized Beurling LASSO” terminology comes from the *Beurling LASSO* (BLASSO), which is used in the Dirac recovery literature [47]. Problem (g-BLASSO) is a generalization of the BLASSO due to the presence of a regularization operator D^2 , which is not present in the latter problem. By Theorem 3.3, the extreme-point solutions of the (g-BLASSO) are

piecewise-linear splines of the form

$$f_{\text{opt}}(x) = b_0 + b_1x + \sum_{k=1}^K a_k(x - \tau_k)_+, \tag{7.1}$$

where $x_+ = \max(0, x)$ is the ReLU, $b_0, b_1, a_k, \tau_k \in \mathbb{R}$, and the number of spline knots K is bounded by $K \leq M - 2$. This representer theorem has two important components:

- the (g-BLASSO) has solutions of the prescribed form, *i.e.*, piecewise-linear splines. This stems from the choice of the regularization, *i.e.*, the TV norm of the second derivative;
- the sparsity is bounded by the number of training samples by $K \leq M - 2$.

In terms of model simplicity, the bound $K \leq M - 2$ is typically uninformative in machine learning problems: in Figure 7.1, it yields $K \leq M - 2 = 198$, which is clearly much higher than the desired sparsity. However, this bound is independent of the regularization parameter λ , which has a major effect on the sparsity of the reconstructed signal. Indeed, $\lambda \rightarrow 0$ will roughly lead to a learned function f that interpolates all the data points, with typically close to $K = M - 2$ knots. At the other extreme, the limit $\lambda \rightarrow +\infty$ leads to linear regression and thus sparsity $K = 0$ due to the fact that linear functions are not penalized by the regularization. Therefore, the interesting case is the intermediate regime (as illustrated by the solid curve in Figure 7.1), in which the overall trend is that the sparsity K decreases as λ increases. Hence, λ controls the universality *versus* simplicity trade-off.

7.1.2 Summary of Contributions and Outline

The above purely qualitative observation is far from telling the whole story. In particular, it does not explain how λ should be chosen in practice. We attempt to overcome this impediment by giving a full description of the solution set of the (g-BLASSO). The basis of our analysis is the classical observation that when E is strictly convex, there exists a unique vector $\mathbf{y}_\lambda = (y_{\lambda,1}, \dots, y_{\lambda,M}) \in \mathbb{R}^M$ such that the (g-BLASSO) is equivalent to the constrained problem

$$\arg \min_{\substack{f \in \mathcal{M}_{\mathbb{D}^2}(\mathbb{R}): \\ f(x_m) = y_{\lambda,m}, \\ m \in \{1, \dots, M\}}} \|D^2 f\|_{\mathcal{M}} \tag{g-BPC}$$

(see for instance [45, Theorem 5]). We refer to this constrained problem as the *generalized basis pursuit in the continuum* (g-BPC)¹. Our terminology is inspired by the (discrete) basis pursuit (BP) [6], which is also a constrained problem; as for the (g-BLASSO), the

¹A similar terminology, the “continuous basis pursuit”, is used in a different context in [67, 175].

“generalization” is due to the presence of a regularization operator D^2 , which is absent in the BP. We therefore carry out our theoretical analysis on the simpler (g-BPC) problem, and we attest that these results apply to the (g-BLASSO) as well, provided that \mathbf{y}_λ is known. For this analysis, we use mathematical tools based on duality theory, and we exploit the very specific form of the so-called dual certificate for our regularization operator D^2 . We describe systematically the form of the solution set and identify the set of sparsest solutions. The fact that optimization problems with sparsity-promoting regularization sometimes have multiple solutions is often sidestepped in the literature by identifying specific cases of uniqueness [47, 49, 62, 65]. When it is not, existing works typically provide the form of certain solutions [14, 87], but they do not characterize cases of uniqueness nor do they give a *complete* description of the solution set as we do here. Concerning our specific problem, it is known that the function that simply connects the points $(x_1, y_{0,1}), \dots, (x_M, y_{0,M})$ is always a solution to the (g-BPC) (see [176, Theorem 1] and [177, Proposition 7]). We refer to it as the *canonical solution*. Building on this result, our contributions on the theoretical and algorithmic sides concerning the (g-BLASSO) are summarized below.

1. Theory

Our main theoretical contributions are the following.

- In Section 7.3, we fully describe the solution set of the (g-BPC) by specifying the intervals in which all solutions follow the canonical solution, and those in which they do not (Theorem 7.1). This allows us to characterize the cases where the (g-BPC) admits a unique solution. When they differ, we give a geometrical description of the set in which the graph of all solutions lies in Theorem 7.2.
- When there are multiple solutions, the canonical solution can be made sparser in certain regions, which is the topic of Section 7.4. More precisely, in Theorem 7.3, we express the minimum achievable sparsity of a solution to the (g-BPC) as a simple function of $\mathbf{x} \triangleq (x_1, \dots, x_M)$ and \mathbf{y}_0 , which we denote by $K_{\min}(\mathbf{x}, \mathbf{y}_0)$. Concerning the solution set, we fully describe the set of sparsest solutions of the (g-BPC). In particular, we characterize the cases of uniqueness, and provide a description of the sparsest solutions together with the number of degrees of freedom $n_{\text{free}}(\mathbf{x}, \mathbf{y}_0)$, that we characterize and show to be finite.
- In Section 7.5.1, we extend the results of the first two items to the (g-BLASSO). This is a consequence of the aforementioned equivalence between the (g-BLASSO) and the (g-BPC) problems, which is proved in Proposition 7.7. We also specify the limit value λ_{\max} , for which any $\lambda \geq \lambda_{\max}$ amounts to linear regression in Proposition 7.10.

2. Algorithm

These theoretical findings warrant our simple and fast algorithm, presented in Section 7.5.2, for reaching (one of) the sparsest solution(s) to the (g-BLASSO). The algorithm, which is agnostic to uniqueness, is divided in two parts: first, we compute the \mathbf{y}_λ vector for the (g-BPC) problem by solving a standard discrete ℓ_1 -regularized problem. Next, we find a sparsest solution to the (g-BLASSO) (with sparsity $K_{\min}(\mathbf{x}, \mathbf{y}_\lambda)$) by sparsifying its canonical solution in some prescribed regions that are determined by our theoretical results. This sparsification step is detailed in Algorithm 7.1 and has complexity $\mathcal{O}(M)$.

This two-step algorithm provides a simple and fast way for the user to judiciously choose λ by evaluating the data fidelity loss $\sum_{m=1}^M E(f(x_m), y_m)$ versus the optimal sparsity $K_{\min}(\mathbf{x}, \mathbf{y}_0)$ —which depends on λ —as a proxy for the universality versus simplicity trade-off. We illustrate this in our experiments in Section 7.6. We may select $\lambda \in (0, \lambda_{\max}]$, where the limit $\lambda \rightarrow 0$ amounts to the (g-BPC) problem and λ_{\max} is the upper bound mentioned above. Note that existing algorithms that solve the (g-BLASSO) such as that introduced in [177] are a lot more complex and computationally expensive. Moreover, to the best of our knowledge, no existing algorithm has the guarantee of reaching a sparsest solution of the (g-BPC) or the (g-BLASSO).

7.1.3 Related Works

Previous Studies of the (g-BLASSO)

In [178], Pinkus proved that the canonical solution—that simply connects the data points—is the unique solution to the (g-BPC) in some special cases, a result that we recover in our analysis. Later, Koenker *et al.* [176, Theorem 1] and Mammen and van de Geer [177, Proposition 7] proved that the canonical solution is indeed a solution to the (g-BPC). These works also propose algorithms to solve the (g-BLASSO) for any value of λ . However, contrary to this chapter, none of the aforementioned works describe the full solution set of the (g-BPC), nor identifies its sparsest solutions. Recently, there has been a promising new surge of works on related problems, both on the theoretical and the algorithmic sides [61, 89, 90, 92, 95, 179, 180]. Several very general theories that incorporate the (g-BLASSO) and the (g-BPC) and that deal with optimization in Banach spaces with various differential regularization operators, have also been recently developed; we refer to our literature review in Section 1.2.4.

ReLU Networks, Piecewise-Linear Splines, and the (g-BLASSO)

A modern approach to supervised learning is neural networks, which in recent years have become the gold standard for an impressive number of applications [181]. Many recent papers have highlighted the property that today’s state-of-the-art convolutional neural

networks (CNNs) with rectified linear unit (ReLU) activations specify an input-output relation $f : \mathbb{R}^d \rightarrow \mathbb{R}$, where d is the number of dimensions, that is continuous and piecewise-linear (CPWL) [182, 183]. This result stems from the fact that the ReLU nonlinearity is itself a CPWL function, as well as, for instance, the widespread max-pooling operation. In fact, there are indications that using more general piecewise-linear splines as activation functions could be more effective than restricting to the ReLU or leaky ReLU [179, 184, 185]. In the one-dimensional case $d = 1$, it follows that the learned function of a ReLU network is a piecewise-linear spline [186], just like the solutions to the (g-BLASSO) given by (7.1). The trade-off between universality and Occam’s razor is then determined by the network size and architecture. Many recent papers in the literature have investigated this connection between ReLU networks and piecewise-linear splines [187, 188], including universality properties [186, 189, 190]. We also mention [191], which considers more general spline activation functions.

Moreover, several works have specifically underscored the relevance of the (g-BLASSO)—or related problems [192]—in machine learning by showing that it is equivalent to the training of a one-dimensional ReLU network with standard weight decay [193, 194]. Therefore, although the current trend of overparametrizing neural networks is somewhat antagonistic to our paradigm of sparsity, our full description of the solution set of the (g-BLASSO) (including its nonsparse solutions) could be relevant to the neural network community. Others recent works have designed multidimensional ($d > 1$) equivalents of the regularization term $\|D^2 f\|_{\mathcal{M}}$ and derive similar connections to neural networks [195, 196].

7.2 Mathematical Preliminaries

The task of recovering a continuous-domain function from finitely many samples is obviously ill-posed; this issue is commonly addressed by adding a regularization term. As a regularization norm, we consider $\|\cdot\|_{\mathcal{M}}$, which is the continuous-domain counterpart of the ℓ_1 -norm in the sense that it promotes sparse solutions, as demonstrated in Theorem 3.3.

7.2.1 The Native Space $\mathcal{M}_{D^2}(\mathbb{R})$

In this section, we recall the definition of the search space $\mathcal{M}_{D^2}(\mathbb{R})$ of the (g-BPC) and the (g-BLASSO) problems, and some of its important properties. This recollection is a summary of Section 2.4 in the special case of the second-derivative operator $L = D^2$. Following (2.33), the native space of D^2 is defined as

$$\mathcal{M}_{D^2}(\mathbb{R}) \triangleq \{f \in \mathcal{S}'(\mathbb{R}) : D^2\{f\} \in \mathcal{M}(\mathbb{R})\}, \quad (7.2)$$

with $D^2 : \mathcal{S}'(\mathbb{R}) \rightarrow \mathcal{S}'(\mathbb{R})$ the second-derivative operator and $\mathcal{M}(\mathbb{R})$ the space of Radon measures defined in (2.18). The space $\mathcal{M}_{D^2}(\mathbb{R})$ has been specifically studied in [179,

Section 2.2], and is sometimes denoted as $BV^{(2)}(\mathbb{R})$ in the literature. It is the second-order generalization of the well-known space of functions with bounded variation.

Next, we introduce the right-inverse operator of $D^2 : \mathcal{M}_{D^2}(\mathbb{R}) \rightarrow \mathcal{M}(\mathbb{R})$ defined in Theorem 2.1, which we call D_0^{-2} . For the latter, we use the canonical biorthogonal system for \mathcal{N}_{D^2} introduced in Proposition 2.4. The Schwartz kernel (2.36) of D_0^{-2} , which we call g_0 , is then given by

$$g_0(x, y) \triangleq (x - y)_+ - (-y)_+ + x((-y)_+ - (1 - y)_+). \quad (7.3)$$

By (2.43), any $f \in \mathcal{M}_{D^2}(\mathbb{R})$ can be uniquely decomposed as

$$\forall x \in \mathbb{R}, \quad f(x) = D_0^{-2}\{w\}(x) + \beta_0 + \beta_1 x, \quad (7.4)$$

where $w \in \mathcal{M}(\mathbb{R})$ and $\beta_0, \beta_1 \in \mathbb{R}$ satisfy

$$w = D^2 f, \quad \beta_0 = f(0), \quad \text{and} \quad \beta_1 = f(1) - f(0). \quad (7.5)$$

We call the measure w the *innovation* of f .

Finally, we recall the definition of L-splines, introduced in Section 3.1, for the special case $L = D^2$. Such splines f are called *piecewise-linear splines*, and are given by

$$f(x) = b_0 + b_1 x + \sum_{k=1}^K a_k (x - \tau_k)_+, \quad (7.6)$$

where $b_0, b_1 \in \mathbb{R}$, $K \geq 1$ is the number of knots, $x_+ \triangleq \max(x, 0)$, the $\tau_k \in \mathbb{R}$ are the knot locations, and the $a_k \in \mathbb{R}$ are the amplitudes. Note that this representation is different from that of (7.4) (in general, $(\beta_0, \beta_1) \neq (b_0, b_1)$); however we favor the representation (7.6) for splines due to its simplicity.

7.2.2 Problem Formulation and Representer Theorem for $\mathcal{M}_{D^2}(\mathbb{R})$

We now recall our inverse-problem formulations of interest—the (g-BPC) and the (g-BLASSO)—and the representer theorem from Chapter 3.3 applied to the special case when $L = D^2$ and the measurement operator ν is of the form $\nu(f) = (f(x_1), \dots, f(x_M))$.

Let $\mathbf{x} = (x_1, \dots, x_M) \in \mathbb{R}^M$ be a collection of distinct $M \geq 2$ ordered sampling locations and $\mathbf{y}_0 \in \mathbb{R}^M$. The noiseless (g-BPC) problem is formulated as

$$\mathcal{V}_0 \triangleq \arg \min_{\substack{f \in \mathcal{M}_{D^2}(\mathbb{R}) \\ f(x_m) = y_{0,m}, \quad m=1, \dots, M}} \|D^2 f\|_{\mathcal{M}}. \quad (\text{g-BPC})$$

Next, we fix $\lambda > 0$ and $\mathbf{y} \in \mathbb{R}^M$, together with a cost function $E : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}^+$ such that $E(\cdot, y)$ is strictly convex, coercive, and differentiable for any $y \in \mathbb{R}$ and $\lambda > 0$. The noisy (g-BLASSO) problem is then formulated as

$$\mathcal{V}_\lambda \triangleq \arg \min_{f \in \mathcal{M}_{\mathbb{D}^2}(\mathbb{R})} \left(\sum_{m=1}^M E(f(x_m), y_m) + \lambda \|D^2 f\|_{\mathcal{M}} \right). \quad (\text{g-BLASSO})$$

Then, Theorem 3.3 states that for any $\lambda \geq 0$ (including 0), \mathcal{V}_λ is nonempty, convex, and weak* compact in $\mathcal{M}_{\mathbb{D}^2}(\mathbb{R})$, and is the weak* closure of the convex hull of its extreme points. The latter are all piecewise-linear splines of the form

$$f_{\text{extreme}}(x) = b_0 + b_1 x + \sum_{k=1}^K a_k (x - \tau_k)_+, \quad (7.7)$$

where $b_0, b_1 \in \mathbb{R}$, the weights a_k are nonzero, the knots locations $\tau_k \in \mathbb{R}$ are distinct, and $K \leq M - 2$.

Remark 7.1. *There is a slight difference in the problem formulations (g-BLASSO) and (3.31). Indeed, (g-BLASSO) is a special case of (3.31) when the data-fidelity cost function $E : \mathbb{R}^M \times \mathbb{R}^M \rightarrow \mathbb{R}$ in (3.31) is separable and is applied component-wise, i.e., $E(\mathbf{z}, \mathbf{y}) = \sum_{m=1}^M \tilde{E}(z_m, y_m)$ where $\tilde{E} : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$. Then, if $\tilde{E}(\cdot, y)$ is strictly convex, coercive, and differentiable for any $y \in \mathbb{R}$, E satisfies the conditions of Theorem 3.3. In this Chapter, we restrict to the separable case for the sake of clarity, and the notation E (instead of \tilde{E}) refers to the component-wise loss. However, all our results still apply under the more general conditions of Theorem 3.3. Finally, as discussed in Example 3.1, the (g-BPC) is a special case of (3.31) when $E : \mathbb{R}^M \times \mathbb{R}^M \rightarrow \mathbb{R}^+ \cup \{+\infty\}$ is an indicator function.*

Remark 7.2. *The applications of Theorem 3.3 to Problems (g-BPC) and (g-BLASSO) require that the point evaluation $f \mapsto f(x_0)$ be weak* continuous on $\mathcal{M}_{\mathbb{D}^2}(\mathbb{R})$ for any $x_0 \in \mathbb{R}$, which is proved in Proposition 2.3.*

7.2.3 Dual Certificates

This section presents the main tools from duality theory for the study of the (g-BPC) problem, which are at the core of our contributions. Our strategy consists in studying a particular class of continuous functions, called *dual certificates*, which can be used individually to certify that an element $f \in \mathcal{M}_{\mathbb{D}^2}(\mathbb{R})$ is a solution of the optimization Problem (g-BPC). More interestingly, from the properties of a given dual certificate, it is possible to precisely describe the whole structure of the set of solutions (see Theorem 7.1) and, in particular, to determine whether the sparse solution given by Theorem 3.3 is the unique solution of the problem (see Proposition 7.6).

Before giving the main results of this section (Propositions 7.1 and 7.2), we first define dual pre-certificates.

Definition 7.1 (Dual Pre-Certificate). *We say that a function $\eta \in \mathcal{C}_0(\mathbb{R})$ is a dual pre-certificate (for the Problem (g-BPC)) if its norm satisfies $\|\eta\|_\infty \leq 1$ and if η is of the form*

$$\eta = \sum_{m=1}^M c_m (x_m - \cdot)_+ \tag{7.8}$$

for some vector $\mathbf{c} = (c_1, \dots, c_M) \in \mathbb{R}^M$ such that $\langle \mathbf{c}, \mathbf{1} \rangle = \langle \mathbf{c}, \mathbf{x} \rangle = 0$ (with $\mathbf{1} \triangleq (1, \dots, 1) \in \mathbb{R}^M$).

A dual pre-certificate is therefore a piecewise-linear spline. The conditions $\langle \mathbf{c}, \mathbf{1} \rangle = \langle \mathbf{c}, \mathbf{x} \rangle = 0$ ensure that η is compactly supported, and is thus an element of $\mathcal{C}_0(\mathbb{R})$ (indeed, we have $\eta(x) = -\langle \mathbf{c}, \mathbf{1} \rangle x + \langle \mathbf{c}, \mathbf{x} \rangle = 0$ for any $x \leq x_1$). We shall present an explicit construction of such a pre-certificate in Proposition 7.4 with the piecewise-linear spline η_{cano} . A dual certificate is a pre-certificate that satisfies an additional condition that ensures that the vector $\mathbf{c} \in \mathbb{R}^M$ in Definition 7.1 is a solution of the dual problem of (g-BPC) (see Proposition 7.1).

From (7.4), we know we can parametrize any $f \in \mathcal{M}_{D^2}(\mathbb{R})$ with a unique element $(w, \beta_0, \beta_1) \in \mathcal{M}(\mathbb{R}) \times \mathbb{R}^2$ via the relation

$$\forall x \in \mathbb{R}, \quad f(x) = D_0^{-2}\{w\}(x) + \beta_0 + \beta_1 x. \tag{7.9}$$

Dual certificates determine the localization of the support of w when f is a solution of (g-BPC). To formulate this property mathematically, we require the following definition which introduces the signed support of a measure (see [65, Section 1.4]) and signed saturation set of a pre-certificate (see [65, Definition 3]).

Definition 7.2 (Signed Support and Signed Saturation Set). *Let $w \in \mathcal{M}(\mathbb{R})$ and $\eta \in \mathcal{C}_0(\mathbb{R})$ be a dual pre-certificate in the sense of Definition 7.1. The signed support of w is defined as*

$$\text{supp}_\pm(w) \triangleq \text{supp}(w_+) \times \{1\} \cup \text{supp}(w_-) \times \{-1\}, \tag{7.10}$$

where w_+ and w_- are positive measures coming from the Jordan decomposition of $w = w_+ - w_-$. Moreover from the positive and negative saturation sets of η , defined as

$$\text{sat}_+(\eta) \triangleq \{x \in \mathbb{R} : \eta(x) = 1\} \quad \text{and} \quad \text{sat}_-(\eta) \triangleq \{x \in \mathbb{R} : \eta(x) = -1\}, \tag{7.11}$$

respectively, we define the signed saturation set of η by

$$\text{sat}_\pm(\eta) \triangleq \text{sat}_+(\eta) \times \{1\} \cup \text{sat}_-(\eta) \times \{-1\}. \tag{7.12}$$

Note that the sets $\text{supp}_\pm(w)$, $\text{sat}_+(\eta)$, $\text{sat}_-(\eta)$, $\text{sat}_\pm(\eta)$ are all closed. A dual pre-certificate η is a piecewise-linear spline in $\mathcal{C}_0(\mathbb{R})$ with norm $\|\eta\|_{L_\infty} \leq 1$. Hence, its signed saturation set is necessarily a union of closed intervals (that can be singletons).

We can now state the first main result of this section. It characterizes the solutions of (g-BPC) via the signed support of their innovation using the signed saturation set of some dual pre-certificate.

Proposition 7.1 (Dual Certificate). *Let $\mathbf{x} \in \mathbb{R}^M$ be the ordered sampling locations, and $\mathbf{y}_0 \in \mathbb{R}^M$. An element $f_{\text{opt}} \in \mathcal{M}_{\text{D}^2}(\mathbb{R})$ is a solution of (g-BPC) if and only if f_{opt} satisfies the interpolation conditions $f_{\text{opt}}(x_m) = y_{0,m}$ for all $m \in \{1, \dots, M\}$ and one can find a dual pre-certificate η (Definition 7.1) such that*

$$\|w\|_{\mathcal{M}} = \langle w, \eta \rangle, \quad (7.13)$$

where $w \triangleq \text{D}^2\{f_{\text{opt}}\}$ is the innovation of f_{opt} . Moreover, the condition (7.13) is equivalent to the inclusion

$$\text{supp}_\pm(w) \subset \text{sat}_\pm(\eta). \quad (7.14)$$

The dual pre-certificate η is then called a dual certificate (for Problem (g-BPC)).

The proof of Proposition 7.1 is given in Appendix 7.A.

Remark 7.3. *When $f_{\text{opt}} \in \mathcal{M}_{\text{D}^2}(\mathbb{R})$ is a piecewise-linear spline of the form $f_{\text{opt}}(x) = \sum_{k=1}^K a_k(x - \tau_k)_+ + b_0 + b_1x$ for all $x \in \mathbb{R}$, the condition (7.14) is equivalent to the following interpolation requirements on the dual pre-certificate η*

$$\forall k \in \{1, \dots, K\}, \quad \eta(\tau_k) = \text{sign}(a_k). \quad (7.15)$$

From Proposition 7.1, a dual certificate η is thus a dual pre-certificate that certifies that a given $f_{\text{opt}} \in \mathcal{M}_{\text{D}^2}(\mathbb{R})$ is a solution of (g-BPC), *i.e.*, f_{opt} satisfies $f_{\text{opt}}(x_m) = y_{0,m}$ for all $m \in \{1, \dots, M\}$ and $\text{supp}_\pm(\text{D}^2 f_{\text{opt}}) \subset \text{sat}_\pm(\eta)$ (or equivalently $\|\text{D}^2 f_{\text{opt}}\|_{\mathcal{M}} = \langle \text{D}^2 f_{\text{opt}}, \eta \rangle$). Once we know that some η is a dual certificate, it can be used to check whether *any* $f \in \mathcal{M}_{\text{D}^2}(\mathbb{R})$ is a solution of (g-BPC). In other words, contrary to what is seemingly implied in Proposition 7.1, there is no need to find a new dual pre-certificate for each candidate solution f . This is formulated in the following proposition.

Proposition 7.2. *Let $\mathbf{x} \in \mathbb{R}^M$ be the ordered sampling locations, $\mathbf{y}_0 \in \mathbb{R}^M$, and let $\eta \in \mathcal{C}_0(\mathbb{R})$ be a dual certificate as defined in Proposition 7.1 for the (g-BPC) problem. Then, an element $f_{\text{opt}} \in \mathcal{M}_{\text{D}^2}(\mathbb{R})$ is a solution of (g-BPC) if and only if f_{opt} satisfies the interpolation conditions $f_{\text{opt}}(x_m) = y_{0,m}$ for all $m \in \{1, \dots, M\}$ and*

$$\text{supp}_\pm(w) \subset \text{sat}_\pm(\eta), \quad (7.16)$$

or equivalently $\|w\|_{\mathcal{M}} = \langle w, \eta \rangle$, where $w \triangleq D^2 f_{\text{opt}}$ is the innovation of f_{opt} .

Proof. The proof of Proposition 7.2 is very similar to the proof of Proposition 7.1, and is derived from the optimality conditions given in Lemma 7.5.

Let η be a dual certificate in the sense of Proposition 7.2. By definition of η (it is in particular a dual pre-certificate in the sense of Definition 7.1) and by Lemma 7.7, there exists $\mathbf{c} \in \mathbb{R}^M$ such that $\eta = \nu_{\mathcal{M}}^*(\mathbf{c})$ and $\langle \mathbf{c}, \mathbf{1} \rangle = \langle \mathbf{c}, \mathbf{x} \rangle = 0$. Since η is a dual certificate, Proposition 7.1 implies that there exists a $\tilde{f} \in \mathcal{M}_{D^2}(\mathbb{R})$ satisfying the interpolation conditions and such that $\|D^2 \tilde{f}\|_{\mathcal{M}} = \langle D^2 \tilde{f}, \eta \rangle$. This implies that \mathbf{c} and $(\tilde{w}, (\tilde{\beta}_0, \tilde{\beta}_1)) \in \mathcal{M}(\mathbb{R}) \times \mathbb{R}^2$, where $\tilde{f} = D_0^{-2}\{\tilde{w}\} + \tilde{\beta}_0 + \tilde{\beta}_1(\cdot)$, satisfy (7.63) and (7.64) *i.e.*, in particular \mathbf{c} is a solution of the dual Problem (7.62) by Lemma 7.5. Using this fixed vector $\mathbf{c} \in \mathbb{R}^M$ and the decomposition of any $f \in \mathcal{M}_{D^2}(\mathbb{R})$ as $f = D_0^{-2}\{w\} + \beta_0 + \beta_1(\cdot)$ (see (7.4)), the equivalence in Lemma 7.5 directly yields that f_{opt} is a solution of (g-BPC) if and only if f_{opt} satisfies the interpolation conditions $f_{\text{opt}}(x_m) = y_{0,m}$ and $\|D^2 f_{\text{opt}}\|_{\mathcal{M}} = \langle D^2 f_{\text{opt}}, \eta \rangle$, which concludes the proof. \square

To end this section, we illustrate how the concept of dual certificates can be used to describe the solution set of (g-BPC). Suppose that we know that some η is a dual certificate (we prove in Proposition 7.5 that this is the case of the dual pre-certificate η_{cano} introduced in Proposition 7.4), then the condition $\text{supp}_{\pm}(w) \subset \text{sat}_{\pm}(\eta)$ of Proposition 7.2 enforces strong constraints on any candidate solution of (g-BPC). This is all the more true when $\text{sat}_{\pm}(\eta)$ is a discrete set, which we consider in the next definition and proposition.

Definition 7.3 (Nondegeneracy). *Let $\mathbf{x} \in \mathbb{R}^M$ be the ordered sampling locations, $\mathbf{y}_0 \in \mathbb{R}^M$ and let $\eta \in \mathcal{C}_0(\mathbb{R})$ be any dual certificate as defined in Proposition 7.1. We say that η is nondegenerate if its signed saturation set $\text{sat}_{\pm}(\eta)$ defined in Definition 7.2 is a discrete set. Otherwise, we say that it is degenerate.*

Proposition 7.3 (General Uniqueness Result for (g-BPC)). *Let $\mathbf{x} \in \mathbb{R}^M$ be the ordered sampling locations and $\mathbf{y}_0 \in \mathbb{R}^M$. If there exists a nondegenerate dual certificate in the sense of Definition 7.1, then the optimization Problem (g-BPC) has a unique solution, which is a piecewise-linear spline of the form (7.6) with $K \leq M - 2$ knots τ_k that form a subset of the sampling points $\{x_2, \dots, x_{M-1}\}$.*

Proof. Let $f_{\text{opt}} \in \mathcal{M}_{D^2}(\mathbb{R})$ be a solution of Problem (g-BPC) of the form (7.7) (which exists by Theorem 3.3). By (7.4), there exist $w \in \mathcal{M}(\mathbb{R})$ and $(\beta_0, \beta_1) \in \mathbb{R}^2$ such that $f_{\text{opt}} = D_0^{-2}\{w\} + \beta_0 + \beta_1(\cdot)$. By the assumption of Proposition 7.3, there exists a nondegenerate dual certificate η , so that by applying Proposition 7.2, we obtain $\text{supp}_{\pm}(w) \subset \text{sat}_{\pm}(\eta)$. Moreover, we have that $\text{sat}_{\pm}(\eta) \subset \{x_2, \dots, x_{M-1}\}$ due to the two following facts

- $\eta = \sum_{m=1}^M c_m (x_m - \cdot)_+$ (as a dual pre-certificate, see Lemma 7.7),
- $\text{sat}_{\pm}(\eta)$ is a discrete set (as η is nondegenerate).

This implies that η must be equal to ± 1 at the points $\{x_2, \dots, x_{M-1}\}$, which yields

$$w = \sum_{k=2}^{M-1} a_k \delta(\cdot - x_k), \quad (7.17)$$

where the $a_k \in \mathbb{R}$ are (possibly zero) weights. In particular, this implies that f_{opt} is a piecewise-linear spline with at most $(M - 2)$ knots that are a subset of $\{x_2, \dots, x_{M-1}\}$. It remains to prove that the coefficients $a_2, \dots, a_{M-1}, \beta_0, \beta_1$ are uniquely determined to conclude that f_{opt} is the unique solution of (g-BPC).

Since f_{opt} is a solution of (g-BPC), we have that $\nu(f_{\text{opt}}) = \mathbf{y}_0$. This implies that

$$\sum_{k=2}^{M-1} a_k \mathbf{g}_k + \beta_0 \mathbf{1} + \beta_1 \mathbf{x} = \mathbf{y}_0 \quad \text{with} \quad \mathbf{g}_k \triangleq \nu_{\mathcal{M}}(\delta(\cdot - x_k)) = (g_0(x_m, x_k))_{1 \leq m \leq M} \in \mathbb{R}^M. \quad (7.18)$$

We now prove that this equation uniquely determines the coefficients $a_2, \dots, a_{M-1}, \beta_0, \beta_1$ by showing that the family $(\mathbf{1}, \mathbf{x}, \mathbf{g}_2, \dots, \mathbf{g}_{M-1})$ is a basis of \mathbb{R}^M . Indeed, by definition of g (see (7.3)), we have that

$$\mathbf{g}_k = ((x_m - x_k)_+)_{1 \leq m \leq M} - (-x_k)_+ \mathbf{1} + \left((-x_k)_+ - (1 - x_k)_+ \right) \mathbf{x} \quad (7.19)$$

for all $k \in \{2, \dots, M-1\}$. Hence, by writing the matrix of the family $(\mathbf{1}, \mathbf{x}, \mathbf{g}_2, \dots, \mathbf{g}_{M-1})$ in the canonical basis of \mathbb{R}^M , subtracting thanks to (7.19) appropriate linear combinations of the first two columns (given by the vectors $\mathbf{1}$ and \mathbf{x}) to all of the other columns and finally subtracting x_1 times the first column to the second one, we end up with the following matrix

$$\begin{pmatrix} 1 & 0 & 0 & 0 & \dots & 0 \\ 1 & (x_2 - x_1) & 0 & 0 & \dots & 0 \\ 1 & (x_3 - x_1) & (x_3 - x_2) & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & (x_M - x_1) & (x_M - x_2) & (x_M - x_3) & \dots & (x_M - x_{M-1}) \end{pmatrix}. \quad (7.20)$$

The latter is a lower triangular matrix with nonzero coefficients on the diagonal (as the sampling points x_m are pairwise distinct), and is thus invertible, which proves the desired result. \square

7.3 The Solutions of the (g-BPC)

In this section, we consider the optimization Problem (g-BPC) where the x_m for $m \in \{1, \dots, M\}$ are distinct and ordered sampling locations and $\mathbf{y}_0 \in \mathbb{R}^M$ is a fixed measurement vector. This setting is especially relevant when the measurements $y_{0,m}$ are exactly the values of the input signal at locations x_m (noiseless case). The solution set is

$$\mathcal{V}_0 \triangleq \underset{\substack{f \in \mathcal{M}_{D^2}(\mathbb{R}) \\ f(x_m) = y_{0,m}, m \in \{1, \dots, M\}}}{\arg \min} \|D^2 f\|_{\mathcal{M}}, \quad (\text{g-BPC})$$

and is known to admit at least one piecewise-linear solution of the form (7.7) due to Theorem 3.3.

7.3.1 Canonical Solution and Canonical Dual Certificate

Thereafter, we identify the complete set of solutions (g-BPC). This allows us to fully determine in which cases this optimization problem admits a unique solution. Our analysis is based on the construction of a pair $(f_{\text{cano}}, \eta_{\text{cano}}) \in \mathcal{M}_{D^2}(\mathbb{R}) \times \mathcal{C}_0(\mathbb{R})$ that satisfies Proposition 7.1, which we call the canonical solution and canonical dual certificate respectively. The former is simply the function that connects the points $\mathbf{P}_{0,m} = \begin{bmatrix} x_m & y_{0,m} \end{bmatrix}^T$.

Definition 7.4 (Canonical Interpolant). *Let $\mathbf{x} \in \mathbb{R}^M$ be the ordered sampling locations and $\mathbf{y}_0 \in \mathbb{R}^M$ with $M \geq 2$. We define f_{cano} as the unique piecewise-linear spline that interpolates the data points with the minimum number of knots, i.e., such that*

- $f_{\text{cano}}(x_m) = y_{0,m}$ for any $m \in \{1, \dots, M\}$, and
- f_{cano} has at most $M - 2$ knots which form a subset of $\{x_m : 2 \leq m \leq M - 1\}$.

We refer to f_{cano} as the canonical interpolant.

The existence and uniqueness of f_{cano} in Definition 7.4 simply follows from the number of degrees of freedom of a piecewise-linear spline whose knots are known. The canonical interpolant is of the form

$$f_{\text{cano}}(x) = a_1 x + a_M + \sum_{m=2}^{M-1} a_m (x - x_m)_+ \quad (7.21)$$

with $\mathbf{a} = (a_1, \dots, a_M) \in \mathbb{R}^M$. By definition, f_{cano} is linear on the interval (x_m, x_{m+1}) for $m \in \{2, \dots, M - 1\}$. The interpolatory conditions $f_{\text{cano}}(x_m) = y_m$ and $f_{\text{cano}}(x_{m+1}) = y_{m+1}$ then imply that its slope is $s_m = \frac{y_{0,m+1} - y_{0,m}}{x_{m+1} - x_m}$. Yet from (7.21) we get that $s_m = a_1 + \dots + a_m$. This implies that $a_1 = s_1$ and that $a_m = s_m - s_{m-1}$ for $m \in \{2, \dots, M - 1\}$.

Finally, the equation $f_{\text{cano}}(x_1) = y_{0,1}$ yields $a_M = y_{0,1} - a_1 x_1$. Consequently, the vector $\mathbf{a} \in \mathbb{R}^M$ in (7.21) is given by

$$\begin{cases} a_1 = \frac{y_{0,2} - y_{0,1}}{x_2 - x_1} \\ a_m = \frac{y_{0,m+1} - y_{0,m}}{x_{m+1} - x_m} - \frac{y_{0,m} - y_{0,m-1}}{x_m - x_{m-1}} \quad \forall m \in \{2, \dots, M-1\} \\ a_M = y_{0,1} - \frac{y_{0,2} - y_{0,1}}{x_2 - x_1} x_1 \end{cases} \quad (7.22)$$

In order to prove that f_{cano} is always a solution of (g-BPC), we construct a particular dual pre-certificate η_{cano} .

Proposition 7.4 (Canonical Pre-Certificate). *Let $\mathbf{x} \in \mathbb{R}^M$ be the ordered sampling locations, $\mathbf{y}_0 \in \mathbb{R}^M$, and $\mathbf{a} \in \mathbb{R}^M$ be the vector defined by (7.22). Then, there exists a unique piecewise-linear spline η_{cano} given by*

$$\eta_{\text{cano}} \triangleq \sum_{m=1}^M c_m (x_m - \cdot)_+ \quad \text{with } \mathbf{c} = (c_1, \dots, c_M) \in \mathbb{R}^M, \quad (7.23)$$

$$\text{s.t. } \langle \mathbf{c}, \mathbf{1} \rangle = \langle \mathbf{c}, \mathbf{x} \rangle = 0, \quad \eta_{\text{cano}}(x_m) = \text{sign}(a_m) \quad \forall m \in \{2, \dots, M-1\}. \quad (7.24)$$

with the convention $\text{sign}(0) = 0$. Moreover, since $\eta_{\text{cano}}(x) = 0$ for $x \leq x_1$ and $x \geq x_M$, we have $\eta_{\text{cano}} \in \mathcal{C}_0(\mathbb{R})$ and $\|\eta_{\text{cano}}\|_\infty = 1$. Hence, η_{cano} is a dual pre-certificate in the sense of Definition 7.1.

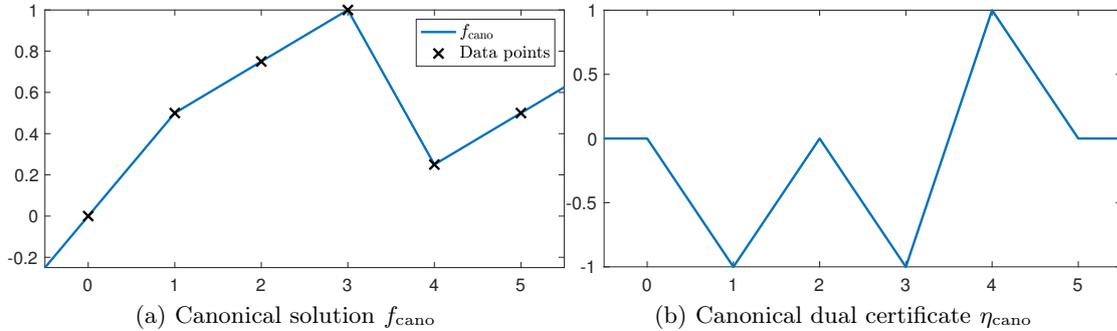


Figure 7.2 – Example of a canonical solution and canonical dual certificate for $M = 6$ with $x_m = m - 1$. We have $a_2 < 0$, $a_3 = 0$, $a_4 < 0$, and $a_5 > 0$, where the a_m are defined in (7.22).

Proof. The existence and uniqueness of such a spline follows the same argument as for f_{cano} , applied to the data points $(x_1 - 1, 0)$, $(x_1, 0)$, $(x_m, \text{sign}(a_m))$ for $m \in \{2, \dots, M-1\}$, $(x_M, 0)$ and $(x_M + 1, 0)$. Note that the points $(x_1 - 1, 0)$ and $(x_M + 1, 0)$ at the boundaries add two additional interpolation constraints to (7.24). Moreover, they imply that η_{cano} does not have a linear term and is thus of the form (7.23).

Next, we notice that for $x \leq x_1$, we have $\eta_{\text{cano}}(x) = -\langle \mathbf{c}, \mathbf{x} \rangle x + \langle \mathbf{c}, \mathbf{1} \rangle = 0$, due to $\langle \mathbf{c}, \mathbf{x} \rangle = \langle \mathbf{c}, \mathbf{1} \rangle = 0$. For $x \geq x_M$, $(x_m - x)_+ = 0$ for every $m \in \{1, \dots, M\}$, hence $\eta_{\text{cano}}(x) = 0$. Then, as a piecewise-linear spline with compact support, we must have $\eta_{\text{cano}} \in \mathcal{C}_0(\mathbb{R})$, and η_{cano} clearly reaches its maximum and minimum values at its knots. Hence, we have that $\|\eta_{\text{cano}}\|_\infty = \max_{m \in \{1, \dots, M\}} |\eta_{\text{cano}}(x_m)| = 1$.

□

We now prove that the pair $(f_{\text{cano}}, \eta_{\text{cano}}) \in \mathcal{M}_{\text{D}^2}(\mathbb{R}) \times \mathcal{C}_0(\mathbb{R})$ satisfies Proposition 7.1. Although the fact that f_{cano} is a solution to (g-BPC) is known [176, 177] and is significant in its own right, the key element of this result is the construction of the dual certificate η_{cano} . The latter will be essential to fully describe the solution set \mathcal{V}_0 .

Proposition 7.5. *Let $\mathbf{x} \in \mathbb{R}^M$ be the ordered sampling locations and $\mathbf{y}_0 \in \mathbb{R}^M$. The canonical interpolant f_{cano} defined in Definition 7.4 is a solution of (g-BPC) and η_{cano} , defined in Proposition 7.4, is a dual certificate in the sense of Proposition 7.1.*

Proof. By construction, the interpolation conditions $f_{\text{cano}}(x_m) = y_{0,m}$ for all $m \in \{1, \dots, M\}$ are satisfied. Moreover, by Proposition 7.4, η_{cano} is a dual pre-certificate. By Proposition 7.1, it remains to prove that

$$\text{supp}_\pm(\text{D}^2 f_{\text{cano}}) \subset \text{sat}_\pm(\eta_{\text{cano}}), \tag{7.25}$$

from which we deduce both that f_{cano} is a solution of (g-BPC) and that η_{cano} is a dual certificate. Since by construction, we have $\eta_{\text{cano}}(x_m) = \text{sign}(a_m)$ for all $m \in \{1, \dots, M\}$ and $\text{D}^2 f_{\text{cano}} = \sum_{m=2}^{M-1} a_m \delta(\cdot - x_m)$, this proves (7.25). □

Due to Proposition 7.5, we call f_{cano} the *canonical solution* and η_{cano} the *canonical dual certificate* of the optimization Problem (g-BPC). We show an example of such functions for given data points $(x_m, y_{0,m})_{m \in \{1, \dots, 6\}}$ in Figure 7.2. Notice that the points $P_{0,2}$, $P_{0,3}$, and $P_{0,4}$ are aligned, which implies that $a_3 = 0$ (defined in (7.22)).

7.3.2 Characterization of the Solution Set

Although identifying a solution f_{cano} to (g-BPC) is an important first step, this solution is not unique in general. We characterize the case of uniqueness in Proposition 7.6, and then provide a complete description of the solution set when the solution is not unique in Theorem 3.3. We shall see that the canonical dual certificate η_{cano} plays an essential role regarding these issues.

Proposition 7.6 (Uniqueness Result for (g-BPC)). *Let $\mathbf{x} \in \mathbb{R}^M$ be the ordered sampling locations and $\mathbf{y}_0 \in \mathbb{R}^M$. Then, the following conditions are equivalent.*

1. (g-BPC) has a unique solution.
2. The canonical dual certificate η_{cano} (defined in Proposition 7.4) is nondegenerate (see Definition 7.3).
3. For all $m \in \{2, \dots, M-2\}$, $a_m a_{m+1} \leq 0$, where $\mathbf{a} \in \mathbb{R}^M$ is given by (7.22).

Proof. The equivalence 2. \Leftrightarrow 3. comes from the fact that η_{cano} is nondegenerate if and only if it never saturates at 1 or -1 between two consecutive knots. This is equivalent to Item 3 because for all $m \in \{2, \dots, M-1\}$, $\eta_{\text{cano}}(x_m) = \text{sign}(a_m)$.

The implication 2. \Rightarrow 1. is given by Proposition 7.3. We now prove the contraposition of the reverse implication 1. \Rightarrow 2. We thus assume that η_{cano} is degenerate, and wish to prove that (g-BPC) has multiple solutions. Using Item 3, there exists an index $m \in \{2, \dots, M-2\}$ such that $a_m a_{m+1} > 0$. We now invoke the following lemma (illustrated in Figure 7.3) that plays an important role throughout this chapter.

Lemma 7.1. *Let $\mathbf{x} \in \mathbb{R}^M$ be the ordered sampling locations, and $\mathbf{y}_0 \in \mathbb{R}^M$ with $M \geq 4$. Let $m \in \{2, \dots, M-2\}$ be an index such that $a_m a_{m+1} > 0$, where $\mathbf{a} \in \mathbb{R}^M$ is defined as in (7.22). Then, the lines $(P_{0,m-1}, P_{0,m})$ and $(P_{0,m+1}, P_{0,m+2})$ are intersecting at a point $\tilde{\mathbf{P}} = \begin{bmatrix} \tilde{\tau} & \tilde{y} \end{bmatrix}^T$ such that $x_m < \tilde{\tau} < x_{m+1}$. Moreover, the piecewise-linear spline f_{opt} defined by*

$$f_{\text{opt}}(x) \triangleq \begin{cases} \frac{y_{0,m} - y_{0,m-1}}{x_m - x_{m-1}}(x - x_{m-1}) + y_{0,m-1} & x_m < x \leq \tilde{\tau} \\ \frac{y_{0,m+2} - y_{0,m+1}}{x_{m+2} - x_{m+1}}(x - x_{m+1}) + y_{0,m+1} & \tilde{\tau} < x < x_{m+1} \\ f_{\text{cano}}(x) & x \notin (x_m, x_{m+1}) \end{cases}, \quad (7.26)$$

which has no knot at x_m or x_{m+1} , is a solution of (g-BPC).

Proof. Let $I_0 = \{2, \dots, M-1\} \setminus \{m, m+1\}$. We then define

$$f_{\text{opt}}(x) \triangleq a_1 x + a_M + \sum_{m' \in I_0} a_{m'}(x - x_{m'})_+ + \tilde{a}(x - \tilde{\tau})_+, \quad (7.27)$$

where $\tilde{a} \triangleq a_m + a_{m+1}$ and $\tilde{\tau} \triangleq \frac{a_m x_m + a_{m+1} x_{m+1}}{\tilde{a}}$. By definition, $\tilde{\tau}$ is a barycenter of x_m and x_{m+1} with weights $\frac{a_m}{\tilde{a}}$ and $\frac{a_{m+1}}{\tilde{a}}$. Yet a_m and a_{m+1} have the same (nonzero) signs, which implies that these weights are in the interval $(0, 1)$ and thus that $\tilde{\tau} \in (x_m, x_{m+1})$. Yet f_{opt} has no knot at x_m and x_{m+1} , so it must follow the line $(P_{0,m-1}, P_{0,m})$ in the interval $[x_m, \tilde{\tau}]$, and the line $(P_{0,m+1}, P_{0,m+2})$ in the interval $[\tilde{\tau}, x_{m+1}]$, which conforms with the first two first lines in (7.26). Due to the continuity of f_{opt} , these lines are therefore intersecting at the point $\tilde{\mathbf{P}} = \begin{bmatrix} \tilde{\tau} & \tilde{y} \end{bmatrix}^T = \begin{bmatrix} \tilde{\tau} & f_{\text{opt}}(\tilde{\tau}) \end{bmatrix}^T$.

Next, for $x \leq x_m$, we have $a_m(x - x_m)_+ + a_{m+1}(x - x_{m+1})_+ = \tilde{a}(x - \tilde{\tau})_+ = 0$. Similarly, for $x \geq x_{m+1}$, we have $a_m(x - x_m)_+ + a_{m+1}(x - x_{m+1})_+ = \tilde{a}(x - \tilde{\tau})_+ = \tilde{a}(x - \tilde{\tau})$ since

$x \geq \tilde{\tau}$. Therefore, for any $x \notin (x_m, x_{m+1})$, we have $f_{\text{cano}}(x) = f_{\text{opt}}(x)$, which conforms with the third line in (7.26). This also implies that $f_{\text{opt}}(x_m) = f_{\text{cano}}(x_m) = y_{0,m}$ for all $m \in \{1, \dots, M\}$. Moreover, we have $\|D^2 f_{\text{cano}}\|_{\mathcal{M}} = \sum_{m=2}^{M-1} a_m = \sum_{m \in I_0} a_m + \tilde{a} = \|D^2 f_{\text{opt}}\|_{\mathcal{M}}$. Therefore, f_{opt} has the same measurements and regularization cost as f_{cano} , which implies that it is also a solution of (g-BPC). \square

Since f_{opt} defined in Lemma 7.1 is a solution to the (g-BPC) such that $f_{\text{opt}} \neq f_{\text{cano}}$, the (g-BPC) has multiple solutions, which concludes the proof. \square

To the best of our knowledge, Proposition 7.6 is a new result. A similar uniqueness result is presented in [178, Theorem 4.2], but with more restrictive conditions than Item 3. It follows from Proposition 7.6 that when $M = 3$, the solution of the (g-BPC) is always unique because the certificate is always nondegenerate, and is given by f_{cano} . We go much further in Theorem 7.1 by providing the full characterization of the solution set when $M \geq 4$.

Theorem 7.1 (Characterization of the Solution Set of the (g-BPC)). *Let $\mathbf{x} \in \mathbb{R}^M$ be the ordered sampling locations and $\mathbf{y}_0 \in \mathbb{R}^M$ with $M \geq 4$, and let f_{cano} and η_{cano} be the functions defined in Definition 7.4 and Proposition 7.4 respectively. A function $f_{\text{opt}} \in \mathcal{M}_{D^2}(\mathbb{R})$ is a solution of the (g-BPC) if and only if $f_{\text{opt}}(x_m) = y_{0,m}$ for $m \in \{1, \dots, M\}$, and the following conditions are satisfied for $m \in \{2, \dots, M-2\}$*

1. $f_{\text{opt}} = f_{\text{cano}}$ in $[x_m, x_{m+1}]$ if $|\eta_{\text{cano}}| < 1$ in (x_m, x_{m+1}) ;
2. f_{opt} is convex in $[x_{m-1}, x_{m+2}]$ if $\eta_{\text{cano}} = 1$ in $[x_m, x_{m+1}]$;
3. f_{opt} is concave in $[x_{m-1}, x_{m+2}]$ if $\eta_{\text{cano}} = -1$ in $[x_m, x_{m+1}]$;
4. $f_{\text{opt}} = f_{\text{cano}}$ in $(-\infty, x_2)$ and $(x_{M-1}, +\infty)$.

Proof. Let f_{opt} be a solution of the (g-BPC). According to Proposition 7.5, η_{cano} is a dual certificate. According to Proposition 7.2, we therefore have that $\text{supp}_{\pm}(D^2 f_{\text{opt}}) \subset \text{sat}_{\pm}(\eta_{\text{cano}})$, meaning that $D^2 f_{\text{opt}} = 0$ on the complement $\text{sat}_{\pm}(\eta_{\text{cano}})^c$ of $\text{sat}_{\pm}(\eta_{\text{cano}})$. In particular, we have that $(-\infty, x_2] \subset \text{sat}_{\pm}(\eta_{\text{cano}})^c$, hence f_{opt} is linear on this interval. The interpolation constraints $f_{\text{opt}}(x_1) = f_{\text{cano}}(x_1)$ and $f_{\text{opt}}(x_2) = f_{\text{cano}}(x_2)$ then imply that $f_{\text{opt}} = f_{\text{cano}}$ on $(-\infty, x_2]$. The same argument holds for the interval $[x_{M-1}, +\infty)$ and any interval (x_m, x_{m+1}) on which η_{cano} does not saturate.

Assume now that $[x_m, x_{m+1}] \subset \text{sat}_{+}(\eta_{\text{cano}})$; that is, $\eta_{\text{cano}} = 1$ on $[x_m, x_{m+1}]$. We use the Jordan decomposition of $D^2 f_{\text{opt}} = w = w_+ - w_-$ where w_+ and w_- are positive measures. By (7.14), we know that $w_- = 0$ on $[x_m, x_{m+1}]$ because its support is included in $\text{sat}_{-}(\eta_{\text{cano}})$. Hence, on this interval, $D^2 f_{\text{opt}} = w = w_+$ is a positive measure, implying that Df_{opt} is increasing and therefore that f_{opt} is convex over $[x_m, x_{m+1}]$. Now, if

$(x_{m-1}, x_m) \subset \text{sat}_+(\eta)^c \cap \text{sat}_-(\eta)^c$ then, as above, $D^2 f_{\text{opt}}|_{(x_{m-1}, x_m)} = 0$. Otherwise, by continuity of η_{cano} , we have $(x_{m-1}, x_m) \subset \text{sat}_+(\eta_{\text{cano}})$ hence $D^2 f_{\text{opt}}|_{(x_{m-1}, x_m)} \geq 0$. As a result f_{opt} is convex over $(x_{m-1}, x_{m+1}]$. The same argument proves that f_{opt} is convex over $[x_m, x_{m+2})$, and therefore on the whole interval (x_{m-1}, x_{m+2}) .

Conversely, suppose that f_{opt} satisfies all the conditions of Theorem 7.1. We now prove that it is a solution of the (g-BPC). By Proposition 7.2, it only remains to prove that f_{opt} satisfies $\text{supp}_\pm(D^2 f_{\text{opt}}) \subset \text{sat}_\pm(\eta_{\text{cano}})$ since, by construction, we have $f_{\text{opt}}(x_m) = y_{0,m}$. By definition of η_{cano} , we have $D^2 f_{\text{opt}} = 0$ on $\text{sat}_+(\eta_{\text{cano}})^c \cap \text{sat}_-(\eta_{\text{cano}})^c$ (because $D^2 f_{\text{opt}}$ is equal to f_{cano} which is linear on that set). Moreover, we have $D^2 f_{\text{opt}} \geq 0$ on $\text{sat}_+(\eta_{\text{cano}})$ (because by assumption, f_{opt} is convex over intervals where $\eta_{\text{cano}} = 1$) and $D^2 f_{\text{opt}} \leq 0$ on $\text{sat}_-(\eta_{\text{cano}})$ (because f_{opt} is concave on intervals where $\eta_{\text{cano}} = -1$). This implies that $\text{supp } w_+ \subset \text{sat}_+(\eta_{\text{cano}})$ and $\text{supp } w_- \subset \text{sat}_+(\eta_{\text{cano}})$ where $D^2 f_{\text{opt}} = w_+ - w_-$ is again the Jordan decomposition of $D^2 f_{\text{opt}}$. Finally, as expected, we have that

$$\begin{aligned} \text{supp}_\pm(D^2 f_{\text{opt}}) &= \text{supp } w_+ \times \{1\} \cup \text{supp } w_- \times \{-1\} \\ &\subset \text{sat}_+(\eta_{\text{cano}}) \times \{1\} \cup \text{sat}_-(\eta_{\text{cano}}) \times \{-1\} \\ &= \text{sat}_\pm(\eta_{\text{cano}}), \end{aligned} \tag{7.28}$$

hence f_{opt} is a solution of the (g-BPC). □

To illustrate Theorem 7.1, a simple example with $M = 4$ data points for which the solution is not unique is given in Figure 7.1. Indeed, the canonical dual certificate saturates at -1 in the interval $[1, 2]$. Therefore, by Theorem 7.1, any function that coincides with f_{cano} in $\mathbb{R} \setminus [1, 2]$ and that is concave in the interval $[0, 3]$ is a solution. This includes the sparsest solution (with a single knot), as well as nonsparse solutions, *e.g.*, with a quadratic regime in $[1, 2]$ as in Figure 7.3.

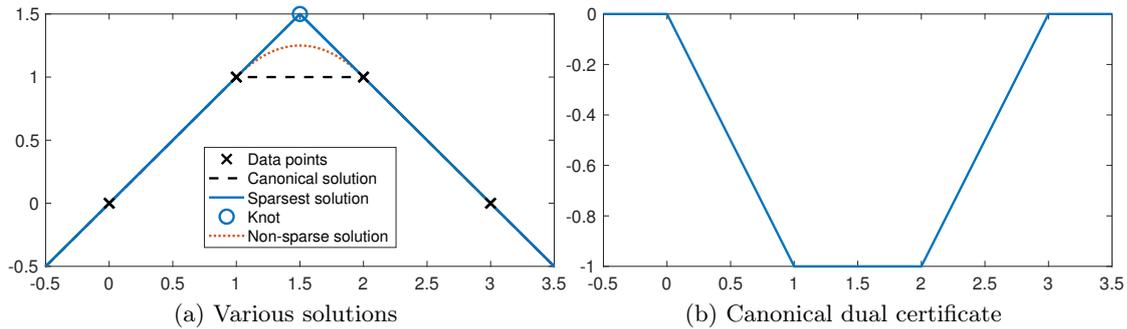


Figure 7.3 – Example with $M = 4$ of a nonunique solution (η_{cano} saturates at -1). An example of a nonsparse solution with a quadratic regime in $[1, 2]$ is given.

Even when the (g-BPC) has infinitely many solutions, we are able to delimit the geometric domain that contains the graphs of all solutions by exploiting the local convex/concavity.

We recall that $P_{0,m} = [x_m \ y_{0,m}]^T$ for $m \in \{1, \dots, M\}$, and that for $A, B \in \mathbb{R}^2$, we denote by (A, B) the line joining A and B . Then, for $M \geq 4$, we consider the set of indices

$$\mathcal{X} \triangleq \mathcal{X}(\mathbf{x}, \mathbf{y}_0) \triangleq \left\{ m \in \{2, \dots, M-2\}; a_m a_{m+1} > 0 \right\}, \quad (7.29)$$

where we recall that $a_m = \frac{y_{0,m+1} - y_{0,m}}{x_{m+1} - x_m} - \frac{y_{0,m} - y_{0,m-1}}{x_m - x_{m-1}}$ (see (7.22)). The slope condition $a_m a_{m+1} > 0$ in (7.29) is equivalent to the fact that the lines $(P_{0,m-1}, P_{0,m})$ and $(P_{0,m+1}, P_{0,m+2})$ are not parallel (otherwise we would have that $a_m = -a_{m+1}$, hence $a_m a_{m+1} \leq 0$) and that their intersection point, that we denote by $\tilde{P}_m = [\tilde{\tau}_m \ \tilde{y}_m]^T$, is such that $x_m \leq \tilde{\tau}_m \leq x_{m+1}$ according to Lemma 7.1. We can thus introduce the triangles Δ_m , whose vertices are the points $P_{0,m}$, \tilde{P}_m , and $P_{0,m+1}$. Theorem 7.2 makes the link between the graph of any solution $f_{\text{opt}} \in \mathcal{M}_{\mathbb{D}^2}(\mathbb{R})$ of the (g-BPC), the graph of f_{cano} and the triangles Δ_m .

Theorem 7.2 (Geometric Domain of the Graph of Solutions of the (g-BPC)). *Let $\mathbf{x} \in \mathbb{R}^M$ be the ordered sampling locations and $\mathbf{y}_0 \in \mathbb{R}^M$ with $M \geq 4$. Then, we have*

$$\cup_{f_{\text{opt}} \in \mathcal{V}_0} \mathcal{G}(f_{\text{opt}}) = \mathcal{G}(f_{\text{cano}}) \cup (\cup_{m \in \mathcal{X}} \Delta_m), \quad (7.30)$$

where f_{cano} is defined in Definition 7.4, \mathcal{X} is defined in (7.29), and the Δ_m triangles are defined in the above paragraph.

Proof. Let $f_{\text{opt}} \in \mathcal{V}_0$. We fix $m \in \{2, M-2\}$. First, as we have seen in the proof of Theorem 7.1, if $a_m a_{m+1} \leq 0$, then $f_{\text{opt}} = f_{\text{cano}}$ on $[x_m, x_{m+1}]$, and the graph of f_{opt} in this interval is equal to the one of f_{cano} . Assume now that $a_m a_{m+1} > 0$. We now show that $\{(x, f_{\text{opt}}(x)) : x \in [x_m, x_{m+1}]\} \subset \Delta_m$. The slope condition $a_m a_{m+1} > 0$ implies that η_{cano} is degenerate and that $\eta_{\text{cano}} = \pm 1$ is constant over $[x_m, x_{m+1}]$. Assume for instance that the value is 1, in which case f_{opt} is convex over $[x_{m-1}, x_{m+2}]$ according to Theorem 7.1.

We shall use the following well-known fact on convex functions. Fix $a < b < c$ and assume that f is convex over $[a, c]$. Then, f is below its arc between a and b on (a, b) , that is, $f(x) \leq \frac{f(b) - f(a)}{b - a}(x - a) + f(a)$ for any $x \in (a, b)$. Moreover, f is above the same arc over (b, c) , that is, $f(x) \geq \frac{f(b) - f(a)}{b - a}(x - a) + f(a)$ for any $x \in (b, c)$.

Let $x^* \in [x_m, x_{m+1}]$. By convexity, f_{opt} is below its arc between x_m and x_{m+1} . Hence, we have that

$$f_{\text{opt}}(x^*) \leq \frac{y_{0,m+1} - y_{0,m}}{x_{m+1} - x_m}(x^* - x_m) + y_{0,m}. \quad (7.31)$$

Moreover, the convexity over $[x_{m-1}, x_{m+1}]$ implies that $f_{\text{opt}}(x^*)$ is above the arc of f_{opt}

between x_{m-1} and x_m . This implies that

$$f_{\text{opt}}(x^*) \geq \frac{y_{0,m} - y_{0,m-1}}{x_m - x_{m-1}}(x^* - x_{m-1}) + y_{0,m-1}. \quad (7.32)$$

A similar argument over $[x^*, x_{m+2}]$ implies that

$$f_{\text{opt}}(x^*) \geq \frac{y_{0,m+2} - y_{0,m+1}}{x_{m+2} - x_{m+1}}(x^* - x_{m+1}) + y_{0,m+1}. \quad (7.33)$$

The conditions (7.31), (7.32), and (7.33) are precisely equivalent to $(x^*, f_{\text{opt}}(x^*)) \in \Delta_m$, since the three linear equations delineate this domain in this case. The same proof applies when $\eta_{\text{cano}} = -1$ over $[x_m, x_{m+1}]$ by using concavity instead of convexity. This proves that $\mathcal{G}(f_{\text{opt}}) \subset \mathcal{G}(f_{\text{cano}}) \cup (\cup_{m \in \mathcal{X}} \Delta_m)$ for every $f_{\text{opt}} \in \mathcal{V}_0$, and hence the direct inclusion in (7.30).

For the reverse inclusion, we already know that $f_{\text{cano}} \in \mathcal{V}_0$, therefore it suffices to show that, for any $m \in \mathcal{X}$ and any $(x^*, y^*) \in \Delta_m$, there exists a solution $f_{\text{opt}} \in \mathcal{V}_0$ such that $f_{\text{opt}}(x^*) = y^*$. As before, since $m \in \mathcal{X}$, we know that $\eta_{\text{cano}} = \pm 1$ on $[x_m, x_{m+1}]$ and we can assume without loss of generality that the value is 1. Then, any solution is convex and satisfies the relations (7.31), (7.32), and (7.33). By convexity of \mathcal{V}_0 , it suffices to show the result for (x^*, y^*) in the boundary of Δ_m , which is delimited by the relations

$$\frac{y_{0,m+1} - y_{0,m}}{x_{m+1} - x_m}(x^* - x_m) + y_{0,m} = y^*, \text{ or} \quad (7.34)$$

$$\frac{y_{0,m} - y_{0,m-1}}{x_m - x_{m-1}}(x^* - x_{m-1}) + y_{0,m-1} = y^*, \text{ or} \quad (7.35)$$

$$\frac{y_{0,m+2} - y_{0,m+1}}{x_{m+2} - x_{m+1}}(x^* - x_{m+1}) + y_{0,m+1} = y^*. \quad (7.36)$$

The solution f_{cano} is such that $f_{\text{cano}}(x^*) = \frac{y_{0,m+1} - y_{0,m}}{x_{m+1} - x_m}(x^* - x_m) + y_{0,m} = y^*$, hence any (x^*, y^*) satisfying (7.34) is attained by a solution (the canonical one) in \mathcal{V}_0 . Assume that (x^*, y^*) satisfies (7.35) (the case of (7.36) follows the same argument). We construct f_{opt} as follows. First, $f_{\text{opt}}(x) = f_{\text{cano}}(x)$ for any $x \notin (x_m, x_{m+1})$. Then, we set

$$f_{\text{opt}}(x) = \frac{y_{0,m} - y_{0,m-1}}{x_m - x_{m-1}}(x - x_{m-1}) + y_{0,m-1} \quad (7.37)$$

for $x \in (x_m, x^*)$. In particular, $f(x^*) = y^*$, and f_{opt} is linear on $[x_m, x^*]$. Finally, we impose that f_{opt} is linear on $[x^*, x_{m+1}]$, which is equivalent to the relation

$$f_{\text{opt}}(x) = \frac{y_{0,m+1} - y^*}{x_{m+1} - x^*}(x - x^*) + y^* \quad (7.38)$$

for any $x \in [x^*, x_{m+1}]$. We then claim that $f_{\text{opt}} \in \mathcal{V}_0$, the argument being very similar to the one of Lemma 7.1. Indeed, to show this, it suffices to remark that f_{opt} , which

is piecewise-constant and coincides with f_{cano} outside of (x_m, x_{m+1}) , is convex over $[x_{m-1}, x_{m+2}]$ (this is guaranteed by the slope condition $a_m a_{m+1} > 0$ and the construction of f_{opt}). According to Theorem 7.1, this implies that $f_{\text{opt}} \in \mathcal{V}_0$, with $f_{\text{opt}}(x^*) = y^*$. This finally shows that $(x^*, y^*) \in \cup_{f_{\text{opt}} \in \mathcal{V}_0} \mathcal{G}(f_{\text{opt}})$, which proves (7.30). □

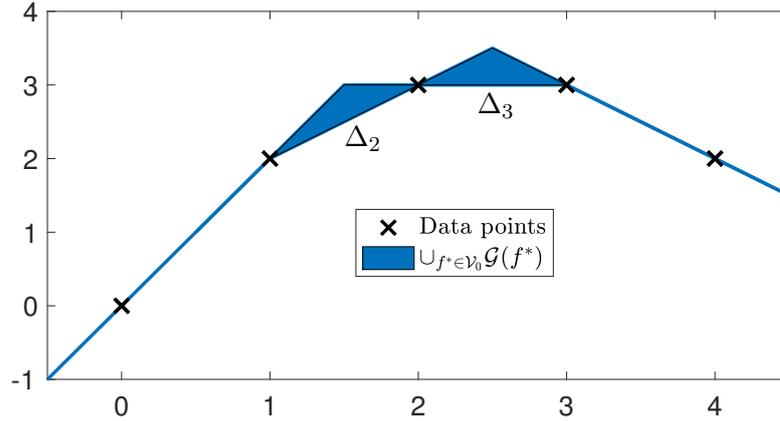


Figure 7.4 – Example with $M = 5$ of the geometric domain $\cup_{f_{\text{opt}} \in \mathcal{V}_0} \mathcal{G}(f_{\text{opt}})$ containing all the solutions to the (g-BPC). We have $\mathcal{X} = \{2, 3\}$ and thus two triangles Δ_m ; all solutions follow f_{cano} everywhere else.

The relation (7.30) reveals the smallest possible geometric domain containing all the graphs of the solutions of the (g-BPC). To obtain a solution of the (g-BPC), one just needs to follow the graph of f_{cano} outside the triangles Δ_m and take a convex or concave function inside them. An example of this domain is given in Figure 7.4 with $M = 5$ and $\#\mathcal{X} = 2$ triangles (this same example is treated further later in Figure 7.6). Next, Section 7.4 is dedicated to the study of the sparsest piecewise-linear solutions of the (g-BPC).

7.4 The Sparsest Solution(s) of the (g-BPC)

7.4.1 Characterization of the Sparsest Solution(s)

We have already identified the situations where the (g-BPC) admits a unique solution, in which case it is the canonical solution introduced in Definition 7.4. When the solution is not unique, Theorem 3.3 ensures that the extreme-point solutions are piecewise-linear functions with at most $K - 2$ knots, and Theorem 7.1 gives a complete description of the solution set. In this section, we go further by providing a complete answer to the following questions:

- what is the minimal number of knots of a solution of the (g-BPC)?

- what are the sparsest solutions, *i.e.*, the ones reaching this minimum number of knots?

These questions are addressed in Theorem 7.3. Let η_{cano} be defined as in Proposition 7.4 for fixed values of $\mathbf{x}, \mathbf{y}_0 \in \mathbb{R}^M$, and let

$$\begin{aligned} I_{\text{sat}} &\triangleq \left\{ m \in \{2, \dots, M-1\} : \eta_{\text{cano}}(x_m) = \pm 1 \text{ and } \eta_{\text{cano}}(x_m) \neq \eta_{\text{cano}}(x_{m-1}) \right\} \\ &= \{s_1, \dots, s_{N_s}\} \quad \text{with } s_1 < \dots < s_{N_s}. \end{aligned} \quad (7.39)$$

In other words, $N_s = \#I_{\text{sat}}$ corresponds to the number of times η_{cano} reaches ± 1 . Next, let $\alpha_n \in \mathbb{N}$ for $n \in \{1, \dots, N_s\}$ be the number of consecutive intervals starting from x_{s_n} in which η_{cano} saturates at ± 1 , *i.e.*,

$$\alpha_n \triangleq \min \left\{ k \in \mathbb{N} : \eta_{\text{cano}}(x_{s_n+k+1}) \neq \eta_{\text{cano}}(x_{s_n}) \right\}. \quad (7.40)$$

In what follows, $\lceil x \rceil$ is the smallest integer larger or equal to $x \in \mathbb{R}$.

Theorem 7.3 (Sparsest Solutions of the (g-BPC)). *Let $\mathbf{x} \in \mathbb{R}^M$ be the ordered sampling locations, $\mathbf{y}_0 \in \mathbb{R}^M$ with $M \geq 4$. Concerning the minimum sparsity of a solution of the (g-BPC), the following hold.*

1. *The lowest possible sparsity (i.e., number of knots) of a piecewise-linear solution of the (g-BPC) is*

$$K_{\min}(\mathbf{x}, \mathbf{y}_0) = \sum_{n=1}^{N_s} \left\lceil \frac{\alpha_n + 1}{2} \right\rceil, \quad (7.41)$$

where the α_n are defined in (7.40), and $N_s = \#I_{\text{sat}}$ where I_{sat} is defined in (7.39).

2. *There is a unique sparsest solution of the (g-BPC) if and only if none of the α_n are nonzero even numbers.*
3. *If one or more $\alpha_n > 0$ are even, then there are uncountably many sparsest solutions to the (g-BPC). The number of degrees of freedom $n_{\text{free}}(\mathbf{x}, \mathbf{y}_0)$ of the set of sparsest solutions is equal to the number of even α_n coefficients, that is,*

$$n_{\text{free}}(\mathbf{x}, \mathbf{y}_0) = \sum_{n=1}^{N_s} \mathbb{1}_{2\mathbb{N}_{\geq 1}}(\alpha_n). \quad (7.42)$$

More precisely, for each saturation region of η_{cano} , fixing a single knot within a certain admissible segment uniquely determines the other knots within the saturation region.

Proof. Using Theorem 7.1, for any $f_{\text{opt}} \in \mathcal{V}_0$, we have $f_{\text{opt}}(x) = f_{\text{cano}}(x)$ for any x such that $\eta_{\text{cano}}(x) \neq \pm 1$. We now focus on regions where $\eta_{\text{cano}}(x) = \pm 1$. For all $n \in \{1, \dots, N_s\}$, f_{cano} has $\alpha_n + 1$ knots in the interval $[x_{s_n}, x_{s_n + \alpha_n}]$. In order to construct one of the sparsest solutions, we must therefore replace these $\alpha_n + 1$ knots with as little knots as possible in each saturation region, since all solutions must coincide with f_{cano} outside these regions. In order to lighten the notations, in what follows, we focus on a single saturation region determined by a fixed $n \in \{1, \dots, N_s\}$ and we write $\alpha \triangleq \alpha_n$ and $s \triangleq s_n$.

Similarly to the proof of Proposition 7.6, a piecewise-linear spline f that coincides with f_{cano} outside the interval $[x_s, x_{s+\alpha}]$ must be of the form

$$f(x) = f_{\text{cano}}(x) - \sum_{n'=0}^{\alpha} a_{s+n'}(x - x_{s+n'})_+ + \sum_{p=1}^P \tilde{a}_p(x - \tilde{\tau}_p)_+, \quad (7.43)$$

where $\tilde{a}_p \in \mathbb{R}$, $\tilde{\tau}_p \in [x_s, x_{s+\alpha}]$ such that $\tilde{\tau}_1 < \dots < \tilde{\tau}_P$ and P is the number of knots of f in this interval. We then prove the following lemma.

Lemma 7.2. *If f in (7.43) satisfies the constraints $f(x_m) = y_{0,m}$ for all $m \in \{1, \dots, M\}$, then the number of knots P in $[x_s, x_{s+\alpha}]$ satisfies $P \geq \lceil \frac{\alpha+1}{2} \rceil$.*

Proof. Lemma 7.2 is trivially true for $\alpha = 0$, since we must have $f = f_{\text{cano}}$ and thus $P = 1$. Assume now that $\alpha > 0$. Firstly, we show that we must have $\tilde{\tau}_1 \in [x_s, x_{s+1})$. Assume by contradiction that $\tilde{\tau}_1 \geq x_{s+1}$: then, f has no knot in the interval (x_{s-1}, x_{s+1}) . Yet f must satisfy the interpolation constraints $f(x_m) = y_{0,m}$ for all $m \in \{1, \dots, M\}$, which implies that the points $P_{0,s-1}$, $P_{0,s}$, and $P_{0,s+1}$ are aligned. Therefore, f_{cano} has a weight $a_s = 0$ (defined in (7.22)) which implies that $\eta_{\text{cano}}(x_s) = 0$, which contradicts the assumption $\eta_{\text{cano}}(x_s) = \pm 1$. We can then prove in a similar fashion that $\tilde{\tau}_P \in (x_{s+\alpha-1}, x_{s+\alpha}]$ when $\alpha > 1$.

Next, we show that for $\alpha \geq 2$, we have

$$\forall n' \in \{1, \dots, \alpha - 1\}, \exists p \in \{1, \dots, P\} \text{ such that } \tilde{\tau}_p \in (x_{s+n'-1}, x_{s+n'+1}), \quad (7.44)$$

i.e., there must be a knot in all blocks of two consecutive saturation intervals. We assume by contradiction that this is not the case. Similarly to above, this implies that $P_{0,s+n'-1}$, $P_{0,s+n'}$, and $P_{0,s+n'+1}$ are aligned and thus that $\eta_{\text{cano}}(x_{s+n'}) = 0$, which yields a contradiction.

Lemma 7.2 immediately follows from the constraints $\tilde{\tau}_1 \in [x_s, x_{s+1})$ and $\tilde{\tau}_P \in [x_{s+\alpha-1}, x_{s+\alpha}]$ for $\alpha \leq 2$. For $\alpha > 2$, by the two aforementioned constraints, f must have at least two knots in the first and last saturation intervals $[x_s, x_{s+1})$ and $(x_{s+\alpha-1}, x_{s+\alpha}]$ respectively. Next, consider the interval $[x_{s+1}, x_{s+\alpha-1}]$, which consists of the central $\alpha - 2$ consecutive saturations. Using (7.44), this interval must contain at least $\lfloor \frac{\alpha-2}{2} \rfloor$ knots, which yields

the lower bound $P \geq 2 + \lfloor \frac{\alpha-2}{2} \rfloor = \lceil \frac{\alpha+1}{2} \rceil$ (the last equality can easily be verified for every $\alpha \in \mathbb{N}$). \square

The following lemma then states that the bound in Lemma 7.2 is tight.

Lemma 7.3. *The lower bound in Lemma 7.2 is always reached, i.e., there exists a piecewise-linear spline $f_{\text{opt}} \in \mathcal{V}_0$ of the form (7.43) with $P = \lceil \frac{\alpha+1}{2} \rceil$ knots in $[x_s, x_{s+\alpha}]$. If α is odd or $\alpha = 0$, then f_{opt} is unique. If $\alpha > 0$ is even, then there are uncountably many such functions f_{opt} .*

Proof. Lemma 7.3 is trivially true for $\alpha = 0$, i.e., when no saturation occurs. Indeed, the saturation interval is then reduced to the point $\{x_s\}$, and the only solution $f_{\text{opt}} \in \mathcal{V}_0$ of the form (7.43) is $f_{\text{opt}} = f_{\text{cano}}$ for which $P = 1$.

Assume now that $\alpha = 2k + 1$ is odd. The bound in Lemma 7.2 then reads $P \geq k + 1$. Similarly to the proof of Proposition 7.6, we construct a function f_{opt} of the form (7.43) with $P = k + 1$ and

$$\begin{cases} \tilde{a}_1 \triangleq a_s + a_{s+1} \text{ and } \tilde{\tau}_1 \triangleq \frac{a_s x_s + a_{s+1} x_{s+1}}{\tilde{a}_1}; \\ \tilde{a}_2 \triangleq a_{s+2} + a_{s+3} \text{ and } \tilde{\tau}_2 \triangleq \frac{a_{s+2} x_{s+2} + a_{s+3} x_{s+3}}{\tilde{a}_2}; \\ \vdots \\ \tilde{a}_{k+1} \triangleq a_{s+2k} + a_{s+2k+1} \text{ and } \tilde{\tau}_k \triangleq \frac{a_{s+2k} x_{s+2k} + a_{s+2k+1} x_{s+2k+1}}{\tilde{a}_{k+1}}. \end{cases} \quad (7.45)$$

Since the $a_s, \dots, a_{s+\alpha}$ all have the same (nonzero) sign, the $\tilde{\tau}_i$, $i = 1, \dots, k + 1$, are all barycenters with positive weights, which implies that $\tilde{\tau}_i \in (x_{s+2i}, x_{s+2i+1})$. Then, as in the proof of Proposition 7.6, replacing the knots at x_{s+2i} and x_{s+2i+1} in f_{cano} by a single knot at $\tilde{\tau}_i$ does not change the expression of f_{opt} outside the interval (x_{s+2i}, x_{s+2i+1}) , which implies that all the constraints $f_{\text{opt}}(x_m) = y_{0,m}$ for all $m \in \{1, \dots, M\}$ are satisfied.

Next, let $I_s = \{1, \dots, M\} \setminus \{s, \dots, s + \alpha\}$ be the set of indices outside our interval of interest. Since $a_s, \dots, a_{s+\alpha}$ and thus $\tilde{a}_1, \dots, \tilde{a}_{k+1}$ all have the same sign, we have $\|\mathbf{D}^2 f_{\text{opt}}\|_{\mathcal{M}} = \sum_{m \in I_s} |a_m| + |\sum_{i=1}^{k+1} \tilde{a}_i| = \sum_{m \in I_s} |a_m| + |\sum_{n=0}^{\alpha} a_{s+n}| = \|\mathbf{D}^2 f_{\text{cano}}\|_{\mathcal{M}}$, which together with the interpolation constraints implies that $f_{\text{opt}} \in \mathcal{V}_0$.

To show the uniqueness, consider once again a function f_{opt} of the form (7.43) with $P = k + 1$ and $\tilde{\tau}_1 < \dots < \tilde{\tau}_{k+1}$. We then invoke Lemma 7.2, which stipulates that there must be knots in the first and last saturation intervals as well as every two consecutive saturation intervals. The only way to achieve this is to have $\tilde{\tau}_i \in (x_{s+2i}, x_{s+2i+1})$, $i = 0, \dots, k$. The intervals (x_{s+2i-1}, x_{s+2i}) for all $i \in \{1, \dots, k\}$ thus have no knot, which implies that in these intervals, f_{opt} must follow the line $(P_{0,s+2i-1}, P_{0,s+2i})$. The knots are then necessarily the intersection of these lines, which yields the solution given in (7.45). The latter is therefore the unique function in \mathcal{V}_0 with $P = k + 1$ knots in the interval

$[x_s, x_{s+\alpha}]$. An example of such a sparsest solution is shown in Figure 7.5 with $M = 6$ and $\alpha = 3$ consecutive saturation intervals.

Assume now that $\alpha = 2k$ is even, with $k > 0$. The bound in Lemma 7.2 then reads $P \geq k + 1$. By Lemma 7.1, the intersection $\tilde{P} = \begin{bmatrix} \tilde{\tau} & \tilde{y} \end{bmatrix}^T$ between the lines $(P_{0,s-1}, P_{0,s})$ and $(P_{0,s+1}, P_{0,s+2})$ exists and satisfies $\tilde{\tau} \in (x_s, x_{s+1})$. Then, let $\tilde{P}_1 = \begin{bmatrix} \tilde{\tau}_1 & \tilde{y}_1 \end{bmatrix}^T$ be any point on the line segment $[P_{0,s}, \tilde{P}]$, *i.e.*, with $\tilde{\tau}_1 \in [x_s, \tilde{\tau}]$. Then, we define \tilde{P}_2 as the intersection between the lines $(\tilde{P}_1, P_{0,s+1})$ and $(P_{0,s+2}, P_{0,s+3})$. Similarly, if $\alpha \geq 4$, for every $i \in \{3, \dots, k+1\}$, we define $\tilde{P}_i = \begin{bmatrix} \tilde{\tau}_i & \tilde{y}_i \end{bmatrix}^T$ as the intersection between the lines $(P_{0,s+2i-4}, P_{0,s+2i-3})$ and $(P_{0,s+2i-2}, P_{0,s+2i-1})$. Due to a similar barycenter argument as in (7.45), these intersections are well defined and satisfy $\tilde{\tau}_i \in (x_{s+2i-3}, x_{s+2i-2})$. Let f_{opt} be the piecewise-linear spline that coincides with f_{cano} outside the interval $(x_s, x_{s+\alpha})$, and that connects the points $P_{0,s-1}$, \tilde{P}_1 , \dots , \tilde{P}_{k+1} , and $P_{0,s+\alpha}$ in that interval. By construction, f_{opt} satisfies the constraints $f_{\text{opt}}(x_m) = y_{0,m}$, $m \in \{1, \dots, M\}$. Moreover, once again in a similar manner to (7.45), we have that $\|f_{\text{opt}}\|_{\mathcal{M}} = \|f_{\text{cano}}\|_{\mathcal{M}}$, which implies that $f_{\text{opt}} \in \mathcal{V}_0$. Finally, f_{opt} is of the form (7.43) with the lowest possible sparsity $P = k + 1$ in the interval $[x_s, x_{s+\alpha}]$ (by Lemma 7.2). Yet there are uncountably many possible choices of \tilde{P}_1 (it can be any point on a non-singleton line segment). All of these choices lead to a different solution $f_{\text{opt}} \in \mathcal{V}_0$ that is uniquely defined, since the choice of \tilde{P}_1 specifies $\tilde{P}_2, \dots, \tilde{P}_{k+1}$. This proves that there are uncountably many solutions of the (g-BPC) with sparsity $k + 1$ in $[x_s, x_{s+\alpha}]$, and that there is a single degree of freedom for the choice of these $k + 1$ knots. An example of such a sparsest solution is shown in Figure 7.6 with $M = 5$ and $\alpha = 2$ consecutive saturation intervals. In our algorithm, we simply choose $\tilde{P}_1 = P_{0,s}$, which yields a function f_{opt} of the form (7.43) with

$$\begin{cases} \tilde{a}_1 \triangleq a_s \text{ and } \tilde{\tau}_1 \triangleq x_s; \\ \tilde{a}_2 \triangleq a_{s+1} + a_{s+2} \text{ and } \tilde{\tau}_2 \triangleq \frac{a_{s+1}x_{s+1} + a_{s+2}x_{s+2}}{\tilde{a}_2}; \\ \vdots \\ \tilde{a}_{k+1} \triangleq a_{s+2k-1} + a_{s+2k} \text{ and } \tilde{\tau}_k \triangleq \frac{a_{s+2k-1}x_{s+2k-1} + a_{s+2k}x_{s+2k}}{\tilde{a}_{k+1}}. \end{cases} \quad (7.46)$$

□

Theorem 7.3 then directly derives from Lemma 7.3 applied independently to each saturation interval $[x_{s_n}, x_{s_n+\alpha_n}]$ for $n \in \{1, \dots, N_s\}$. Note that Lemma 7.3 also applies when no saturation occurs, *i.e.*, $\alpha_n = 0$. A sparsest solution of the (g-BPC) thus coincides with a function of the form (7.43) constructed in Lemma 7.3 in each of these intervals, and with f_{cano} outside these intervals. Finally, since the behavior of a solution in each saturation interval does not affect its behavior outside of it, the number of degrees of freedom in the set of sparsest solutions of the (g-BPC) is simply the sum of the number

of degrees of freedom in each saturation interval. Yet by Lemma 7.3, there are no degrees of freedom in intervals such that α_n is odd (a sparsest solution is uniquely determined on that interval), and there is one when α_n is even. Therefore, the total number of degrees of freedom of the set of sparsest solutions of the (g-BPC) is equal to the number of even values of α_n for $n \in \{1, \dots, N_s\}$. \square

Illustrations of its Items 2. and 3. with a single saturation region (*i.e.*, $N_s = 1$) are given in Figures 7.5 and 7.6 respectively. In Figure 7.5, the unique sparsest solution is shown. In Figure 7.6, any point \tilde{P}_1 in the segment that connects the points $P_{0,2}$ and \tilde{P} yields one of the sparsest solutions, with a uniquely determined second knot \tilde{P}_2 . In the latter example, there is thus a single degree of freedom $n_{\text{free}}(\mathbf{x}, \mathbf{y}_0)$ in the set of sparsest solutions to the (g-BPC).

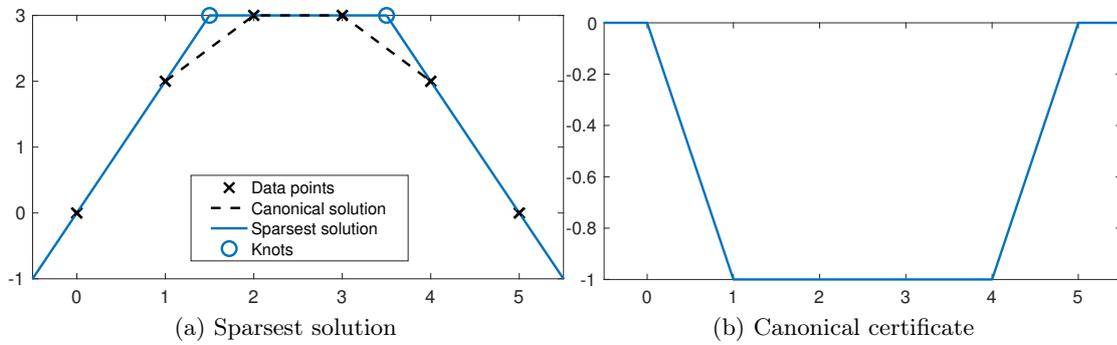


Figure 7.5 – Example with $M = 6$ and $\alpha = 3$ consecutive saturation intervals of η_{cano} at -1 . The unique sparsest solution has $P = 2$ knots.

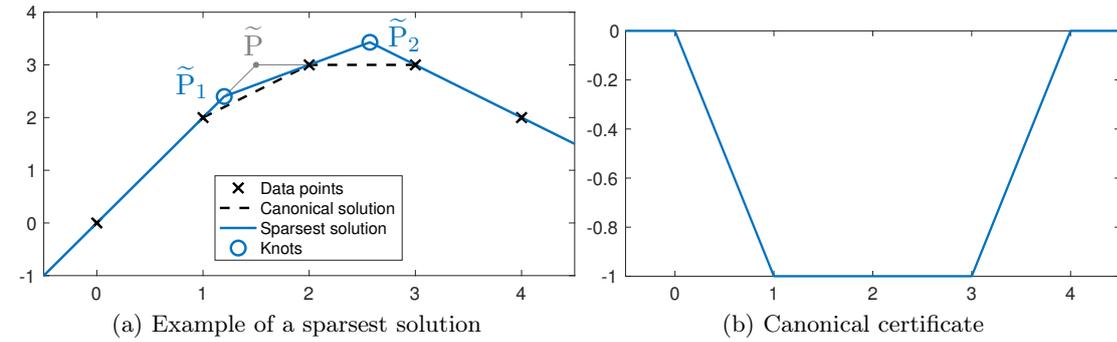


Figure 7.6 – Example with $M = 5$ and $\alpha = 2$ consecutive saturation intervals of η_{cano} at -1 . The sparsest solutions have $P = 2$ knots.

7.4.2 Algorithm for Reaching a Sparsest Solution

The results of Theorem 7.3 suggest a simple yet elegant algorithm for constructing a sparsest solution of the (g-BPC) for given sampling locations $\mathbf{x} = (x_1, \dots, x_M)$ and

data $\mathbf{y}_0 = (y_{0,1}, \dots, y_{0,M})$. The pseudocode is given in Algorithm 7.1, which applies the sparsifying procedure described in Lemma 7.3 in every saturation interval. The proof of Theorem 7.3 guarantees that the output f^* of Algorithm 7.1 is indeed a sparsest solution to the (g-BPC), with sparsity $K_{\min}(\mathbf{x}, \mathbf{y}_0)$ as defined in Theorem 7.3. The following observations can be made concerning Algorithm 7.1.

Input: \mathbf{x}, \mathbf{y}_0

compute a_1, \dots, a_M defined in (7.22);

$[\eta_{\text{cano}}(x_1), \dots, \eta_{\text{cano}}(x_M)] = [0, \text{sign}(a_2), \dots, \text{sign}(a_{M-1}), 0]$;

compute N_s, s_1, \dots, s_{N_s} and $\alpha_1, \dots, \alpha_{N_s}$ defined in (7.39) and (7.40);

$\hat{\boldsymbol{\tau}} = []$; $\hat{\mathbf{a}} = []$;

for $n \leftarrow 1$ **to** N_s **do**

$P \leftarrow \lceil \frac{\alpha_n + 1}{2} \rceil$;

compute $\tilde{\tau}_1, \dots, \tilde{\tau}_P$ and $\tilde{a}_1, \dots, \tilde{a}_P$ using (7.45) or (7.46);

$\hat{\boldsymbol{\tau}} \leftarrow [\hat{\boldsymbol{\tau}}, \tilde{\tau}_1, \dots, \tilde{\tau}_P]$;

$\hat{\mathbf{a}} \leftarrow [\hat{\mathbf{a}}, \tilde{a}_1, \dots, \tilde{a}_P]$;

end

return $f_{\text{opt}} \leftarrow \sum_{k=1}^K \hat{a}_k (\cdot - \hat{\tau}_k)_+$

Algorithm 7.1: Pseudocode of our algorithm to find a sparsest solution of the (g-BPC).

- In the cases where the sparsest solution is not unique, the choice of solution specified by (7.46) (which is *not* the one shown in Figure 7.6) is guided by simplicity. However, it is an arbitrary choice that can be adapted depending on the application.
- Notice that the x_m such that $\eta_{\text{cano}}(x_m) = 0$ need not be included in the vector of knots \mathbf{x}' built in the algorithm, since we have $a_m = 0$. Therefore, there is in fact no knot at x_m in the canonical solution, which implies that the sparsity of f_{cano} is strictly less than $M - 2$. This corresponds to alignment cases of the data points, *i.e.*, the points $P_{0,m-1}$, $P_{0,m}$, and $P_{0,m+1}$ are aligned, as illustrated in Figure 7.2.
- Algorithm 7.1 can be translated into an online algorithm, *i.e.*, an updated solution can be computed efficiently if a new input data point is added. More precisely, when a new data point $P_{0,M+1}$ is added, the reconstructed signal is at worst only modified in the saturation interval $I = [x_{s_n-1}, x_{s_n+\alpha_n}]$ if $x_{M+1} \in I$. Since in practice, we usually have $\alpha_n \ll M$, the computational complexity of updating the solution is typically much smaller than rerunning the complete offline algorithm.

7.4.3 Computational Complexity

Algorithm 7.1 is very fast and memory-efficient; it requires at most two passes through the data points, and thus has linear time and space complexity $\mathcal{O}(M)$ with respect to the number of data points. More precisely, computing the canonical interpolant (*i.e.*, the a_m coefficients using (7.22)) requires about $3M$ subtractions and M divisions, and storing

two arrays of size M . Next, in the worst-case scenario where $\text{sign}(a_2) = \dots = \text{sign}(a_{M-1})$, computing the sparsest interpolant (*i.e.*, the \tilde{a}_k and \tilde{x}_k coefficients using (7.45) or (7.46)) requires approximately M multiplications, M additions, $\frac{M}{2}$ divisions and storing two arrays of size $\frac{M}{2}$. Hence, the complete worst-case time complexity for Algorithm 7.1 requires $4M$ additions, M multiplications and $\frac{3M}{2}$ divisions, and its space complexity is $3M$.

7.5 The Solutions of the (g-BLASSO)

We now focus on the (g-BLASSO) problem, in which the interpolation of the data is no longer required to be exact as in Section 7.3, but is formulated as a penalized problem with a regularization parameter $\lambda > 0$. In practice, such problems are typically formulated when we have access to noise-corrupted measurements $\mathbf{y} = \mathbf{y}_0 + \mathbf{n}$ where $\mathbf{n} \in \mathbb{R}^M$ is a noise term. In this case, we solve the following optimization problem

$$\mathcal{V}_\lambda \triangleq \arg \min_{f \in \mathcal{M}_{\mathbb{D}^2}(\mathbb{R})} \left(\sum_{m=1}^M E(f(x_m), y_m) + \lambda \|D^2 f\|_{\mathcal{M}} \right), \quad (\text{g-BLASSO})$$

where $E(\cdot, y)$ is a strictly convex, coercive, and differentiable cost function (typically quadratic, *i.e.*, $E(z, y) = \frac{1}{2}(z - y)^2$) for any $y \in \mathbb{R}$, and $\lambda > 0$ is a regularization parameter. The latter controls the weight between the data fidelity term $\sum_{m=1}^M E(f(x_m), y_m)$ and the regularization term $\|D^2 f\|_{\mathcal{M}}$, and should therefore be adapted to the noise level.

7.5.1 From the (g-BPC) to the (g-BLASSO): Reduction to the Noiseless Case

We now show, using the strict convexity of $E(\cdot, y)$, that the (g-BLASSO) can be reduced to an optimization problem of the form (g-BPC). The proof is a straightforward generalization of a well-known result in finite-dimensional optimization (see Lemma 4.2 or [21, Lemma 1]); we include it for the sake of completeness.

Proposition 7.7 (Reformulation of the (g-BLASSO) as a (g-BPC) Problem). *Let $\mathbf{x} \in \mathbb{R}^M$ be the ordered sampling locations, and $\mathbf{y} \in \mathbb{R}^M$ with $M \geq 2$. Let $E : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}^+$ be a cost function such that $E(\cdot, y)$ is strictly convex, coercive, and differentiable for every $y \in \mathbb{R}$. Then, there exists a unique $\mathbf{y}_\lambda \in \mathbb{R}^M$ such that, for any $f_{\text{opt}} \in \mathcal{V}_\lambda$, $f_{\text{opt}}(x_m) = y_{\lambda, m}$ for all $m \in \{1, \dots, M\}$. Moreover, we have that the (g-BLASSO) is equivalent to the (g-BPC) with the measurement vector $\mathbf{y}_0 = \mathbf{y}_\lambda$, *i.e.*,*

$$\mathcal{V}_\lambda = \arg \min_{\substack{f \in \mathcal{M}_{\mathbb{D}^2}(\mathbb{R}) \\ f(x_m) = y_{\lambda, m}, \quad m=1, \dots, M}} \|D^2 f\|_{\mathcal{M}}. \quad (7.47)$$

Proof. Assume by contradiction that there exist $f_1, f_2 \in \mathcal{V}_\lambda$ and $m_0 \in \{1, \dots, M\}$ such

that $f_1(x_{m_0}) \neq f_2(x_{m_0})$, and let $f_\gamma = \gamma f_1 + (1 - \gamma)f_2$, where $0 < \gamma < 1$. We then have

$$\begin{aligned}
 & \sum_{m=1}^M E(f_\gamma(x_m), y_m) + \lambda \|D^2 f_\gamma\|_{\mathcal{M}} \\
 & < \gamma \sum_{m=1}^M E(f_1(x_m), y_m) + (1 - \gamma) \sum_{m=1}^M E(f_2(x_m), y_m) + \lambda \left(\gamma \|D^2 f_1\|_{\mathcal{M}} + (1 - \gamma) \|D^2 f_2\|_{\mathcal{M}} \right) \\
 & = \gamma \mathcal{J}_\lambda^0 + (1 - \gamma) \mathcal{J}_\lambda^0 = \mathcal{J}_\lambda^0,
 \end{aligned} \tag{7.48}$$

where \mathcal{J}_λ^0 is the optimal cost of the (g-BLASSO). The inequality is due to the convexity of the $\|\cdot\|_{\mathcal{M}}$ norm and of $E(\cdot, y)$ for any $y \in \mathbb{R}$. The fact that it is strict is due to the strict convexity of $E(\cdot, y_{m_0})$ and the fact that $f_1(x_{m_0}) \neq f_2(x_{m_0})$. Yet since \mathcal{V}_λ is a convex set, we have $f_\gamma \in \mathcal{V}_\lambda$: this implies that $\mathcal{J}_\lambda^0 = \sum_{m=1}^M E(f_\gamma(x_m), y_m) + \lambda \|D^2 f_\gamma\|_{\mathcal{M}} < \mathcal{J}_\lambda^0$, which yields a contradiction.

Therefore, there exists a unique vector $\mathbf{y}_\lambda \in \mathbb{R}^M$ such that for any $f_{\text{opt}} \in \mathcal{V}_\lambda$, $f_{\text{opt}}(x_m) = y_{\lambda, m}$ for all $m \in \{1, \dots, M\}$. This implies that $\mathcal{V}_\lambda \subset \{f \in \mathcal{M}_{D^2}(\mathbb{R}) : f(x_m) = y_{\lambda, m}, 1 \leq m \leq M\}$. Moreover, we have that for any $f_{\text{opt}} \in \mathcal{V}_\lambda$, $E(f_{\text{opt}}(x_m), y_m) = E(y_{\lambda, m}, y_m)$, and thus that the data fidelity is constant in the constrained space $\{f \in \mathcal{M}_{D^2}(\mathbb{R}) : f(x_m) = y_{\lambda, m}, 1 \leq m \leq M\}$. This proves the equality between the solution sets of the (g-BLASSO) and (7.47). \square

The implications of Proposition 7.7 for our problem are huge: it implies that all the results of Section 7.3—in particular, uniqueness, form the solutions, and sparsest solutions—can be applied to the penalized Problem (g-BLASSO). The only—but crucial—catch is that the samples $\mathbf{y}_\lambda \in \mathbb{R}^M$ are unknown. Fortunately, the following proposition enables us to compute them through a standard ℓ_1 -regularized discrete optimization.

Proposition 7.8. *Assume that the hypotheses of Proposition 7.7 are met. Then, the vector $\mathbf{y}_\lambda \in \mathbb{R}^M$ defined in Proposition 7.7 is the unique solution of the discrete minimization problem*

$$\mathbf{y}_\lambda = \arg \min_{\mathbf{z} \in \mathbb{R}^M} \left(\sum_{m=1}^M E(z_m, y_m) + \lambda \|\mathbf{L}\mathbf{z}\|_1 \right), \tag{7.49}$$

where $\mathbf{L} \in \mathbb{R}^{(M-2) \times M}$ is given by

$$\mathbf{L} \triangleq \begin{pmatrix} v_1 & -(v_1 + v_2) & v_2 & 0 & \cdots & 0 \\ 0 & v_2 & -(v_2 + v_3) & v_3 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & v_{M-2} & -(v_{M-2} + v_{M-1}) & v_{M-1} \end{pmatrix}, \tag{7.50}$$

and $\mathbf{v} \triangleq (v_1, \dots, v_{M-1}) \in \mathbb{R}^{M-1}$ is defined as $v_m \triangleq \frac{1}{x_{m+1} - x_m}$ for $m \in \{1, \dots, M-1\}$.

Proof. In this proof, we denote by $f_{\mathbf{z}}$ the canonical solution (defined in Definition 7.4) of the (g-BPC) with sampling locations \mathbf{x} and data point $\mathbf{y}_0 = \mathbf{z}$. We first prove that if $\mathbf{z}_{\text{opt}} \in \mathbb{R}^M$ is a solution of Problem (7.49), then $f_{\mathbf{z}_{\text{opt}}} \in \mathcal{M}_{\mathbb{D}^2}(\mathbb{R})$ is a solution of the (g-BLASSO). We then deduce that for all $m \in \{1, \dots, M\}$, $z_m = f_{\mathbf{z}_{\text{opt}}}(x_m) = y_{\lambda, m}$ (where the last equality is true thanks to Proposition 7.7), which proves the desired result, *i.e.*, $\mathbf{y}_{\lambda} = \mathbf{z}_{\text{opt}}$ is the unique solution of Problem (7.49).

Let $\mathbf{z} \in \mathbb{R}^M$. Using Equations (7.21) and (7.22), we have that $\|\mathbb{D}^2 f_{\mathbf{z}}\|_{\mathcal{M}} = \sum_{m=2}^{M-1} |a_m|$, where $a_m = \frac{z_{m+1} - z_m}{x_{m+1} - x_m} - \frac{z_m - z_{m-1}}{x_m - x_{m-1}}$. Therefore, we have $\|\mathbb{D}^2 f_{\mathbf{z}}\|_{\mathcal{M}} = \|\mathbf{L}\mathbf{z}\|_1$, where \mathbf{L} is given by Equation (7.50). This yields $\sum_{m=1}^M E(f_{\mathbf{z}}(x_m), y_m) + \lambda \|\mathbf{L}\mathbf{z}\|_1 = \sum_{m=1}^M E(z_m, y_m) + \lambda \|\mathbf{L}\mathbf{z}\|_1$. Applied to the particular case $\mathbf{z} = \mathbf{y}_{\lambda}$, we obtain the equality $\sum_{m=1}^M E(y_{\lambda, m}, y_m) + \lambda \|\mathbf{L}\mathbf{y}_{\lambda}\|_1 = \mathcal{J}_{\lambda}^0$, where \mathcal{J}_{λ}^0 is the optimal cost of the (g-BLASSO), since by Proposition 7.5, $f_{\mathbf{y}_{\lambda}} \in \mathcal{V}_{\lambda}$. This proves that the optimal value of Problem (7.49) is lower or equal than \mathcal{J}_{λ}^0 .

Next, let \mathbf{z}_{opt} be a solution of Problem (7.49) (which exists due to the coercivity of $E(\cdot, y)$ for any $y \in \mathbb{R}$). We thus have from before that

$$\mathcal{J}_{\lambda}^0 \leq \sum_{m=1}^M E(f_{\mathbf{z}_{\text{opt}}}(x_m), y_m) + \lambda \|\mathbb{D}^2 f_{\mathbf{z}_{\text{opt}}}\|_{\mathcal{M}} = \sum_{m=1}^M E(z_m, y_m) + \lambda \|\mathbf{L}\mathbf{z}_{\text{opt}}\|_1 \leq \mathcal{J}_{\lambda}^0, \quad (7.51)$$

which yields the desired result $f_{\mathbf{z}_{\text{opt}}} \in \mathcal{V}_{\lambda}$. \square

7.5.2 Algorithm for Reaching a Sparsest Solution of the (g-BLASSO)

By combining results from the previous sections, we now formulate the following simple algorithmic pipeline to reach a sparsest solution of the (g-BLASSO).

Proposition 7.9. *Let $\mathbf{x} \in \mathbb{R}^M$ be the ordered sampling locations and $\mathbf{y} \in \mathbb{R}^M$ with $M \geq 2$, and let $E : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}^+$ be a cost function such that $E(\cdot, y)$ is strictly convex, coercive, and differentiable for any $y \in \mathbb{R}$. Let the function f_{opt} be obtained through the following two-step procedure:*

1. Compute $\mathbf{y}_{\lambda} \in \mathbb{R}^M$ (defined in Proposition 7.7) by solving Problem (7.49);
2. Apply Algorithm 7.1 with the measurement vector $\mathbf{y}_0 = \mathbf{y}_{\lambda}$ to compute a sparsest solution f_{opt} of the (g-BPC) given by Equation (7.47).

Then, f_{opt} is one of the sparsest solutions to the (g-BLASSO), with sparsity $K_{\min}(\mathbf{x}, \mathbf{y}_{\lambda})$ as defined in Theorem 7.3.

Proof. Proposition 7.7 guarantees that the (g-BLASSO) is equivalent to the (g-BPC) with

the measurement vector $\mathbf{y}_0 = \mathbf{y}_\lambda$. Proposition 7.8 then specifies that \mathbf{y}_λ can be computed by solving Problem (7.49). Finally, as demonstrated in the proof of Theorem 7.3, the output f_{opt} of Algorithm 7.4.2 reaches a sparsest solution of the corresponding (g-BPC) problem, which thus has sparsity $K_{\min}(\mathbf{x}, \mathbf{y}_\lambda)$. \square

Proposition 7.9 proposes a simple but very powerful algorithm. It reaches a sparsest solution of the (g-BLASSO) - a challenging task a priori - in two simple steps. The first consists in solving a standard ℓ_1 -regularized discrete problem, for which many off-the-shelf solvers such as ADMM [19] are available (see Appendix A). The second is our proposed sparsifying procedure, which converges in finite time. The following remarks can be made concerning Proposition 7.9.

Remark 7.4. *Algorithm 7.1 still converges to a solution of the (g-BLASSO) when E is only a convex function, and not strictly convex as assumed in Propositions 7.7 and 7.8. The difference is that Proposition 7.7 no longer holds true in that there is no unique vector of measurements \mathbf{y}_λ . The solution set of the constrained Problem (7.47) is thus in general a strict subset of \mathcal{V}_λ . Hence, the obtained solution is not necessarily the sparsest solution of the full solution set \mathcal{V}_λ , but only of this subset.*

As for the assumption that E is differentiable, it is not a requirement for Proposition 7.9. However, as it is needed later on in Proposition 7.10, we include it in order to have consistent assumptions concerning E throughout this chapter.

7.5.3 Range of the Regularization Parameter λ

In practice, the choice of the regularization parameter λ is the critical element that determines the performance of our algorithm. Although this choice is highly data-dependant, in this section, we show that the search can be restricted to a bounded interval. The lower bound is $\lambda \rightarrow 0$, which corresponds at the limit to exact interpolation, that is the (g-BPC). The upper bound $\lambda \rightarrow +\infty$ corresponds to the linear regression regime, which is described in the following proposition.

Proposition 7.10 (Linear Regression Regime of the (g-BLASSO)). *Let $\mathbf{x} \in \mathbb{R}^M$ be the ordered sampling locations and $\mathbf{y} \in \mathbb{R}^M$ with $M \geq 2$. Let $E : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}^+$ be a cost function such that $E(\cdot, y)$ is strictly convex, coercive, and differentiable for any $y \in \mathbb{R}$. Then, the following properties hold.*

1. *There is a unique solution $(\beta_{\text{opt}_0}, \beta_{\text{opt}_1}) \in \mathbb{R}^2$ to the linear regression problem*

$$(\beta_{\text{opt}_0}, \beta_{\text{opt}_1}) \triangleq \arg \min_{(\beta_0, \beta_1) \in \mathbb{R}^2} \left(\sum_{m=1}^M E(\beta_0 + \beta_1 x_m, y_m) \right). \quad (7.52)$$

We can thus define the value

$$\lambda_{max} \triangleq \left\| \mathbf{L}^{T\dagger} \begin{pmatrix} \partial_1 E(\beta_{opt_0} + \beta_{opt_1} x_m, y_1) \\ \vdots \\ \partial_1 E(\beta_{opt_0} + \beta_{opt_1} x_M, y_M) \end{pmatrix} \right\|_{\infty}, \quad (7.53)$$

where $\partial_1 E$ denotes the partial derivative with respect to the first variable of E , the matrix $\mathbf{L}^{T\dagger}$ denotes the pseudoinverse of \mathbf{L}^T , and \mathbf{L} is defined as in (7.50).

2. For any $\lambda \geq \lambda_{max}$, the solution to the discrete Problem (7.49) is given by $\mathbf{y}_\lambda = \beta_{opt_0} \mathbf{1} + \beta_{opt_1} \mathbf{x}$, where $\mathbf{1} \triangleq (1, \dots, 1) \in \mathbb{R}^M$.
3. For any $\lambda \geq \lambda_{max}$, the solution to the (g-BLASSO) is unique and is the linear function f_{max} given by $f_{max}(x) \triangleq \beta_{opt_0} + \beta_{opt_1} x$.

Proof. Item 1: Let $\mathcal{J}(\beta_0, \beta_1) = \sum_{m=1}^M E(\beta_0 + \beta_1 x_m, y_m)$ be the objective function of Problem (7.52). We show that Problem (7.52) indeed has a unique solution by proving that \mathcal{J} is strictly convex and coercive when $M \geq 2$ and the x_m are pairwise distinct.

Concerning the coercivity, let $\|(\beta_0, \beta_1)\|_2 \rightarrow +\infty$. Assume by contradiction that $\beta_0 + \beta_1 x_m$ is bounded for every $m \in \{1, \dots, M\}$. Then, since $M \geq 2$, $\beta_0 + \beta_1 x_1 - (\beta_0 + \beta_1 x_2) = \beta_1(x_1 - x_2)$ must also be bounded, which implies that β_1 is bounded since the x_m are pairwise distinct. Therefore, we must have $|\beta_0| \rightarrow +\infty$, which implies that $|\beta_0 + \beta_1 x_1| \rightarrow +\infty$ which yields a contradiction. Therefore, there exists a $m_0 \in \{1, \dots, M\}$ such that $|\beta_0 + \beta_1 x_{m_0}| \rightarrow +\infty$. The coercivity of \mathcal{J} then directly follows from that of $E(\cdot, y_{m_0})$.

Next, to prove the strict convexity of \mathcal{J} , let $(\beta_0, \beta_1), (\beta'_0, \beta'_1) \in \mathbb{R}^2$ with $(\beta_0, \beta_1) \neq (\beta'_0, \beta'_1)$, and $0 < s < 1$. For any m , we have $s\beta_0 + (1-s)\beta'_0 + (s\beta_1 + (1-s)\beta'_1)x_m = s(\beta_0 + \beta_1 x_m) + (1-s)(\beta'_0 + \beta'_1 x_m)$. Since $(\beta_0, \beta_1) \neq (\beta'_0, \beta'_1)$ and the x_m are distinct, the equation $\beta_0 + \beta_1 x_m = \beta'_0 + \beta'_1 x_m$ can only be satisfied for at most a single $m \in \{1, \dots, M\}$. Yet $M \geq 2$, which implies that $\exists m_0, \beta_0 + \beta_1 x_{m_0} \neq \beta'_0 + \beta'_1 x_{m_0}$. Therefore, due to the strict convexity of $E(\cdot, y_{m_0})$, we have

$$E\left(\left(s\beta_0 + (1-s)\beta'_0\right) + \left(s\beta_1 + (1-s)\beta'_1\right)x_{m_0}, y_{m_0}\right) \quad (7.54)$$

$$< sE\left(\beta_0 + \beta_1 x_{m_0}, y_{m_0}\right) + (1-s)E\left(\beta'_0 + \beta'_1 x_{m_0}, y_{m_0}\right). \quad (7.55)$$

It then follows from the convexity of $E(\cdot, y_m)$ for all m that $\mathcal{J}(s(\beta_0, \beta_1) + (1-s)(\beta'_0, \beta'_1)) < s\mathcal{J}(\beta_0, \beta_1) + (1-s)\mathcal{J}(\beta'_0, \beta'_1)$, which proves the strict convexity of \mathcal{J} . Together with the fact that \mathcal{J} is coercive, this proves that (7.52) has a unique solution.

Item 2: Assume that $\lambda \geq \lambda_{max}$. By Fermat's rule, a vector \mathbf{z}_{opt} is a solution of Problem (7.49) if and only if the zero vector belongs to the subdifferential of the objective

function evaluated at \mathbf{z}_{opt} . We thus have $\mathbf{z}_{\text{opt}} = \mathbf{y}_\lambda$ if and only if

$$\mathbf{0} \in \underbrace{\begin{pmatrix} \partial_1 E(z_{\text{opt},1}, y_1) \\ \vdots \\ \partial_1 E(z_{\text{opt},M}, y_M) \end{pmatrix}}_{\triangleq \mathbf{v}(\mathbf{z}_{\text{opt}})} + \lambda \partial \|\mathbf{L} \cdot\|_1(\mathbf{z}_{\text{opt}}), \quad (7.56)$$

where ∂_1 denotes the partial derivative with respect to the first variable, and ∂ the subdifferential. The chain rule for subdifferentials [197, Theorem 23.9] yields $\partial \|\mathbf{L} \cdot\|_1(\mathbf{z}) = \{\mathbf{L}^T \mathbf{g} : \mathbf{g} \in \partial \|\cdot\|_1(\mathbf{Lz}) \subset \mathbb{R}^{M-2}\}$, where $\partial \|\cdot\|_1(\mathbf{a}) = \{\mathbf{g} \in \mathbb{R}^{M-2} : \|\mathbf{g}\|_\infty \leq 1, \mathbf{a}^T \mathbf{g} = \|\mathbf{a}\|_1\}$. The vector \mathbf{Lz}_{opt} lists the weights a_m associated to the knots of the canonical solution $f_{\mathbf{z}_{\text{opt}}}$ (see the proof of Proposition 7.8). Therefore, the linear regression case (in which $f_{\mathbf{z}_{\text{opt}}}$ has no knot) corresponds to $\mathbf{Lz}_{\text{opt}} = \mathbf{0}$. In this case, since $\partial \|\cdot\|_1(\mathbf{0}) = \{\mathbf{g} \in \mathbb{R}^{M-2} : \|\mathbf{g}\|_\infty \leq 1\}$, the optimality condition (7.56) now reads

$$\exists \mathbf{g} \in \mathbb{R}^{M-2}, \quad \|\mathbf{g}\|_\infty \leq 1, \quad \text{s.t.} \quad \mathbf{v}(\mathbf{z}_{\text{opt}}) + \lambda \mathbf{L}^T \mathbf{g} = \mathbf{0}. \quad (7.57)$$

We now prove that $\mathbf{z}_{\text{opt}} = \beta_{\text{opt}_0} \mathbf{1} + \beta_{\text{opt}_1} \mathbf{x}$ satisfies the optimality conditions (7.56), and thus that $\mathbf{y}_\lambda = \beta_{\text{opt}_0} \mathbf{1} + \beta_{\text{opt}_1} \mathbf{x}$. To achieve this, we prove that $\mathbf{g} = -\frac{1}{\lambda} \mathbf{L}^{T\dagger} \mathbf{v}(\mathbf{z}_{\text{opt}})$ satisfies $\mathbf{v}(\mathbf{z}_{\text{opt}}) + \lambda \mathbf{L}^T \mathbf{g} = \mathbf{0}$. Firstly, since $\lambda \geq \lambda_{\text{max}}$, we have that $\|\mathbf{g}\|_\infty \leq 1$ by definition of λ_{max} . Next, let V be the orthogonal complement of $\ker \mathbf{L} \subset \mathbb{R}^M$. A known property of the pseudoinverse operator [198, Corollary 7] is that $\mathbf{L}^T \mathbf{L}^{T\dagger}$ is the orthogonal projection operator onto V . By decomposing $\mathbf{v}(\mathbf{z}_{\text{opt}}) = \mathbf{v}_1 + \mathbf{v}_2$, where $\mathbf{v}_1 \in V$ and $\mathbf{v}_2 \in \ker \mathbf{L}$, we thus get $\mathbf{v}(\mathbf{z}_{\text{opt}}) + \lambda \mathbf{L}^T \mathbf{g} = \mathbf{v}_2$. Yet $\ker \mathbf{L} = \text{span}\{\mathbf{1}, \mathbf{x}\}$, since the canonical solutions $f_{\mathbf{1}}$ and $f_{\mathbf{x}}$ (that satisfy $f_{\mathbf{1}}(x_m) = 1$ and $f_{\mathbf{x}}(x_m) = x_m$ for every $m \in \{1, \dots, M\}$ respectively) are linear functions that are thus not penalized by the regularization. The optimality conditions of Problem (7.52) (*i.e.*, setting the gradient to zero) then yield $\mathbf{v}(\mathbf{z}_{\text{opt}}) \perp \ker \mathbf{L}$, which implies that $\mathbf{v}_2 = \mathbf{0}$ and thus that $\mathbf{v}(\mathbf{z}_{\text{opt}}) + \lambda \mathbf{L}^T \mathbf{g} = \mathbf{0}$. This proves that \mathbf{z}_{opt} satisfies the optimality condition of Problem (7.49), and thus that $\mathbf{z}_{\text{opt}} = \mathbf{y}_\lambda = \beta_{\text{opt}_0} \mathbf{1} + \beta_{\text{opt}_1} \mathbf{x}$.

Item 3: Due to Item 2, we have $\mathbf{y}_\lambda = \beta_{\text{opt}_0} \mathbf{1} + \beta_{\text{opt}_1} \mathbf{x}$ which implies that the points $\begin{bmatrix} x_m & y_{\lambda,m} \end{bmatrix}^T$ are aligned. Hence, the canonical dual certificate of the constrained Problem (7.47) is $\eta_{\text{cano}} = 0$, which is nondegenerate. By Proposition 7.6, this implies that the unique solution to Problem (7.47) is the canonical solution $f_{\mathbf{z}_{\text{opt}}} = f_{\text{max}} = \beta_{\text{opt}_0} + \beta_{\text{opt}_1}(\cdot)$. Due to the equivalence between problems (7.47) and the (g-BLASSO) proved in Proposition 7.7, this concludes the proof. \square

Proposition 7.10 guarantees that the range of λ can be restricted to the interval $(0, \lambda_{\text{max}}]$: indeed, all values $\lambda \geq \lambda_{\text{max}}$ lead to linear regression. Moreover, the value of λ_{max} given in (7.53) only depends on the data $\mathbf{x}, \mathbf{y} \in \mathbb{R}^M$ and is easy to compute numerically - the most costly step being the computation of the pseudoinverse $\mathbf{L}^{T\dagger}$. Note that Item 2 in Proposition 7.10, which stems from duality theory, is a generalization of a well-

known result for the LASSO Problem [199, Proposition 1.3], which plays a crucial role in the homotopy method [200]. The difference here is the presence of a noninvertible regularization matrix \mathbf{L} in Problem (7.49), which requires additional arguments in the proof.

7.6 Experimental Results

In this section, we describe the implementation of our two-step algorithm presented in Section 7.5.2 and show our experimental results. The first step of our algorithm - which consists in solving problem (7.49) with ADMM (see Appendix A) - is implemented using GlobalBioIm, a Matlab inverse-problem library developed by the Biomedical Imaging Group at EPFL [118]. In all our experiments, we choose the standard quadratic data fidelity loss $E(z, y) = \frac{1}{2}(z - y)^2$. This choice leads to $\partial_1 E(z, y) = z - y$, which enables the simple computation of λ_{\max} using (7.53).

We present an illustrative example with $M = 30$ simulated data points in Figure 7.7. A small number is chosen for visualization purposes; an application of our algorithm with a larger number of $M = 200$ data points was shown in Figure 7.1. The sampling locations x_m are generated following a uniform distribution in the $[\frac{m-1}{M}, \frac{m}{M}]$ intervals for $m = 1, \dots, M$. Next, the ground-truth signal, a piecewise-linear spline f_0 with 2 knots, is generated, with random knot locations τ_m within the interval $[0, 1]$, and i.i.d. Gaussian amplitudes a_m ($\sigma_a^2 = 1$). We then have $y_m = f_0(x_m) + n_m$ for $m = 1, \dots, M$, where $\mathbf{n} \in \mathbb{R}^M$ is i.i.d. Gaussian noise ($\sigma_n^2 = 4 \times 10^{-4}$).

7.6.1 Extreme Values of λ

The reconstructions using our algorithm for extreme values of λ - *i.e.*, $\lambda \rightarrow 0$ which leads to exact interpolation of the data, and $\lambda = \lambda_{\max}$ which leads to linear regression - are shown in Figure 7.7a. Clearly, none of these solutions are satisfactory: on one hand, linear regression is too simple to model the data adequately. On the other hand, the exact interpolator suffers from overfitting. Although thanks to the sparsification procedure in Algorithm 7.1, its sparsity $K_{\min}(\mathbf{x}, \mathbf{y}_\lambda) = 20$ is smaller than the theoretical bound $M - 2 = 28$ given by Theorem 3.3, it is still clearly much larger than the desired outcome.

7.6.2 Sparsity *versus* Data Fidelity Loss Trade-Off

Next, we show the sparsity $K_{\min}(\mathbf{x}, \mathbf{y}_\lambda)$ *versus* error $\|\mathbf{y} - \mathbf{y}_\lambda\|$ trade-off curve in Figure 7.7b. The latter was obtained by applying our algorithm with 20 values of λ (equispaced on a logarithmic scale) within the range $[\lambda_{\min}, \lambda_{\max}]$, with $\lambda_{\max} = 0.1713$ (as defined in (7.53)) and $\lambda_{\min} \triangleq 10^{-5} \times \lambda_{\max}$. We thus observe the evolution from exact interpolation to linear regression as λ increases.

Ideally, one would like to choose a value of λ that minimizes $\|\mathbf{y}_0 - \mathbf{y}_\lambda\|$, *i.e.*, the error with respect to the noiseless data \mathbf{y}_0 . However, in practice, the noiseless data is unknown, and one must use the noisy data \mathbf{y} . Depending on the noise level, solely minimizing $\|\mathbf{y} - \mathbf{y}_\lambda\|$ might not be a desirable objective, since it leads to overfitting. Hence, we consider the trade-off between data fidelity loss and sparsity as a proxy for the standard universality *versus* simplicity trade-off in machine learning. Note that we choose the data fidelity loss $\|\mathbf{y} - \mathbf{y}_\lambda\|$ instead of λ as the x -axis metric, since it is an increasing function of the latter, and the former is easier to interpret.

This trade-off curve does not specify a single optimal value of the regularization parameter λ . Instead, it helps the user choose an appropriate balance by giving quantitative, interpretable data about the possible trade-offs. A key observation is that this curve is not necessarily monotonous: the sparsity can increase as $\|\mathbf{y} - \mathbf{y}_\lambda\|$ increases, as shown in Figure 7.7b. This lack of monotonicity is rather counter-intuitive, since the overall trend as λ increases is to go from sparsity $K_{\min}(\mathbf{x}, \mathbf{y}) = 20$ to $K_{\min}(\mathbf{x}, \mathbf{y}_{\lambda_{\max}}) = 0$. Note that a similar behavior has been known to occur in the context of the homotopy method [199], although it is far from being systematic. However, the interesting feature is that, in the sparsity *versus* error trade-off, some values of λ are sometimes strictly better than others for both metrics, such as the star point over the square point in Figure 7.7b. Having access to the full trade-off curve such as Figure 7.7b is very helpful to judiciously select a suitable value of λ . This holds true as well when the curve is monotonic: indeed, the user should select the value of λ such that the data fidelity is lowest for the desired level of sparsity, *i.e.*, the leftmost point of every plateau.

7.6.3 Example Reconstructions

To illustrate the nonmonotonicity of the sparsity *versus* error curve, examples of reconstructions for two specific values of λ are shown in Figures 7.7c and 7.7d. Indeed, the former reconstruction has a lower value of λ , and thus lower data-fidelity loss. Nevertheless, the reconstruction in Figure 7.7c is sparser, with $K_{\min}(\mathbf{x}, \mathbf{y}_\lambda) = 3$ *versus* 6 in Figure 7.7d. Note that this gap is not a numerical artefact, since the magnitude of the weights \tilde{a}_k associated to the knots in Figure 7.7d is much greater than numerical precision. This indicates that the value of λ for Figure 7.7c should be preferred to that of 7.7d.

7.7 Conclusion

In this chapter, we fully described the solution set of the (g-BPC), which consists in interpolating data points by minimizing the TV norm of the second derivative. More precisely, we specified the cases in which it has a unique solution, the form of all the solutions, and the subset of sparsest solutions. We also proposed a simple and fast algorithm to reach (one of) the sparsest solution(s). We then extended these results to

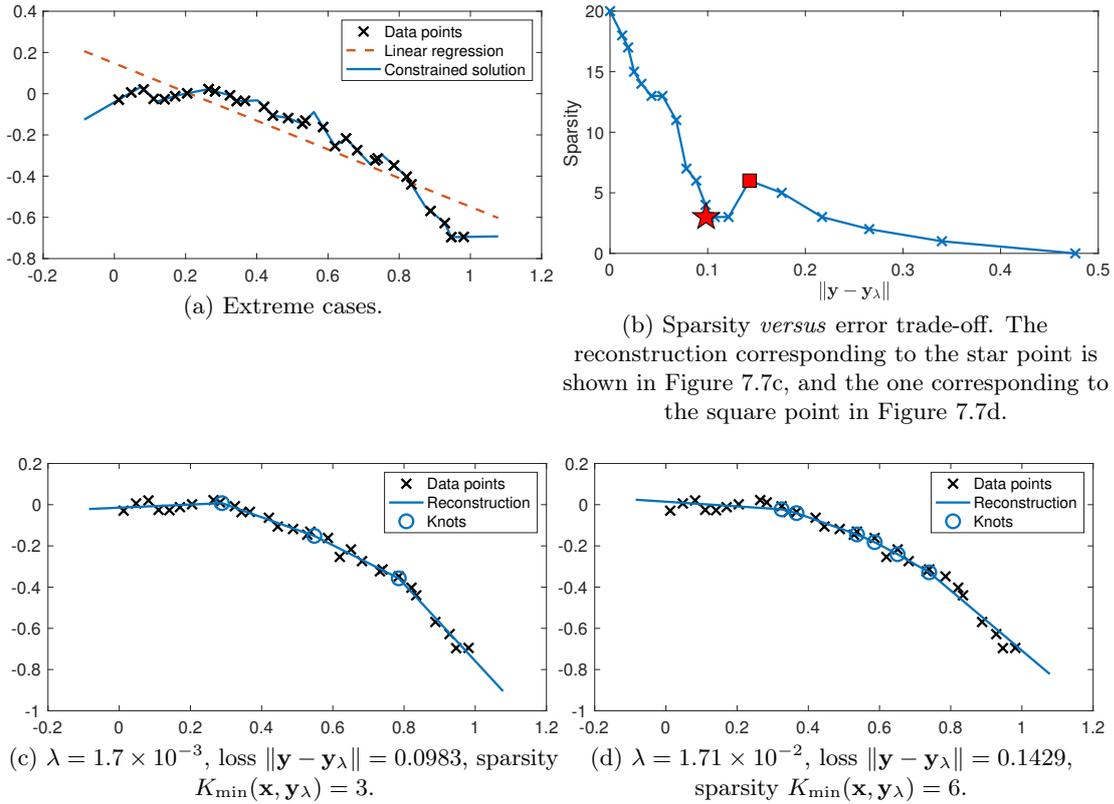


Figure 7.7 – Example of reconstruction for varying regularization $0 \leq \lambda \leq \lambda_{\max} = 0.1713$ with $M = 30$ simulated data points.

the (g-BLASSO), by showing that it can be reformulated as a (g-BPC) problem. Next, we introduced a two-step algorithm to solve the (g-BLASSO), the first step of which consists in solving a discrete ℓ_1 -regularized problem, and the second in applying our algorithm to solve a (g-BPC) problem. Finally, we applied our algorithm to some simulated data, and suggested plotting the sparsity *versus* data fidelity error plot in order to judiciously select a suitable value of the regularization parameter. This chapter paves the way for the study of supervised learning problems through the formulation of variational inverse problems with TV-based regularization, by completely describing the one-dimensional scenario. A future exciting—albeit much more challenging—prospect would be to achieve similar results in higher dimensions, *i.e.*, to reconstruct functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ with $d > 1$. This would be a major milestone to better understand ReLU networks and deep learning in general, whose practical outstanding performances are yet to be fully explained.

7.A Proof of Proposition 7.1

The forward operator considered in this paper is a sampling operator (the functions $f \in \mathcal{M}_{D^2}(\mathbb{R})$ are sampled at the locations $x_m \in \mathbb{R}$ for $m \in \{1, \dots, M\}$). For the

convenience of the proof and following the notation of the whole thesis, we denote it as the linear operator $\nu : \mathcal{M}_{D^2}(\mathbb{R}) \rightarrow \mathbb{R}^M$ such that

$$\forall f \in \mathcal{M}_{D^2}(\mathbb{R}), \quad \nu(f) \triangleq (f(x_m))_{1 \leq m \leq M}. \quad (7.58)$$

The proof of Proposition 7.1 is divided in several steps. First, we reformulate the (g-BPC) into an equivalent optimization problem thanks to the decomposition of any $f \in \mathcal{M}_{D^2}(\mathbb{R})$ given by (7.4). This is stated in the following lemma.

Lemma 7.4. *The Problem (g-BPC) is equivalent to*

$$\min_{(w, \beta_0, \beta_1) \in \mathcal{M}(\mathbb{R}) \times \mathbb{R}^2} \left(\iota_{\{\mathbf{y}_0\}} \left(\nu_{\mathcal{M}}(w) + \beta_0 \mathbf{1} + \beta_1 \mathbf{x} \right) + \|w\|_{\mathcal{M}} \right), \quad (7.59)$$

where $\iota_{\{\mathbf{y}_0\}}$ is the indicator of the convex set $\{\mathbf{y}_0\}$, which is zero at \mathbf{y}_0 and $+\infty$ elsewhere, and

$$\nu_{\mathcal{M}} \triangleq \nu \circ D_0^{-2} : \mathcal{M}(\mathbb{R}) \rightarrow \mathbb{R}^M \quad (7.60)$$

is the modified forward operator. More precisely, there exists a bijection given by the unique decomposition of any $f \in \mathcal{M}_{D^2}(\mathbb{R})$ as $f = D_0^{-2}\{w\} + \beta_0 + \beta_1(\cdot)$ with $(w, \beta_0, \beta_1) \in \mathcal{M}(\mathbb{R}) \times \mathbb{R}^2$ (see(7.4)) between the solution sets of both optimization problems.

Proof. The result immediately follows by plugging in the decomposition $f = D_0^{-2}\{w\} + \beta_0 + \beta_1(\cdot)$ with $(w, \beta_0, \beta_1) \in \mathcal{M}(\mathbb{R}) \times \mathbb{R}^2$ into the objective function of Problem (g-BPC), and by reformulating the interpolation constraints $f(x_m) = y_{0,m}$ using the indicator function $\iota_{\{\mathbf{y}_0\}}$. \square

From now on, we consider the equivalent Problem (7.59) and analyze it using tools from duality theory. The search space $\mathcal{M}(\mathbb{R}) \times \mathbb{R}^2$ of this optimization problem is endowed with the weak* topology, which is defined in terms of its predual space $\mathcal{C}_0(\mathbb{R}) \times \mathbb{R}^2$. By definition of D_0^{-2} via its Schwartz kernel g_0 (see Theorem 2.1), the modified operator $\nu_{\mathcal{M}}$ can be expressed as $\nu_{\mathcal{M}}(w) = (\langle w, g_0(x_m, \cdot) \rangle)_{1 \leq m \leq M}$, where $g_0(x_m, \cdot) \in \mathcal{C}_0(\mathbb{R})$ for all $m \in \{1, \dots, M\}$. Since $\mathcal{M}(\mathbb{R})$ is the dual of $\mathcal{C}_0(\mathbb{R})$, this implies that the linear functional $\nu_{\mathcal{M}} : \mathcal{M}(\mathbb{R}) \rightarrow \mathbb{R}^M$ is weak* continuous [113, Theorem IV.20]. The adjoint operator $\nu_{\mathcal{M}}^* : \mathbb{R}^M \rightarrow \mathcal{C}_0(\mathbb{R})$ of $\nu_{\mathcal{M}}$ is thus uniquely defined and is given by

$$\forall \mathbf{c} \in \mathbb{R}^M, \quad \nu_{\mathcal{M}}^*(\mathbf{c}) = \sum_{m=1}^M c_m g_0(x_m, \cdot), \quad (7.61)$$

since $\langle w, \nu_{\mathcal{M}}^*(\mathbf{c}) \rangle = \langle \nu_{\mathcal{M}}(w), \mathbf{c} \rangle = \left\langle \left(\langle w, g_0(x_m, \cdot) \rangle \right)_{1 \leq m \leq M}, \mathbf{c} \right\rangle = \left\langle w, \sum_{m=1}^M c_m g_0(x_m, \cdot) \right\rangle$, for all $w \in \mathcal{M}(\mathbb{R})$ and $\mathbf{c} \in \mathbb{R}^M$.

The second part of the proof of Proposition 7.1 consists in determining the dual problem of (7.59), proving that strong duality between the primal and dual problem holds (*i.e.*, that the optimal values of both problems are equal and finite), and then deriving the optimality conditions which characterize the solutions of Problem (7.59). This is done in the next lemma.

Lemma 7.5. *The dual problem of (7.59) is given by*

$$\sup_{\mathbf{c} \in \mathcal{C}} \langle \mathbf{y}_0, \mathbf{c} \rangle, \quad \text{with } \mathcal{C} \triangleq \left\{ \mathbf{c} \in \mathbb{R}^M : \langle \mathbf{c}, \mathbf{1} \rangle = \langle \mathbf{c}, \mathbf{x} \rangle = 0, \|\boldsymbol{\nu}_{\mathcal{M}}^*(\mathbf{c})\|_{\infty} \leq 1 \right\}. \quad (7.62)$$

Moreover, it has at least one solution and strong duality holds between Problems (7.59) and (7.62). Finally, for any $(w, \beta_0, \beta_1) \in \mathcal{M}(\mathbb{R}) \times \mathbb{R}^2$ and $\mathbf{c} \in \mathbb{R}^M$, the following statements are equivalent:

1. (w, β_0, β_1) is a solution of (7.59) and \mathbf{c} is a solution of (7.62).
2. (w, β_0, β_1) and \mathbf{c} satisfy the following conditions:

$$\boldsymbol{\nu}_{\mathcal{M}}(w) + \beta_0 \mathbf{1} + \beta_1 \mathbf{x} = \mathbf{y}_0, \quad (7.63)$$

$$\langle \mathbf{c}, \mathbf{1} \rangle = \langle \mathbf{c}, \mathbf{x} \rangle = 0, \quad \|w\|_{\mathcal{M}} = \langle w, \boldsymbol{\nu}_{\mathcal{M}}^*(\mathbf{c}) \rangle, \quad \text{and} \quad \|\boldsymbol{\nu}_{\mathcal{M}}^*(\mathbf{c})\|_{\infty} \leq 1. \quad (7.64)$$

Proof. Let us first derive the dual Problem (7.62). The proof follows the technique of perturbed problems detailed in [26, Chapter 3].

Dual problem Let us write the (primal) Problem (7.59) as

$$\min_{(w, \beta_0, \beta_1) \in \mathcal{M}(\mathbb{R}) \times \mathbb{R}^2} \left(F(w, \beta_0, \beta_1) + G\left(\Lambda(w, \beta_0, \beta_1)\right) \right), \quad (7.65)$$

where $F(w, \beta_0, \beta_1) \triangleq \|w\|_{\mathcal{M}}$, $G(\mathbf{c}) \triangleq \iota_{\{\mathbf{y}_0\}}(\mathbf{c})$ for all $\mathbf{c} \in \mathbb{R}^M$, and $\Lambda(w, \beta_0, \beta_1) \triangleq \boldsymbol{\nu}_{\mathcal{M}}(w) + \beta_0 \mathbf{1} + \beta_1 \mathbf{x}$.

The functions F and G are convex, lower semi-continuous and not identically equal to $\pm\infty$. By [26, Equation (4.18)], the dual problem of (7.65) is thus given by $\sup_{\mathbf{c} \in \mathbb{R}^M} (-F^*(\Lambda^*(\mathbf{c})) - G^*(-\mathbf{c}))$, where F^* and G^* are the Fenchel conjugates of F and G respectively, and $\Lambda^* : \mathbb{R}^M \rightarrow \mathcal{C}_0(\mathbb{R}) \times \mathbb{R}^2$ is the adjoint operator of Λ . One can verify that for all $\mathbf{c} \in \mathbb{R}^M$, $G^*(\mathbf{c}) = \langle \mathbf{c}, \mathbf{y}_0 \rangle$, for all $\eta \in \mathcal{C}_0(\mathbb{R})$ and $\beta_0, \beta_1 \in \mathbb{R}$, $F^*(\eta, \beta_0, \beta_1) = \iota_{\|\cdot\|_{\infty} \leq 1}(\eta) + \iota_{\{(0,0)\}}(\beta_0, \beta_1)$ (with $\iota_{\|\cdot\|_{\infty} \leq 1}$ the indicator function of the closed unit ball in $\mathcal{C}_0(\mathbb{R})$ for the supremum norm), and for all $\mathbf{c} \in \mathbb{R}^M$, $\Lambda^*(\mathbf{c}) = (\boldsymbol{\nu}_{\mathcal{M}}^*(\mathbf{c}), \langle \mathbf{c}, \mathbf{1} \rangle, \langle \mathbf{c}, \mathbf{x} \rangle)$. Therefore, the dual problem can be rewritten as

$$- \inf_{\mathbf{c} \in \mathbb{R}^M} \left(\iota_{\mathcal{C}}(\mathbf{c}) + \langle -\mathbf{c}, \mathbf{y}_0 \rangle \right), \quad (7.66)$$

where $\mathcal{C} \subset \mathbb{R}^M$ is the convex set defined in (7.62). Problem (7.66) is clearly the same as Problem (7.62), which proves the first statement of the Lemma 7.5.

Strong duality To prove strong duality between Problems (7.59) and (7.62) (*i.e.*, that they have the same optimal value), we start by showing strong duality between

$$\inf_{\mathbf{c} \in \mathbb{R}^M} \left(\iota_{\mathcal{C}}(\mathbf{c}) + \langle -\mathbf{c}, \mathbf{y}_0 \rangle \right), \quad (7.67)$$

and its dual problem. We then conclude by observing that the optimal value of the dual problem of (7.67) is equal to the optimal value of Problem (7.59) up to a sign. Indeed, this last statement proves that both Problems (7.59) and (7.62) have the same optimal value since Problem (7.67) is, up to a sign, the dual Problem (7.62) (which rewrites as in (7.66)).

We first start by proving that strong duality holds between Problem (7.67) and its dual problem. The aim is to apply [26, Proposition 2.3, Chapter 3]. With the notations of [26], let us denote the map $\Phi : \mathbb{R}^M \times \mathcal{C}_0(\mathbb{R}) \rightarrow \mathbb{R} \cup \{+\infty\}$ as

$$\Phi(\mathbf{c}, \eta) \triangleq \langle -\mathbf{c}, \mathbf{y}_0 \rangle + \iota_{\{(0,0)\}} \left(\langle \mathbf{c}, \mathbf{1} \rangle, \langle \mathbf{c}, \mathbf{x} \rangle \right) + \iota_{\|\cdot\|_{\infty} \leq 1} \left(\boldsymbol{\nu}_{\mathcal{M}}^*(\mathbf{c}) - \eta \right) \quad (7.68)$$

for all $(\mathbf{c}, \eta) \in \mathbb{R}^M \times \mathcal{C}_0(\mathbb{R})$. This map Φ defines a perturbed version of Problem (7.67), since by definition, for all $\mathbf{c} \in \mathbb{R}^M$,

$$\Phi(\mathbf{c}, 0) = \iota_{\mathcal{C}}(\mathbf{c}) + \langle -\mathbf{c}, \mathbf{y}_0 \rangle \quad (7.69)$$

is the objective function of Problem (7.67). We now verify that the assumptions of [26, Proposition 2.3] are satisfied for Φ and Problem (7.67):

- Φ is convex,
- the optimal value of Problem (7.67) is finite due to the weak duality (primal-dual inequality given below) between Problems (7.65) and (7.66), which yields

$$\begin{aligned} -\infty &< - \inf_{\mathbf{c} \in \mathbb{R}^M} \left(\iota_{\mathcal{C}}(\mathbf{c}) + \langle -\mathbf{c}, \mathbf{y}_0 \rangle \right) \\ &\leq \inf_{(w, \beta_0, \beta_1) \in \mathcal{M}(\mathbb{R}) \times \mathbb{R}^2} \left(\|w\|_{\mathcal{M}} + \iota_{\{\mathbf{y}_0\}} \left(\boldsymbol{\nu}_{\mathcal{M}}(w) + \beta_0 \mathbf{1} + \beta_1 \mathbf{x} \right) \right) \\ &< +\infty, \end{aligned} \quad (7.70)$$

- the map $\eta \in \mathcal{C}_0(\mathbb{R}) \mapsto \Phi(\mathbf{0}, \eta) = \iota_{\|\cdot\|_{\infty} \leq 1}(-\eta)$ is finite and continuous at $\eta = 0 \in \mathcal{C}_0(\mathbb{R})$.

Therefore, we deduce that strong duality holds between Problem (7.67) and its dual

problem given by

$$\sup_{w \in \mathcal{M}(\mathbb{R})} \left(-\Phi^*(\mathbf{0}, w) \right), \quad (7.71)$$

and that this last optimization problem has at least one solution. By writing the map Φ as $\Phi(\mathbf{c}, \eta) = \tilde{F}(\mathbf{c}) + \tilde{G}(\tilde{\Lambda}(\mathbf{c}) - \eta)$ with $\tilde{F}(\mathbf{c}) \triangleq \langle -\mathbf{c}, \mathbf{y}_0 \rangle + \iota_{V^\perp}(\mathbf{c})$, $V \triangleq \text{span}(\mathbf{1}, \mathbf{x}) \subset \mathbb{R}^M$, $\tilde{G} \triangleq \iota_{\|\cdot\|_\infty \leq 1}(\cdot)$, and $\tilde{\Lambda} = \nu_{\mathcal{M}}^*$, we get that $\Phi^*(\mathbf{c}, w) = \tilde{F}^*(\tilde{\Lambda}^*(w) + \mathbf{c}) + \tilde{G}^*(-w)$ for any $(\mathbf{c}, w) \in \mathbb{R}^M \times \mathcal{M}(\mathbb{R})$, and thus that Problem (7.71) becomes

$$- \min_{w \in \mathcal{M}(\mathbb{R})} \left(\iota_V(\nu_{\mathcal{M}}(w) + \mathbf{y}_0) + \|w\|_{\mathcal{M}} \right). \quad (7.72)$$

We now verify that the optimal value of

$$\min_{w \in \mathcal{M}(\mathbb{R})} \left(\iota_V(\nu_{\mathcal{M}}(w) + \mathbf{y}_0) + \|w\|_{\mathcal{M}} \right), \quad (7.73)$$

i.e., minus the optimal value of the dual problem of (7.67), is equal to the optimal value of Problem (7.59)

$$\min_{(w, \beta_0, \beta_1) \in \mathcal{M}(\mathbb{R}) \times \mathbb{R}^2} \left(\iota_{\{\mathbf{y}_0\}}(\nu_{\mathcal{M}}(w) + \beta_0 \mathbf{1} + \beta_1 \mathbf{x}) + \|w\|_{\mathcal{M}} \right). \quad (7.74)$$

Let $w \in \mathcal{M}(\mathbb{R})$ be a solution of Problem (7.73), which we know to exist by [26, Proposition 2.3]. Since the objective function of Problem (7.73) is finite at w , we obtain that $\nu_{\mathcal{M}}(w) + \mathbf{y}_0 \in V$, *i.e.*, there exists $(\beta_0, \beta_1) \in \mathbb{R}^2$ such that $\mathbf{y}_0 = \nu_{\mathcal{M}}(-w) + \beta_0 \mathbf{1} + \beta_1 \mathbf{x}$. Assume by contradiction that there exist $(\tilde{w}, \tilde{\beta}_0, \tilde{\beta}_1) \in \mathcal{M}(\mathbb{R}) \times \mathbb{R}^2$ that achieve a lower cost than (w, β_0, β_1) in (7.59), *i.e.*,

$$\iota_{\{\mathbf{y}_0\}}(\nu_{\mathcal{M}}(-w) + \beta_0 \mathbf{1} + \beta_1 \mathbf{x}) + \|-w\|_{\mathcal{M}} > \iota_{\{\mathbf{y}_0\}}(\nu_{\mathcal{M}}(\tilde{w}) + \tilde{\beta}_0 \mathbf{1} + \tilde{\beta}_1 \mathbf{x}) + \|\tilde{w}\|_{\mathcal{M}}. \quad (7.75)$$

Since the left term of this inequality is finite, we must have $\mathbf{y}_0 = \nu_{\mathcal{M}}(\tilde{w}) + \tilde{\beta}_0 \mathbf{1} + \tilde{\beta}_1 \mathbf{x}$ and

$$\|w\|_{\mathcal{M}} > \|\tilde{w}\|_{\mathcal{M}}. \quad (7.76)$$

Since $\nu_{\mathcal{M}}(-\tilde{w}) + \mathbf{y}_0 = \tilde{\beta}_0 \mathbf{1} + \tilde{\beta}_1 \mathbf{x} \in V$, we deduce using (7.76) that $-\tilde{w}$ achieves a lower cost than w for Problem (7.73), which contradicts the assumption on w . Hence, for all $w \in \mathcal{M}(\mathbb{R})$, $(\beta_0, \beta_1) \in \mathbb{R}^2$, we have

$$\iota_{\{\mathbf{y}_0\}}(\nu_{\mathcal{M}}(-w) + \beta_0 \mathbf{1} + \beta_1 \mathbf{x}) + \|-w\|_{\mathcal{M}} \leq \iota_{\{\mathbf{y}_0\}}(\nu_{\mathcal{M}}(w) + \beta_0 \mathbf{1} + \beta_1 \mathbf{x}) + \|w\|_{\mathcal{M}}, \quad (7.77)$$

i.e., $(-w, \beta_0, \beta_1) \in \mathcal{M}(\mathbb{R}) \times \mathbb{R}^2$ is a solution of Problem (7.59). Therefore, we get that

the optimal values of Problems (7.73) and (7.59) are equal since

$$\iota_V(\nu_{\mathcal{M}}(w) + \mathbf{y}_0) + \|w\|_{\mathcal{M}} = \iota_{\{\mathbf{y}_0\}}(\nu_{\mathcal{M}}(-w) + \beta_0 \mathbf{1} + \beta_1 \mathbf{x}) + \|-w\|_{\mathcal{M}}. \quad (7.78)$$

Optimality conditions To derive the optimality conditions given in (7.63) and (7.64), we apply [26, Proposition 2.4, Chapter 3]. To this end, we must first prove that the assumptions of this proposition hold true, which requires three intermediate results. Firstly, we have already proved that strong duality holds, and that the primal Problem (7.59) has at least one solution. Secondly, we require the following lemma.

Lemma 7.6. *The set $\mathcal{D} \triangleq \{\mathbf{c} \in \mathbb{R}^M : \|\nu_{\mathcal{M}}^*(\mathbf{c})\|_{\infty} \leq 1\}$ is a compact set.*

Proof. Consider the map $F : \mathbb{R}^M \rightarrow \mathcal{F}$ given by

$$\forall \mathbf{c} \in \mathbb{R}^M, \quad F(\mathbf{c}) \triangleq \sum_{m=1}^M c_m g_0(x_m, \cdot) = \nu_{\mathcal{M}}^*(\mathbf{c}) \quad (7.79)$$

(using (7.61) for the last equality), where $\mathcal{F} \triangleq \text{span}(\{g_0(x_m, \cdot) : 1 \leq m \leq M\})$. The family $(g_0(x_m, \cdot))_{1 \leq m \leq M}$ is linearly independent due to the fact that each $g_0(x_m, \cdot)$ is a piecewise-linear splines with finitely many knots, and so there exists a nonempty interval I in which all the $g_0(x_m, \cdot)$ are linear functions. The function F is therefore a bijective linear map between finite-dimensional spaces. Hence, F^{-1} is continuous when $\mathcal{F} \subset \mathcal{C}_0(\mathbb{R})$ is endowed with the supremum norm $\|\cdot\|_{\infty}$ due the continuity of $g_0(x_m, \cdot)$ for that norm. Moreover, we have that $\mathcal{E} \triangleq \{f \in \mathcal{F} : \|f\|_{\infty} \leq 1\}$ is bounded and closed, and is thus compact (since $\mathcal{F} = \text{Im}(\nu_{\mathcal{M}}^*)$ is finite dimensional). This proves that $\mathcal{D} = F^{-1}(\mathcal{E})$ is compact. \square

Thirdly, the two following statements hold true:

- the objective function of Problem (7.62) is a continuous linear form over the convex set \mathcal{C} ,
- the convex set $\mathcal{C} = V^{\perp} \cap \mathcal{D} \subset \mathbb{R}^M$ is compact as the intersection of the closed set V^{\perp} and the compact set $\mathcal{D} \triangleq \{\mathbf{c} \in \mathbb{R}^M : \|\nu_{\mathcal{M}}^*(\mathbf{c})\|_{\infty} \leq 1\}$.

Hence, the convexity and the compactness of \mathcal{C} imply that there is at least one extreme point of \mathcal{C} that is a solution of the dual Problem (7.62).

Hence, we have proved that the assumptions of [26, Proposition 2.4, Chapter 3] are satisfied, which implies that any solution $(w, \beta_0, \beta_1) \in \mathcal{M}(\mathbb{R}) \times \mathbb{R}^2$ of (the primal) Problem (7.59)

and $\mathbf{c} \in \mathbb{R}^M$ of (the dual) Problem (7.62) are linked by the optimality conditions

$$\boldsymbol{\nu}_{\mathcal{M}}(w) + \beta_0 \mathbf{1} + \beta_1 \mathbf{x} = \mathbf{y}_0, \quad (7.80)$$

$$\langle \mathbf{c}, \mathbf{1} \rangle = \langle \mathbf{c}, \mathbf{x} \rangle = 0, \quad \|w\|_{\mathcal{M}} = \langle w, \boldsymbol{\nu}_{\mathcal{M}}^*(\mathbf{c}) \rangle \quad \text{and} \quad \|\boldsymbol{\nu}_{\mathcal{M}}^*(\mathbf{c})\|_{\infty} \leq 1. \quad (7.81)$$

Conversely, if any $(w, \beta_0, \beta_1) \in \mathcal{M}(\mathbb{R}) \times \mathbb{R}^2$ and $\mathbf{c} \in \mathbb{R}^M$ satisfy the optimality conditions given above, then by [26, Proposition 2.4, Chapter 3] we obtain that $(w, \beta_0, \beta_1) \in \mathcal{M}(\mathbb{R}) \times \mathbb{R}^2$ and $\mathbf{c} \in \mathbb{R}^M$ are solutions of the primal and dual problems, respectively. This proves the last statement of the Lemma 7.5. \square

The last intermediate result needed for the proof of Proposition 7.1 is given in the next lemma, where we prove that any continuous function $\boldsymbol{\nu}_{\mathcal{M}}^*(\mathbf{c}) \in \mathcal{C}_0(\mathbb{R})$ with $\mathbf{c} \in \mathbb{R}^M$ satisfying the orthogonality conditions given in (7.64) is a piecewise-linear spline whose knots are located at the sampling points $\mathbf{x} = (x_m)_{1 \leq m \leq M}$.

Lemma 7.7. *Let $\mathbf{c} \in \mathbb{R}^M$ such that $\langle \mathbf{c}, \mathbf{1} \rangle = \langle \mathbf{c}, \mathbf{x} \rangle = 0$. Then, we have $\boldsymbol{\nu}_{\mathcal{M}}^*(\mathbf{c}) = \sum_{m=1}^M c_m(x_m - \cdot)_+$.*

Proof. We know by (7.61) and (7.3) that

$$\begin{aligned} \boldsymbol{\nu}_{\mathcal{M}}^*(\mathbf{c}) &= \left\langle \mathbf{c}, (g_0(x_m, \cdot))_{1 \leq m \leq M} \right\rangle, \\ &= \left\langle \mathbf{c}, \left((x_m - \cdot)_+ - (-\cdot)_+ + x_m \left((-\cdot)_+ - (1 - \cdot)_+ \right) \right)_{1 \leq m \leq M} \right\rangle, \\ &= \left\langle \mathbf{c}, ((x_m - \cdot)_+)_{1 \leq m \leq M} \right\rangle - (-\cdot)_+ \underbrace{\langle \mathbf{c}, \mathbf{1} \rangle}_{=0} + \left((-\cdot)_+ - (1 - \cdot)_+ \right) \underbrace{\langle \mathbf{c}, \mathbf{x} \rangle}_{=0}, \end{aligned} \quad (7.82)$$

which proves that $\boldsymbol{\nu}_{\mathcal{M}}^*(\mathbf{c}) = \sum_{m=1}^M c_m(x_m - \cdot)_+$. \square

We now have all the necessary elements to prove Proposition 7.1.

Proof of Proposition 7.1. Suppose that $f_{\text{opt}} \in \mathcal{M}_{\text{D}^2}(\mathbb{R})$ is a solution of (g-BPC). Then, f_{opt} satisfies the interpolation conditions $f_{\text{opt}}(x_m) = y_{0,m}$ for all $m \in \{1, \dots, M\}$, and $(w, \beta_0, \beta_1) \in \mathcal{M}(\mathbb{R}) \times \mathbb{R}^2$ is a solution of Problem (7.59), where $f_{\text{opt}} = \text{D}_0^{-2}\{w\} + \beta_0 + \beta_1(\cdot)$. By Lemma 7.5, there exists a solution $\mathbf{c} \in \mathbb{R}^M$ of Problem (7.62) which satisfies $\langle \mathbf{c}, \mathbf{1} \rangle = \langle \mathbf{c}, \mathbf{x} \rangle = 0$ with $\|\boldsymbol{\nu}_{\mathcal{M}}^*(\mathbf{c})\|_{\infty} \leq 1$. Let us denote $\eta \triangleq \boldsymbol{\nu}_{\mathcal{M}}^*(\mathbf{c}) \in \mathcal{C}_0(\mathbb{R})$. By Lemma 7.7, we have $\eta = \sum_{m=1}^M c_m(x_m - \cdot)_+$ i.e., η is a dual pre-certificate (Definition 7.1). Moreover, again by Lemma 7.5, we know that $\|w\|_{\mathcal{M}} = \langle w, \eta \rangle$ which gives the direct implication in Proposition 7.1.

For the reverse implication, the dual pre-certificate η given by the statement satisfies $\eta = \nu_{\mathcal{M}}^*(\mathbf{c})$ by Lemma 7.7, and since f_{opt} satisfies the interpolation conditions, we deduce that $\nu_{\mathcal{M}}(w) + \beta_0 \mathbf{1} + \beta_1 \mathbf{x} = \mathbf{y}_0$ where β_0 and β_1 are defined via the relation $f_{\text{opt}} = D_0^{-2}\{w\} + \beta_0 + \beta_1(\cdot)$. Hence, by Lemma 7.5, $(w, \beta_0, \beta_1) \in \mathcal{M}(\mathbb{R}) \times \mathbb{R}^2$ is a solution of Problem (7.59) (and \mathbf{c} is a solution of Problem (7.62)), *i.e.*, f_{opt} is a solution of (g-BPC).

Let us now prove that the relation $\|w\|_{\mathcal{M}} = \langle w, \eta \rangle$ is equivalent to $\text{supp}_{\pm}(w) \subset \text{sat}_{\pm}(\eta)$ when η is a dual pre-certificate (see Definition 7.2 for the definition of the signed support and signed saturation set). First, we have that $\|w\|_{\mathcal{M}} = \|w|_{\text{sat}_+(\eta)}\|_{\mathcal{M}} + \|w|_{\text{sat}_-(\eta)}\|_{\mathcal{M}} + \|w|_{S^c}\|_{\mathcal{M}}$ (see [99, Theorem 6.2]), where $S \triangleq \text{sat}_+(\eta) \cup \text{sat}_-(\eta)$, hence

$$\left(\|w|_{\text{sat}_+(\eta)}\|_{\mathcal{M}} - \langle w|_{\text{sat}_+(\eta)}, \eta \rangle \right) + \left(\|w|_{\text{sat}_-(\eta)}\|_{\mathcal{M}} - \langle w|_{\text{sat}_-(\eta)}, \eta \rangle \right) \quad (7.83)$$

$$+ \left(\|w|_{S^c}\|_{\mathcal{M}} - \langle w|_{S^c}, \eta \rangle \right) = 0. \quad (7.84)$$

Each of the three terms in the sum is nonnegative by definition of $\|\cdot\|_{\mathcal{M}}$ and the fact that $\|\eta\|_{\infty} \leq 1$, so that the equality $\|w\|_{\mathcal{M}} = \langle w, \eta \rangle$ is equivalent to

$$\|w|_{\text{sat}_+(\eta)}\|_{\mathcal{M}} = \langle w|_{\text{sat}_+(\eta)}, \eta \rangle, \quad (7.85)$$

$$\|w|_{\text{sat}_-(\eta)}\|_{\mathcal{M}} = \langle w|_{\text{sat}_-(\eta)}, \eta \rangle, \quad (7.86)$$

$$\|w|_{S^c}\|_{\mathcal{M}} = \langle w|_{S^c}, \eta \rangle. \quad (7.87)$$

Consider the Jordan decomposition $w = w_+ - w_-$ of w . Then, $\|w|_{\text{sat}_+(\eta)}\|_{\mathcal{M}} = w_+(\text{sat}_+(\eta)) + w_-(\text{sat}_+(\eta))$ and $\langle w|_{\text{sat}_+(\eta)}, \eta \rangle = \int_{\text{sat}_+(\eta)} dw = w_+(\text{sat}_+(\eta)) - w_-(\text{sat}_+(\eta))$, so that (7.85) is equivalent to $w_-(\text{sat}_+(\eta)) = 0$ *i.e.*,

$$\text{supp}(w_-) \cap \text{sat}_+(\eta) = \emptyset. \quad (7.88)$$

Similarly, we can prove that (7.86) is equivalent to

$$\text{supp}(w_+) \cap \text{sat}_-(\eta) = \emptyset, \quad (7.89)$$

since $\langle w|_{\text{sat}_-(\eta)}, \eta \rangle = - \int_{\text{sat}_-(\eta)} dw$. As a result, to obtain the desired equivalence, it remains to prove that (7.87) is the same as $w|_{S^c} = 0$. The arguments can be found for example in [47] (see the proof of Lemma A.1), but we reproduce the reasoning here for the sake of completeness. Consider the closed sets for all $k > 0$

$$\Omega_k \triangleq \mathbb{R} \setminus \left(S + \left(-\frac{1}{k}, \frac{1}{k} \right) \right) \subset S^c. \quad (7.90)$$

Suppose by contradiction that there exists $k > 0$ such that $\|w|_{\Omega_k}\|_{\mathcal{M}} > 0$. Since $|\eta| < 1$ on the closed set Ω_k (because it is true on the bigger open set S^c), we deduce that

$\langle w|_{\Omega_k}, \eta \rangle < \|w|_{\Omega_k}\|_{\mathcal{M}}$ and then

$$\|w\|_{\mathcal{M}} = \langle w|_{\Omega_k}, \eta \rangle + \langle w|_{\Omega_k^c}, \eta \rangle < \|w|_{\Omega_k}\|_{\mathcal{M}} + \|w|_{\Omega_k^c}\|_{\mathcal{M}} = \|w\|_{\mathcal{M}}, \quad (7.91)$$

which is a contradiction. Hence, we have $\|w|_{\Omega_k}\|_{\mathcal{M}} = 0$ for all $k > 0$, which yields $\|w|_{S^c}\|_{\mathcal{M}} = 0$ since $S^c = \cup_{k>0} \Omega_k$, *i.e.*, $w|_{S^c} = 0$. \square

8 Interpolation of Data Under Lipschitz Constraint

This chapter is based on the following publication [201]:

S. Aziznejad*, T. Debarre*, and M. Unser, “Sparsest Univariate Learning Models Under Lipschitz Constraint”, *IEEE Open Journal of Signal Processing*, vol. 3, pp. 140–154, Mar. 2022.

8.1 Introduction

The goal of a regression model is to learn a mapping $f : \mathcal{X} \rightarrow \mathcal{Y}$ from a collection of data points $(x_m, y_m) \in \mathcal{X} \times \mathcal{Y}$ with $m = 1, \dots, M$ such that $y_m \approx f(x_m)$, while avoiding the problem of overfitting [43, 202, 203]. Here, \mathcal{X} denotes the input domain and \mathcal{Y} is the set of possible outcomes. A common way of carrying out this task is to solve a minimization problem of the form

$$\min_{f \in \mathcal{F}} \left(\sum_{m=1}^M E(f(x_m), y_m) + \mathcal{R}(f) \right), \quad (8.1)$$

where \mathcal{F} is the underlying search space, the convex loss function $E : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}^+$ enforces the consistency of the learned mapping with the given data points, and the regularization functional $\mathcal{R} : \mathcal{F} \rightarrow \mathbb{R}_{\geq 0}$ injects prior knowledge on the form of the mapping f , which is designed to alleviate the problem of overfitting.

8.1.1 Nonparametric Regression

In some cases, the optimization can be performed over an infinite-dimensional function space—see for example [204–206] for applications in signal and image processing. A prominent example is the family of reproducing-kernel Hilbert spaces (RKHS) $\mathcal{F} = \mathcal{H}(\mathbb{R}^d)$,

* Equal contribution.

$\mathcal{X} = \mathbb{R}^d$, $\mathcal{Y} = \mathbb{R}$ [207, 208], in which the regression problem is formulated as

$$\min_{f \in \mathcal{H}(\mathbb{R}^d)} \left(\sum_{m=1}^M E(f(\mathbf{x}_m), y_m) + \lambda \|f\|_{\mathcal{H}}^2 \right). \quad (8.2)$$

The fundamental result in RKHS theory is the kernel representer theorem [42, 209], which states that the unique solution of (8.2) admits the kernel expansion

$$f(\cdot) = \sum_{m=1}^M a_m k(\cdot, \mathbf{x}_m), \quad (8.3)$$

where $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is the unique reproducing kernel of $\mathcal{H}(\mathbb{R}^d)$ and $a_m \in \mathbb{R}$ with $m = 1, \dots, M$ are learnable parameters. The expansion (8.3) allows one to recast the infinite-dimensional problem (8.2) as a finite-dimensional one and to use standard computational tools of convex optimization to solve it. Many classical kernel-based schemes follow this approach, including support-vector machines and radial basis functions [210–212].

8.1.2 Parametric Regression

In cases when Problem (8.1) cannot be recast as a finite-dimensional optimization problem, another common approach is to restrict \mathcal{F} to a subspace that admits a parametric representation. This approach is used in deep neural networks (DNNs), which have become a prominent tool in machine learning and data science in recent years [181, 213]. They outperform classical kernel-based methods for various image-processing tasks. In particular, they have become state-of-the-art for image classification [214], inverse problems [215], and image segmentation [216]. However, most published works are empirical, and the outstanding performance of DNNs is yet to be fully understood. To this end, many recent works are directed towards studying DNNs from a theoretical perspective. Unsurprisingly, *stability* and *interpretability*, which are key principles in machine learning, play a central role in these works. For example, the stability of state-of-the-art deep-learning-based methods has been dramatically challenged in image classification [217, 218] and image reconstruction [219]. Attempts have also been made to understand and interpret DNNs from different perspectives, such as rate-distortion theory [220, 221]. However, the community is still far from reaching a global understanding and these questions remain active areas of research.

8.1.3 Our Contributions

In this chapter, we introduce two variational formulations for regressing one-dimensional data that favor stable and simple regression models. Similarly to RKHS theory, the latter are nonparametric continuous-domain problems in the sense that \mathcal{F} in (8.1) is

an infinite-dimensional function space. Inspired by the stability principle, we focus on the development of regression schemes with controlled Lipschitz regularity. This is motivated by the observation that many analyses in deep learning require assumptions on the Lipschitz constant of the learned mapping [222–224]. Likewise, in the context of so-called “plug-and-play” methods—*i.e.*, when a trainable module is inserted into an iterative-reconstruction framework—the rate of convergence of the overall scheme often depends on the Lipschitz constant of this module [225–230].

In our first formulation, we use the Lipschitz constant of the learned mapping as a regularization term. Specifically, we consider the minimization problem

$$\min_{f \in \text{Lip}(\mathbb{R})} \left(\sum_{m=1}^M E(f(x_m), y_m) + \lambda L(f) \right), \quad (8.4)$$

where $\text{Lip}(\mathbb{R})$ is the space of Lipschitz-continuous real functions and $L(f)$ denotes the Lipschitz constant of $f \in \text{Lip}(\mathbb{R})$. In this formulation, one can implicitly control the Lipschitz regularity of the learned function by varying the regularization parameter λ . Our first contribution is the proof of a representer theorem that characterizes the solution set of (8.4). In particular, we prove that the global minimum is achieved by a continuous and piecewise-linear (CPWL) mapping. Next, motivated by our quest for simplicity, we find the mapping with the minimal number of linear regions. Note that many previous works study problems similar to (8.4) in more general settings, typically using the Lipschitz constant of the n th derivative $\mathcal{R}(f) = L(f^{(n)})$ with $n \geq 0$ as the regularization term [88, 178, 231–234]. More recently, [235] studied the classification problem over metric spaces and derived a parametric form of a solution of this problem. Our work complements this interesting line of research by providing an in-depth analysis of the $n = 0$ case which is related to second-order total-variation (TV) minimization, and by focusing more on computational aspects of (8.4). More precisely, we propose a two-step algorithm to reach the sparsest CPWL solution of (8.4). The first step consists in solving a discrete problem with ℓ_∞ regularization, and the second is a sparsification step using Algorithm 7.1 that reaches the sparsest solution.

In the second scenario, we explicitly control the Lipschitz constant of the learned mapping by imposing a hard constraint. Inspired by the theoretical insights of the first problem, we add a second-order TV regularization term that is known to promote sparse CPWL functions, as demonstrated in the previous chapter and in [103]. This leads to the minimization problem

$$\min_{f \in \mathcal{M}_{\mathbb{D}^2}(\mathbb{R})} \left(\sum_{m=1}^M E(f(x_m), y_m) + \lambda \|D^2 f\|_{\mathcal{M}} \right) \quad \text{s.t.} \quad L(f) \leq \bar{L}, \quad (8.5)$$

where $\|\cdot\|_{\mathcal{M}}$ is the TV norm for measures, $\mathcal{M}_{\mathbb{D}^2}(\mathbb{R})$ is the space of functions with bounded second-order TV (see Chapter 2), and \bar{L} is the user-defined upper bound for the desired

Lipschitz regularity of the learned mapping. An interesting feature of (8.5) is that the simplicity and stability of the learned mapping can be adjusted by tuning the parameters $\lambda > 0$ and $\bar{L} > 0$, respectively. In this case as well, we prove a representer theorem that guarantees the existence of CPWL solutions. We also propose a two-step algorithm, similar to that of the first scenario, to find the sparsest CPWL solution. The main difference is the first step, where the discrete problem has a ℓ_1 regularization term and a ℓ_∞ constraint.

8.1.4 Connection to Neural Networks

Another major motivation for this work is to further elucidate the tight connection between CPWL functions and neural networks. It is well known that the input-output mapping of any feed-forward DNN with linear spline activations is a CPWL function [182]. Prominent examples are rectified linear unit (ReLU) activation functions, which are used in many state-of-the-art DNN architectures. This CPWL property is due to the fact that these mappings are compositions of affine transformations and pointwise activations. Hence, since the activation functions are CPWL and the CPWL property of functions is preserved by composition, the full input-output mapping is CPWL. Conversely, any CPWL function can be represented *exactly* by a DNN with linear-spline activations [236]. This establishes a direct link with spline theory, as first highlighted by Poggio *et al.* [187] and then further explored in various works [103, 183, 194, 195, 237].

When it comes to shallow networks, the connection with our framework becomes even more explicit. It is well known in the literature that the standard training (*i.e.*, with weight decay) of a two-layer univariate ReLU network is equivalent to solving a TV-based variational problem such as (8.5) without the Lipschitz constraint [193, 194]. Specifically, the weight-decay penalty can be shown to be equal to the second-order TV of the input-output mapping of the full network at the optimum [194, Proposition 18]. As we demonstrate, these results can be readily extended to prove the equivalence between the training of a Lipschitz-constrained two-layer univariate ReLU network and our formulation (8.5). Our description of the solution set of Problem (8.5) thus provides insights into the training of Lipschitz-aware neural networks.

8.1.5 Outline

The chapter is organized as follows: we review the required mathematical background in Section 8.2. In Section 8.3, we introduce our supervised-learning formulations and we state their corresponding representer theorems. We then propose algorithms for finding the corresponding sparsest CPWL solution in Section 8.4. Finally, we provide numerical illustrations and discussions in Section 8.5.

8.2 Mathematical Preliminaries

8.2.1 Banach Spaces

A Banach space is a normed topological vector space that is complete in its norm topology. The prototypical examples of Banach spaces are the Lebesgue spaces $L_p(\mathbb{R})$ for $p \in [1, +\infty]$, which are the spaces of Lebesgue-measurable functions with finite L_p norm. For $p \neq +\infty$, this reads as

$$L_p(\mathbb{R}) \triangleq \{f : \mathbb{R} \rightarrow \mathbb{R} \text{ measurable: } \|f\|_{L_p} < +\infty\}, \quad (8.6)$$

where $\|f\|_{L_p} \triangleq \left(\int_{\mathbb{R}} |f(x)|^p dx\right)^{\frac{1}{p}}$. Alternatively, one can define $L_p(\mathbb{R}) = \overline{(\mathcal{S}(\mathbb{R}), \|\cdot\|_{L_p})}$ as the completion of the Schwartz space $\mathcal{S}(\mathbb{R})$ (defined in (2.5)) with respect to the L_p norm for $p \in [1, +\infty)$. The case $p = +\infty$ is special: indeed, the L_∞ norm is defined as $\|f\|_{L_\infty} = \text{ess sup}_{x \in \mathbb{R}} |f(x)|$, where the essential supremum extracts an upper bound that is valid almost everywhere. Moreover, contrary to the other $L_p(\mathbb{R})$ spaces (with $1 \leq p < +\infty$), the space $\mathcal{S}(\mathbb{R})$ is not dense in $L_\infty(\mathbb{R})$. In fact, the completion of $\mathcal{S}(\mathbb{R})$ with respect to the L_∞ norm is the space $\mathcal{C}_0(\mathbb{R})$ of continuous functions that vanish at infinity [15].

Finally, following Section 2.4, we introduce the Banach space $\mathcal{M}_{D^2}(\mathbb{R})$ (see Theorem 2.2) of functions with finite second-order TV norm (see (2.17)):

$$\|D^2\{f\}\|_{\mathcal{M}} \triangleq \sup_{\varphi \in \mathcal{S}(\mathbb{R}): \|\varphi\|_{L_\infty} = 1} \langle D^2 f, \varphi \rangle. \quad (8.7)$$

8.2.2 Lipschitz Constant

We denote by $\text{Lip}(\mathbb{R})$ the space of Lipschitz-continuous functions $f : \mathbb{R} \rightarrow \mathbb{R}$ with a finite Lipschitz constant

$$L(f) \triangleq \sup_{x_1 \neq x_2} \frac{|f(x_1) - f(x_2)|}{|x_1 - x_2|} < +\infty. \quad (8.8)$$

Following Rademacher's theorem, any Lipschitz-continuous function $f \in \text{Lip}(\mathbb{R})$ is differentiable almost everywhere with a measurable and essentially bounded derivative. The Lipschitz constant of the function then corresponds to the essential supremum of its derivative, so that

$$L(f) = \|D\{f\}\|_{L_\infty} = \text{ess sup}_{x \in \mathbb{R}} |f'(x)|. \quad (8.9)$$

Conversely, any distribution $f \in \mathcal{S}'(\mathbb{R})$ whose weak derivative lies in $L_\infty(\mathbb{R})$ is a Lipschitz-continuous function [238, Theorem 1.36]. In other words, we have that

$$\text{Lip}(\mathbb{R}) \triangleq \{f \in \mathcal{S}'(\mathbb{R}) : D\{f\} \in L_\infty(\mathbb{R})\}, \quad (8.10)$$

which allows us to view $\text{Lip}(\mathbb{R})$ as the native Banach space associated to the pair $(L_\infty(\mathbb{R}), D)$ in the sense of [91].

8.3 Lipschitz-Aware Formulations for Supervised Learning

A way of indirectly controlling the Lipschitz constant of a function is to use a second-order TV-type regularizer within a variational formulation. Indeed, the two seminorms are connected, as demonstrated in Theorem 8.1.

Theorem 8.1. *Any function with second-order bounded variation is Lipschitz continuous. Moreover, for any $f \in \mathcal{M}_{D^2}(\mathbb{R})$, we have the upper bound*

$$L(f) \leq \|D^2 f\|_{\mathcal{M}} + \ell(f) \quad (8.11)$$

for the Lipschitz constant of f , where

$$\ell(f) \triangleq \inf_{x_1 \neq x_2} \frac{|f(x_1) - f(x_2)|}{|x_1 - x_2|} \geq 0. \quad (8.12)$$

Finally, (8.11) holds with equality if and only if f is monotone and convex/concave.

Proof. For any $h > 0$ and $\mathbf{p} = (p_1, p_2) \in \mathbb{R}^2$ with $p_1 < p_2$, we first define the test function $\varphi_h(\cdot; \mathbf{p}) \in \mathcal{C}_0(\mathbb{R})$ as

$$\varphi_h(x; \mathbf{p}) \triangleq h^{-1} \left((x - (p_1 - h))_+ - (x - p_1)_+ + (x - (p_2 + h))_+ - (x - p_2)_+ \right).$$

This function will be used on several occasions throughout the proof. In particular, we use the explicit form of its second-order derivative, given by

$$D^2 \varphi_h(\cdot; \mathbf{p}) = h^{-1} \left(\delta(\cdot - (p_1 - h)) - \delta(\cdot - p_1) + \delta(\cdot - (p_2 + h)) - \delta(\cdot - p_2) \right). \quad (8.13)$$

Upper Bound: Similarly to (8.9), we have that $\ell(f) = \text{ess inf}_{x \in \mathbb{R}} |f'(x)|$. For a fixed $\epsilon > 0$, by definition of essential extrema, there exist $\bar{x}, \underline{x} \in \mathbb{R}$ at which f is differentiable with $|f'(\bar{x})| \geq (L(f) - \epsilon)$ and $|f'(\underline{x})| \leq (\ell(f) + \epsilon)$. Without loss of generality, we assume that $\bar{x} < \underline{x}$. Following the limit definition of the derivative, we then consider a small

radius $h > 0$ such that

$$\begin{aligned} \left| \frac{f(\bar{x} + h) - f(\bar{x})}{h} \right| &\geq |f'(\bar{x})| - \epsilon \geq L(f) - 2\epsilon, \\ \left| \frac{f(\underline{x} + h) - f(\underline{x})}{h} \right| &\leq |f'(\underline{x})| + \epsilon \leq \ell(f) + 2\epsilon. \end{aligned}$$

Next, we consider the test function $\varphi = \varphi_h(\cdot; (\bar{x} + h, \underline{x}))$, which satisfies $\|\varphi\|_\infty = 1$. Following the definition of the TV norm (8.7), we deduce that $\|\mathbf{D}^2 f\|_{\mathcal{M}} \geq |\langle \mathbf{D}^2 f, \varphi \rangle| = |\langle f, \mathbf{D}^2 \varphi \rangle|$, where the last equality follows from the self-adjointness of the second-order derivative. Using (8.13), we thus have that

$$\begin{aligned} \|\mathbf{D}^2 f\|_{\mathcal{M}} &\geq h^{-1} \left| f(\bar{x}) - f(\bar{x} + h) + f(\underline{x} + h) - f(\underline{x}) \right| \\ &\geq \frac{|f(\bar{x} + h) - f(\bar{x})|}{h} - \frac{|f(\underline{x} + h) - f(\underline{x})|}{h} \\ &\geq L(f) - 2\epsilon - \ell(f) - 2\epsilon = L(f) - \ell(f) - 4\epsilon. \end{aligned}$$

Finally, by letting $\epsilon \rightarrow 0$, we deduce the desired upper bound.

Equality—Sufficient Conditions: Assume that $f \in \mathcal{M}_{\mathbb{D}^2}(\mathbb{R})$ is convex and increasing; we denote its second-order derivative by $w = \mathbf{D}^2 f$. Note that, in this case, the functions $(-f(\cdot))$, $f(-\cdot)$, and $(-f(-\cdot))$ are concave/decreasing, convex/decreasing, and concave/increasing, respectively. Hence, it only remains to prove the equality for f and the other cases immediately follow.

Let $\epsilon > 0$, and let $\mathcal{D}(\mathbb{R})$ be the space of infinitely smooth functions with compact support. Due to the denseness of $\mathcal{D}(\mathbb{R})$ in $\mathcal{S}(\mathbb{R})$ with respect to the supremum norm [15], by (8.7), there exists a test function $\psi \in \mathcal{S}(\mathbb{R})$ with compact support $K \triangleq \text{supp}(\psi)$ such that $\|\psi\|_{L_\infty} = 1$ and $\langle w, \psi \rangle \geq (\|\mathbf{D}^2 f\|_{\mathcal{M}} - \epsilon)$. For any $T > 0$, we consider the test function $\psi_T = \varphi_1(\cdot; (-T, T))$. From (8.13), we obtain that

$$\begin{aligned} \langle w, \psi_T \rangle &= \langle f, \mathbf{D}^2 \psi_T \rangle \\ &= (f(T+1) - f(T)) - (f(-T) - f(-T-1)) \\ &\leq L(f) - \ell(f), \end{aligned}$$

where we have used the increasing assumption to deduce that $f(T+1) \geq f(T)$ and $f(-T) \geq f(-T-1)$. By choosing T large enough so that $K \subseteq [-T, T]$, we ensure that $(\psi_T - \psi)$ is a nonnegative function, since for all $x \in K$, we will have that $\psi_T(x) = 1 = \|\psi\|_{L_\infty} \geq \psi(x)$. Next, the convexity of f implies that $w = \mathbf{D}^2 f$ is a positive measure. Hence,

$$0 \leq \langle w, \psi_T - \psi \rangle \leq L(f) - \ell(f) - \|\mathbf{D}^2 f\|_{\mathcal{M}} + \epsilon. \quad (8.14)$$

By letting $\epsilon \rightarrow 0$, we deduce that $\|\mathbf{D}^2 f\|_{\mathcal{M}} \leq (L(f) - \ell(f))$, which implies that (8.11)

holds with equality.

Equality—Necessary Conditions: Let $f \in \mathcal{M}_{D^2}(\mathbb{R})$ be a function for which (8.11) holds with equality. We now prove that f must be monotone and convex/concave.

Monotonicity: Assume by contradiction that f is not monotone. Hence, there exists $x_n \in \mathbb{R}$ such that $f'(x_n) < 0$. Indeed, if f' were a positive distribution, then for any $a, b \in \mathbb{R}$ with $a < b$, we would have that $(f(b) - f(a)) = \langle f', \mathbb{1}_{[a,b]} \rangle \geq 0$, which contradicts the assumption of nonmonotonicity. Similarly, there exists $x_p \in \mathbb{R}$ such that $f'(x_p) > 0$.

Next, consider a point $x_L \in \mathbb{R}$, distinct from x_n and x_p , such that $|f'(x_L)| > (L(f) - \epsilon) > 0$, where $0 < \epsilon < \frac{\min(-f'(x_n), f'(x_p))}{3}$ is a small constant. We assume without loss of generality that $f'(x_L) > 0$. There exists a small radius $h \in (0, \frac{|x_L - x_n|}{2})$ such that

$$\begin{aligned} \frac{f(x_n + h) - f(x_n)}{h} &\leq f'(x_n) + \epsilon < 0, \\ \frac{f(x_L + h) - f(x_L)}{h} &\geq f'(x_L) - \epsilon > 0. \end{aligned}$$

By considering the test function

$$\varphi \triangleq \begin{cases} \varphi_h(\cdot; (x_n + h, x_L)) & \text{if } x_n < x_L \\ \varphi_h(\cdot; (x_L, x_n + h)) & \text{if } x_n > x_L \end{cases}$$

and using (8.7) once again, we deduce that

$$\begin{aligned} \|D^2 f\|_{\mathcal{M}} &\geq h^{-1} |f(x_n) - f(x_n + h) + f(x_L + h) - f(x_L)| \\ &= \frac{f(x_L + h) - f(x_L)}{h} - \frac{f(x_n + h) - f(x_n)}{h} \\ &\geq f'(x_L) - \epsilon - f'(x_n) - \epsilon \geq L(f) - f'(x_n) - 3\epsilon \\ &> L(f), \end{aligned}$$

which contradicts the original assumption that (8.11) holds with equality. For the case $f'(x_L) < 0$, the same arguments can be applied to x_p instead of x_n . This proves that f is monotone. In the following, we consider the case where f is an increasing function; the decreasing case can be deduced by symmetry.

Convexity/Concavity: We first consider the canonical decomposition $f = D_0^{-2}w + p$, where $w \triangleq D^2 f$, D_0^{-2} is a right inverse of the second-order derivative whose Schwartz kernel is given in (7.3), and $p(x) = ax + b$ is an affine term (see Theorem 2.1). We then use the Jordan decomposition of $w = D^2 f$ as $w = (w_+ - w_-)$, where $w_+, w_- \in \mathcal{M}(\mathbb{R})$ are positive measures such that $\|w\|_{\mathcal{M}} = \|w_+\|_{\mathcal{M}} + \|w_-\|_{\mathcal{M}}$. This allows us to form the decomposition $f = (f_+ - f_-)$, where $f_s = D_0^{-2}w_s + p_s$, $s \in \{+, -\}$, $p_+(x) = (A + a)x + b$, and $p_-(x) = Ax$ with $A > 0$ being a sufficiently large constant so that the functions f_+

and f_- are both convex and strictly increasing. Hence, they both satisfy the sufficient conditions for equality, which implies that $\|\mathbf{D}^2 f_s\|_{\mathcal{M}} = (L(f_s) - \ell(f_s))$ for $s \in \{+, -\}$.

Assume by contradiction that $w_s \neq 0$ for $s \in \{+, -\}$ and let $\epsilon < \frac{\min(\|\mathbf{D}^2 f_+\|_{\mathcal{M}}, \|\mathbf{D}^2 f_-\|_{\mathcal{M}})}{2}$ be a small positive constant. Let $\bar{x}, \underline{x} \in \mathbb{R}$ such that $f'(\bar{x}) \geq (L(f) - \epsilon)$ and $f'(\underline{x}) \leq (\ell(f) + \epsilon)$. Using these inequalities, we deduce that

$$\begin{aligned} \|\mathbf{D}^2 f\|_{\mathcal{M}} &= L(f) - \ell(f) \\ &\leq f'(\bar{x}) - f'(\underline{x}) + 2\epsilon \\ &= (f'_+(\bar{x}) - f'_-(\bar{x})) - (f'_+(\underline{x}) - f'_-(\underline{x})) + 2\epsilon \\ &= A_+ - A_- + 2\epsilon, \end{aligned} \tag{8.15}$$

where $A_s = (f'_s(\bar{x}) - f'_s(\underline{x}))$ for $s \in \{+, -\}$. We consider the two following exhaustive cases.

Case I: $\bar{x} > \underline{x}$. The convexity of f_- implies that $A_- \geq 0$. Moreover, we have that $A_+ = (f'_+(\bar{x}) - f'_+(\underline{x})) \leq (L(f_+) - \ell(f_+)) = \|\mathbf{D}^2 f_+\|_{\mathcal{M}}$. Using (8.15), this yields that $\|\mathbf{D}^2 f\|_{\mathcal{M}} \leq \|\mathbf{D}^2 f_+\|_{\mathcal{M}} + 2\epsilon$, which can be rewritten as $2\epsilon \geq \|\mathbf{D}^2 f_-\|_{\mathcal{M}}$. However, our original choice of ϵ implies that $\epsilon < \|\mathbf{D}^2 f_-\|_{\mathcal{M}}/2$, which leads to a contradiction.

Case II: $\bar{x} \leq \underline{x}$. Similarly to the first case, we deduce that $A_+ \leq 0$ and $-A_- \leq \|\mathbf{D}^2 f_-\|_{\mathcal{M}}$. Hence, we get that $2\epsilon \geq \|\mathbf{D}^2 f_+\|_{\mathcal{M}}$, which contradicts the assumption $\epsilon < \|\mathbf{D}^2 f_+\|_{\mathcal{M}}/2$. Since both cases lead to a contradiction, we have $w_- = 0$ or $w_+ = 0$. This implies that f is either convex or concave. \square

A weaker version of Theorem 8.1 is proven in [185], where $\ell(f)$ is replaced with the clear upper bound $|f(1) - f(0)|$. The importance of the updated bound is due to its sharpness, in the sense that it is an equality for monotone and convex/concave functions.

A weaker version of (8.11) motivated the authors of [185] to provide a global bound for the Lipschitz constant of DNNs and to regularize it during training. Although this is an interesting approach to control the Lipschitz constant of the learned mapping, the obtained guarantee is too conservative. This is due to the fact that, as soon as f has some oscillations, the difference between the two sides of (8.11) dramatically increases and the bound becomes loose. Here, by contrast, we ensure the global stability of the learned mapping by directly controlling the Lipschitz constant itself.

8.3.1 Lipschitz Regularization

We first consider the Lipschitz constant as a regularizer and study the minimization problem

$$\mathcal{V}_{\text{Lip}} \triangleq \arg \min_{f \in \text{Lip}(\mathbb{R})} \left(\sum_{m=1}^M E(f(x_m), y_m) + \lambda L(f) \right), \quad (8.16)$$

where $E : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ is a strictly convex and coercive function and $\lambda > 0$ is the regularization parameter. We also assume, without loss of generality, that the data points x_m are sorted in the increasing order $x_1 < x_2 < \dots < x_M$. In Theorem 8.2, we state our main theoretical contributions regarding Problem (8.16).

Theorem 8.2. *Regarding the minimization Problem (8.16), the following statements hold true.*

1. *The solution set \mathcal{V}_{Lip} is a nonempty, convex and weak*-compact subset of $\text{Lip}(\mathbb{R})$.*
2. *There exists a unique vector $\mathbf{z} = (z_m) \in \mathbb{R}^M$ such that*

$$\mathcal{V}_{\text{Lip}} = \arg \min_{f \in \text{Lip}(\mathbb{R})} L(f) \quad \text{s.t.} \quad f(x_m) = z_m \quad 1 \leq m \leq M. \quad (8.17)$$

3. *The optimal Lipschitz constant has the closed-form expression*

$$L_{\min} \triangleq \max_{2 \leq m \leq M} \left| \frac{z_m - z_{m-1}}{x_m - x_{m-1}} \right|. \quad (8.18)$$

Consequently, any L_{\min} -Lipschitz function f that satisfies $f(x_m) = z_m$ for $m = 1, \dots, M$ is a solution of (8.16).

4. *Let $\mathcal{E} \subseteq \mathbb{R}^2$ be the union of the graphs of all solutions of (8.16), defined as*

$$\mathcal{E} \triangleq \{(x, y) \in \mathbb{R}^2 : \exists f \in \mathcal{V}_{\text{Lip}}, y = f(x)\}. \quad (8.19)$$

We also define the right and left planar cones $\mathcal{R}, \mathcal{L} \subseteq \mathbb{R}^2$ as

$$\mathcal{R} \triangleq \{\alpha_1(1, L_{\min}) + \alpha_2(1, -L_{\min}) : \alpha_1, \alpha_2 \geq 0\}, \quad (8.20)$$

and $\mathcal{L} \triangleq -\mathcal{R}$. With the convention that $\mathcal{R}_0 \triangleq \mathcal{L}_{M+1} \triangleq \mathbb{R}^2$, we have that

$$\mathcal{E} = \bigcup_{m=1}^{M+1} (\mathcal{R}_{m-1} \cap \mathcal{L}_m), \quad (8.21)$$

where the \mathcal{R}_m and \mathcal{L}_m are shifted versions of \mathcal{R} and \mathcal{L} , with

$$\mathcal{R}_m \triangleq (x_m, z_m) + \mathcal{R} \quad \mathcal{L}_m \triangleq (x_m, z_m) + \mathcal{L} \quad \forall 1 \leq m \leq M. \quad (8.22)$$

5. Any solution of the constrained minimization problem

$$\min_{f \in \mathcal{M}_{D^2}(\mathbb{R})} \|D^2 f\|_{\mathcal{M}} \quad \text{s.t.} \quad f(x_m) = z_m \quad 1 \leq m \leq M \quad (8.23)$$

is included in \mathcal{V}_{Lip} . In particular, the solution set of (8.16) always includes a CPWL function.

The proof of Theorem 8.2 is given in Appendix B.3. Items 1 and 2 are classical results that hold for a general class of variational problems (see [112] for a generic result). Their practical implication is Item 3, which provides a way of identifying solutions of (8.16). The properties of the solution set \mathcal{V}_{Lip} are further explored in Item 4, where a geometrical insight is given (see Figure 8.1).

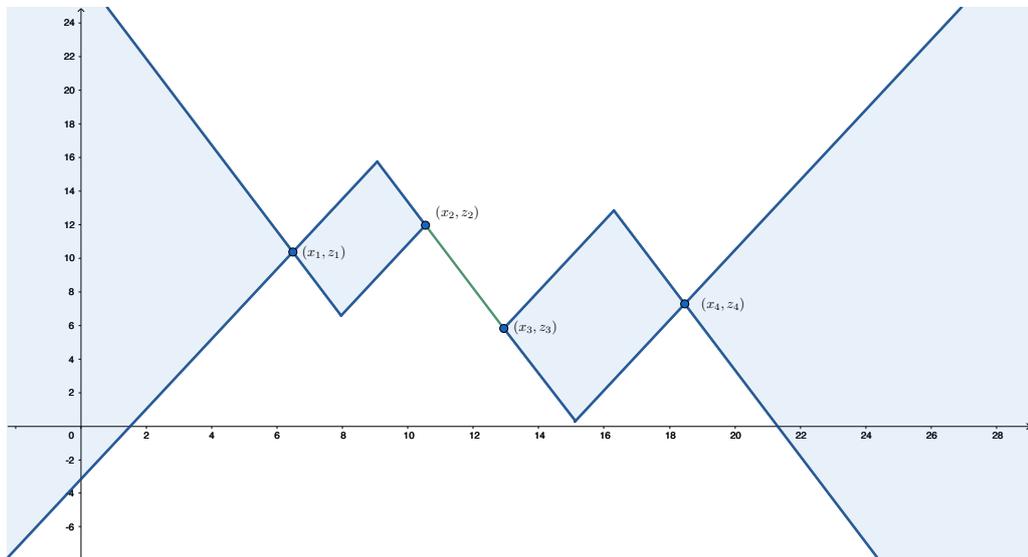


Figure 8.1 – The union of the graphs of all solutions in a simple example with four data points. Note that all solutions must directly connect (x_2, z_2) to (x_3, z_3) , since the slope of this segment is L_{\min} whose formula is given in (8.18).

Finally, the result that has the greatest practical relevance is stated in Item 5, which creates an interesting link with second-order TV minimization problems and hence guarantees the existence of CPWL solutions.

8.3.2 Lipschitz Constraint

While the first formulation is interesting on its own right and results in learning CPWL mappings with tunable Lipschitz constants, it does not necessarily yield a sparse (and, hence, interpretable) solution. In fact, the learned mapping can have undesirable oscil-

lations as illustrated in Figure 8.4. This observation motivates us to propose a second formulation that combines second-order TV regularization with a constraint over the Lipschitz constant, as expressed by

$$\mathcal{V}_{\text{hyb}} \triangleq \arg \min_{f \in \mathcal{M}_{\mathbb{D}^2}(\mathbb{R})} \left(\sum_{m=1}^M \mathbb{E}(f(x_m), y_m) + \lambda \|D^2 f\|_{\mathcal{M}} \right) \quad \text{s.t.} \quad L(f) \leq \bar{L}. \quad (8.24)$$

The quantity \bar{L} is the upper bound for the Lipschitz constant of the learned mapping. In this way, the stability is directly controlled by the user, while the regularization term removes undesired oscillations. The strength of the regularization can be adapted with the tunable parameter $\lambda > 0$. The solution set \mathcal{V}_{hyb} is characterized in Theorem 8.3, from which we also deduce the existence of CPWL solutions.

Theorem 8.3. *The solution set \mathcal{V}_{hyb} of Problem (8.24) is a nonempty, convex, and weak*-compact subset of $\mathcal{M}_{\mathbb{D}^2}(\mathbb{R})$ whose extreme points are linear splines with at most $(M - 1)$ linear regions. Moreover, there exists a unique vector $\mathbf{z} = (z_m)$ such that*

$$\mathcal{V}_{\text{hyb}} = \arg \min_{f \in \mathcal{M}_{\mathbb{D}^2}(\mathbb{R})} \|D^2 f\|_{\mathcal{M}} \quad \text{s.t.} \quad f(x_m) = z_m \quad 1 \leq m \leq M. \quad (8.25)$$

Finally, the optimal second-order TV cost has the closed-form expression

$$\mathcal{J}_{\text{hyb}}^0 \triangleq \sum_{m=2}^{M-1} \left| \frac{z_m - z_{m-1}}{x_m - x_{m-1}} - \frac{z_m - z_{m+1}}{x_m - x_{m+1}} \right|. \quad (8.26)$$

The proof of Theorem 8.3, which is given in Appendix B.4, involves the weak*-closedness of the constraint box $L(f) \leq \bar{L}$, which is essential to prove existence. Once the existence of a minimizer is guaranteed, we can apply the results of the previous chapter for second-order TV minimization to deduce the remaining parts. We also remark that the Lipschitz constraint only affects the vector \mathbf{z} in (8.25), which imposes that its entries satisfy the inequalities

$$\left| \frac{z_m - z_{m-1}}{x_m - x_{m-1}} \right| \leq \bar{L} \quad m = 2, \dots, M. \quad (8.27)$$

8.3.3 Connection to Neural Networks

In this part, we show that our second formulation (8.24) is equivalent to training a two-layer neural network with weight decay and a Lipschitz constraint. We recall that a univariate ReLU network with two layers and skip connections is a mapping $f_{\boldsymbol{\theta}} : \mathbb{R} \rightarrow \mathbb{R}$

of the form

$$f_{\boldsymbol{\theta}}(x) = c_0 + c_1 x + \sum_{k=1}^K v_k (w_k x - b_k)_+, \quad (8.28)$$

where $c_1 \in \mathbb{R}$ is the weight of the skip connection, $K \in \mathbb{N}$ is the width of the network, $v_k, w_k \in \mathbb{R}, k = 1, \dots, K$ are the linear weights, and $b_k \in \mathbb{R}, k = 1, \dots, K$ and $c_0 \in \mathbb{R}$ are the bias terms of the first and second layers, respectively. These parameters are concatenated into a single vector $\boldsymbol{\theta} \triangleq (K, \mathbf{v}, \mathbf{w}, \mathbf{b}, \mathbf{c})$, and we denote by Θ the set of all possible parameter vectors $\boldsymbol{\theta}$. Thus, the training problem with Lipschitz constraint and weight decay is formulated as

$$\mathcal{V}_{\text{NN}} \triangleq \arg \min_{\boldsymbol{\theta} \in \Theta} \left(\sum_{m=1}^M \mathbb{E}(f_{\boldsymbol{\theta}}(x_m), y_m) + \lambda \mathcal{R}_{\text{NN}}(\boldsymbol{\theta}) \right) \quad \text{s.t.} \quad L(f_{\boldsymbol{\theta}}) \leq \bar{L}, \quad (8.29)$$

where $\mathcal{R}_{\text{NN}}(\boldsymbol{\theta}) \triangleq \sum_{k=1}^K \left(\frac{|v_k|^2 + |w_k|^2}{2} \right)$ is the regularization term corresponding to weight decay. In Proposition 8.1, we show the equivalence between this training problem and our Lipschitz-constrained formulation (8.24).

Proposition 8.1. *For any solution $\boldsymbol{\theta}^*$ of (8.29), $f_{\boldsymbol{\theta}^*}$ is a CPWL solution of (8.24). Moreover, any CPWL solution of (8.24) can be expressed as a two-layer ReLU network $f_{\boldsymbol{\theta}^*}$ with skip connections whose parameter vector is optimal in the sense of (8.29), i.e., $\boldsymbol{\theta}^* \in \mathcal{V}_{\text{NN}}$.*

Proof. We start by proving a useful lemma.

Lemma 8.1. *For any $\boldsymbol{\theta}^* = (K^*, \mathbf{v}^*, \mathbf{w}^*, \mathbf{b}^*, \mathbf{c}^*) \in \mathcal{V}_{\text{NN}}$, we have that $|v_k^*| = |w_k^*|$ for any $k = 1, \dots, K$.*

Proof. Let $\boldsymbol{\theta}^* = (K^*, \mathbf{v}^*, \mathbf{w}^*, \mathbf{b}^*, \mathbf{c}^*) \in \mathcal{V}_{\text{NN}}$ and $1 \leq k \leq K$. For any $\epsilon \in (-1, 1)$, we define a perturbed parameter vector $\boldsymbol{\theta}_{\epsilon} = (K^*, \mathbf{v}_{\epsilon}, \mathbf{w}_{\epsilon}, \mathbf{b}_{\epsilon}, \mathbf{c}^*)$, where for any $k' = 1, \dots, K$ we have that

$$v_{\epsilon, k'} \triangleq \begin{cases} v_{k'}^* & k' \neq k \\ (1 + \epsilon)^{\frac{1}{2}} v_k^* & k' = k \end{cases}, \quad (8.30)$$

$$w_{\epsilon, k'} \triangleq \begin{cases} w_{k'}^* & k' \neq k \\ (1 + \epsilon)^{-\frac{1}{2}} w_k^* & k' = k \end{cases}, \quad (8.31)$$

$$b_{\epsilon, k'} \triangleq \begin{cases} b_{k'}^* & k' \neq k \\ (1 + \epsilon)^{-\frac{1}{2}} b_k^* & k' = k \end{cases}. \quad (8.32)$$

Due to the positive homogeneity of the ReLU, one readily deduces from (8.28) that $f_{\boldsymbol{\theta}^*} = f_{\boldsymbol{\theta}_{\epsilon}}$ for any $\epsilon \in (-1, 1)$. Together with the optimality of $\boldsymbol{\theta}^*$ in Problem (8.29), this

implies that

$$v_k^{*2} + w_k^{*2} \leq (1 + \epsilon)v_k^{*2} + (1 + \epsilon)^{-1}w_k^{*2} \quad \forall \epsilon \in (-1, 1).$$

Multiplying both sides of the above inequality by $(1 + \epsilon) > 0$ yields

$$\epsilon w_k^{*2} \leq \epsilon(1 + \epsilon)v_k^{*2} \quad \forall \epsilon \in (-1, 1).$$

Letting $\epsilon \rightarrow 0^+$ yields $w_k^{*2} \leq v_k^{*2}$ and $\epsilon \rightarrow 0^-$ yields $w_k^{*2} \geq v_k^{*2}$, which proves that $|w_k^*| = |v_k^*|$. \square

Using Lemma 8.1, we observe that for any $\theta^* \in \mathcal{V}_{\text{NN}}$, we have that

$$\mathcal{R}_{\text{NN}}(\theta^*) = \frac{1}{2} \sum_{k=1}^K (v_k^{*2} + w_k^{*2}) = \sum_{k=1}^K |v_k^*| |w_k^*| = \|\mathbb{D}^2 f_{\theta^*}\|_{\mathcal{M}},$$

where the last inequality comes from the simple observation that $\|\mathbb{D}^2\{v(w(\cdot) - b)_+\}\| = |v||w|$ for any $v, w, b \in \mathbb{R}$. Hence, one can rewrite the solution set \mathcal{V}_{NN} as

$$\mathcal{V}_{\text{NN}} = \arg \min_{\theta \in \Theta_{\text{red}}} \left(\sum_{m=1}^M \mathbb{E}(f_{\theta}(x_m), y_m) + \lambda \|\mathbb{D}^2 f_{\theta}\|_{\mathcal{M}} \right) \quad \text{s.t.} \quad L(f_{\theta}) \leq \bar{L},$$

where $\Theta_{\text{red}} = \{\theta \in \Theta : \mathcal{R}_{\text{NN}}(\theta) = \|\mathbb{D}^2 f_{\theta}\|_{\mathcal{M}}\}$ is the reduced parameter space. To prove the announced equivalence, it remains to prove that the mapping $\Theta_{\text{red}} \rightarrow \mathcal{M}_{\mathbb{D}^2}(\mathbb{R}) : \theta \mapsto f_{\theta}$ is a bijection onto the CPWL members of $\mathcal{M}_{\mathbb{D}^2}(\mathbb{R})$ with finitely many linear regions.

For any $\theta \in \Theta_{\text{red}}$, the function f_{θ} is a CPWL member of $\mathcal{M}_{\mathbb{D}^2}(\mathbb{R})$ with finitely many linear regions. To prove the converse, let $f \in \mathcal{M}_{\mathbb{D}^2}(\mathbb{R})$ be a CPWL function with finitely many linear regions. Using the canonical representation of f , there exist $c_0, c_1 \in \mathbb{R}$, $K \in \mathbb{N}$ and $a_k, \tau_k \in \mathbb{R}$ with $a_k \neq 0$ for $k = 1, \dots, K$ such that

$$f(x) = c_0 + c_1 x + \sum_{k=1}^K a_k (x - \tau_k)_+.$$

Next, by defining $v_k \triangleq \frac{a_k}{\sqrt{|a_k|}}$, $w_k \triangleq \sqrt{|a_k|}$ and $b_k \triangleq \sqrt{|a_k|}\tau_k$ for $k = 1, \dots, K$, the homogeneity of the ReLU function $(\cdot)_+$ yields $f = f_{\theta}$ with $\theta = (K, \mathbf{c}, \mathbf{v}, \mathbf{w}, \mathbf{b}) \in \Theta_{\text{red}}$, where the latter inclusion is due to the equalities $|v_k| = |w_k|$ for $k = 1, \dots, K$. \square

Proposition 8.1 is an extension of the results of [193, 194], where this equivalence is proved in the absence of a Lipschitz constraint. These works rely on a result that describes the energy propagation in the training of feed-forward neural networks with weight decay (*e.g.*, [193, Corollary C.2]), which can easily be extended to the Lipschitz-constrained case (Lemma 8.1). Proposition 8.1 provides a functional framework to study the training of

Lipschitz-aware neural networks, which is a nontrivial task. To this end, Proposition 8.1 allows us to deploy our sparsification Algorithm 7.1, as demonstrated in the following section.

8.4 Finding the Sparsest CPWL Solution

Using the theoretical results of Section 8.3, we propose an algorithm to find the sparsest CPWL solution of Problems (8.16) and (8.24). To that end, we first compute the vector \mathbf{z} of the value of the optimal function at the data points x_1, \dots, x_m . Using this vector, we then deploy our sparsification Algorithm 7.1, whose use in the present method is motivated by the following theorem.

Theorem 8.4. *Let $(x_m, z_m) \in \mathbb{R}^2$ with $m = 1, \dots, M$ be a collection of ordered data points with $x_1 < \dots < x_M$. Then, the output f_{sparse} of the sparsification Algorithm 7.1 is the sparsest linear-spline interpolator of the data points. In other words, f_{sparse} is the CPWL interpolator with the fewest number of linear regions.*

Proof. Let f^* be the output of Algorithm 7.1. It is thus a CPWL solution of Problem (8.17) with the minimum number of linear regions. We prove that *any* CPWL interpolant f of the data points $P_m = (x_m, z_m)$ with $m = 1, \dots, M$ —not necessarily a minimizer of $\|\mathbf{D}^2\{\cdot\}\|_{\mathcal{M}}$ —has at least as many linear regions as f^* . Our proof is based on induction over the number M of data points. The initialization $M = 2$ trivially holds, since f^* then has a single linear region—it is simply the line connecting the two data points. Next, let $M > 2$, and assume that Theorem 8.4 holds for $(M - 1)$ or less data points (the induction hypothesis). The canonical interpolant f_{cano} introduced in Definition B.1 can be expressed as

$$f_{\text{cano}}(x) = \alpha_1 x + \alpha_2 + \sum_{m=2}^{M-1} a_m (x - x_m)_+ \quad (8.33)$$

for some coefficients $\alpha_1, \alpha_2, a_m \in \mathbb{R}$. There are three possible scenarios:

1. all a_m s are positive (or negative);
2. at least one of them is zero;
3. there are two consecutive coefficients with opposite signs, so that $a_m a_{m+1} < 0$ for some m .

We analyze each case separately and use the induction hypothesis to deduce the desired result. In this proof, we refer to singularities of CPWL functions (*i.e.*, the boundary points between linear regions) as *knots*.

Case 1: In this case, Theorem 7.3 states that f^* has $K = (\lceil \frac{M}{2} \rceil - 1)$ knots. Assume by contradiction that there exists a CPWL interpolant f with fewer knots and consider the K disjoint intervals (x_{2k-1}, x_{2k+1}) for $1 \leq k \leq (\lceil \frac{M}{2} \rceil - 1) = K$. We deduce that there exists an interval (x_{2k-1}, x_{2k+1}) in which f has no knot. This in turn implies that the data points P_{2k-1} , P_{2k} , and P_{2k+1} are aligned so that $a_{2k} = 0$, which yields a contradiction.

Case 2: Let $m \in \{2, M-1\}$ be such that $a_m = 0$, and consider the collection of $m < M$ data points $(P_{m'})_{1 \leq m' \leq m}$. By the induction hypothesis, f^* interpolates these points with the minimal number K_1 of knots. The same applies to the collection of $(M-m+1) < M$ points $(P_{m'})_{m \leq m' \leq M}$ with K_2 knots. Let f be a CPWL interpolant of all the M data points with the minimal number of knots. By definition of the K_i , f must have at least K_1 knots in the interval (x_1, x_m) and K_2 knots in the interval (x_m, x_M) . Since these intervals are disjoint, f must have at least $(K_1 + K_2)$ knots in total. Yet, f^* has exactly $(K_1 + K_2)$ knots: indeed, f^* follows f_{cano} in the interval $[x_{m-1}, x_{m+1}]$, which has no knot at x_m since $a_m = 0$ (the points P_{m-1} , P_m , and P_{m+1} are aligned). This concludes that f^* has the minimum number of knots.

Case 3: Let $m \in \{2, M-2\}$ be such that $a_m a_{m+1} < 0$. Consider the collection of $(m+1) < M$ data points $(P_{m'})_{1 \leq m' \leq m+1}$; by the induction hypothesis, f^* interpolates them with the minimal number K_1 of knots. Similarly, f^* interpolates the $(M-m+1) < M$ points $(P_{m'})_{m \leq m' \leq M}$ with the minimal number K_2 of knots. We now state a useful lemma whose proof is given below.

Lemma 8.2. *Let $m \in \{2, \dots, M-2\}$ be such that $a_m a_{m+1} < 0$. Then, any CPWL interpolant f of the data points $(P_{m'})_{1 \leq m' \leq M}$ can be modified to become another CPWL interpolant \tilde{f} with as many (or fewer) knots such that \tilde{f} has no knot in the interval (x_m, x_{m+1}) .*

Proof. Let f be a CPWL interpolant of the data points $(P_{m'})_{1 \leq m' \leq M}$ with P knots. In what follows, we consider a CPWL function \tilde{f} that follows f outside this interval and (x_{m-1}, x_{m+2}) , and we modify it inside this interval in order to remove all knots in (x_m, x_{m+1}) without increasing the total number of knots.

We consider the case $a_m > 0$ and $a_{m+1} < 0$ without loss of generality. Let $s^- \triangleq f'(x_{m-1}^-)$ and $s^+ \triangleq f'(x_{m+2}^+)$ be the slopes of f before and after the interval of interest (x_{m-1}, x_{m+2}) , respectively, and let $s_{\text{cano}}^- \triangleq f'_{\text{cano}}(x_{m-1}^-)$ and $s_{\text{cano}}^+ \triangleq f'_{\text{cano}}(x_{m+2}^+)$ be those of f_{cano} . We also introduce the linear functions $f^-(x) \triangleq z_{m-1} + s^-(x - x_{m-1})$ and $f^+(x) \triangleq z_{m+2} + s^+(x - x_{m+2})$. They prolong f in a straight line after P_{m-1} and before P_{m+2} , respectively. We now distinguish cases based on s^- and s^+ .

Case I: $s^- \leq s_{\text{cano}}^-$ and $s^+ \leq s_{\text{cano}}^+$. Graphically, this corresponds to f lying in none of the gray regions in Figure 8.2. In this case, the line $(P_m P_{m+1})$ intersects the linear function f^- at some point $P^- \triangleq (x^-, z^-)$ where $x^- \in (x_{m-1}, x_m)$, and with f^+ at some

point $P^+ \triangleq (x^+, z^+)$ with $x^+ \in (x_{m+1}, x_{m+2})$. This is obvious graphically (see Figure 8.2 as an illustration for P^-), and is due to the fact that $a_m > 0$ and $a_{m+1} < 0$. Hence, by taking an \tilde{f} that connects the points P_{m-1} , P^- , P^+ , and P_{m+2} , then \tilde{f} has two knots in $[x_{m-1}, x_{m+2}]$ and its knots satisfy $x^-, x^+ \notin (x_m, x_{m+1})$. Since f clearly cannot have fewer than two knots in this interval, this proves the desired result.

Case II: $s^+ > s_{\text{cano}}^+$ and $s^- > s_{\text{cano}}^-$. In this case, f lies in both gray regions in Figure 8.2. To pass through P_m , f must have at least one knot in $[x_{m-1}, x_m]$; let $P^- \triangleq (x^-, z^-)$ be the first of those knots (with $x^- < x_m$). Similarly, to pass through P_{m+1} , f must have a knot in $(x_{m+1}, x_{m+2}]$; let $P^+ \triangleq (x^+, z^+)$ be the last of those knots (with $x^+ > x_{m+1}$). Then, f must pass through the points P^- , P_m , P_{m+1} , P^+ . Yet, the lines (P^-P_m) and $(P_{m+1}P^+)$ clearly cannot intersect in the interval $[x_m, x_{m+1}]$, which implies that at least two knots are needed in the interval (x^-, x^+) . We conclude that f must have at least four knots in the interval $[x_{m-1}, x_{m+2}]$. Hence, we take an \tilde{f} that simply connects the points P_{m-1} , P_m , P_{m+1} , and P_{m+2} and follows f elsewhere; the latter has four knots in $[x_{m-1}, x_{m+2}]$, which is no more than \tilde{f} and thus fulfills the requirements of the proof.

Case III: $s^+ > s_{\text{cano}}^+$ and $s^- \leq s_{\text{cano}}^-$. This case is illustrated in Figure 8.2: f is outside the gray region on the left, and inside the one on the right. With a similar argument as in Case II, f must have at least three knots in the interval $[x_{m-1}, x_{m+2}]$. The fact that $a_m > 0$ implies that the line (P_mP_{m+1}) intersects the linear function f^- at some point $P^- \triangleq (x^-, z^-)$ where $x^- \in (x_{m-1}, x_m)$. We then take an \tilde{f} that connects the points P_{m-1} , P^- , P_{m+1} , and P_{m+2} and follows f elsewhere. The interpolant \tilde{f} has three knots at x^- , x_{m+1} , and x_{m+2} in $[x_{m-1}, x_{m+2}]$ and thus satisfies the requirements of the proof.

Case IV: $s^+ \leq s_{\text{cano}}^+$ and $s^- > s_{\text{cano}}^-$. This is similar to Case III, and can be readily deduced by symmetry, thus completing the proof of Lemma 8.2.

□

Let f be a CPWL interpolant of all the M data points with the minimal number of knots. By Lemma 8.2, it can be modified to become another interpolant \tilde{f} with the same total number of knots and none in the interval (x_m, x_{m+1}) . By definition of the K_i , \tilde{f} must have at least K_1 knots in the interval (x_1, x_{m+1}) and K_2 knots in the interval (x_m, x_M) . Yet, \tilde{f} has no knot in the interval (x_m, x_{m+1}) , so it must have at least K_1 knots in $(x_1, x_m]$ and K_2 knots in $[x_{m+1}, x_M)$. Since these intervals are disjoint, \tilde{f} must have at least $(K_1 + K_2)$ knots in total. Yet, f^* follows f_{cano} in the interval $[x_{m-1}, x_{m+2}]$ and thus also has no knot in the interval (x_1, x_{m+1}) . Therefore, by the induction hypothesis, f^* has K_1 knots in $(x_1, x_m]$ and K_2 knots in $[x_{m+1}, x_M)$, for a total of $(K_1 + K_2)$ knots. Since this is no more than \tilde{f} , f^* has the minimal number of knots, which proves the induction. □

Theorem 8.4 is a strong enhancement of Theorem 7.3, where it is merely established that

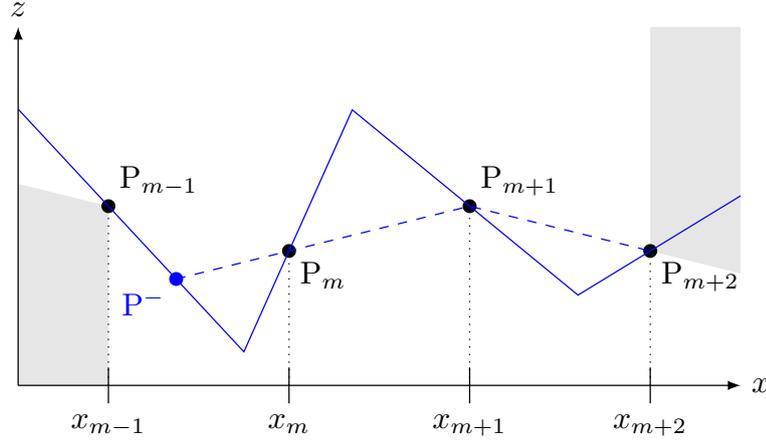


Figure 8.2 – Illustration of Lemma 8.2 in the case $a_m > 0$ and $a_{m+1} < 0$. The interpolant f (solid line) satisfies $s^+ > s_{\text{cano}}^+$ and $s^- \leq s_{\text{cano}}^-$. The modified interpolant \tilde{f} (dashed line) also has three knots P^- , P_{m+1} , and P_{m+2} , but none in (x_m, x_{m+1}) .

f_{sparse} is the sparsest CPWL solution of (8.23). In Theorem 8.4, we prove that f_{sparse} is in fact the sparsest of *all* CPWL interpolants of the data points (x_m, z_m) , without restricting the search to the solutions of (8.23). This is a remarkable result in its own right, as it gives a nontrivial answer to the seemingly simple question: how to interpolate data points with the minimum number of lines? Here, we apply Theorem 8.4 to deduce that, with the vector \mathbf{z} defined in Item 2 of Theorem 8.2, f_{sparse} is the sparsest CPWL solution of (8.17). Similarly, with the vector \mathbf{z} defined in Theorem 8.3, f_{sparse} is the sparsest CPWL solution of (8.24).

In the remaining part of this section, we detail our computation of the vectors \mathbf{z} defined in Theorems 8.2 and 8.3. We define the empirical loss function $F : \mathbb{R}^M \rightarrow \mathbb{R}_{\geq 0}$ as

$$F(\mathbf{z}) \triangleq \sum_{m=1}^M E(z_m, y_m). \quad (8.34)$$

For simplicity, we assume that F is differentiable; the prototypical example is the quadratic loss $F(\mathbf{z}) = \frac{1}{2} \sum_{m=1}^M (z_m - y_m)^2$. Following this notation and using (8.18), the vector \mathbf{z} in Problem (8.17) is the solution to the minimization problem

$$\min_{\mathbf{z} \in \mathbb{R}^M} \left(F(\mathbf{z}) + \lambda \|\mathbf{L}_{\text{inf}} \mathbf{z}\|_{\infty} \right), \quad (8.35)$$

where the matrix $\mathbf{L}_{\text{inf}} \in \mathbb{R}^{(M-1) \times M}$ is given by

$$[\mathbf{L}_{\text{inf}}]_{m,n} \triangleq \begin{cases} -v_{m+1} & n = m \\ v_{m+1} & n = m + 1, \\ 0 & \text{otherwise} \end{cases} \quad (8.36)$$

where $v_m \triangleq \frac{1}{x_m - x_{m-1}}$ for $m = 2, \dots, M$. To solve (8.35), we use the well-known ADMM algorithm [19] (see Appendix A).

Similarly and using (8.26), we formulate the search for the vector \mathbf{z} associated to the Problem (8.24) as

$$\min_{\mathbf{z} \in \mathbb{R}^M} \left(F(\mathbf{z}) + \lambda \|\mathbf{L}_1 \mathbf{z}\|_1 + i_{\|\mathbf{L}_{\text{inf}} \mathbf{z}\|_\infty \leq \bar{L}} \right), \quad (8.37)$$

where i_E denotes the indicator function of the set E and $\mathbf{L}_1 \in \mathbb{R}^{(M-2) \times M}$ with

$$[\mathbf{L}_1]_{m,n} \triangleq \begin{cases} -v_{m+1} & n = m \\ (v_{m+1} + v_{m+2}) & n = m + 1 \\ -v_{m+2} & n = m + 2 \\ 0 & \text{otherwise} \end{cases}, \quad (8.38)$$

for all $m = 1, \dots, M - 2$ and $n = 1, \dots, M$.

8.5 Experimental Results

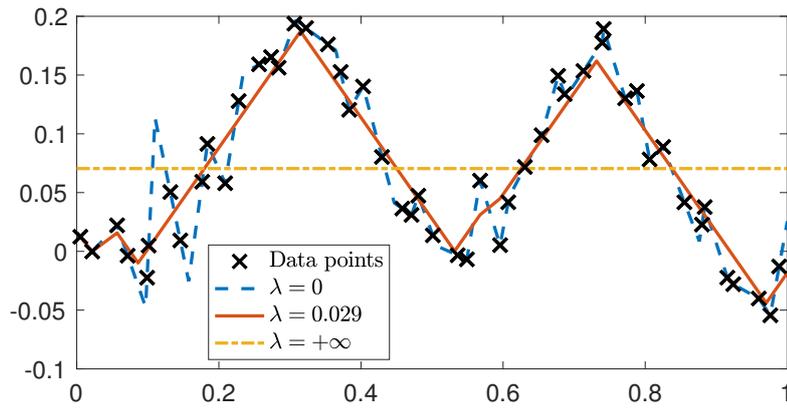
8.5.1 Experimental Setup

In all our experiments, we consider the standard quadratic loss $E(y, z) = \frac{1}{2}(y - z)^2$. We draw the data-point locations x_m randomly in the interval $[0, 1]$. The values y_m are then generated as $y_m = f_0(x_m) + n_m$, where f_0 is some known CPWL function (gold standard) and n_m is drawn i.i.d. from a zero-mean normal distribution with variance σ^2 .

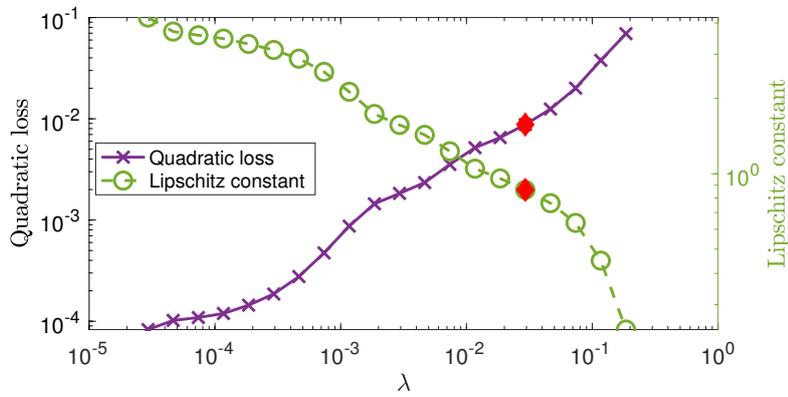
8.5.2 Example of Lipschitz Regularization

In this first experiment, we illustrate our first formulation (8.16). We take $M = 50$ data points, a CPWL ground-truth f_0 with 6 linear regions, and a noise level $\sigma = 0.02$.

The results are shown in Figure 8.3. In Figure 8.3a, we show the reconstructions for extreme values of λ . On one hand, $\lambda \rightarrow 0$ corresponds to the exact interpolation Problem (8.17). On the other hand, $\lambda = +\infty$ corresponds to constant regression. Obviously, neither is very satisfactory: interpolation leads to overfitting (the reconstruction has 37 linear regions), and the constant regression to underfitting. We show an example of a more satisfactory reconstruction for $\lambda = 0.029$ (10 linear regions), which is visually acceptable. In Figure 8.3b, we show the evolution of the quadratic loss $\frac{1}{2} \sum_{m=1}^M (f^*(x_m) - y_m)^2$ and the Lipschitz constant $L(f^*)$, for various values of λ . With the aid of such curves, the user can choose what is considered acceptable for either of these costs and select a suitable value of λ .



(a) Reconstructions for different values of λ . Number of linear regions: 10 for $\lambda = 0.029$ versus 37 for $\lambda = +\infty$.



(b) Evolution of the training error and the Lipschitz regularity with respect to λ . The diamond corresponds to $\lambda = 0.029$ (shown in Figure 8.3a).

Figure 8.3 – Example of our first formulation (8.16) for $M = 50$ data points.

8.5.3 Limitations of Lipschitz-Only Regularization

Despite its interesting theoretical properties, Problem (8.16) does not always yield satisfactory reconstructions. This is because it does not enforce a sparse reconstruction, despite the fact that our algorithm reconstructs (one of) the sparsest elements of \mathcal{V}_{lip} . This leads to learned mappings with too many linear regions and, consequently, poor interpretability.

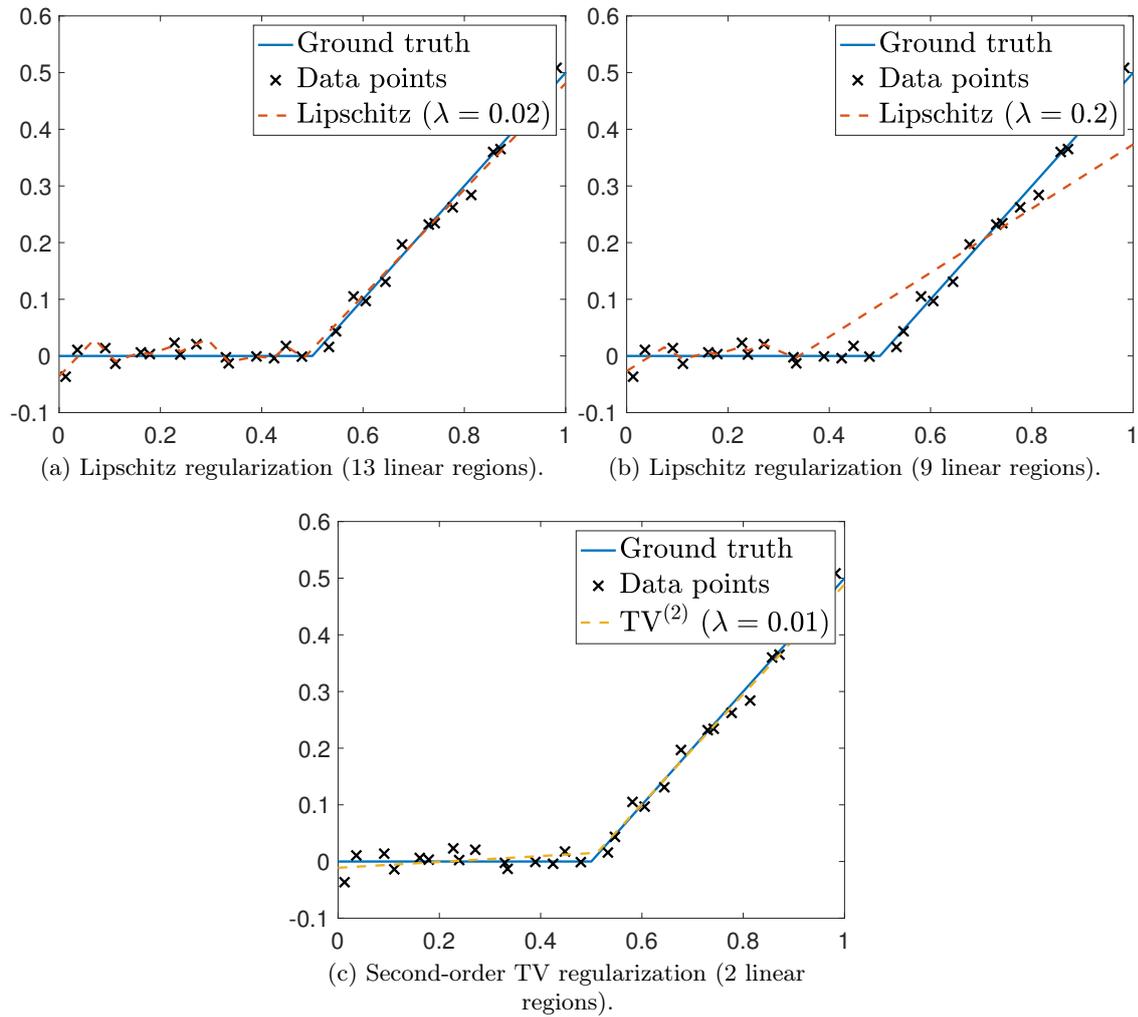
One such example is shown in Figure 8.4, where we consider the shifted ReLU function $f_0(\cdot) = (\cdot - \frac{1}{2})_+$ as the ground-truth mapping. We also fix the standard deviation of the noise to $\sigma = 0.02$. Figure 8.4a shows a reconstruction that solves Problem (8.16) with the regularization parameter $\lambda = 0.02$. Although the reconstruction is satisfactory in the active section ($x > 1/2$), it has many linear regions in the flat section ($x < 1/2$) that are not present in f_0 . This is due to the fact that the active section forces the Lipschitz constant of the reconstruction to be around 1, while oscillations with a slope smaller than 1 in the flat section are not penalized by the regularization. This problem clearly cannot be fixed by a simple increase in the regularization parameter: with $\lambda = 0.2$ (Figure 8.4b), not only there are still too many linear regions in the flat section (the reconstruction has 9 linear regions in total), but also the active section is poorly reconstructed because the Lipschitz constant is penalized too heavily by the regularization.

Hence, to reconstruct such a ground truth accurately, it is necessary to enforce the sparsity of the reconstruction, which is exactly the purpose of TV regularization. The reconstruction result of the second-order TV-regularized problem (*i.e.*, Problem (8.24) with a relatively large Lipschitz bound) with $\lambda = 0.01$ is also shown in Figure 8.4c; it is clearly much more satisfactory than any of the Lipschitz-penalized reconstructions since it is very close to the ground truth and has the same sparsity (two linear regions).

8.5.4 Robustness to Outliers of the Lipschitz-Constrained Formulation

In this final experiment, we demonstrate the pertinence of our second formulation (Problem (8.24)). More precisely, we examine the increased robustness to outliers of our second formulation (8.24) with respect to second-order TV regularization. To that end, we generate the CPWL ground truth f_0 with 6 linear regions and $M = 50$ data points. We then consider an additive Gaussian-noise model with low standard deviation $\sigma = 10^{-3}$ for 90% of the data, and a much stronger $\sigma' = 3.5 * 10^{-2}$ for the remaining 10%, which can be considered outliers.

We show in Figure 8.5 the reconstruction results using our second formulation with $\lambda = 10^{-4}$ and $\bar{L} = 0.66$. The latter is quite satisfactory despite the presence of a strong outlier around $x_m = 0.22$. This is due to the fact that the Lipschitz constant is constrained. When using second-order TV regularization alone with the same regularization parameter, the reconstruction is very similar in most regions but is much more sensitive to this outlier,

Figure 8.4 – Reconstructions with a ReLU ground truth and $M = 30$ data points.

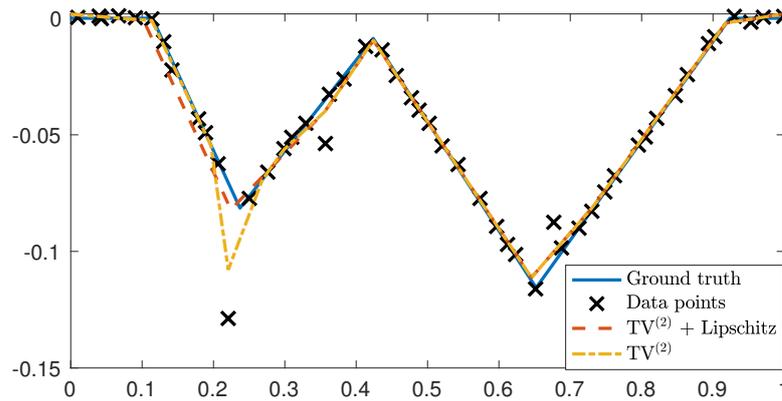


Figure 8.5 – Reconstruction of $M = 50$ data points for $\lambda = 10^{-4}$. Our second formulation with $\bar{L} = 0.66$ produces 9 linear regions. We compare it to that of second-order TV regularization, which produces 12 linear regions.

which leads to an unwanted sharp peak and to a high Lipschitz constant $L(f^*) = 2.21$. Moreover, our reconstruction is more satisfactory in terms of sparsity (9 linear regions compared to 12, which is closer to the 6 linear regions of the target function f_0).

8.6 Conclusion

We have proposed two schemes for the learning of one-dimensional CPWL mappings with tunable Lipschitz constant. In the first scheme, we directly use the Lipschitz constant as a regularization term. We establish a representer theorem that allows us to deduce the existence of a CPWL solution for this continuous-domain optimization problem. In the second scheme, we use the second-order TV seminorm as the regularization term to which we add a Lipschitz constraint. Again, we proved the existence of a CPWL solution for this problem. Finally, we proposed an efficient algorithm to find the sparsest CPWL solution of each problem. We illustrated the outcome of each scheme via numerical examples. A potential application of the proposed algorithm is to design stable CPWL activation functions with a minimum number of linear regions in DNNs. This can, for example, be useful to train a denoising module in the context of plug-and-play methods for image reconstruction, whose convergence rates typically depend on the Lipschitz constant of the trainable denoising module [239].

9 Reconstruction of Periodic Signals from Fourier-Domain Measurements

This chapter is based on the following work [240]:

T. Debarre, Q. Denoyelle, and J. Fageot, “TV-Based Spline Reconstruction with Fourier Measurements: Uniqueness and Convergence of Grid-Based Methods”, *arXiv preprint arXiv:2202.05059*, Feb. 2022.

We also recall a result of our companion work [241]:

T. Debarre, Q. Denoyelle, and J. Fageot, “On the Uniqueness of Solutions for the Basis Pursuit in the Continuum”, *arXiv preprint arXiv:2009.11855*, Feb. 2022.

9.1 Introduction

In this chapter, we study the reconstruction of an unknown periodic real function $f_0 : \mathbb{T} \rightarrow \mathbb{R}$ from the knowledge of its possibly noise-corrupted low-frequency Fourier series coefficients, where $\mathbb{T} \triangleq [0, 2\pi]$ is the one-dimensional torus. Let $K_c \geq 0$ be the cutoff frequency; we therefore have access to

$$\mathbf{y} = (y_0, y_1, \dots, y_{K_c}) \in \mathbb{R} \times \mathbb{C}^{K_c}, \quad (9.1)$$

such that $y_k \approx \widehat{f}_0[k]$, where $\widehat{f}_0[k]$ is the k th Fourier-series coefficient of f_0 . Note that, since f_0 is a real function, $y_0 \in \mathbb{R}$ is approximately the mean $\widehat{f}_0[0] = \langle f_0, 1 \rangle$ of f_0 , while $y_k \in \mathbb{C}$ for $k \neq 0$. Moreover, the Fourier series of f_0 is Hermitian symmetric, meaning that $\widehat{f}_0[-k] = \overline{\widehat{f}_0[k]} \in \mathbb{C}$ for every $k \in \mathbb{Z}$. The observation vector \mathbf{y} in (9.1) has $(2K_c + 1)$ (real) degrees of freedom: one for the real mean y_0 and two for each other complex Fourier series coefficients in \mathbb{C} .

9.1.1 Reconstruction via TV-Based Optimization

We choose to formulate the reconstruction task as an optimization problem with generalized total-variation (gTV) of the form (3.31). More precisely, the reconstruction f^* is the solution of

$$f^* \in \arg \min_{f \in \mathcal{M}_L(\mathbb{T})} \left(E(\boldsymbol{\nu}(f), \mathbf{y}) + \lambda \|Lf\|_{\mathcal{M}} \right), \quad (9.2)$$

where $\mathcal{M}_L(\mathbb{T})$ is the native space defined in (2.33), $\mathbf{y} \in \mathbb{R} \times \mathbb{C}^{K_c}$ is the observation vector, $\boldsymbol{\nu}(f) \in \mathbb{R} \times \mathbb{C}^{K_c}$ is the measurement vector

$$\boldsymbol{\nu}(f) = \left(\widehat{f}[0], \widehat{f}[1], \dots, \widehat{f}[K_c] \right) \in \mathbb{R} \times \mathbb{C}^{K_c}, \quad (9.3)$$

$E(\cdot, \mathbf{y}) : \mathbb{R} \times \mathbb{C}^{K_c} \rightarrow \mathbb{R}^+ \cup \{\infty\}$ is a data-fidelity cost functional (see Example 3.1 for classical examples), $\|\cdot\|_{\mathcal{M}}$ is the TV norm for periodic Radon measures defined in (2.17), and L is a regularization operator acting on periodic functions. As in the rest of this thesis, we focus on N_d th-order derivative operators $L = D^{N_d}$ with $N_d \geq 1$.

Remark 9.1. *Contrary to the hypotheses of Theorem 3.4 where we have $\boldsymbol{\nu} : \mathcal{M}_L(\mathbb{T}) \rightarrow \mathbb{R}^M$, the measurement operator $\boldsymbol{\nu}$ defined in (9.3) has complex output values. This allows us to use the Fourier series, which leads to more convenient calculations. However, this does not prevent us from applying Theorem 3.4 to Problem (9.2) by separating the real and imaginary parts into two separate real outputs as in Example 3.5, so that there is a total of $M = 2K_c + 1$ real outputs.*

9.1.2 Contributions

Problems of the form (9.2) have previously been studied in [22] in a more general setting; the representer theorem from that paper (Theorem 3.4) guarantees the existence of a solution without adjudicating on its uniqueness. Moreover, it gives the form of the extreme-point solution(s) as periodic L-splines, *i.e.*, functions f such that

$$Lf = \sum_{k=1}^K a_k \mathbb{I}\mathbb{I}(\cdot - x_k) \quad (9.4)$$

is a finite sum of shifted Dirac combs, the distinct Dirac locations x_k being the knots of the spline. Moreover, known proof techniques [14, 22, 87] allow us to show that the number of knots K is bounded by $K \leq 2K_c + 1$. Our contributions can be detailed as follows.

(i) *Uniqueness of the Solution.* Our main result is Theorem 9.1, in which we prove that the solution to Problem (9.2) is always unique. Moreover, we slightly improve the upper

bound on the number of knots to $K \leq 2K_c$ in (9.4). To the best of our knowledge, Theorem 9.1 is the first systematic uniqueness result for the analysis of TV-based variational problems such as (9.2). This result has both theoretical and algorithmic implications, which we leverage in our other contributions.

(ii) *Uniform Convergence of Grid-Based Methods.* We study the grid-based discretization of Problem (9.2). More precisely, we restrict its search space to the finite-dimensional space of uniform L-splines, *i.e.*, L-splines whose knots lie on a uniform grid. We show that as the grid gets finer, any sequence of solutions of the discretized problems converges in uniform norm towards the unique solution of Problem (9.2). This form of convergence is remarkably strong: in particular, it implies convergence for any L_p norm with $1 \leq p \leq \infty$.

(iii) *Grid-Based Algorithm.* We propose a periodic adaptation of the B-spline-based algorithm developed in Chapter 4 to solve Problem (9.2). Thanks to our aforementioned uniform convergence result, the reconstructed signal is guaranteed to be uniformly close to the gridless solution when the grid is sufficiently fine. We provide some experimental results of our algorithm on some simulated data that demonstrate its numerical feasibility.

In our setting, contrary to most gTV-based optimization problems, the solution to Problem (9.2) is always unique, as we show in Theorem 9.1. In the literature, the closest work in this direction is that presented in Chapter 7, where we provide a full description of the solution set of nonperiodic TV-based optimization problems with a regularization operator $L = D^2$ (which leads to piecewise-linear reconstructions), and spatial sampling measurements ν . That study includes the characterization of the cases of uniqueness (Proposition 7.6 and Theorem 7.1), which, contrary to Problem (9.2), is not systematic.

9.1.3 Related Works

The convergence of discretized optimization schemes to the solutions of continuous-domain TV-regularized problems has been studied in multiple works, such as [12, 51, 65, 78, 85]. Grid-based methods have specifically been considered in [65–67]. In these works, the authors prove convergence results in the weak* sense, which is adapted to the space of Radon measures, in a setting where no systematic uniqueness results are known. To the best of our knowledge, our work is the first to prove the convergence of solutions of discretized gTV-based problems towards the solution of the original problem, let alone in a strong sense such as the uniform norm. To achieve this, we leverage our uniqueness result of Theorem 9.1. On the algorithmic side, grid-based methods to solve optimization problems with TV-based regularization have been proposed in [45] and the works of Part II of this thesis.

9.1.4 Outline

The chapter is organized as follows. Section 9.2 introduces the necessary mathematical background. In Section 9.3, we present our optimization problem of interest (9.2) and prove that it always has a unique solution. In Section 9.4, we present the grid-based discretization of Problem (9.2), and prove that its solutions converge uniformly to that of the original problem when the grid size goes to zero. We present our proposed method for solving this discretized problem using a B-spline basis in Section 9.5. Finally, we exemplify our results on simulations in Section 9.6.

9.2 Mathematical Preliminaries

9.2.1 Periodic Functions and Periodic Splines

We first introduce some notations and recall some basic facts concerning the periodic setting introduced in Chapters 2 and 3, namely periodic functions, the generalized Fourier series (Section 2.1.6), and periodic splines (Section 3.1).

The Schwartz space of infinitely smooth periodic function (2.6) is denoted by $\mathcal{S}(\mathbb{T})$, and its dual is the space of periodic generalized functions $\mathcal{S}'(\mathbb{T})$. Their complex-valued counterparts are denoted by $\mathcal{S}(\mathbb{T}, \mathbb{C})$ (defined in (2.7)) and $\mathcal{S}'(\mathbb{T}, \mathbb{C})$ (defined in (2.8)), respectively. For $k \in \mathbb{Z}$, let $e_k : \mathbb{T} \rightarrow \mathbb{C}$ be the complex exponential function $e_k \triangleq e^{jk \cdot} \in \mathcal{S}(\mathbb{T}, \mathbb{C})$. The Fourier series coefficients of any $f \in \mathcal{S}'(\mathbb{T})$ are given by $\hat{f}[k] = \langle f, e_k \rangle \in \mathbb{C}$. Following (2.16), we have the Fourier series expansion $f = \sum_{k \in \mathbb{Z}} \hat{f}[k] e_k$ for any $f \in \mathcal{S}'(\mathbb{T})$. For example, the Dirac comb $\text{III} \triangleq \sum_{n \in \mathbb{Z}} \delta(\cdot - 2\pi n)$ has Fourier series coefficients $\widehat{\text{III}}[k] = 1$ for any $k \in \mathbb{Z}$.

The weak-derivative operator (Section 2.3) is denoted by $D : \mathcal{S}'(\mathbb{T}) \rightarrow \mathcal{S}'(\mathbb{T})$. More generally, we consider the N_d th-order derivative operator¹ $L = D^{N_d}$ for a fixed integer $N_d \geq 1$, whose null space is given by $\mathcal{N}_L \triangleq \text{span}\{x \mapsto 1\}$ (defined in (2.31)), and whose Fourier multiplier is $\widehat{L}[k] = (jk)^{N_d}$. We then have that $Lf = \sum_{k \in \mathbb{Z}} (jk)^{N_d} \hat{f}[k] e_k$ for any $f \in \mathcal{S}'(\mathbb{T})$.

Following (2.32), the *periodic Green's function* g_L of $L = D^{N_d}$ is given by

$$g_L \triangleq \sum_{k \neq 0} \frac{e_k}{(jk)^{N_d}}. \quad (9.5)$$

¹As in most chapters of the thesis, all the results and algorithms from this chapter can be extended to more general operators L , e.g., rational operators [102].

A L-spline is a periodic function $f \in \mathcal{S}'(\mathbb{T})$ that satisfies

$$Lf = w = \sum_{k=1}^K a_k \text{III}(\cdot - x_k) \quad (9.6)$$

where $K \geq 0$, $a_k \in \mathbb{R} \setminus \{0\}$, and the knots $x_k \in \mathbb{T}$ are pairwise distinct. By Proposition 3.2, a L-spline f that satisfies (9.6) can be represented as

$$f = a_0 + \sum_{k=1}^K a_k g_L(\cdot - x_k) \quad (9.7)$$

for some $a_0 \in \mathbb{R}$. In this case, we necessarily have that $\sum_{k=1}^K a_k = 0$.

9.2.2 Periodic Radon Measures and Native Spaces

We now briefly recall the necessary elements from the periodic setting concerning periodic Radon measures (Section 2.2) and native spaces (Section 2.4). Let $\mathcal{M}(\mathbb{T})$ be the space of periodic Radon measures defined in (2.18). The space $\mathcal{M}(\mathbb{T})$ is a Banach space endowed with the TV norm $\|\cdot\|_{\mathcal{M}}$ (defined in (2.17)).

We denote by $\mathcal{M}_0(\mathbb{T})$ the set of Radon measures with zero mean, *i.e.*, $\mathcal{M}_0(\mathbb{T}) \triangleq \{w \in \mathcal{M}(\mathbb{T}) : \widehat{w}[0] = 0\}$. It is the dual of the space $\mathcal{C}_0(\mathbb{T}) \triangleq \{f \in \mathcal{C}(\mathbb{T}) : \widehat{f}[0] = 0\}$ of continuous functions with zero mean endowed with the supremum norm $\|\cdot\|_{L_\infty}$.

Let $L = D^{N_d}$ for some $N_d \geq 1$. The *native space* of L (defined in (2.33)) is given by

$$\mathcal{M}_L(\mathbb{T}) \triangleq \{f \in \mathcal{S}'(\mathbb{T}) : Lf \in \mathcal{M}(\mathbb{T})\}. \quad (9.8)$$

We introduce the pseudo-inverse operator L^\dagger defined in Proposition 2.6 as

$$L^\dagger f \triangleq \sum_{k \neq 0} \frac{\widehat{f}[k]}{(jk)^{N_d}} e_k \quad (9.9)$$

for any $f \in \mathcal{S}'(\mathbb{T})$. By Theorem 2.4, the native space $\mathcal{M}_L(\mathbb{T})$ then admits the direct-sum decomposition

$$\mathcal{M}_L(\mathbb{T}) = L^\dagger \mathcal{M}_0(\mathbb{T}) \oplus \text{span}\{1\}, \quad (9.10)$$

and any $f \in \mathcal{M}_L(\mathbb{T})$ has a unique decomposition

$$f = L^\dagger w + a \quad (9.11)$$

where $w \in \mathcal{M}_0(\mathbb{T})$ and $a \in \mathbb{R}$ are given by $w = Lf$ and $a = \widehat{f}[0]$.

9.2.3 Dual Certificates for the Basis Pursuit in the Continuum

In this section, we introduce the material from [170], which concerns the *basis pursuit in the continuum*² (BPC) problem

$$\min_{w \in \mathcal{M}(\mathbb{T}), \nu(w)=\mathbf{y}} \|w\|_{\mathcal{M}}. \quad (\text{BPC})$$

Note that the results of [170] are more general than the ones presented here; however, since our focus in this thesis is on the generalized BPC (gBPC)—the BPC with a regularization operator L —we restrict here to the minimum required results from [170] from which the results of [241] (that concern the gBPC) can be deduced.

The analysis of (BPC) benefits from the theory of duality for infinite-dimensional convex optimization, as exposed for instance by Ekeland and Temam in [26]. This line of research has proven to be fruitful to study optimization on measure spaces [47, 49, 62, 65, 66]. We mostly rely on the concepts and results exposed in [65], but very similar tools can be found elsewhere [47, 49]. Considering the dual problem to (BPC) and writing the optimality conditions that link the solutions of both problems³ leads to the notion of *dual certificates*, which are continuous functions on \mathbb{T} satisfying some conditions (see Proposition 9.1 below). In particular, dual certificates enable to *certify* that some $w \in \mathcal{M}(\mathbb{T})$ is a solution of (BPC) and to localize its support. In the next definition, we introduce some notations that allow us to formulate these statements mathematically. We recall that a Radon measure $w \in \mathcal{M}(\mathbb{T})$ can be uniquely decomposed as $w = w_+ - w_-$, where w_+ and w_- are nonnegative measures (Jordan decomposition). In the following definition, we define signed supports and saturation sets; it is the periodic equivalent of Definition 7.2 from Chapter 7.

Definition 9.1. *Let $w \in \mathcal{M}(\mathbb{T})$. We define the signed support of w as*

$$\text{supp}_{\pm}(w) \triangleq \text{supp}(w_+) \times \{1\} \cup \text{supp}(w_-) \times \{-1\}, \quad (9.12)$$

where $\text{supp}(\tilde{w})$ is the support of $\tilde{w} \in \mathcal{M}(\mathbb{T})$.

Let $\eta \in \mathcal{C}(\mathbb{T})$. The positive and negative saturation sets of η are given by

$$\text{sat}_+(\eta) \triangleq \eta^{-1}(\{1\}) \quad \text{and} \quad \text{sat}_-(\eta) \triangleq \eta^{-1}(\{-1\}), \quad (9.13)$$

respectively. Finally, we define the signed saturation set of η as

$$\text{sat}_{\pm}(\eta) \triangleq \text{sat}_+(\eta) \times \{1\} \cup \text{sat}_-(\eta) \times \{-1\}. \quad (9.14)$$

The following proposition introduces formally the notion of dual certificates for the (BPC)

²A similar terminology, the “continuous basis pursuit”, is used in a different context in [67, 175].

³Note that dual certificates always exist for the (BPC) with Fourier measurements [13].

problem.

Proposition 9.1. *There exists a function $\eta \in \mathcal{C}(\mathbb{T})$, which is a real trigonometric polynomial of degree at most K_c , that satisfies $\|\eta\|_{L^\infty} \leq 1$, and is such that for any solution $w_0 \in \mathcal{M}(\mathbb{T})$ of (BPC), we have one of the following equivalent conditions:*

- $\langle w_0, \eta \rangle = \|w_0\|_{\mathcal{M}}$;
- $\text{supp}_\pm(w_0) \subset \text{sat}_\pm(\eta)$.

Such a function η is called a dual certificate.

Proposition 9.1 is stated in an equivalent form in [65, Proposition 3], where dual certificates of Problem (BPC) are studied⁴. Moreover, the proof of Proposition 9.1 is a seamless adaptation of that of Proposition 7.1 to the periodic setting in the absence of a regularization operator and with a specific forward model ν of the form (9.3). Hence, we do not provide the proof.

Concerning the study the uniqueness of (BPC), the following proposition is an important consequence of Proposition 9.1. Proposition 9.2 can also be deduced from [65] (see also [47]).

Proposition 9.2. *If there exists a nonconstant⁵ dual certificate for the Problem (BPC), then it has a unique solution of the form $w^* = \sum_{k=1}^K a_k \delta(\cdot - x_k)$ with $K \leq 2K_c$ and $a_k \in \mathbb{R}$, $x_k \in \mathbb{T}$.*

Proof. Let η be a nonconstant dual certificate for (BPC) and $w^* \in \mathcal{M}(\mathbb{T})$ a solution. Firstly, η is a trigonometric polynomial of degree at most K_c and so is its derivative η' . Since η is nonconstant, η' has at most $2K_c$ roots [242, p. 150]. By Proposition 9.1, we have $\text{supp}_\pm(w^*) \subset \text{sat}_\pm(\eta)$, hence any point in the support of w^* is a root of η' . Consequently, w^* is composed of at most $2K_c$ Dirac masses. Let $\boldsymbol{\tau} = (\tau_1, \dots, \tau_P) \in \mathbb{T}^P$ be the pairwise distinct roots of η' , with $P \leq 2K_c$. Then, we have $w^* = \sum_{p=1}^P a_p \delta(\cdot - \tau_p)$, with $\mathbf{a} = (a_1, \dots, a_P) \in \mathbb{R}^P$ (note that some weights may be equal to 0). Moreover, any other solution of (BPC) must be of the form $w_{\tilde{\mathbf{a}}, \boldsymbol{\tau}}^* = \sum_{p=1}^P \tilde{a}_p \delta(\cdot - \tau_p)$, with $\tilde{\mathbf{a}} = (\tilde{a}_1, \dots, \tilde{a}_P) \in \mathbb{R}^P$ (once again some weights may be equal to 0), where

$$\nu(w^*) = \mathbf{y} = \nu(w_{\tilde{\mathbf{a}}, \boldsymbol{\tau}}^*). \quad (9.15)$$

⁴Duval and Peyré consider more general measurement operators whose image can lie in a Hilbert space and exemplify their results for low-frequency measurements.

⁵In the more general case studied in [65], this corresponds to the nondegeneracy condition of the dual certificate.

Consider the matrix

$$\mathbf{M}_\tau \triangleq \begin{bmatrix} e^{jK_c\tau_1} & \dots & e^{jK_c\tau_P} \\ \vdots & \ddots & \\ e^{j\tau_1} & \dots & e^{j\tau_P} \\ 1 & \dots & 1 \\ e^{-j\tau_1} & \dots & e^{-j\tau_P} \\ \vdots & \ddots & \\ e^{-jK_c\tau_1} & \dots & e^{-jK_c\tau_P} \end{bmatrix} = \left[e^{jk\tau_p} \right]_{\substack{1 \leq p \leq P \\ -K_c \leq k \leq K_c}} \in \mathbb{C}^{(2K_c+1) \times P}. \quad (9.16)$$

By definition of $\boldsymbol{\nu}$ and by Equation (9.15), we get that $\boldsymbol{\nu}(\mathbf{a} - \tilde{\mathbf{a}}) = \mathbf{M}_\tau(\mathbf{a} - \tilde{\mathbf{a}}) = \mathbf{0}$. The matrix \mathbf{M}_τ is a Vandermonde-type matrix, which is therefore of full rank P , since $P \leq 2K_c$ and τ_1, \dots, τ_P are pairwise distinct. Hence the null space of \mathbf{M}_τ is trivial and $\mathbf{a} = \tilde{\mathbf{a}}$, which prove the uniqueness of the solution $w^* = w_{\tilde{\mathbf{a}}, \tau}^*$. \square

From Proposition 9.2, we can deduce the following corollary. The latter is a special case of [240, Corollary 1]; the specialization leads to a much simpler proof.

Corollary 9.1. *If $y_0 = 0$, then Problem (BPC) has a unique solution of the form $w^* = \sum_{k=1}^K a_k \delta(\cdot - x_k)$ with $K \leq 2K_c$ and $a_k \in \mathbb{R}$, $x_k \in \mathbb{T}$.*

Proof. The result trivially holds true for $\mathbf{y} = \mathbf{0}$, since the unique solution of (BPC) is then $w^* = 0$. Assume that $\mathbf{y} \neq \mathbf{0}$, and let η be a dual certificate for Problem (BPC). Assume by contradiction that η is constant, *i.e.*, $\eta = \pm 1$. Then, by Proposition 9.1, we have $\langle w_0, \eta \rangle = \pm \|w^*\|_{\mathcal{M}}$. Yet we have that $\langle w^*, \eta \rangle = \pm \widehat{w^*}[0] = \pm \nu_1(w^*) = \pm y_0 = 0$ and hence that $w^* = 0$, which contradicts the assumption $\mathbf{y} = \mathbf{0}$ since $\boldsymbol{\nu}(w^*) = \mathbf{y}$.

Hence, η must be nonconstant; the desired result then directly follows from Proposition 9.2. \square

9.3 Uniqueness of TV-Based Penalized Problems

It is well known that gTV-based optimization problems lead to splines solutions (see the representer theorems from Section 3.3). In addition to being existence results, these representer theorems provide the form of the (extreme-point) solutions of the optimization task. Here, we focus on problems of the form (9.2), whose specificity compared to the generic formulation in Problem 3.31 is the periodic setting ($\mathbb{K} = \mathbb{T}$) and the Fourier-domain measurement operator $\boldsymbol{\nu}$.

We now state our main result, which guarantees the uniqueness of the solution to Problem (9.2). The proof relies on Corollary 9.1 by reformulating Problem (9.2) over the space of Radon measures; interestingly, the regularization operator $L = D^{N_d}$ leads

to systematic uniqueness, which is not true of generic problems formulated over Radon measures.

Theorem 9.1. *Let $L = D^{N_d}$ with $N_d \geq 1$, $K_c \geq 0$ be the cutoff frequency of the low-pass filter $\nu : \mathcal{M}_L(\mathbb{T}) \rightarrow \mathbb{R} \times \mathbb{C}^{K_c}$ defined in (9.3), $\mathbf{y} \in \mathbb{R} \times \mathbb{C}^{K_c}$, $E(\cdot, \mathbf{y}) : \mathbb{R} \times \mathbb{C}^{K_c} \rightarrow \mathbb{R}^+$ be a functional that is a proper, convex, strictly convex over its effective domain⁶, lower semi-continuous, and coercive, and $\lambda > 0$. Then, the optimization problem*

$$\mathcal{V} \triangleq \arg \min_{f \in \mathcal{M}_L(\mathbb{T})} \left(E(\nu(f), \mathbf{y}) + \lambda \|Lf\|_{\mathcal{M}} \right), \quad (9.17)$$

admits a unique solution that is a L-spline whose number of knots is bounded by $2K_c$.

Proof. Using a classical argument based on the strict convexity of $E(\cdot, \mathbf{y})$ (see Proposition 7.7), we deduce that all solutions f^* of Problem (9.17) share an identical observation vector $\mathbf{y}_\lambda \in \mathbb{R}^{N_d}$, that is, $\forall f^* \in \mathcal{V}$, we have $\nu(f^*) = \mathbf{y}_\lambda$. Hence, Problem (9.17) is equivalent to

$$\mathcal{V} = \arg \min_{f \in \mathcal{M}_L(\mathbb{T}), \nu(f) = \mathbf{y}_\lambda} \|Lf\|_{\mathcal{M}}. \quad (9.18)$$

By (9.11), any $f \in \mathcal{M}_L(\mathbb{T})$ admits a unique decomposition $f = L^\dagger w + a \in \mathcal{M}_L(\mathbb{T})$ with $(w, a) \in \mathcal{M}_0(\mathbb{T}) \times \mathbb{R}$. By plugging in this expansion into the cost functional of Problem (9.18), we get that the latter is equivalent to

$$\arg \min_{(w, a) \in \mathcal{M}_0(\mathbb{T}) \times \mathbb{R}, \nu(L^\dagger w + a) = \mathbf{y}_\lambda} \|w\|_{\mathcal{M}} \iff \arg \min_{w \in \mathcal{M}_0(\mathbb{T}), \nu(w) = \mathbf{z}} \|w\|_{\mathcal{M}}, \quad (9.19)$$

where $\mathbf{z} \in \mathbb{R} \times \mathbb{C}^{K_c}$ is defined as $z_0 = 0$ and $z_k = \widehat{L}[k](\mathbf{y}_\lambda)_k$ for $k \neq 0$. The equivalence in (9.19) comes from the fact that $\nu(L^\dagger w + a) = \left(a, \frac{\widehat{w}[1]}{\widehat{L}[1]}, \dots, \frac{\widehat{w}[K_c]}{\widehat{L}[K_c]} \right)$. Any $f^* \in \mathcal{V}$ can thus be decomposed as $f^* = L^\dagger w^* + (\mathbf{y}_\lambda)_0$ where w^* is a solution of Problem (9.19).

Note that z_0 in Problem (9.19) satisfies the assumptions of Corollary 9.1. Hence, Problem (9.19) has a unique solution w^* that is a sum of at most $2K_c$ Dirac impulses. This in turn implies the uniqueness of the solution $f^* = L^\dagger w^* + (\mathbf{y}_\lambda)_0$ as well as the fact that it is a L-spline with at most $2K_c$ knots. \square

Remark 9.2. *Theorem 9.1 remains valid for more general operators L , namely any spline-admissible operator in the sense of [22, Definition 2] whose null space includes constant functions, i.e., $L\{1\} = 0$.*

Theorem 9.1 has three components: i) it guarantees the uniqueness of the solution, ii) it provides the form of the solution, and iii) it gives an upper bound on the number of

⁶The effective domain of a convex function $g : X \rightarrow \mathbb{R}^+ \cup \{\infty\}$ is the set $\{x \in X, g(x) < \infty\}$ [197].

knots of the solution. The first item, arguably the most striking one, is completely new; existing results typically provide the form of extreme-point solutions of the problem. We are not aware of any other systematic uniqueness result concerning inverse problems with TV-based regularization. The second item is already known (Theorem 3.4). Finally, concerning the third item, known proof techniques (Theorem 3.4 and [14, 87]) allow us to reach the bound $2K_c + 1$, which we improve to $2K_c$.

One can actually be slightly more precise and show that the mean of the solution is known under very mild conditions on the cost functional E . Under this assumption, we also provide a reformulation of Problem (9.17) over the space of Radon measures.

Proposition 9.3. *We assume that we are under the conditions of Theorem 9.1 and that the data-fidelity cost functional E is such that for any fixed $(z_1, \dots, z_{K_c}) \in \mathbb{C}^{K_c}$, we have*

$$y_0 = \arg \min_{z_0 \in \mathbb{R}} E(\mathbf{z}, \mathbf{y}) \quad (9.20)$$

where $\mathbf{y} = (y_0, y_1, \dots, y_{K_c}) \in \mathbb{R} \times \mathbb{C}^{K_c}$ and $\mathbf{z} = (z_0, z_1, \dots, z_{K_c}) \in \mathbb{R} \times \mathbb{C}^{K_c}$. Then, the unique solution f^* to (9.17) admits the decomposition $f^* = y_0 + L^\dagger w^*$ where

$$w^* \triangleq \arg \min_{w \in \mathcal{M}_0(\mathbb{T})} \left(E(\boldsymbol{\nu}(L^\dagger w + y_0), \mathbf{y}) + \lambda \|w\|_{\mathcal{M}} \right). \quad (9.21)$$

In particular, this implies that $\widehat{f^*}[0] = y_0$.

Proof. Similarly to our manipulation in (9.19), Problem (9.17) is equivalent to

$$(w^*, a^*) = \arg \min_{(w, a) \in \mathcal{M}_0(\mathbb{T}) \times \mathbb{R}} \left(E(\boldsymbol{\nu}(L^\dagger w + a), \mathbf{y}) + \lambda \|w\|_{\mathcal{M}} \right), \quad (9.22)$$

with $f^* = L^\dagger w^* + a^*$. Problem (9.22) has a unique solution due to that of Problem (9.17) (proved in Theorem 9.1), and to the uniqueness of the decomposition of f^* (Equation (9.11)). Then, we have that $\boldsymbol{\nu}(a^* + L^\dagger w^*) = (a^*, \widehat{L^\dagger w^*}[1], \dots, \widehat{L^\dagger w^*}[K_c])$, which by (9.20) implies that $E(\boldsymbol{\nu}(a^* + L^\dagger w^*), \mathbf{y}) \geq E(\boldsymbol{\nu}(y_0 + L^\dagger w^*), \mathbf{y})$, with equality if and only if $a^* = y_0$. Hence, since the constant a does not impact the regularization in (9.21), we must have that $a^* = \widehat{f^*}[0] = y_0$. Problem (9.22) can thus be rewritten as (9.21). \square

Remark 9.3. *The relation (9.20) holds for virtually all classical cost functionals, including any ℓ_p norm-based cost such as the quadratic data fidelity $E(\mathbf{z}, \mathbf{y}) = \frac{1}{2} \|\mathbf{z} - \mathbf{y}\|_2^2$, or any separable cost whose minimum over each component is reached when $y_m = z_m$, such as indicator functions. Proposition 9.3 ensures that the mean of the solution of Problem (9.17) is given by $\widehat{f^*}[0] = y_0$.*

9.4 Uniform Convergence of Grid-Based Methods

A common way to solve infinite-dimensional continuous-domain problems such as (9.23) algorithmically is to discretize them using a uniform finite grid [66, 67]. In this section, we propose such a discretization method of the problem

$$f^* = \arg \min_{f \in \mathcal{M}_L(\mathbb{T})} \left(\frac{1}{2} \|\boldsymbol{\nu}(f) - \mathbf{y}\|_2^2 + \lambda \|Lf\|_{\mathcal{M}} \right), \quad (9.23)$$

that is, Problem (9.17) with a quadratic data-fidelity cost $E(\mathbf{z}, \mathbf{y}) = \frac{1}{2} \|\mathbf{z} - \mathbf{y}\|_2^2$. Note that we no longer denote the solution of Problem (9.23) as a set but as a function f^* , since Theorem 9.1 guarantees that this solution is unique. We restrict to the case of the quadratic data fidelity for the sake of simplicity, although our results hereafter hold for more general choices of E . Our choice clearly satisfies the assumption of Proposition 9.3, hence the solution f^* of (9.23) satisfies $\widehat{f^*}[0] = y_0$.

Our discretization method, which was introduced for similar problems in Chapter 4, consists in restricting the search space of Problem (9.23) to the space of uniform L-splines $\mathcal{M}_{L,h}(\mathbb{T})$, *i.e.*, L-splines that satisfy (9.6) with knots x_k on a uniform grid. The space $\mathcal{M}_{L,h}(\mathbb{T})$ is defined as follows for a grid size $h \triangleq \frac{2\pi}{P}$, where $P \geq 1$ is the number of grid points:

$$\mathcal{M}_{L,h}(\mathbb{T}) \triangleq \left\{ f \in \mathcal{S}'(\mathbb{T}), Lf = \sum_{p=0}^{P-1} a[p] \text{III}(\cdot - hp) \right\} \quad \text{with} \quad h \triangleq \frac{2\pi}{P}. \quad (9.24)$$

Our choice of restricting the search space of Problem (9.23) to $\mathcal{M}_{L,h}(\mathbb{T})$ is guided by Theorem 9.1, which states that the unique solution to this problem is a L-spline. Hence, this choice of space is compatible with the sparsity-promoting regularization $\|L \cdot\|_{\mathcal{M}}$. Although in general, the solution of our problem does not have knots on a uniform grid, it can be approximated arbitrary closely with an element of $\mathcal{M}_{L,h}(\mathbb{T})$ when P is large. The other main feature of our method is that the computations are exact in the continuous domain, both those of the forward model and of the regularization term. Our restricted optimization problem then becomes

$$\mathcal{V}_{\text{res}}(h) \triangleq \arg \min_{f \in \mathcal{M}_{L,h}(\mathbb{T})} \left(\frac{1}{2} \|\boldsymbol{\nu}(f) - \mathbf{y}\|_2^2 + \lambda \|Lf\|_{\mathcal{M}} \right) \quad \text{with} \quad h \triangleq \frac{2\pi}{P}. \quad (9.25)$$

Note that contrary to the original Problem (9.23), the solution set $\mathcal{V}_{\lambda,P}(\mathbf{y})$ of the discretized Problem (9.25) is not necessarily unique.

As we shall demonstrate in Section 9.5, Problem (9.23) can be solved algorithmically with standard finite-dimensional solvers. However, the important question of how well it approximates the original Problem (9.23) still remains. We answer this question in

Theorem 9.2 by proving that any sequence of elements of $\mathcal{V}_{\text{res}}(h)$ converge in a strong sense—namely, uniform convergence—towards f^* when $P \rightarrow \infty$.

Theorem 9.2. *Let $L = D^{N_d}$ with $N_d \geq 2$, $\mathbf{y} \in \mathbb{R}^+ \times \mathbb{C}^{K_c}$, and $\lambda > 0$. We denote by f^* the unique solution to (9.23). For any $P \geq 1$, we set $f_P^* \in \mathcal{V}_{\text{res}}(\frac{2\pi}{P})$. Then, we have that*

$$\|f^* - f_P^*\|_{L^\infty} \xrightarrow{P \rightarrow \infty} 0. \quad (9.26)$$

Remark 9.4. *Despite the fact that the solutions to (9.34) may not be unique, Theorem 9.2 ensures that the convergence (9.26) holds for any choice of the f_P^* .*

Remark 9.5. *Uniform convergence implies convergence with respect to any L_p norm for $1 \leq p \leq \infty$, since we have $\|f\|_p \leq (2\pi)^{1/p} \|f\|_{L^\infty}$ for any $f \in \mathcal{M}_L(\mathbb{T})$.*

Remark 9.6. *Theorem 9.2 holds for more general settings than Problem (9.23). More specifically, our proof seamlessly extends to the more general setting of Theorem 9.1 for any cost functional E that is continuous with respect to its second argument, such as ℓ_p losses of the form $E(\mathbf{z}, \mathbf{y}) = \|\mathbf{z} - \mathbf{y}\|_p^p$. Compared to the setting of Theorem 9.1, this notably excludes indicator functions, i.e., the constrained optimization Problem (9.18). Concerning the regularization operator L , Theorem 9.2 readily extends to any operator L such that $L\{1\} = \{0\}$ and whose periodic Green's function is Lipschitz. This notably excludes the case $L = D$, i.e., $N_d = 1$.*

Proof. We first introduce $\mathcal{M}_{0,h}(\mathbb{T}) = \{w \in \mathcal{M}_0(\mathbb{T}), w = \sum_{p=0}^{P-1} a[p] \text{III}(\cdot - hp)\}$, where $h = \frac{2\pi}{P}$, the uniform discretization of $\mathcal{M}_0(\mathbb{T})$ using Dirac impulses. Then, using Proposition 9.3 (with a restriction of the search space that does not affect the proof), we have that $f_P^* = y_0 + L^\dagger w_P^*$, where

$$w_P^* \in \arg \min_{w \in \mathcal{M}_{0,h}(\mathbb{T})} \left(\frac{1}{2} \|\boldsymbol{\nu}(L^\dagger w + y_0) - \mathbf{y}\|_2^2 + \lambda \|w\|_{\mathcal{M}} \right). \quad (9.27)$$

We now prove that the Radon measures w_P^* converge towards the unique solution of

$$w^* \triangleq \arg \min_{w \in \mathcal{M}_0(\mathbb{T})} \left(\frac{1}{2} \|\boldsymbol{\nu}(L^\dagger w + y_0) - \mathbf{y}\|_2^2 + \lambda \|w\|_{\mathcal{M}} \right) \quad (9.28)$$

for the weak* topology when $P \rightarrow \infty$, where the uniqueness of w^* follows from (9.21) in Proposition 9.3. This convergence is proved by following [66, Proposition 4]; the fact that the search space in (9.28) is $\mathcal{M}_0(\mathbb{T})$ rather than $\mathcal{M}(\mathbb{T})$ does not impact the proof. Then, the operator L^\dagger is linear and continuous between $\mathcal{M}_0(\mathbb{T})$ and $\mathcal{M}_L(\mathbb{T})$ for their respective weak* topologies. This implies that f_P^* converges to f^* for the weak* topology over $\mathcal{M}_L(\mathbb{T})$. According to [22, Proposition 9], $L = D^{N_d}$ is sampling-admissible for $N_d \geq 2$, which implies in particular that III is in the predual of $\mathcal{M}_L(\mathbb{T})$. Equivalently, this implies that $f \mapsto f(x)$ is weak*-continuous over $\mathcal{M}_L(\mathbb{T})$, which then implies that $f_P^*(x) \rightarrow f^*(x)$ for any $x \in \mathbb{T}$ (pointwise convergence).

We now prove that the family $(f_P^*)_{P \in \mathbb{N}}$ is equicontinuous. By [243, Theorem 15, Chapter 7], pointwise and uniform convergences are equivalent, which will then conclude the proof. Since f_P^* is a L-spline, using the expansion (9.7), we have

$$f_P^* = y_0 + L^\dagger w_P^* = y_0 + \sum_{p=0}^{P-1} a_P[p] g_L(\cdot - hp) \quad (9.29)$$

for some coefficients $a_P[p]$, $0 \leq p \leq P-1$, where g_L is the Green's function of L defined in (9.5). Moreover, g_L is a periodic Lipschitz function for $L = D^{N_d}$ and $N_d \geq 2$, hence $\|g_L\|_{\text{Lip}} \triangleq \sup_{x,y \in \mathbb{R}, x \neq y} \frac{|g_L(x) - g_L(y)|}{|x-y|} < \infty$. For any $x, y \in \mathbb{R}$, we have that

$$\begin{aligned} |f_P^*(x) - f_P^*(y)| &\leq \sum_{p=0}^{P-1} |a_P[p]| |g_L(x - hp) - g_L(y - hp)| \\ &\leq \left(\sum_{p=0}^{P-1} |a_P[p]| \right) \|g_L\|_{\text{Lip}} |x - y| \\ &= \|w_P^*\|_{\mathcal{M}} \|g_L\|_{\text{Lip}} |x - y|. \end{aligned} \quad (9.30)$$

We have seen that $w_P^* \rightarrow w^*$ when $P \rightarrow \infty$ for the weak* topology. It is therefore bounded for the TV norm, thanks to the uniform boundedness principle. We therefore deduce from (9.30) that the f_P^* are uniformly Lipschitz, and therefore equicontinuous, which proves the desired result. \square

The first part of the proof of Theorem 9.2, dealing with the pointwise convergence, mostly relies on the generalization of the weak* convergence studied in [66, Proposition 4]. Duval and Peyré use tools from Γ -convergence (see [244] for an introduction) and are themselves inspired by [245].

Theorem 9.2 shows that our grid-based discretization yields spline solutions that are arbitrarily close to the unique solution f^* of (9.23) in the uniform sense when the discretization step $h = \frac{2\pi}{P}$ vanishes. It leverages the uniqueness of the spline reconstruction from Fourier measurements ensured by Theorem 9.1.

9.5 B-spline-Based Algorithm

We now introduce our proposed algorithm to solve the discretized Problem (9.25) in an exact way, *i.e.*, without any discretization error. The algorithm is based on Algorithm 4.1 and uses the B-spline basis to represent the space of uniform splines $\mathcal{M}_{L,h}(\mathbb{T})$. The main difference here with Chapter 4 is the periodic setting, which actually simplifies the treatment of the boundary conditions. Moreover, for the sake of conciseness, we focus here on discretizing for a fixed grid; we do not present the multiresolution aspect of

Algorithm 4.1, although it can seamlessly be adapted to our setting.

9.5.1 Preliminaries on Uniform Periodic Polynomial Splines

We first recall some background information on B-splines and their periodized versions; we refer to Section 3.1 for more information.

For any integer $P \geq 1$, the periodized B-spline matched to L with grid size $h \triangleq \frac{2\pi}{P}$ is then defined as

$$\beta_{L,h}^{\text{per}}(x) \triangleq \sum_{k \in \mathbb{Z}} \beta_{L,h}(x - 2\pi k), \quad \forall x \in \mathbb{T}, \quad (9.31)$$

where $\beta_{L,h}$ is the (nonperiodic) scaled B-spline of the operator $L = D^{N_d}$ defined in (3.13).

As demonstrated in Proposition 3.4, a convenient feature of the space on periodic uniform L-splines $\mathcal{M}_{L,h}(\mathbb{T})$ defined in (9.24) is that it is generated by periodic B-splines:

$$\mathcal{M}_{L,h}(\mathbb{T}) = \left\{ f = \sum_{p=0}^{P-1} c[p] \beta_{L,h}^{\text{per}}(\cdot - hp), \mathbf{c} = (c[0], \dots, c[P-1]) \in \mathbb{R}^P \right\}. \quad (9.32)$$

By (3.27), the innovation of a uniform spline $f = \sum_{p=0}^{P-1} c[p] \beta_{L,h}^{\text{per}}(\cdot - hp) \in \mathcal{M}_{L,h}(\mathbb{T})$ is given by

$$L\{f\} = \frac{1}{h^{N_d-1}} \sum_{p=0}^{P-1} (\mathbf{d}_L * \mathbf{c})[p] \text{III}(\cdot - ph), \quad (9.33)$$

with $\mathbf{c} = (c[0], \dots, c[P-1])$ and $\mathbf{d}_L = (d_L[0], \dots, d_L[P-1])$, and where the $P = \frac{2\pi}{h}$ -periodic sequence d_L is characterized by its discrete Fourier transform (DFT) $D_L[k] = (1 - e^{-jkh})^{N_d}$.

9.5.2 Discrete Problem Formulation

In practice, we use the B-spline representation (9.32) of $\mathcal{M}_{L,h}(\mathbb{T})$ to solve Problem (9.25). The choice of the B-spline representation is guided by numerical considerations: B-splines have the shortest support among any uniform L-spline, and thus lead to well-conditioned optimization tasks. The problem thus consists in optimizing over the $c[0], \dots, c[P-1]$ coefficients, which leads to a computationally feasible finite-dimensional problem, as demonstrated in the following proposition.

Proposition 9.4. *Problem (9.25) is exactly equivalent to solving the finite-dimensional*

problem

$$\mathcal{V}_f(h) \triangleq \left(\arg \min_{\mathbf{c} \in \mathbb{R}^P} \frac{1}{2} \|\mathbf{H}\mathbf{c} - \mathbf{y}\|_2^2 + \frac{\lambda}{h^{N_d-1}} \|\mathbf{d}_L * \mathbf{c}\|_1 \right) \quad \text{with } h = \frac{2\pi}{P}, \quad (9.34)$$

where the matrix $\mathbf{H} \in \mathbb{C}^{(K_c+1) \times P}$ is given by $\mathbf{H}_{n,m} \triangleq \nu_k \left(\beta_{L,h}^{\text{per}}(\cdot - mh) \right) = e^{-jmh} \widehat{\beta}_{L,h}^{\text{per}}[n]$, $\mathbf{d}_L = (d_L[0], \dots, d_L[P-1])$, and $\mathbf{c} = (c[0], \dots, c[P-1])$. The continuous-domain reconstructed signal is then $f^* \triangleq \sum_{p=0}^{P-1} c_p^* \beta_{L,h}^{\text{per}}(\cdot - hp)$, where $\mathbf{c}^* \in \mathcal{V}_f(h)$.

Proof. This equivalence is obtained by plugging in $f = \sum_{p=0}^{P-1} c_p \beta_{L,h}^{\text{per}}(\cdot - hp) \in \mathcal{M}_{L,h}(\mathbb{T})$ into the cost function of Problem (9.23). The expression of the system matrix \mathbf{H} immediately follows. The expression of the regularization term follows from (9.33) and the fact that $\|\sum_{p=0}^{P-1} a_p \mathbb{I}(\cdot - x_p)\|_{\mathcal{M}} = \|\mathbf{a}\|_1$ for pairwise-distinct knot locations x_p . \square

Problem (9.34) is a standard discrete problem with ℓ_1 regularization, and a solution to the latter can be reached using proximal solvers such as ADMM [19] (see Appendix A).

9.6 Experimental Results

We now present some results of our discretization method introduced in the previous section in various experimental settings.

9.6.1 Effects of Gridding

Qualitative Effect

We first present a toy experiment to illustrate the effect of gridding in our discretization method, *i.e.*, restricting the search space to $\mathcal{M}_{L,h}(\mathbb{T})$. We therefore design an experiment in which the solution of the problem is known, in order to observe whether our algorithm is able to reconstruct it. To this end, we take $L = D^2$, and generate a ground-truth signal f_0 that is a periodic D^2 -spline with 2 knots (the locations and amplitudes of the knots are picked at random). We then compute the noiseless data vector $\mathbf{y} = \boldsymbol{\nu}(f_0)$ for $K_c = 3$, and solve the corresponding problem (9.23) with a small regularization parameter $\lambda = 10^{-7}$ in order to enforce the constraints $\boldsymbol{\nu}(f) \approx \mathbf{y}$ with very low error. Since the form of f_0 is compatible with that of the solution given by Theorem 9.1, the hope is that f_0 will be very close to the solution f^* to problem (9.23), which is confirmed by our experiments.

In Figure 9.1a, we show the reconstruction results of our algorithm, using a voluntarily coarse grid with $P = 16$ points for visualization purposes. We observe that since the knot of f_0 are quite far from the grid, it is difficult to approximate f_0 with an element of $\mathcal{M}_{L,h}(\mathbb{T})$. The reconstruction therefore requires several knots on the grid to mimic a

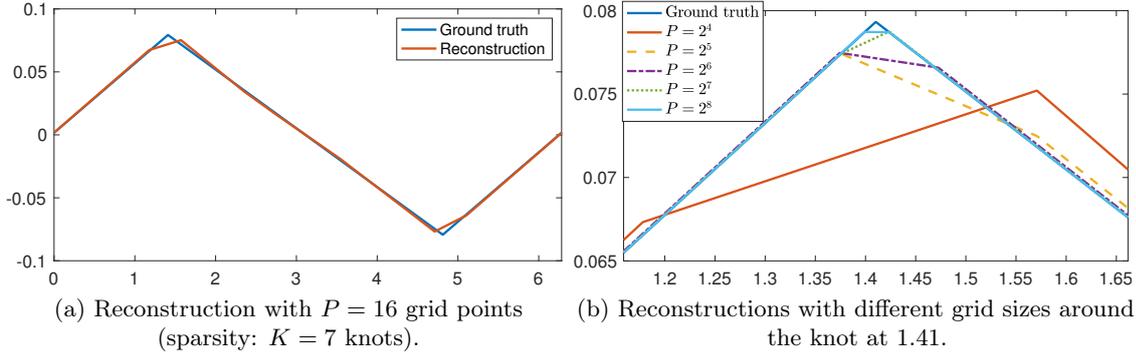


Figure 9.1 – Noiseless reconstruction of a piecewise-linear spline with $N = 2$ knots, $K_c = 3$, and $\lambda = 10^{-7}$.

single knot of f_0 , and thus has a much higher sparsity ($K = 7$ knots versus $K = 2$ for f_0).

However, as we increase the number of grid points, the effect of gridding is greatly reduced, as illustrated in Figure 9.1b: with $P = 512$, the reconstruction using our algorithm is visually indistinguishable from f_0 (which is why we do not show it). However, the knot locations of f_0 still do not exactly lie on the grid, and thus our reconstruction still requires multiple knots to mimic a single knot of f_0 , which leads to a sparsity of $K = 4$. Specifically, our reconstruction has two knots at consecutive grid points 1.3990 and 1.4113 mimicking the knot at 1.4103 of f_0 , and two knots at 4.8106 and 4.8228 mimicking the knot at 4.8122 of f_0 . This effect of knot multiplication due to gridding has already been observed and studied extensively in [66] in the absence of a regularization operator L .

The conclusion is thus that gridding leads to visually near-perfect reconstruction when the number of grid points is very large, which is in line with Theorem 9.2; however, the sparsity of the reconstruction is a poor indicator of the sparsity of the true solution of Problem (9.23), since gridding induces clusters of knots.

Quantitative Effect

In Theorem 9.2, we have proved that any sequence of continuous-domain solutions f_P^* to the grid-restricted problem converges uniformly towards the unique solution f^* of problem (9.23) when P goes to infinity. In order to quantify the speed of this convergence, using the same experimental setting as in Figure 9.1, we compute the error $\|f_P^* - f_0\|_{L^\infty}$ where f_P^* is the reconstructed signal using our grid-based algorithm, and the ground truth f_0 is a proxy for the solution f^* to problem (9.23). As explained earlier, this is a reasonable proxy due to the very small regularization parameter $\lambda = 10^{-7}$. In order to limit the effect of randomness in the choice of the knots of the ground truth, we apply a Monte Carlo-type method by generating 100 different ground-truth signals (following the methodology described in the previous section) and averaging the error over these

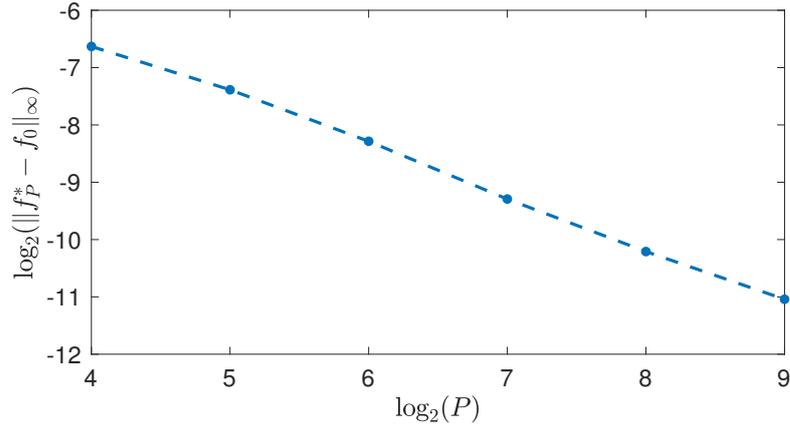


Figure 9.2 – Average error $\|f_P^* - f_0\|_{L_\infty}$ over 100 runs for different grid sizes P (in log-log scale).

100 runs. These average errors for different grid sizes are shown in Figure 9.2. The trend appears to be linear in log-log scale, which indicates an empirical speed of convergence of $\|f_P^* - f_0\|_{L_\infty} \approx (\frac{C}{P^s})$ for some constant $C > 0$ and where $-s < 0$ is the slope of the linear function. We observe here that $s \approx 1$ with $s < 1$. This is consistent with classical approximation theory results, since the approximation power of linear splines with grid size h is in $\mathcal{O}(h)$ for the supremum norm, which corresponds to $s = 1$. There is therefore no hope of having $s > 1$; our observation $s < 1$ can likely be attributed to the fact that we use f_0 as a proxy for f^* and to increased numerical issues when the grid size decreases.

9.6.2 Noisy Recovery of Sparse Splines

We now attempt to recover a ground-truth signal f_0 based on noisy data \mathbf{y} with a regularization operator $L = D$. Once again, the ground-truth signal fits the signal model of problem (9.23), *i.e.*, f_0 is a periodic D-spline (piecewise-constant signal) with $K = 7$ knots. Each knot x_k is chosen at random within consecutive intervals of length $\frac{2\pi}{7}$, and the vector of amplitudes $\mathbf{a} = (a_1, \dots, a_K)$ is an i.i.d. Gaussian random vector projected on the space of zero-mean vectors. The measurements are corrupted by some additive i.i.d. Gaussian noise⁷ $\mathbf{n} \in \mathbb{R} \times \mathbb{C}^{K_c}$ with standard deviation $\sigma = 10^{-3}$, *i.e.*, $\mathbf{y} = \boldsymbol{\nu}(f_0) + \mathbf{n}$.

The signal reconstructed using our algorithm is shown in Figure 9.3. Despite the presence of noise, the reconstruction of the ground truth f_0 is almost perfect. As observed in the previous experiment, the sparsity of the reconstruction ($K = 20$) is higher than that of the ground truth ($K = 7$) due to clusters of knots. We compare our reconstruction to the truncated Fourier series of f_0 up to K_c , *i.e.*, $f_{K_c} = \sum_{k=-K_c}^{K_c} \hat{f}_0[k]e_k$, which solely depends on the noiseless data vector $\boldsymbol{\nu}(f_0)$. Without any prior knowledge, this is the simplest reconstruction one can think of based on the available data $\boldsymbol{\nu}(f_0) = (\hat{f}_0[0], \dots, \hat{f}_0[K_c])$.

⁷For complex entries, both the real and imaginary parts are i.i.d. Gaussian variables with the same σ .

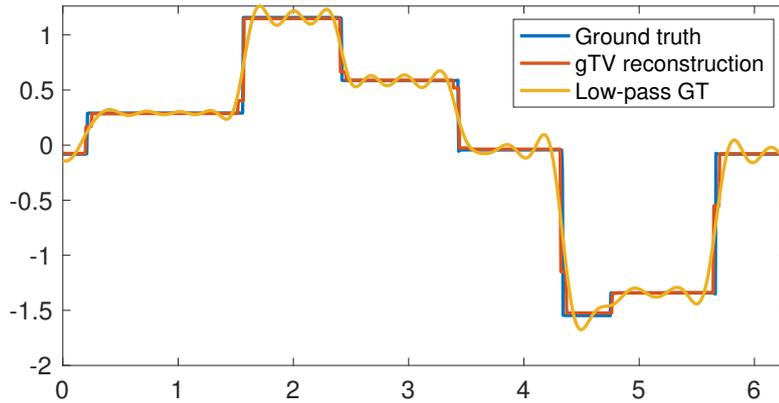


Figure 9.3 – Recovery of piecewise-constant spline with 7 knots with $K_c = 20$. For our reconstruction (gTV), we use $\lambda = 10^{-2}$ and $P = 256$ grid points; the sparsity of the reconstruction is $K = 20$ knots. The data \mathbf{y} is noisy in the gTV case, whereas the noiseless data $\boldsymbol{\nu}(f_0)$ is used for the low-pass reconstruction (partial Fourier series of the ground-truth signal up to the cutoff frequency K_c).

As it turns out, f_{K_c} is also the unique solution to the following constrained L_2 -regularized problem

$$f_{K_c} \triangleq \arg \min_{f: \boldsymbol{\nu}(f) = \boldsymbol{\nu}(f_0)} \|f\|_{L_2}, \quad (9.35)$$

as demonstrated in [45, Theorem 3]. In fact, adding any LSI regularization operator L in (9.35) still yields the same solution, since the basis functions φ_m in [45, Theorem 3] span the same space. This is due to the fact that the measurement functionals ν_m , *i.e.*, complex exponentials, are eigenfunctions of LSI operators.

As expected from the fact that f_{K_c} is a trigonometric polynomial whereas f_0 has sharp jumps, the reconstruction is quite poor and exhibits Gibbs-like oscillations, despite the absence of noise. This clearly demonstrates the superiority of gTV over L_2 regularization for sparse periodic splines reconstruction. Note however that the gap in performance decreases as the order of $L = D^{N_d}$ increases, since Gibbs-like phenomena are less significant for smoother functions.

9.7 Conclusion

This chapter deals with continuous-domain inverse problems where the goal is to recover a periodic function from its low-pass Fourier series measurements. The reconstruction task is formalized as an optimization problem with a TV-based regularization involving a high-order derivative operator. The existence of spline solutions was a known result (representer theorem); our main result further proved that the solution is in fact always unique. We then studied the grid-based discretization of our optimization problem. We

leveraged our uniqueness result to show that any sequence of solutions of the discretized problems converges in uniform norm—a remarkably strong form of convergence—to the solution of the original problem when the grid size vanishes. Finally, we proposed a B-spline-based algorithm to solve the discretized problem, and we illustrated the relevance of our approach on simulations.

Part IV Applications

In this part, we apply our continuous-domain inverse problem frameworks with sparsity-promoting regularization to real-world applications. More precisely, we adapt our B-spline-based exact discretization techniques from Part II. However, contrary to Parts II and III which focused on one-dimensional signals, the works in this part deal with two-dimensional signals, which poses increased computational and theoretical challenges. The applications we consider are the following:

- In Chapter 10, we propose a method for sparse curve fitting based on contour points. We formulate this task as a continuous-domain inverse problem with a novel sparsity-promoting regularization term which we call *rotation-invariant total variation* (RI-TV). We prove a representer theorem for this problem, and we adapt our B-spline-based discretization method from Chapter 4 to solve it computationally. Finally, following Chapter 5, we extend our method to hybrid curve models with varying smoothness properties, for which we also prove a representer theorem. We apply our algorithms to the task of obtaining sparse, stylized representations of fonts.
- In Chapter 11, we propose an image-reconstruction method for scanning transmission X-ray microscopy (STXM). The latter is a two-dimensional imaging modality where the forward model is a nonuniform sampling operator, which justifies our continuous-domain formulation. We use the sparsity-promoting *Hessian-Schatten norm* as a regularization functional, and we deploy B-spline-based techniques inspired from Part II to solve the underlying problem. We apply our method to both simulated and real STXM data.

10 Coupled Splines for Sparse Curve Fitting

This chapter is based on the following work [246]:

I. Lloréns Jover, T. Debarre, S. Aziznejad, and M. Unser, “Coupled Splines for Sparse Curve Fitting”, *arXiv preprint arXiv:2202.01641*, Feb. 2022.

10.1 Introduction

Contour tracing is a common yet rich subject in the image-processing and computer-graphics community. It has numerous applications, such as component labeling [247, 248] or topological structure analysis [249]. The objective is to produce a contour that accurately separates two regions of a given image. This task is, however, not without difficulties. Firstly, the edges suffer from discretization effects and intrinsic image noise. Secondly, the smoothness of the contour may need to be nonuniform, since contours are often made of smooth parts joined by sharp discontinuities.

Our goal is to extract a continuous stylized sparse parametric curve that explains a given set of ordered edge points given by possibly inaccurate two-dimensional coordinates, which is particularly relevant for raster-to-vector conversion, *i.e.*, vectorization [250]. Vectorization consists in converting raster data (e.g. pixel images) into a set of continuous functions representing the contours. This is the principle on which fonts or vector formats like pdf or svg, which allow for zooming into the objects without losing resolution, rely on. It is therefore a problem of great importance for the computer graphics community. Our search for sparsity intimately follows Occam’s razor principle of simplicity. Indeed, it heightens our probability of approaching the true curve, as many real-world signals are sparse. This is the principle on which compressed sensing hinges [8, 10].

Two main approaches come to mind when thinking of contour tracing. The first one consists in joint edge detection and curve fitting. Parametric active contours are popular examples of this approach as these methods provide efficient tools for the extraction of

a contour from an image, for example for point-cloud segmentation [251]. The contour consists in continuous curves that evolve through the optimization of an energy functional and iteratively approximate an image edge [252, 253]. A plethora of parametric snake models can be found in the literature, mostly with model-based energy functionals [254–258], or more recently with learning-based approaches [259, 260]. Of particular relevance to this chapter is a snake model implementation that uses basis functions and that allows for tangent control, a useful property when the smoothness of the contours is nonuniform [261].

The second approach to contour tracing is discrete contour extraction and subsequent curve fitting. In the first approach, the entire image was used to iteratively update the contour, whereas the second approach interpolates a continuous parametric curve from a list of coordinates. This can be achieved using spline curves, which is the method of choice in computer graphics [262–265]. Another popular way to tackle this is through a regularized minimization problem, the regularization enforcing prior knowledge about the curve [266, 267]. The method presented in this work follows the latter paradigm by enforcing a sparsity prior. Other more classical spline-based methods enforce sparsity by simply removing knots from an initially nonsparse curve [268–273]. We also mention other contour-tracing algorithms based on very different techniques [274, 275], as well as recent deep-learning based ones that are applicable to 3D contour tracing [276, 277]; see [278] for a survey.

To attain our goal, we solve a bipartite optimization problem. On one hand, we want that the candidate curve fits the existing contour points exactly. This is achieved through a data-fidelity term. On the other hand, as an infinity of curves could satisfy this fit, we have to enforce prior knowledge into our model. This prior knowledge is introduced as a regularization cost coupled with a regularization operator, the result aiming at the enforcement of desired properties. First, it is likely that the true curve has few variations, which implies that the curve has a sparse representation. Second, it is frequent that variations happen over both the horizontal and the vertical axes simultaneously. Moreover, the recovered curve should be possibly denoised. Finally, the optimization cost should not depend on a rotation of the system of coordinates. We show in this chapter that these specifications lead us to a regularization cost that consists of a mixed (TV- ℓ_2) norm.

In order to sparsify given data, modern regularizers used include structured sparsity [279], namely [280], low rank regularization [281], or deep prior [282]. However, these regularizers assume a discrete setting and thus do not yield a continuous curve as a solution. In addition, the deep prior regularizer does not provide an interpretable model. In this chapter, we explore the continuous setting, as we aim at the recovery of a continuous 2D curve. Similarly to our works in Part II, we explore generalized total-variation (TV) regularization for continuous-domain signal reconstruction using B-splines as basis functions for an exact discretization. Finally, we choose to represent the curve with hybrid splines, which give us the tools to represent curves with nonuniform smoothness. While

Chapter 5 addressed signal reconstruction using hybrid splines, this chapter extends the setting for the handling of curves in 2D, which calls for a new regularizer.

Our main contribution is threefold:

- Firstly, we introduce a continuous rotation-invariant TV (RI-TV) norm as a regularization for the recovery of curves. It effectively reconstructs sparse parametric curves from given contour points while being robust to noise.
- Secondly, we prove a representer theorem according to which there exists a curve with spline components that is a global minimizer of our optimization problem. Building upon this, we propose a curve construction using B-splines, which allows us to discretize the continuous-domain problem exactly with numerical efficiency.
- Finally, we present the combination of such RI-TV norm with a hybrid framework to generate stylized curves with nonuniform smoothness properties.

The remainder of this chapter is organized as follows: in Section 10.2, we present the continuous-domain framework of the optimization problem and introduce our representer theorem. We then show the precise implementation and resolution of this task through the introduction of B-splines in Section 10.3. In Section 10.4, we extend the framework to hybrid splines. Finally, we experimentally verify properties of our contributions and show applications in Section 10.5.

10.2 Continuous-Domain Formulation

Our goal is to recover a 2D parametric curve $\mathbf{r}(t) = (x(t), y(t))$ that best fits a given ordered list of points $\mathbf{p}[m] = (p_x[m], p_y[m])$, $m = 0, \dots, M - 1$. Contours being closed curves, we consider the coordinate functions $x(t)$ and $y(t)$ to be periodic in t . Since we have M data locations, it is convenient to deal with M -periodic functions. We consequently set $t \in \mathbb{T}_M \triangleq [0, M]$. This differs from the periodic setting presented in Chapters 2 and 3 which deals with 2π -periodic functions; however, all the concepts and results can be readily extended to M -periodic functions.

Concurrently, we want to control the sparsity of the fitted curve. This can be achieved by limiting the number of the singularities in the derivatives of its components. This effectively means that $\mathbf{r}(t)$ admits a sparse representation. To that end, we introduce two new elements: a differential operator L and the RI-TV regularization functional.

Like most of this thesis, the mathematical foundations of this chapter rely on Schwartz' theory of distributions [98]. Henceforth, let us denote the Schwartz' space of M -periodic smooth functions by $\mathcal{S}(\mathbb{T}_M)$. Its topological dual, $\mathcal{S}'(\mathbb{T}_M)$ is the space of tempered

distributions over the torus. We refer to Section 2.1 for more background on distribution theory.

10.2.1 Derivative Operators and Splines

The first element we introduce in our problem formulation is $L = D^{N_d}$, the derivative operator whose order N_d , with $N_d \geq 1$, determines the smoothness of the components of the constructed curve¹.

Next, we recall the definition of periodic L-splines from Section 3.1 (Definition 3.3). A periodic L-spline is a function $s : \mathbb{T}_M \rightarrow \mathbb{R}$ that verifies that

$$L\{s\}(t) = \sum_{k=0}^{K-1} a[k] \text{III}_M(t - t_k), \quad (10.1)$$

where $\text{III}_M(t) \triangleq \sum_{k \in \mathbb{Z}} \delta(t - Mk) \in \mathcal{S}'(\mathbb{T}_M)$ is the M -periodic Dirac comb, $K \in \mathbb{N} \setminus \{0\}$ is the number of knots, $a[k] \in \mathbb{R}$ is the amplitude of the k th jump, and $t_k \in \mathbb{R}$ are pairwise-distinct knot locations.

10.2.2 RI-TV Regularization

The second element is \mathcal{R} , a sparsity-promoting regularization functional with key characteristics. Firstly, for 2D curves, the minimization of $\mathcal{R}(L\{\mathbf{r}\})$, where $L\{\mathbf{r}\} \triangleq (L\{x\}, L\{y\})$, should enforce sparsity jointly for the two components of \mathbf{r} . Indeed, we want \mathbf{r} to have few variations, and they often should occur along both components simultaneously. Secondly, if the points $\mathbf{p}[m]$ are rotated by an angle θ , the fitted curve $\mathbf{r}(t)$ should be rotated by the same angle θ . To achieve this, our regularizer \mathcal{R} should be invariant to a rotation of the system of coordinates, meaning that $\mathcal{R}(\mathbf{R}_\theta L\{\mathbf{r}\}) = \mathcal{R}(L\{\mathbf{r}\})$, where \mathbf{R}_θ is a rotation matrix. Similarly, \mathcal{R} should be equivariant to isotropic scaling, meaning that there exists a function A such that $\mathcal{R}(L\{a\mathbf{r}\}) = A(a)\mathcal{R}(L\{\mathbf{r}\})$ for any $a \neq 0$. We now introduce the RI-TV norm, which consists in a mixed continuous (TV- ℓ_2) norm and satisfies our specifications.

Definition 10.1. *Let $p \in [1, +\infty]$. The (TV- ℓ_p) norm of any vector-valued tempered distribution $\mathbf{w} = \begin{bmatrix} w_1 & w_2 \end{bmatrix} \in \mathcal{S}'(\mathbb{T}_M)^2$ is defined as*

$$\|\mathbf{w}\|_{\text{TV-}\ell_p} \triangleq \sup_{\substack{\varphi=(\varphi_1, \varphi_2) \in \mathcal{S}(\mathbb{T}_M)^2 \\ \|\varphi\|_{q, L_\infty} = 1}} (\langle w_1, \varphi_1 \rangle + \langle w_2, \varphi_2 \rangle), \quad (10.2)$$

where $q \in [1, L_\infty]$ is the Hölder conjugate of p with $\frac{1}{p} + \frac{1}{q} = 1$ and $\|\cdot\|_{q, L_\infty}$ is the $(\ell_q - L_\infty)$

¹As in most chapters of the thesis, all the results and algorithms from this chapter can be extended to more general operators L , e.g., rational operators [102].

mixed norm, defined for any $\varphi \in \mathcal{S}(\mathbb{T}_M)^2$ as

$$\|\varphi\|_{q,L_\infty} \triangleq \sup_{t \in \mathbb{T}_M} \|\varphi(t)\|_q. \quad (10.3)$$

The RI-TV norm (or (TV- ℓ_2)) is a well-known quantity; it was previously introduced as the *ground total-variation norm* in [62] in the context of the recovery of Dirac distributions. In fact, in classical measure theory, it is simply known as the total variation of a vector-valued measure [283, Definition 1.4.]; there, it is defined using measure theory instead of by duality as in (10.2). In Theorem 10.1, we compute the (TV- ℓ_p) norm for two general classes of functions or distributions.

Theorem 10.1. *1. For any curve $\mathbf{f} = (f_1, f_2)$ with absolutely integrable components $f_i \in L_1(\mathbb{T}_M)$, $i = 1, 2$, we have that*

$$\left\| \begin{bmatrix} f_1 & f_2 \end{bmatrix} \right\|_{\text{TV-}\ell_p} = \int_0^M \|\mathbf{f}(t)\|_p dt. \quad (10.4)$$

2. Let $\mathbf{w} = (w_1, w_2)$ be a vector-valued distribution of the form $\mathbf{w} = \sum_{k=0}^{K-1} \mathbf{a}[k] \text{III}_M(\cdot - t_k)$ with $\mathbf{a}[k] \in \mathbb{R}^2$, $k = 0, \dots, K-1$. Then, we have that

$$\left\| \begin{bmatrix} w_1 & w_2 \end{bmatrix} \right\|_{\text{TV-}\ell_p} = \sum_{k=0}^{K-1} \|\mathbf{a}[k]\|_p. \quad (10.5)$$

Proof. Item 1: Let $\varphi = (\varphi_1, \varphi_2) \in \mathcal{S}(\mathbb{T}_M)^2$ be an arbitrary smooth curve with $\|\varphi\|_{q,L_\infty} = 1$. On the one hand, the Hölder inequality for vectors implies that, for any $t \in \mathbb{T}_M$,

$$|f_1(t)\varphi_1(t) + f_2(t)\varphi_2(t)| \leq \|\mathbf{f}(t)\|_p \|\varphi(t)\|_q \leq \|\mathbf{f}(t)\|_p, \quad (10.6)$$

where the last inequality is due to $\|\varphi\|_{q,L_\infty} = 1$. On the other hand, the inclusion $f_i \in L_1(\mathbb{T}_M)$ allows us to express the duality product $\langle f_i, \varphi_i \rangle$ as a simple integral of the form

$$\langle f_i, \varphi_i \rangle = \int_0^M f_i(t)\varphi_i(t) dt, \quad i = 1, 2. \quad (10.7)$$

Combining (10.7) with (10.6), we obtain that

$$\langle f_1, \varphi_1 \rangle + \langle f_2, \varphi_2 \rangle = \int_0^M (f_1(t)\varphi_1(t) + f_2(t)\varphi_2(t)) dt \leq \int_0^M \|\mathbf{f}(t)\|_p dt. \quad (10.8)$$

Taking the supremum over all $\varphi \in \mathcal{S}(\mathbb{T}_M)^2$ with $\|\varphi\|_{q,L_\infty} = 1$ then yields $\left\| \begin{bmatrix} f_1 & f_2 \end{bmatrix} \right\|_{\text{TV-}\ell_p} \leq$

$\int_0^M \|\mathbf{f}(t)\|_p dt$. To prove the equality, we first define the functions

$$g_i : \mathbb{T}_M \rightarrow \mathbb{R} : t \mapsto \mathbb{1}_{\mathbf{f} \neq \mathbf{0}} \frac{\text{sgn}(f_i(t)) |f_i(t)|^{(p-1)}}{\|\mathbf{f}(t)\|_p^{(p-1)}}, \quad (10.9)$$

where $\mathbb{1}_A$ denotes the indicator function of the set A . We note that g_i are Borel-measurable with $\|g_i\|_{L_\infty} \leq 1$ for $i = 1, 2$. Further, one readily verifies that

$$\int_0^M (f_1(t)g_1(t) + f_2(t)g_2(t)) dt = \int_0^M \|\mathbf{f}(t)\|_p dt. \quad (10.10)$$

By applying a variant of Lusin's theorem (see [284, Theorem 7.10]) on the space $\mathcal{C}(\mathbb{T}_M)$ of M -periodic continuous functions, we then consider the ϵ -approximations $g_{i,\epsilon} \in \mathcal{C}(\mathbb{T}_M)$ of g_i such that $\|g_{i,\epsilon}\|_{L_\infty} \leq \|g_i\|_{L_\infty} \leq 1$, and $\int_E |f_i(t)| dt \leq \epsilon/8, i = 1, 2$, where $E = \{t \in \mathbb{T}_M : g_{i,\epsilon}(t) \neq g_i(t)\}$. This in effect implies that

$$\begin{aligned} \int_0^M |f_i(t)| \cdot |g_{i,\epsilon}(t) - g_i(t)| dt &= \int_E |f_i(t)| \cdot |g_{i,\epsilon}(t) - g_i(t)| dt \\ &\leq \|\mathbb{1}_E f_i\|_{L_1} \|g_{i,\epsilon} - g_i\|_{L_\infty} \leq \frac{\epsilon}{4}. \end{aligned} \quad (10.11)$$

Then, by the denseness of $\mathcal{S}(\mathbb{T}_M)$ in $\mathcal{C}(\mathbb{T}_M)$, there exist $\varphi_{i,\epsilon} \in \mathcal{S}(\mathbb{T}_M)$ with $\|g_{i,\epsilon} - \varphi_{i,\epsilon}\|_{L_\infty} \leq \frac{\epsilon}{4\|f_i\|_{L_1}}$. This gives us the upper bound

$$\int_0^M |f_i(t)| \cdot |\varphi_{i,\epsilon}(t) - g_{i,\epsilon}(t)| dt \leq \|f_i\|_{L_1} \|\varphi_{i,\epsilon} - g_{i,\epsilon}\|_{L_\infty} \leq \frac{\epsilon}{4}. \quad (10.12)$$

Next, we use the triangle inequality to obtain the lower bound

$$\begin{aligned} \langle f_i, \varphi_{i,\epsilon} \rangle &\geq \int_0^M f_i(t)g_i(t) dt - \int_0^M |f_i(t)| \cdot |g_i(t) - g_{i,\epsilon}(t)| dt - \int_0^M |f_i(t)| \cdot |g_{i,\epsilon}(t) - \varphi_{i,\epsilon}(t)| dt \\ &\geq \int_0^M f_i(t)g_i(t) dt - \frac{\epsilon}{2}, \quad i = 1, 2, \end{aligned} \quad (10.13)$$

where the last inequality follows from the combination of (10.11) and (10.12). Finally, we use (10.10) to conclude that

$$\left\| \begin{bmatrix} f_1 & f_2 \end{bmatrix} \right\|_{\text{TV}-\ell_p} \geq \frac{\langle f_1, \varphi_{1,\epsilon} \rangle + \langle f_2, \varphi_{2,\epsilon} \rangle}{\|(\varphi_{1,\epsilon}, \varphi_{2,\epsilon})\|_{q, L_\infty}} \geq \frac{\int_0^M \|\mathbf{f}(t)\|_p dt - \epsilon}{1 + O(\epsilon)}. \quad (10.14)$$

We conclude the proof by letting $\epsilon \rightarrow 0$.

Item 2: Similarly to the previous part, for any smooth curve $\varphi = (\varphi_1, \varphi_2) \in \mathcal{S}(\mathbb{T}_M)^2$

with $\|\boldsymbol{\varphi}\|_{q,L_\infty} = 1$, we have that

$$\begin{aligned} \langle w_1, \varphi_1 \rangle + \langle w_2, \varphi_2 \rangle &= \sum_{k=0}^{K-1} (a_1[k]\varphi_1(t_k) + a_2[k]\varphi_2(t_k)) \\ &\leq \sum_{k=0}^{K-1} \|\mathbf{a}[k]\|_p \|\boldsymbol{\varphi}(t_k)\|_q \leq \sum_{k=0}^{K-1} \|\mathbf{a}[k]\|_p. \end{aligned} \quad (10.15)$$

Taking the supremum over $\boldsymbol{\varphi} = (\varphi_1, \varphi_2)$ with $\|\boldsymbol{\varphi}\|_{q,L_\infty} = 1$ then yields that $\|\mathbf{w}\|_{\text{TV}-\ell_p} \leq \sum_{k=0}^{K-1} \|\mathbf{a}[k]\|_p$. To prove the equality, we first define a set of vectors $\boldsymbol{\varphi}_k \in \mathbb{R}^2$ such that $\|\boldsymbol{\varphi}_k\|_\infty = 1$ and $\mathbf{a}[k]^T \boldsymbol{\varphi}_k = \|\mathbf{a}[k]\|_p$ for $k = 0, \dots, K-1$. We then consider a smooth curve $\boldsymbol{\varphi}^* \in \mathcal{S}(\mathbb{T}_M)^2$ with $\|\boldsymbol{\varphi}^*\|_{q,L_\infty} = 1$ such that $\boldsymbol{\varphi}^*(t_k) = \boldsymbol{\varphi}_k$. We conclude the proof by verifying that

$$\|\mathbf{w}\|_{\text{TV}-\ell_p} \geq \langle w_1, \boldsymbol{\varphi}_1^* \rangle + \langle w_2, \boldsymbol{\varphi}_2^* \rangle = \sum_{k \in \mathbb{Z}} \|\mathbf{a}[k]\|_p.$$

□

As demonstrated in Theorem 3.3, the outer TV norm is the continuous sparsity-promoting counterpart of the ℓ_1 norm. The inner ℓ_p norm in Item 1 induces a coupling of the f_1 and f_2 components. Indeed, it first aggregates the f_1 and f_2 curve components, which the outer TV norm then jointly sparsifies. This is true of any ℓ_p norm for $p \neq 1$. For $p = 1$, the components are no longer coupled due to the separability of the ℓ_1 norm. For $p = 2$ and for any curve $\mathbf{f} = (f_1, f_2)$, we set

$$\mathcal{R}(\mathbf{f}) \triangleq \left\| \begin{bmatrix} f_1 & f_2 \end{bmatrix} \right\|_{\text{TV}-\ell_2}. \quad (10.16)$$

Proposition 10.1. *The (TV- ℓ_2) norm, denoted by \mathcal{R} , is invariant to rotation, in the sense that $\mathcal{R}(\mathbf{R}_\theta \mathbf{f}) = \mathcal{R}(\mathbf{f})$, where \mathbf{R}_θ is a rotation matrix. Furthermore, the (TV- ℓ_2) norm is the only (TV- ℓ_p) norm that is rotation invariant for $p \in [1, +\infty]$.*

Proof. By substitution of $\mathbf{r}_\theta \mathbf{f}$ in (10.2), we have that

$$\begin{aligned} \mathcal{R}(\mathbf{r}_\theta \mathbf{f}) &= \sup_{\substack{\boldsymbol{\varphi} \in \mathcal{S}(\mathbb{T}_M)^2 \\ \|\boldsymbol{\varphi}\|_{2,L_\infty} = 1}} (\langle \cos(\theta)f_1 - \sin(\theta)f_2, \varphi_1 \rangle + \langle \sin(\theta)f_1 + \cos(\theta)f_2, \varphi_2 \rangle) \\ &= \sup_{\substack{\boldsymbol{\varphi} \in \mathcal{S}(\mathbb{T}_M)^2 \\ \|\boldsymbol{\varphi}\|_{2,L_\infty} = 1}} (\langle f_1, \cos(\theta)\varphi_1 + \sin(\theta)\varphi_2 \rangle + \langle f_2, -\sin(\theta)\varphi_1 + \cos(\theta)\varphi_2 \rangle). \end{aligned} \quad (10.17)$$

We perform the change of variable $\boldsymbol{\psi} = \mathbf{r}_{-\theta} \boldsymbol{\varphi}$. We readily conclude that, since $\mathbf{r}_{-\theta}$ is bijective over $\mathcal{S}(\mathbb{T}_M)^2$, for any $\boldsymbol{\varphi} \in \mathcal{S}(\mathbb{T}_M)^2$, we have that $\boldsymbol{\psi} = \mathbf{r}_{-\theta} \boldsymbol{\varphi} \in \mathcal{S}(\mathbb{T}_M)^2$.

Additionally, and in accordance with (10.3), we have that

$$\|\boldsymbol{\psi}\|_{2,L_\infty} = \sup_{t \in \mathbb{T}_M} \|\boldsymbol{\psi}(t)\|_2 = \sup_{t \in \mathbb{T}_M} \|\mathbf{r}_{-\theta} \boldsymbol{\varphi}(t)\|_2 = \sup_{t \in \mathbb{T}_M} \|\boldsymbol{\varphi}(t)\|_2, \quad (10.18)$$

as $\mathbf{r}_{-\theta}$ is an isometry. Hence, it does not change the ℓ_2 norm of a vector. Consequently, we have that

$$\mathcal{R}(\mathbf{r}_\theta \mathbf{f}) = \sup_{\boldsymbol{\psi} \in \mathcal{S}(\mathbb{T}_M)^2, \|\boldsymbol{\psi}\|_{2,L_\infty} = 1} (\langle f_1, \psi_1 \rangle + \langle f_2, \psi_2 \rangle) = \mathcal{R}(\mathbf{f}). \quad (10.19)$$

Moreover, according to Item 1 of Theorem 10.1, for any curve $\mathbf{f} = (f_1, f_2)$ with absolutely integrable components $f_i \in L_1(\mathbb{T}_M)$, $i = 1, 2$, the TV- ℓ_p norm for $p \in [1, +\infty]$ is

$$\|\mathbf{f}\|_{\text{TV}-\ell_p} = \int_0^M (|f_1(t)|^p + |f_2(t)|^p)^{\frac{1}{p}} dt. \quad (10.20)$$

We take $f_1(t) = 1$, $f_2(t) = 0$, and $\theta = \frac{\pi}{4}$. This gives us

$$\|\mathbf{f}\|_{\text{TV}-\ell_p} = \int_0^M (|1|^p + |0|^p)^{\frac{1}{p}} dt = M. \quad (10.21)$$

When applying the planar rotation \mathbf{r}_θ to the curve \mathbf{f} , we have that

$$\begin{aligned} \|\mathbf{r}_\theta \mathbf{f}\|_{\text{TV}-\ell_p} &= \int_0^M (|f_1(t) \cos(\theta) - f_2(t) \sin(\theta)|^p + |f_1(t) \sin(\theta) + f_2(t) \cos(\theta)|^p)^{\frac{1}{p}} dt \\ &= \int_0^M (|\cos(\theta)|^p + |\sin(\theta)|^p)^{\frac{1}{p}} dt \\ &= \int_0^M \left(2 \left| \frac{\sqrt{2}}{2} \right|^p \right)^{\frac{1}{p}} dt = 2^{\frac{1}{p} - \frac{1}{2}} M. \end{aligned} \quad (10.22)$$

We conclude that $\|\mathbf{f}\|_{\text{TV}-\ell_p} = \|\mathbf{r}_\theta \mathbf{f}\|_{\text{TV}-\ell_p}$ if and only if $p = 2$, which proves that the TV- ℓ_p norm is not rotation invariant for $p \neq 2$. \square

Finally, \mathcal{R} being a norm, it is scale equivariant (homogeneity property).

10.2.3 Continuous-Domain Optimization Problem

The setting we described in this section is typical of a minimization problem with two terms. The first term—the data-fidelity term—ensures that the candidate curve $\mathbf{r}(t)$ is close to the points $\mathbf{p}[m]$. The second term, called regularization, introduces our *a priori* desiderata for the reconstructed curve. The importance of these two terms is weighted by

a parameter $\lambda > 0$. The solution set of the minimization problem is

$$\mathcal{V} \triangleq \arg \min_{\mathbf{r} \in \mathcal{X}_L(\mathbb{T}_M)} \left(\sum_{m=0}^{M-1} \|\mathbf{r}(t)|_{t=m} - \mathbf{p}[m]\|_2^2 + \lambda \mathcal{R}(L\{\mathbf{r}\}) \right), \quad (10.23)$$

where the search space \mathcal{X}_L is defined as

$$\mathcal{X}_L(\mathbb{T}_M) \triangleq \{\mathbf{r} \in \mathcal{S}'(\mathbb{T}_M)^2 : \mathcal{R}(L\{\mathbf{r}\}) < +\infty\}, \quad (10.24)$$

and corresponds to the space of BV^2 curves in [285, Section 2.1] when $L = D^2$. The data-fidelity term in (10.23) penalizes the Euclidean distance between the sample $\mathbf{r}(t)|_{t=m}$ of the curve and the point $\mathbf{p}[m]$ for every $m = 0, \dots, M - 1$. The fact that \mathbf{r} is sampled uniformly along the parameter axis encourages the reconstructed curve to be parametrized by its curvilinear abscissa, promoting the arc length of $\mathbf{r}(t)$ to be a linear function of the parameter t . The underlying assumption behind this statement is that the points $\mathbf{p}[m]$ are spread approximately uniformly along the curve. This is an important assumption, since the regularization term in (10.23) involves the derivatives of $\mathbf{r}(t)$ and thus heavily depends on the choice of parametrization. In that respect, the curvilinear abscissa is a desirable choice. In practice, it often results in rough curves being penalized heavily by our regularization, which other parametrizations may fail to achieve [255].

Our representer theorem (Theorem 10.2) specifies the form of the solution of (10.23).

Theorem 10.2. *Let $L = D^{N_d}$ with $N_d \geq 2$. Then, the global minimizer of (10.23) can be achieved by a periodic L-spline curve \mathbf{r}^* with at most $K \leq 2M + 2$ knots. More precisely, we have that*

$$L\{\mathbf{r}^*\} = \sum_{k=0}^{K-1} \mathbf{a}_k \text{III}_M(\cdot - t_k) \quad (10.25)$$

for some pairwise-distinct knot locations $t_k \in \mathbb{T}_M$ and amplitude vectors $\mathbf{a}_k \in \mathbb{R}^2$.

The proof of Theorem 10.2 is given in Appendix B.5.

Remark 10.1. *Theorem 9.1 remains valid for more general operators L , namely any spline-admissible operator in the sense of [22, Definition 2] whose null space includes constant functions, i.e., $L\{1\} = 0$, and whose Green's function g_L (Definition (2.32)) is continuous.*

Theorem 10.2 states that the solution set \mathcal{V} contains periodic L-splines. Even though our work uses a mixed (TV- ℓ_2) norm as regularization, this result is reminiscent of Theorems 3.3 and 3.4, which prove that inverse problems with TV regularization have spline solutions.

10.3 Exact Discretization

10.3.1 Polynomial B-Splines

Following the works of Part II, Theorem 10.2 motivates our discretization of the continuous-domain Problem (10.23) over the space of periodic cardinal L-splines, *i.e.* with integer knot spacing ($t_{k+1} - t_k = 1$ in (10.1)). When $L = D^{N_d}$, as exposed in Section 3.1, L-splines are piecewise polynomials of degree $N_d - 1$. For simplicity, we consider symmetric N_d th-order—or $(N_d - 1)$ th-degree—B-splines

$$\beta^{(N_d-1)} \triangleq \beta_L \left(\cdot - \frac{N_d}{2} \right), \quad (10.26)$$

where β_L is the causal B-spline of $L = D^{N_d}$ defined in (3.9) (see Section 3.1.4). B-splines are the functions with the smallest support within the space of cardinal periodic L-splines, with a support included in $[-\frac{N_d}{2}, \frac{N_d}{2}]$. This finite-support property is particularly advantageous for numerical efficiency.

For centered B-splines, we use grid points $t_k = k - \frac{N_d}{2}$ for $k \in \mathbb{Z}$. For an even N_d , we then have an integer grid, while an odd N_d gives a half-integer grid. Additionally, the jump amplitudes $a[k]$ in (10.1) for $s = \beta^{(N_d-1)}$ are denoted by $d_L[k]$ (defined in Table 3.1), which is the finite-difference digital filter of order N_d . The support of $d_L[k]$ is $\{0, \dots, N_d\}$, and using (3.23), we have

$$L\{\beta^{(N_d-1)}\} = \sum_{k=0}^{N_d} d_L[k] \text{III} \left(\cdot - \left(k - \frac{N_d}{2} \right) \right). \quad (10.27)$$

10.3.2 Discrete Formulation

As suggested by Theorem 10.2, we take the stance of recasting the continuous-domain problem in (10.23) as a finite-dimensional optimization problem by restricting the search space to periodic L-splines with knots on a uniform grid. This allows us to effectively reduce the complexity of our algorithmic resolution. To do so, we describe our closed curves $\mathbf{r}(t)$ as linear combinations of N shifts of a M -periodic basis function φ_M . Following Section 10.3.1, we choose φ_M to be the M -periodization and h -dilation of the B-spline generator $\varphi \triangleq \beta^{(N_d-1)}$, so that $\varphi_M(t) \triangleq \sum_{k \in \mathbb{Z}} \varphi(\frac{t-Mk}{h})$, where $h \triangleq \frac{M}{N}$ is the grid step size. These basis functions are weighted by two vectors of coefficients $\mathbf{c}_x = (c_x[n])_{n=0}^{N-1}$ and $\mathbf{c}_y = (c_y[n])_{n=0}^{N-1}$. Finally, the weighted functions are shifted by multiples of the grid size h , so that our curves are parametrized as

$$\mathbf{r}(t) = \begin{bmatrix} x(t) \\ y(t) \end{bmatrix} = \begin{bmatrix} \sum_{n=0}^{N-1} c_x[n] \varphi_M(t - nh) \\ \sum_{n=0}^{N-1} c_y[n] \varphi_M(t - nh) \end{bmatrix}. \quad (10.28)$$

10.3.3 Discrete Implementation

Our choice (10.28) of curve parametrization allows us to optimize solely over the coefficients $\mathbf{c}_x, \mathbf{c}_y \in \mathbb{R}^N$ of the two curve components. We implement a system matrix that samples the components $x(t)$ and $y(t)$ of the curve in (10.28) at integer values of t when applied to the spline coefficients \mathbf{c}_x and \mathbf{c}_y . This yields the matrix $\mathbf{H} \in \mathbb{R}^{M \times N}$ with

$$[\mathbf{H}]_{m,n} \triangleq \varphi_M(m - nh). \quad (10.29)$$

The regularization operator L becomes a circulant regularization matrix $\mathbf{L} \in \mathbb{R}^{N \times N}$ composed of shifted versions of the sequence d_L (see Table 3.1). The regularization matrix \mathbf{L} is therefore constructed as

$$[\mathbf{L}]_{m,n} \triangleq \frac{1}{h^\alpha} d_L[(m - n) \bmod N]. \quad (10.30)$$

Our mixed-norm regularization involves, in the discrete setting, an $\ell_1 - \ell_2$ norm given by

$$\left\| \begin{bmatrix} \mathbf{f}_1 & \mathbf{f}_2 \end{bmatrix} \right\|_{\ell_1 - \ell_2} \triangleq \sum_{n=0}^{N-1} \sqrt{(\mathbf{f}_1[n])^2 + (\mathbf{f}_2[n])^2}, \quad (10.31)$$

for $\mathbf{f}_1, \mathbf{f}_2 \in \mathbb{R}^N$. Indeed, we have that (see Theorem 10.1):

$$\left\| \begin{bmatrix} L\{x\} & L\{y\} \end{bmatrix} \right\|_{\text{TV} - \ell_2} = \left\| \mathbf{L} \begin{bmatrix} \mathbf{c}_x & \mathbf{c}_y \end{bmatrix} \right\|_{\ell_1 - \ell_2}. \quad (10.32)$$

Our discrete optimization problem therefore aims at finding \mathbf{c}_x and \mathbf{c}_y such that

$$\arg \min_{\mathbf{c}_x, \mathbf{c}_y \in \mathbb{R}^N} \left(\left\| \begin{bmatrix} \mathbf{H} & \mathbf{0} \\ \mathbf{0} & \mathbf{H} \end{bmatrix} \begin{bmatrix} \mathbf{c}_x \\ \mathbf{c}_y \end{bmatrix} - \begin{bmatrix} \mathbf{p}_x \\ \mathbf{p}_y \end{bmatrix} \right\|_2^2 + \lambda \left\| \mathbf{L} \begin{bmatrix} \mathbf{c}_x & \mathbf{c}_y \end{bmatrix} \right\|_{\ell_1 - \ell_2} \right). \quad (10.33)$$

10.3.4 Algorithmic Resolution

To solve Problem (10.33), we use the ADMM solver [19] (see Appendix A) as implemented in the GlobalBioIm Matlab library [118] dedicated to solving inverse problems.

10.4 Extension to Hybrid B-Spline Contours

10.4.1 Motivation and Continuous Model

In Section 10.2, we have presented a model and its implementation that reconstructs of a sparse curve using splines as basis functions. However, not all curves can be faithfully represented with a single type of spline. We propose to cater to this by modeling our closed function as a sum of two components $\mathbf{r}(t) = \mathbf{r}_1(t) + \mathbf{r}_2(t)$. Similarly to the single-spline

setting, we have M points $\mathbf{p}[m] = (p_x[m], p_y[m])$, $m = 0, \dots, M - 1$. Hence, we again have that \mathbf{r} is M -periodic with $t \in \mathbb{T}_M$. Following the formulation for one-dimensional signals in Chapter 5 and extending it to two dimensions, we consider continuous problems of the form

$$\mathcal{V}_{\text{hyb}} \triangleq \arg \min_{\substack{\mathbf{r}_i \in \mathcal{X}_{L_i}(\mathbb{T}_M) \\ \mathbf{r}_1(0) = \mathbf{0}}} \left(\sum_{m=0}^{M-1} \|\mathbf{r}_1(t)|_{t=m} + \mathbf{r}_2(t)|_{t=m} - \mathbf{p}[m]\|_2^2 + \lambda_1 \|\mathbf{L}_1\{\mathbf{r}_1\}\|_{\text{TV}-\ell_2} + \lambda_2 \|\mathbf{L}_2\{\mathbf{r}_2\}\|_{\text{TV}-\ell_2} \right), \quad (10.34)$$

where $\lambda_1, \lambda_2 > 0$ are the two regularization parameters weighting the two regularization terms, and $\mathbf{L}_1 = \mathbf{D}^{N_{d,1}}$ and $\mathbf{L}_2 = \mathbf{D}^{N_{d,2}}$ are derivative operators of different orders $N_{d,1} \neq N_{d,2}$.

Remark 10.2. *The constraint $\mathbf{r}_1(0) = \mathbf{0}$ is necessary to handle the ill-posedness of the problem. Indeed, without this constraint, for any solution $(\mathbf{r}_1(t), \mathbf{r}_2(t))$ of Problem (10.34), the pair $(\mathbf{r}_1 + \mathbf{v}, \mathbf{r}_2 - \mathbf{v})$, where \mathbf{v} is an arbitrary constant vector, would clearly also be a solution. This implies that the solution set would be unbounded, which can be problematic for numerical implementations. The constraint $\mathbf{r}_1(0) = \mathbf{0}$ resolves this issue without any restriction on the constructed curve, since any constant offset can be included in the \mathbf{r}_2 component. We refer to Chapter 5, where a similar strategy is used, for more details. Note that technically, $\mathbf{L}_1 = \mathbf{D}^{N_{d,1}}$ is sampling-admissible—making the constraint $\mathbf{r}_1(0) = \mathbf{0}$ licit—only when $N_{d,1} \geq 2$. For $N_{d,1} = 1$, using the technique outlined in Remark 5.2, the constraint can be implemented as $\mathbf{r}_1(0^+) = \mathbf{0}$. We skip these technicalities here for the sake of simplicity.*

We now prove a representer theorem that gives a parametric form of the solution of Problem (10.34). It is a multicomponent extension of Theorem 10.2, and is reminiscent of Theorem 5.1.

Theorem 10.3. *There exists a global minimizer \mathbf{r}^* of (10.34) that can be decomposed as $\mathbf{r}^* = \mathbf{r}_1^* + \mathbf{r}_2^*$, where \mathbf{r}_i^* are periodic L_i -splines (see (10.1)) with K_i knots, with $i = 1, 2$. Moreover, the total number of knots of \mathbf{r}^* is bounded by $K_1 + K_2 \leq 2M + 2$.*

The proof of Theorem 10.3 is given in Appendix B.6.

10.4.2 Discretization and Implementation

As in Section 10.3, we derive a discrete setting by using two sets of B-spline basis functions matched to their corresponding regularization operators. Given a grid of step size h , we

consider closed M -periodic curves $\mathbf{r}(t) = \mathbf{r}_1(t) + \mathbf{r}_2(t)$ such that, for $i = 1, 2$, we have

$$\mathbf{r}_i(t) = \begin{bmatrix} x_i(t) \\ y_i(t) \end{bmatrix} = \begin{bmatrix} \sum_{n=0}^{N-1} c_x^i[n] \varphi_M^i(t - nh) \\ \sum_{n=0}^{N-1} c_y^i[n] \varphi_M^i(t - nh) \end{bmatrix}. \quad (10.35)$$

The two regularization operators are set to $L_i = D^{N_{d,i}}$, with $D^{N_{d,i}}$ the derivative operator of order $N_{d,i}$ for $i = 1, 2$. As in Section 10.3, these operators lead to spline solutions (see Chapter 5). We hence set $\varphi_M^i \triangleq \beta^{(N_{d,i}-1)}$ for $i = 1, 2$ and with $N_{d,1} < N_{d,2}$.

This choice of curve allows us to optimize over the coefficients $\mathbf{c}_x^1 = (c_x^1[n])_{n=0}^{N-1}$, $\mathbf{c}_y^1 = (c_y^1[n])_{n=0}^{N-1}$, $\mathbf{c}_x^2 = (c_x^2[n])_{n=0}^{N-1}$, and $\mathbf{c}_y^2 = (c_y^2[n])_{n=0}^{N-1}$. As in Section 10.2, we can define \mathbf{H}_i , the system matrices that samples the continuous curves \mathbf{r}_i , as well as the corresponding regularization matrices \mathbf{L}_i . We construct \mathbf{H}_i and \mathbf{L}_i as in (10.29) and (10.30), respectively. Finally, we cater to the constraint $\mathbf{r}_1(0) = \mathbf{0}$ by enforcing that $(\mathbf{c}_x^1 * \mathbf{b}_{\varphi^1})[0] = 0$ and $(\mathbf{c}_y^1 * \mathbf{b}_{\varphi^1})[0] = 0$, where $b_{\varphi^1}[k] \triangleq \varphi_M^1(kh)$ and $*$ denotes the cyclic discrete convolution.

Our discrete hybrid optimization problem takes the form

$$\begin{aligned} \arg \min_{\substack{\mathbf{c}_x^1, \mathbf{c}_y^1, \\ \mathbf{c}_x^2, \mathbf{c}_y^2 \in \mathbb{R}^N}} & \left(\left\| \mathbf{H}_{\text{hyb}} \begin{bmatrix} \mathbf{c}_x^1 \\ \mathbf{c}_y^1 \\ \mathbf{c}_x^2 \\ \mathbf{c}_y^2 \end{bmatrix} - \begin{bmatrix} \mathbf{p}_x \\ \mathbf{p}_y \end{bmatrix} \right\|_2^2 + \lambda_1 \left\| \mathbf{L}_1 \begin{bmatrix} \mathbf{c}_x^1 & \mathbf{c}_y^1 \end{bmatrix} \right\|_{\ell_1 - \ell_2} + \lambda_2 \left\| \mathbf{L}_2 \begin{bmatrix} \mathbf{c}_x^2 & \mathbf{c}_y^2 \end{bmatrix} \right\|_{\ell_1 - \ell_2} \right), \\ \text{s.t.} & \quad (\mathbf{c}_x^1 * \mathbf{b}_{\varphi^1})[0] = 0, \quad (\mathbf{c}_y^1 * \mathbf{b}_{\varphi^1})[0] = 0, \end{aligned} \quad (10.36)$$

where $\mathbf{H}_{\text{hyb}} \triangleq \begin{bmatrix} \mathbf{H}_1 & \mathbf{0} & \mathbf{H}_2 & \mathbf{0} \\ \mathbf{0} & \mathbf{H}_1 & \mathbf{0} & \mathbf{H}_2 \end{bmatrix}$. As in Section 10.3, we use the ADMM solver to find a solution to Problem (10.36) and GlobalBioIm [118] to implement our algorithms.

10.5 Experimental Results

We evaluate the distance between the constructed curves and the contour points through the quadratic fitting error (QFE) defines as

$$\text{QFE} = \frac{1}{M} \sum_{m=0}^{M-1} \|\mathbf{r}(t)|_{t=m} - \mathbf{p}[m]\|_2^2. \quad (10.37)$$

It is noteworthy that the QFE can be used at the same time in the single-spline setting and in the hybrid setting. Indeed, by replacing the hybrid curve $\mathbf{r} = \mathbf{r}_1 + \mathbf{r}_2$ in (10.37), we obtain a QFE that is consistent with the data-fidelity term in (10.34).

For computational reasons, we chose the lowest resolution, i.e. the largest grid size h , that allowed us to solve the problem in a satisfactory way, thus effectively making h a hyperparameter. In this work, the number of knots N was chosen so that it matched the order of magnitude of the number of data points. It is important to note that increasing N , thus splitting the grid, can only improve the solution in terms of cost.

10.5.1 Rotation Invariance

To verify that our regularization norm is truly rotation-invariant, we apply a planar rotation of angle θ to our data before we reconstruct the curve with the regularization operator $L = D^2$. We have added to the data a Gaussian perturbation with a signal-to-noise ratio (SNR) of 47.28 dB. We compare the curves reconstructed with our regularization to the curves resulting from the $(TV-\ell_1)$ regularization of Definition 10.1. Indeed, ℓ_1 regularization is widely used in the signal-processing community as a sparsifying prior. To do so, we choose λ in the nonrotated $(TV-\ell_1)$ regularized curve (Figure 10.1c) so that the QFE matches the QFE from the nonrotated RI-TV regularized curve (Figure 10.1a). When rotating the measurements, we adjust λ again so that the QFE of the $(TV-\ell_1)$ -regularized curve on the rotated points matches the one of the curve constructed with RI-TV regularization with rotated points. We see in Figure 10.1 that the RI-TV-regularized problem provides the same solution regardless of θ . Indeed, the knot locations do not differ between Figures 10.1a and 10.1b, nor does the number K of knots. This is not the case for the purely $(TV-\ell_1)$ -regularized problem. Not only are the knot locations different when a rotation is applied to the measurements, but the number K of knots varies with θ as well as the QFE of the curve. Additionally, one needs to adapt λ to obtain the same QFE between the constructed curves on rotated and nonrotated measurements.

10.5.2 Resilience to Contour Imprecisions

A beneficial feature derived from the enforcement of joint sparsity in the two curve components is resilience of our reconstructions to imprecisions in the contour points. Indeed, when we expect our data to be imprecise, we can choose to increase the regularization parameter λ at the cost of data fidelity, as the curve cannot rely as much on the data. Particularly, when the regularizer is TV-based, an increase in λ tends to smoothen sharp variations. This is visible in Figure 10.2, where several curves have been reconstructed using linear B-splines β^1 . Figures 10.2a, 10.2b, and 10.2c are reconstructions of increasingly inaccurate measurements using RI-TV. Figures 10.2d, 10.2e, and 10.2f depict reconstructions resulting from a sparsifying regularization without coupling $(TV-\ell_1)$, using a λ tuned so that the QFE matches the QFE of the curves regularized by RI-TV. When TV regularization is used, and as the contour becomes more inaccurate, the number K of knots drastically increases and the angles are deformed. On the contrary, for the reconstructions in Figures 10.2a, 10.2b, even as the noise and λ increase, the

number K of knots remains unchanged and the angles are fairly well preserved.

10.5.3 Hybrid Setting Applications

The single-component framework described in Sections 10.2 and 10.3 only allows for the use of one kind of B-spline per curve. However, when the contour under consideration is composed of smooth sections and kinks, no single type of B-spline can provide a faithful and sparse reconstruction. An example of curve fitting that depicts this problem is given in Figure 10.3. We reconstructed the contour using first β^1 and β^3 as basis functions, giving piecewise-linear and piecewise-cubic curves in Figures 10.3a and 10.3b, respectively. Figure 10.3c contains a reconstruction under the hybrid setting $L_1 = D^2$ and $L_2 = D^4$, thus producing a curve that has both a linear and a cubic component. While all three reconstructions yield the same QFE with respect to the data, the hybrid curve has by far the smallest number of knots. Moreover, upon visual inspection, the hybrid curve in Figure 10.3c portrays the most faithful reconstruction, as it does round neither the angles nor the straight lines, nor does it straighten the smooth sections.

We can observe the effect of the parameters λ_1 and λ_2 on the constructed curve when the hybrid reconstruction setting is applied to real contour points for a constant ratio of knots $\frac{K_1}{K_2} = 0.86$. In Figure 10.4, as λ_1 and λ_2 increase, the total number K of knots decreases and the curve becomes more stylized. In addition, for all values of λ_1 and λ_2 , our algorithm preserves the kinks of the contour while mimicking its smooth segments.

10.6 Conclusion

We have introduced a framework to reconstruct sparse continuous curves from a list of possibly inaccurate contour points using an RI-TV regularization norm. We have proved that an optimal solution to our minimization problem is a curve that uses splines as basis functions, and we have leveraged this result to provide an exact discretization of the continuous-domain framework using B-splines. Furthermore, we have extended our formulation to reconstruct curves with components of distinct smoothness properties using hybrid splines. We have experimental confirmation of the rotation invariance of our regularizer. In addition, our experimental results demonstrate that our formulation yields sparse reconstructions that are close to the data points even when their noise increases, unlike other regularization methods. Finally, our hybrid-curve experiments demonstrate that we are able to faithfully reconstruct contours with varying smoothness properties with a low number of knots.

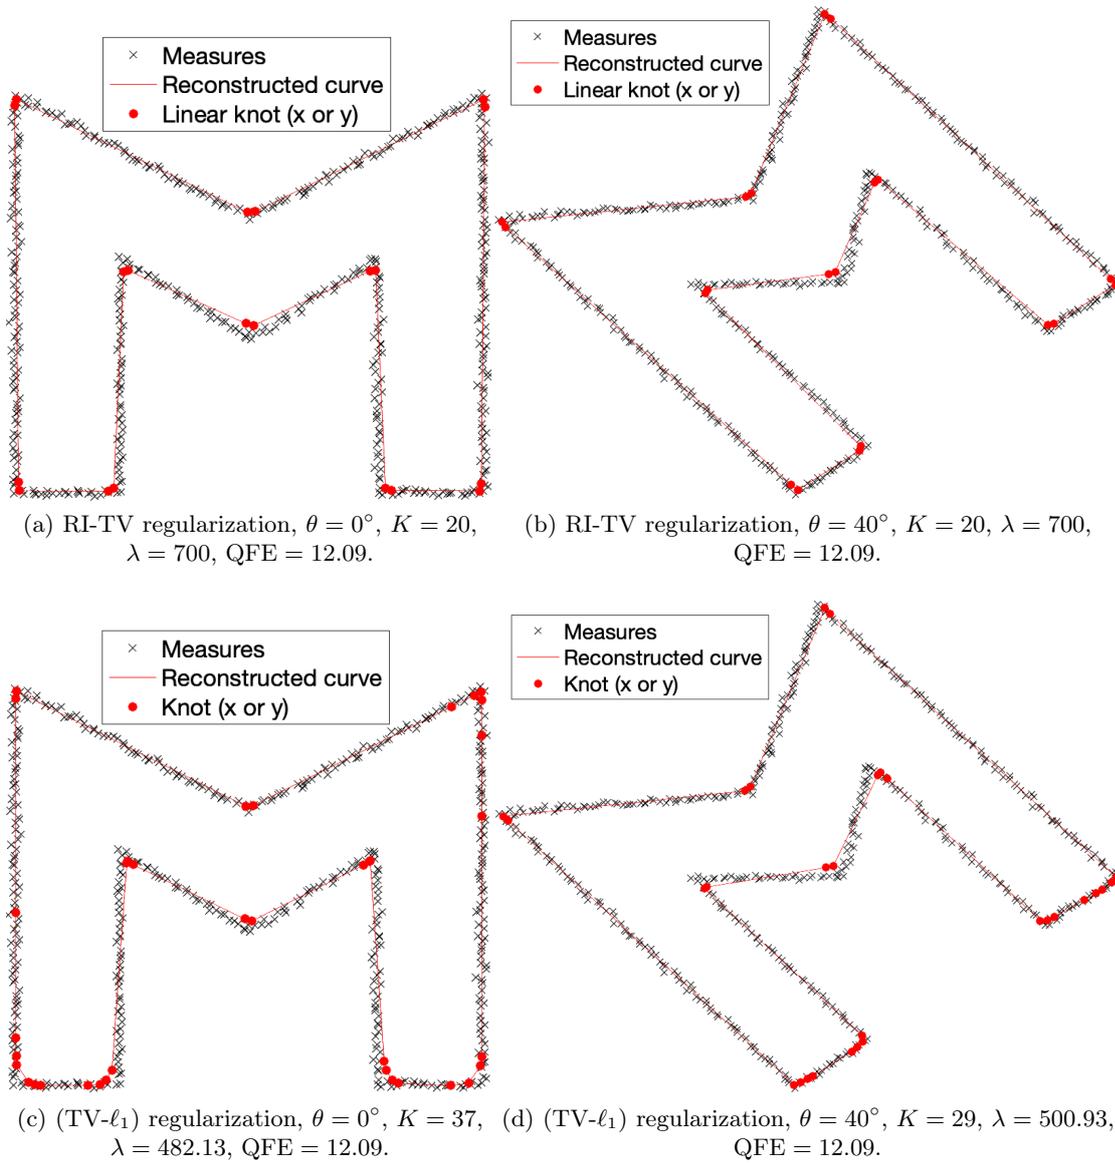


Figure 10.1 – Solutions as a function of the rotation angle θ for RI-TV regularization and $(TV-\ell_1)$ regularization for a same contour. $M = 488$, grid step size $h = 1.9062$, $\varphi = \beta^1$.

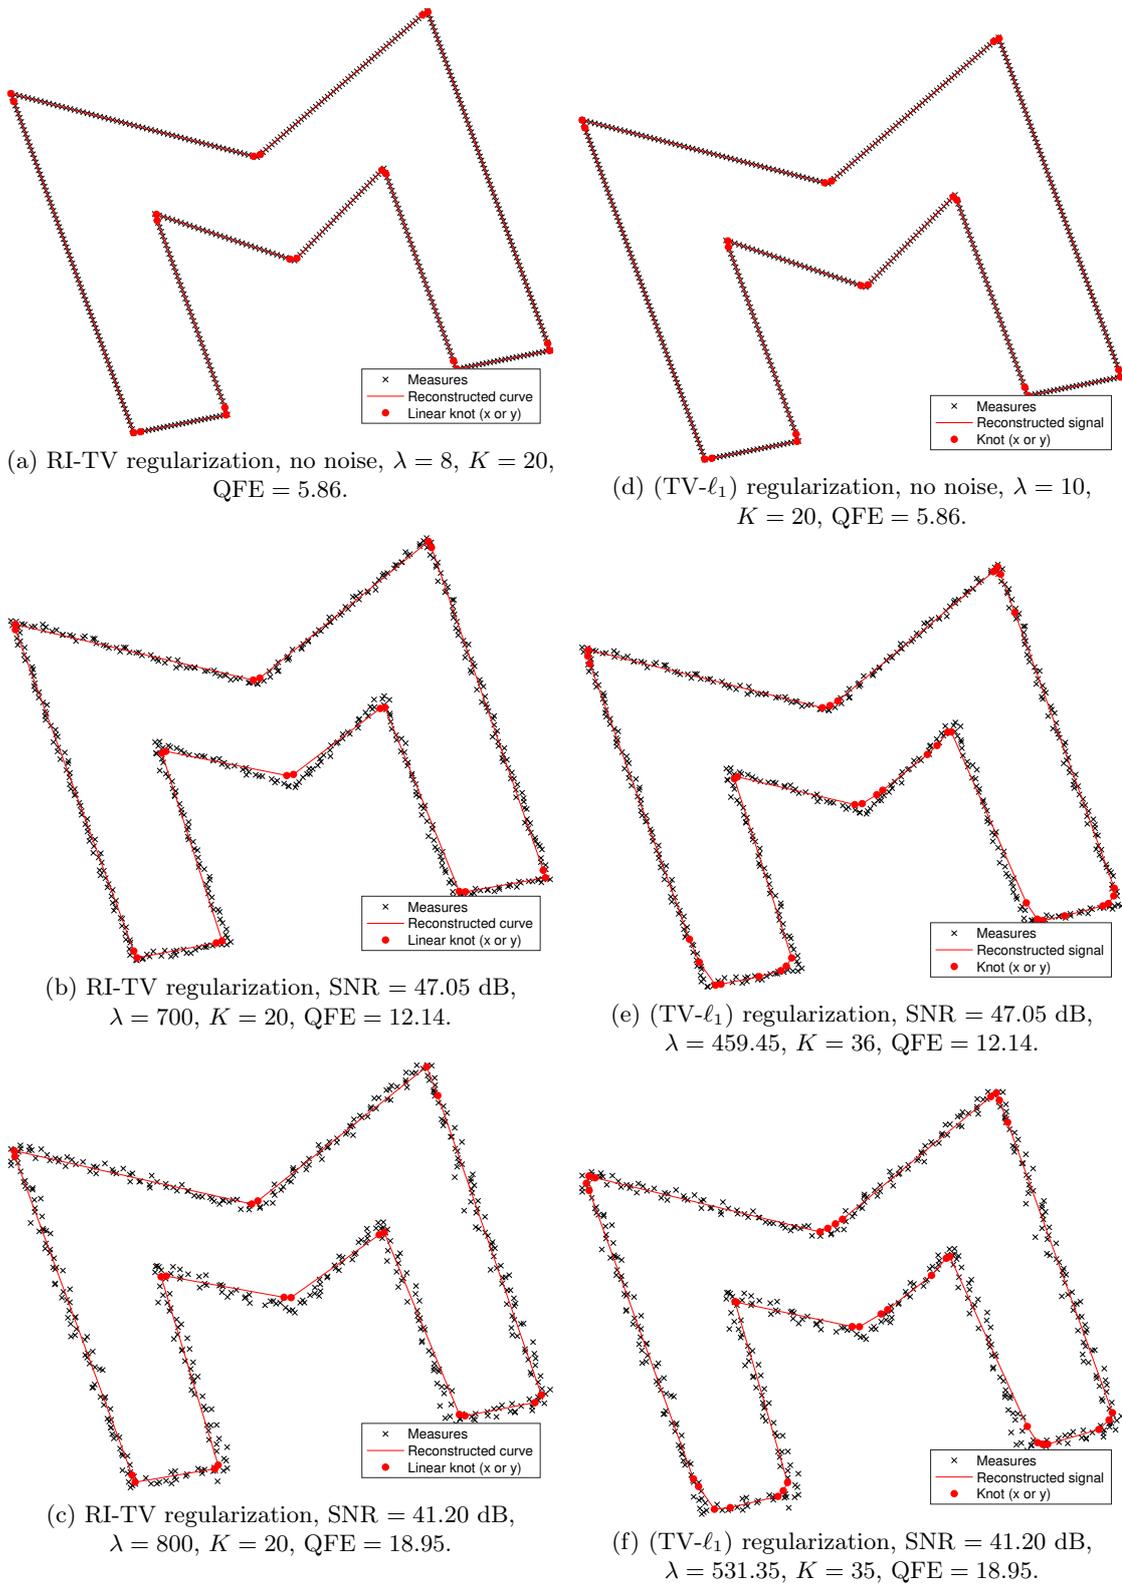


Figure 10.2 – Resilience to noise for RI-TV regularization and $(TV-\ell_1)$ regularization.

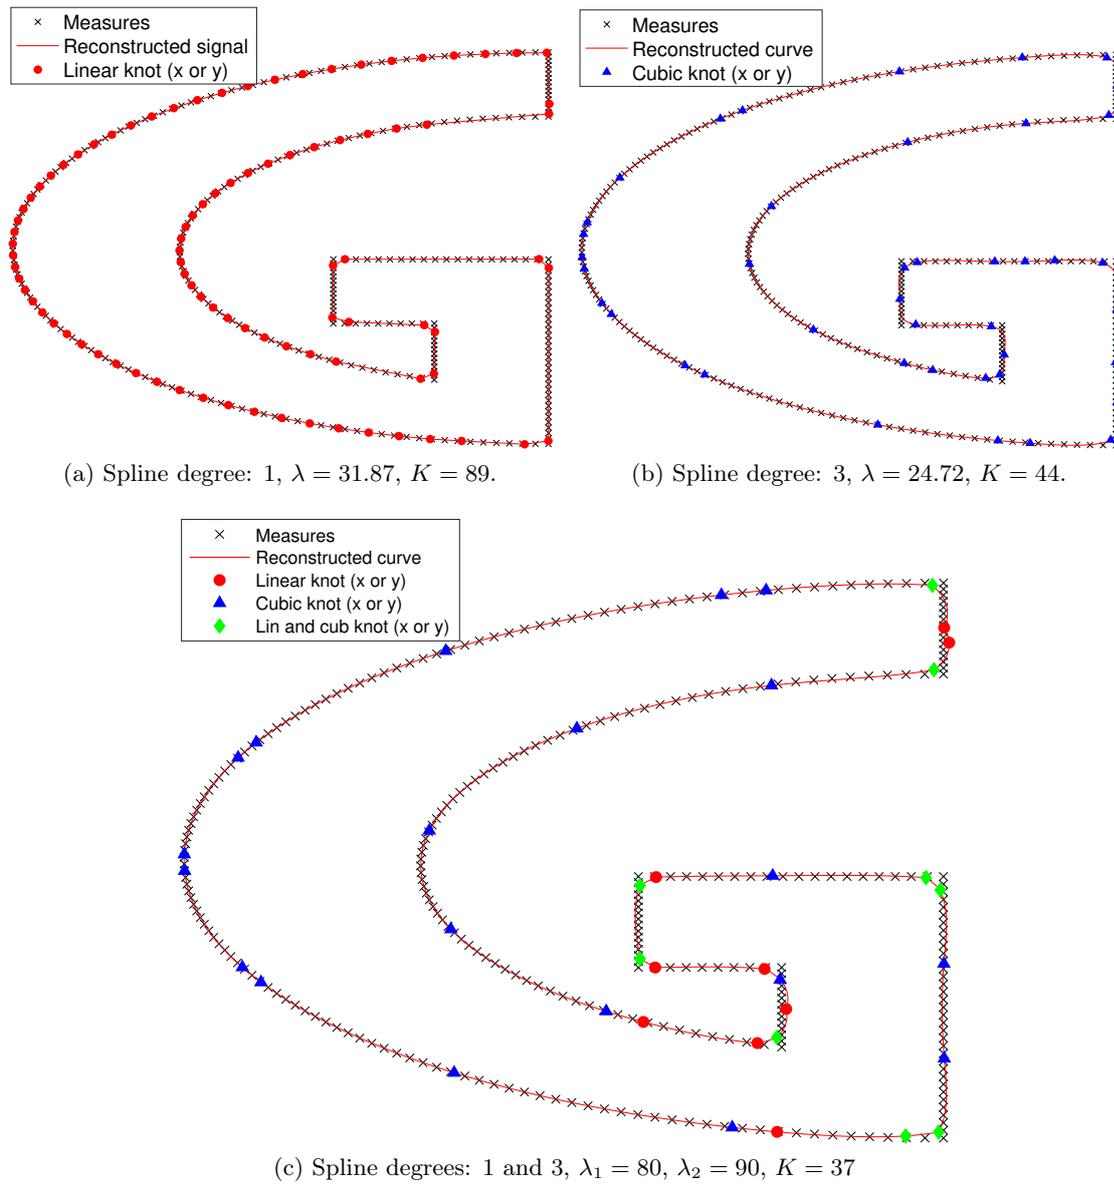


Figure 10.3 – Noiseless curve reconstruction with a single spline, a hybrid setting, and RI-TV regularization. All three reconstructions have a constant QFE = 8.88.

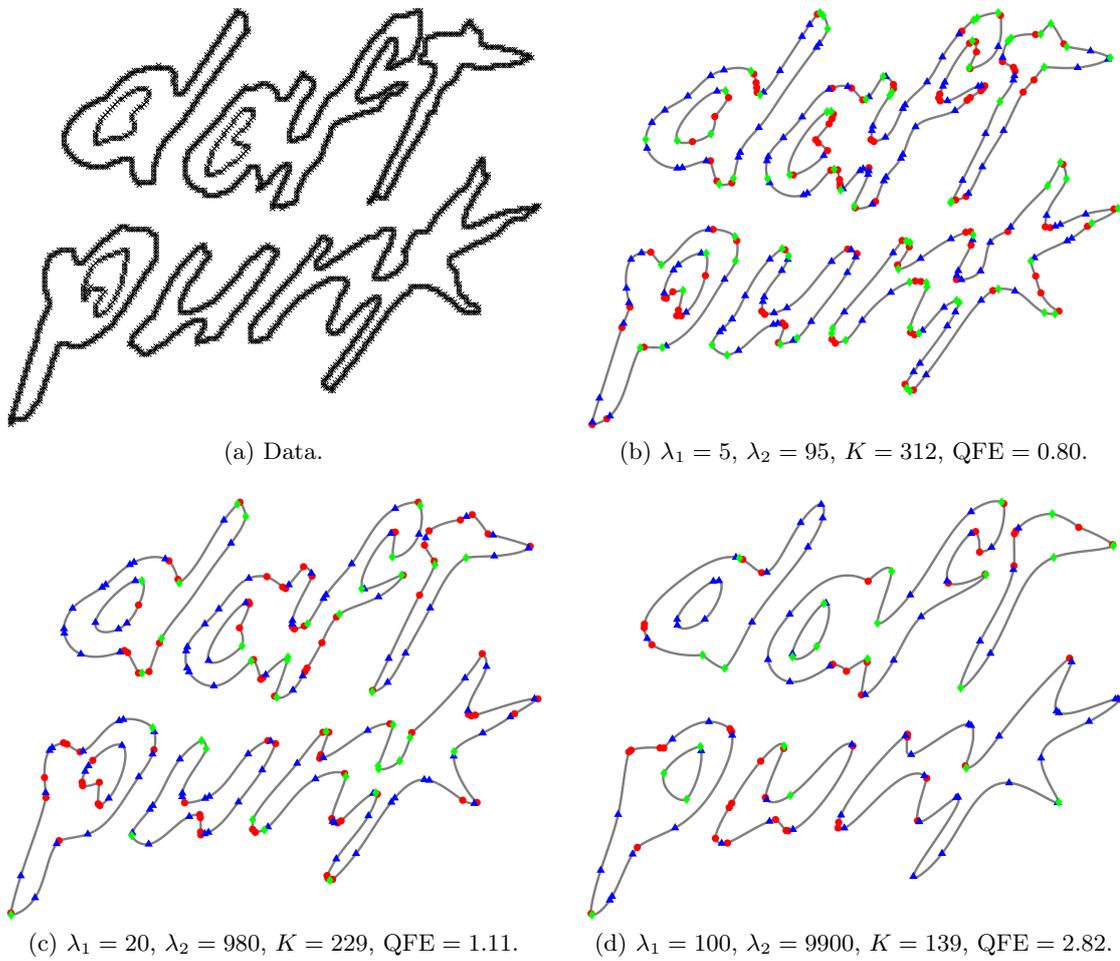


Figure 10.4 – Effect of λ_1 and λ_2 on the reconstructed hybrid curve for $M = 2714$ under RI-TV regularization. The reconstructed curve is represented by the solid line. The round markers and the triangular markers indicate the location of the linear and cubic knots, respectively. The diamond-shaped markers indicate the superimposition of a linear knot and a cubic knot. The data are extracted from the official Daft Punk logo (source: https://en.wikipedia.org/wiki/Daft_Punk).

11 Hessian Splines for Scanning-Transmission X-Ray Microscopy

This chapter is based on the following publication [286]:

T. Debarre, B. Watts, B. Rösner, and M. Unser, “Hessian Splines for Scanning Transmission X-Ray Microscopy”, in *Proceedings of the Seventeenth IEEE International Symposium on Biomedical Imaging (ISBI’20)*, Iowa City IA, USA, Apr. 2020, pp. 199–202.

11.1 Introduction

Scanning transmission X-ray microscopy (STXM) is a non-invasive microscopy technique that uses X-ray spectroscopy to generate contrast based on near-edge X-ray absorption fine structure (NEXAFS) spectroscopy or associated dichroism to quantitatively map material properties such as chemical oxidation state, molecular structure and orientation, and magnetization at the nanoscale [287]. A Fresnel zone plate is used to focus the X-ray beam onto a small region of the sample (a pixel), and the transmitted beam intensity is measured while the sample is raster-scanned in a rectangular array in order to produce a 2D image. With recent advances in the design of the zone plate, X-ray spot sizes well below 10 nm can be achieved [288]. However, demonstrating STXM images with a resolution below 10 nm also requires similarly high precision in positioning the X-ray beam on the sample, which is challenging due to vibrations in the instrument. This imprecision leads to an off-the-grid scanning pattern: in fact, when measuring images close to the resolution limit, the displacement error can easily be larger than the spacing between the array points. This error can be measured using a heterodyne laser interferometer, with spatial resolution of 0.3 nm. The current state-of-the-art resolution of 7 nm for STXM is thus achieved by regriding the measured intensity values using linear interpolation [289].

In this work, we propose a more elaborate interpolation method using a continuous-domain inverse problem formulation. For discretization purposes, we reconstruct the image as a parametric continuous 2D function using a spline-based generalized interpolation model [290]. We then formulate the image-reconstruction task as an optimization problem over

the spline coefficients so as to minimize the discrepancy between the measured data and the reconstructed images.

In order to improve the robustness of the reconstruction, we add a regularization term to the cost functional. In effect, this enables us to reduce the effect of the noise in the measurements, as well as the uncertainty in the interferometer measurements. Our algorithm involves a second-order Hessian nuclear-norm regularization, which has been successfully applied to many imaging problems [27, 291, 292]. The key feature of this regularization is that it enjoys many advantages of the popular total-variation (TV) seminorm [28] such as convexity, or translation and scale invariance, without suffering from the staircasing effect which typically plagues TV-based methods. As opposed to the purely discrete framework of [27], we compute the Hessian in the continuous domain, which yields a new brand of splines that we coin *Hessian splines*. The resulting optimization problem is solved using the standard ADMM algorithm [19]. We illustrate the effectiveness of our approach on a simulated ground-truth image, by showing that it outperforms linear interpolation used in state-of-the-art STXM reconstructions. We also apply our algorithm to real high-resolution STXM data. Note that our approach is pertinent for any imaging modality with nonuniform measurements in which the displacement error is nonnegligible.

11.2 Imaging Model

In this work, we view 2D images as continuous-domain functions $f : \Omega \rightarrow \mathbb{R}$, where $\Omega \subset \mathbb{R}^2$ is the (bounded) image domain. Without loss of generality, we assume that the pixels are located on the integer grid, *i.e.*, $\Omega \cap \mathbb{Z}^2$.

11.2.1 Reconstruction Basis

Following the generalized interpolation approach of [290], we parametrize the reconstructed signal $f : \Omega \rightarrow \mathbb{R}$ in a spline basis as

$$f(\mathbf{x}) = \sum_{\mathbf{k} \in \Omega_\varphi} c[\mathbf{k}] \varphi(\mathbf{x} - \mathbf{k}), \quad (11.1)$$

where $\mathbf{x} = (x_1, x_2) \in \mathbb{R}^2$, $(c[\mathbf{k}])_{\mathbf{k} \in \Omega_\varphi}$ are the spline coefficients and φ is a suitable spline generating function. The domain $\Omega_\varphi \triangleq \{\mathbf{k} \in \mathbb{Z}^2 : \text{Supp}(\varphi(\cdot - \mathbf{k})) \cap \Omega \neq \emptyset\}$ simply selects the integer grid points \mathbf{k} such that the support of the corresponding basis function $\varphi(\cdot - \mathbf{k})$ intersects the image domain Ω . In this work, we choose the tensor-product cubic B-spline

$\varphi(\mathbf{x}) \triangleq \beta^3(x_1)\beta^3(x_2)$, where the univariate centered cubic B-spline is defined as

$$\beta^3(x) \triangleq \begin{cases} \frac{2}{3} - |x|^2 + \frac{|x|^3}{2}, & 0 \leq |x| < 1 \\ \frac{(2-|x|)^3}{6}, & 1 \leq |x| < 2 \\ 0, & \text{otherwise.} \end{cases} \quad (11.2)$$

Our choice of cubic B-splines is motivated by their simplicity and their popularity in applications [24, 25, 105, 293], in part due to their short support. For additional background on B-splines, we refer to Section 3.1.4. Moreover, they are twice differentiable, which is required to compute the Hessian of f . Since the basis function φ is supported in a square of size 4×4 , there is a finite number of spline coefficients $c[\mathbf{k}]$ such that $\mathbf{k} \in \Omega_\varphi$. We denote these coefficients by $\mathbf{c} \in \mathbb{R}^N$, where $N \triangleq \#\Omega_\varphi$ (the cardinality of the set).

11.2.2 Forward Model

In STXM imaging, the task is to reconstruct a continuous-domain function $f_0 : \Omega \rightarrow \mathbb{R}$ (the ground-truth image) based on the measured data $\mathbf{y} \in \mathbb{R}^M$, where M is the number of measurements, *i.e.*, pixels. The data are acquired via the forward model $\boldsymbol{\nu} : f \mapsto (f(\mathbf{t}^1), \dots, f(\mathbf{t}^M)) \in \mathbb{R}^M$, where $\mathbf{t}^m = (t_1^m, t_2^m) \in \mathbb{R}^2$ are the sampling locations measured by the interferometer. The data is corrupted by some additive noise $\mathbf{n} \in \mathbb{R}^M$, so that $\mathbf{y} = \boldsymbol{\nu}(f_0) + \mathbf{n}$.

Using the parametrization (11.1) and due to the linearity of $\boldsymbol{\nu}$, we specify the discrete forward model as the matrix $\mathbf{H} \in \mathbb{R}^{M \times N}$ such that $\mathbf{H}\mathbf{c} \triangleq \boldsymbol{\nu}(f)$, where the vector $\mathbf{c} \in \mathbb{R}^N$ collects the spline coefficients of f . Note that if the pixel grid used for the reconstruction is the same as the one intended by the hardware (which is a natural choice), N is slightly larger than M . This is due to the fact that the length of the support of β^3 is larger than 2, which leads to Ω_φ containing points outside of the image domain near the boundaries, *i.e.*, $\Omega_\varphi \not\subset \Omega$. Since there are more unknowns than data points, the problem is thus ill-posed.

11.2.3 Regularization

In order to increase the robustness of the reconstruction and to handle its aforementioned ill-posedness, we add a Hessian nuclear-norm regularization term to the cost function. This regularization was first introduced in [27], which proposes a discrete version of the functional $\mathcal{R}(f) \triangleq \int_\Omega \|\mathbf{L}_f(\mathbf{x})\|_* d\mathbf{x}$, where $\mathbf{L}_f(\mathbf{x})$ is the Hessian matrix of f at the location \mathbf{x} , and $\|\cdot\|_*$ is the nuclear norm (also known as the trace or 1-Schatten norm). The latter is defined as $\|\mathbf{M}\|_* \triangleq \sum_i |\sigma(\mathbf{M})_i|$, where the $\sigma(\mathbf{M})_i$ are the singular values of the matrix \mathbf{M} . The choice of a second-order differential operator—the Hessian—, as opposed to first-order for TV regularization, is designed to promote piecewise-linear reconstructions.

Indeed, planes (*i.e.*, first-order polynomials) induce no regularization cost as their Hessian is zero.

Here, we adapt the purely discrete setting of [27] to our continuous-domain representation (11.1). More precisely, instead of discretizing the Hessian with finite differences, we compute the continuous Hessian operator in terms of the spline coefficients \mathbf{c} of f . However, due to the nonlinearity of the singular value decomposition, computing the analytical expression of $\mathcal{R}(f)$ as a function of the coefficients \mathbf{c} proves challenging. However, computing the Hessian of f on the integer grid can be done efficiently with digital filtering using the B-spline kernels $b[k] \triangleq (\beta^3(k))_{k \in \mathbb{Z}}$, $b^{(1)}[k] \triangleq \left(\frac{d\beta^3}{dx}(k)\right)_{k \in \mathbb{Z}}$ and $b^{(2)}[k] \triangleq \left(\frac{d^2\beta^3}{dx^2}(k)\right)_{k \in \mathbb{Z}}$, characterized by their z -transforms $B(z) = \frac{z+4+z^{-1}}{6}$, $B^{(1)}(z) = \frac{z+z^{-1}}{2}$, and $B^{(2)}(z) = z - 2 + z^{-1}$ [106]. This yields the Hessian matrix

$$\mathbf{L}_c[\mathbf{k}] \triangleq \begin{pmatrix} \frac{\partial f(\mathbf{k})}{\partial^2 x_1} & \frac{\partial f(\mathbf{k})}{\partial x_1 x_2} \\ \frac{\partial f(\mathbf{k})}{\partial x_2 x_1} & \frac{\partial f(\mathbf{k})}{\partial^2 x_2} \end{pmatrix} = \begin{pmatrix} (c * h_{11})[\mathbf{k}] & (c * h_{12})[\mathbf{k}] \\ (c * h_{12})[\mathbf{k}] & (c * h_{22})[\mathbf{k}] \end{pmatrix}, \quad (11.3)$$

where $\mathbf{k} \in \mathbb{Z}^2$, and the $h_{i,j}[\mathbf{k}]$ are the tensor-product digital filters $h_{1,1}[\mathbf{k}] \triangleq b^{(2)}[k_1]b[k_2]$, $h_{1,2}[\mathbf{k}] \triangleq b^{(1)}[k_1]b^{(1)}[k_2]$, and $h_{2,2}[\mathbf{k}] \triangleq b[k_1]b^{(2)}[k_2]$. Here, $*$ denotes the 2D discrete convolution. We then approximate the integral with a sum over the pixel values, which yields our discretized regularization functional

$$\mathcal{R}_d(\mathbf{c}) \triangleq \sum_{\mathbf{k} \in \Omega \cap \mathbb{Z}^2} \sum_{i=1}^2 |\sigma(\mathbf{L}_c[\mathbf{k}])_i| \quad (11.4)$$

Note that we only consider the pixels inside the image domain Ω instead of Ω_φ , which contains additional grid points near the boundaries. This avoids unwanted boundary effects, due to the fact that the reconstructed signal (11.1) goes to zero near the boundaries. In particular, planes (*i.e.*, first-order polynomials) are not penalized using (11.4), which is a key desired feature of Hessian nuclear-norm regularization.

11.2.4 Positivity

Since STXM image represent photon counts which are positive by nature, the reconstructed image should satisfy $f(\mathbf{x}) \geq 0$ for any $\mathbf{x} \in \Omega$. This positivity constraint does not readily translate into a constraint on the coefficients \mathbf{c} of f —for instance, we do not necessarily have $c_n \geq 0$ for all n . A simple and easily-computable surrogate is to impose the positivity on the pixel locations $\Omega \cap \mathbb{Z}^2$. As before, this can be achieved with digital filtering with $f(\mathbf{k}) = (c * (b \otimes b))[\mathbf{k}]$, where $(b \otimes b)[\mathbf{k}] \triangleq b[k_1]b[k_2]$ is the tensor-product B-spline filter. We thus impose the positivity constraint on the coefficients $\mathbf{k} \in \Omega \cap \mathbb{Z}^2$, which can be

written in terms of the matrix $\mathbf{L}_{\text{pos}} \in \mathbb{R}^{P \times N}$ as

$$\mathbf{L}_{\text{pos}} \mathbf{c} \triangleq (f[\mathbf{k}])_{\mathbf{k} \in \Omega \cap \mathbb{Z}^2} = \left(c * (b \otimes b)[\mathbf{k}] \right)_{\mathbf{k} \in \Omega \cap \mathbb{Z}^2}, \quad (11.5)$$

where $P \triangleq \#(\Omega \cap \mathbb{Z}^2)$ is the number of pixels. Note that we typically have $P = M$, but this is not a requirement of our method.

11.3 Inverse Problem Formulation and Algorithm

We now formulate the image reconstruction task as an optimization problem over the spline coefficients \mathbf{c} . By piecing together the elements of Section 11.2, we get the following problem

$$\mathbf{c}^* \in \arg \min_{\mathbf{c} \in \mathbb{R}^N} \left(\frac{1}{2} \|\mathbf{H}\mathbf{c} - \mathbf{y}\|_2^2 + \lambda \mathcal{R}_d(\mathbf{c}) + \iota_{\geq 0}(\mathbf{L}_{\text{pos}} \mathbf{c}) \right), \quad (11.6)$$

where $\lambda > 0$, and $\iota_{\geq 0}$ is the indicator function defined by

$$\iota_{\geq 0}(\mathbf{x}) \triangleq \begin{cases} 0 & \text{if } \forall n, x_n \geq 0 \\ +\infty & \text{otherwise.} \end{cases} \quad (11.7)$$

The first term in (11.6) is known as the data-fidelity term, and ensures that the reconstructed signal conforms with the measured data \mathbf{y} . The second term is the Hessian nuclear-norm regularization, which tends to promote piecewise-linear reconstructed images. The balance between these two terms is controlled by the regularization parameter $\lambda > 0$, which should be tuned according to the noise level \mathbf{n} and the error in the forward model, *i.e.*, on the sampling locations \mathbf{t}^m . Finally, the last term in (11.6) guarantees that the reconstructed image $f^*(\mathbf{x}) \triangleq \sum_{\mathbf{k} \in \Omega_\varphi} c^*[\mathbf{k}] \varphi(\mathbf{x} - \mathbf{k})$ has positive values at the pixel locations $\mathbf{x} \in \Omega \cap \mathbb{Z}^2$.

Despite its somewhat daunting appearance, Problem (11.6) is a convex problem that can be solved with standard proximal algorithms. This is due to the availability of proximal operators for the sum of nuclear norms [27] and for the indicator function $\iota_{\geq 0}$. We solve it by applying the alternating direction method of multipliers (ADMM) [19] (see Appendix A), which we implemented in Matlab using GlobalBioIm [118], an inverse-problem library developed at the Biomedical Imaging Group at EPFL. The linear step of ADMM is solved using an inner-loop conjugate gradient (CG) algorithm [294] (see Remark A.1). Thanks to the modularity of GlobalBioIm, the only block that required implementation was the continuous Hessian operator $(\mathbf{L}_c[\mathbf{k}])_{\mathbf{k} \in \Omega_\varphi}$ and its adjoint operator.

Although the theoretical convergence speed of ADMM is rather slow— $\mathcal{O}(1/k)$, where k is the number of iteration (see Appendix A.3), in practice, few iterations are necessary to

obtain a decent accuracy in few iterations. This is all the more true in STXM due to the availability of a good warm-start initialization for ADMM. The latter can be obtained by computing the spline coefficients of the uncorrected image (*i.e.*, assuming that the measurements are taken on the desired pixel grid) using inverse filtering [106].

11.4 Experimental Results

11.4.1 Simulated Data

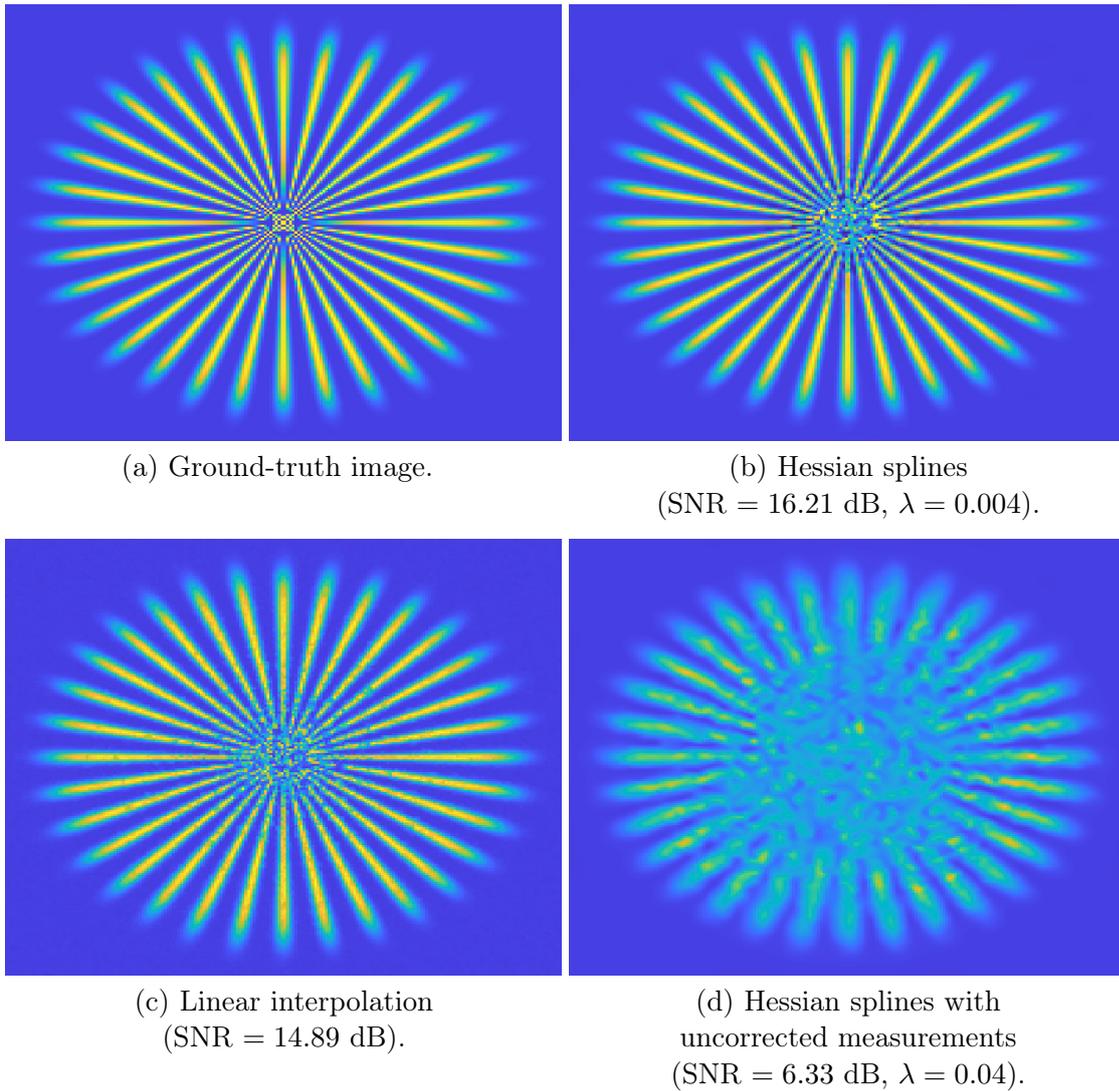


Figure 11.1 – Simulations on a star-like sample.

In order to assess the pertinence of our Hessian-spline framework, we apply our algorithm to a simulated continuous-domain ground-truth image $f_0 : \mathbb{R}^2 \rightarrow [0, 1]$, and we evaluate

its reconstruction performance in terms of signal-to-noise ratio (SNR). For the ground truth, we use the star-like sample shown in Figure 11.1 (a), taken from GlobalBioIm. This sample has high-frequency content at the center of the star and lower frequencies towards the end of the branches, and is thus a good benchmarking example. In order to better conform with real STXM images, we add a constant background of 0.1, and we rescale the image so that it ranges from 0.1 to 1.

For the forward model $\boldsymbol{\nu}$, the sampling pattern \mathbf{t}^m , $m = 1, \dots, M$ is taken from the real STXM data described in Section 11.4.2. Note that this pattern is quite far from being uniform: the average displacement with respect to the desired pixel grid is 2.7 nm, which is larger than the pixel size (1 nm). The image size is 200×200 , which yields $M = 40000$ measurements. The latter are corrupted by a Poisson noise term $\mathbf{n} \in \mathbb{R}^M$, which is applied to the rescaled ground truth image such the maximum pixel value corresponds to the maximum number of photon counts in the real STXM data. We thus have measurements $\mathbf{y} = \boldsymbol{\nu}(f_0) + \mathbf{n}$.

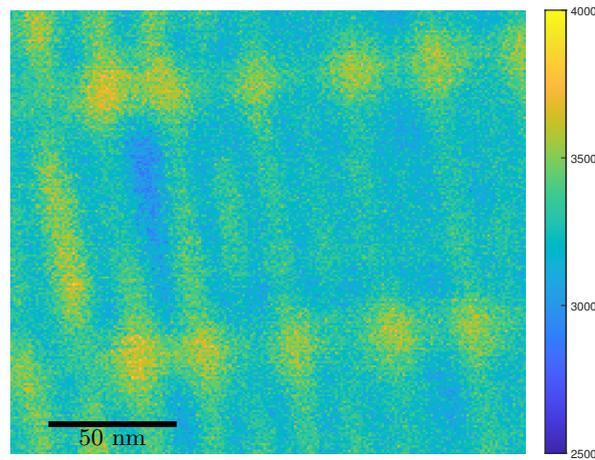
We then apply our Hessian-spline algorithm to solve Problem (11.6). The problem dimension is $N = 42436$, which leads to reasonable computations times (in the order of 30 seconds on commodity hardware). For ADMM, we use 5 inner CG iterations. We optimize the regularization parameter λ using grid search in order to maximize the SNR of the reconstructed image with respect to the ground truth f_0 . The reconstructed image is shown in Figure 11.1 (b), and achieves an SNR of 16.21 dB with a regularization parameter $\lambda = 0.004$.

Next, we compare the performance of Hessian splines with linear interpolation, which is used for achieving the state-of-the-art 7 nm resolution in STXM [289], using the Matlab function “scatteredInterpolant”. This method also uses the knowledge of the sampling locations \mathbf{t}^m , but it does not apply regularization and is therefore more sensitive to noise. The reconstruction image is shown in Figure 11.1 (c), and achieves an SNR of 14.89 dB.

Finally, we show the effect of the sampling location error by running our algorithm with an uncorrected forward model $\boldsymbol{\nu}_{\text{pix}}$, which does not take into account the sampling locations \mathbf{t}^m measured by the interferometer. As before, we optimize the regularization parameter using grid search, which yields the reconstructed image shown in Figure 11.1 (d) and achieves an SNR of 6.33 dB with a regularization parameter $\lambda = 0.04$. We see that the reconstruction fails dramatically, which indicates that using the corrected forward model $\boldsymbol{\nu}$ is critical in order to achieve good reconstruction results.

11.4.2 Real Data

We now apply our Hessian spline algorithm to real high-resolution STXM data provided by the Paul Scherrer Institute. The sample being imaged is a grating structure that is part of a Fresnel zone plate. The image size is 200×200 pixels, and the pixel size is 1 nm.



(a) Uncorrected image.

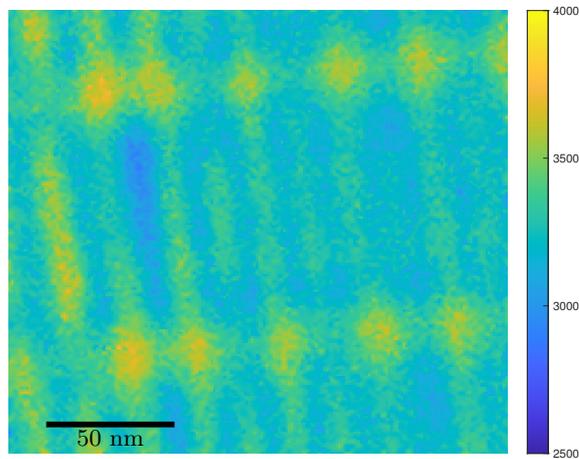
(b) Hessian splines reconstruction with $\lambda = 2.5$.

Figure 11.2 – Real STXM data of a grating structure.

In Figure 11.2 (a), we show the uncorrected image, assuming that the measurements are taken on the desired pixel grid. Although the effect of the displacement errors is not as glaring as in the simulation in Figure 11.1 (d), they lead to noticeable jitter artifacts.

In Figure 11.2 (b), we show our reconstruction results using the corrected sampling locations \mathbf{t}^m and the Hessian spline framework with a regularization parameter $\lambda = 2.5$. Although the reconstruction performance cannot be evaluated quantitatively due to the absence of a ground-truth image, we notice that the jitter artifacts as well as the noise are attenuated, due to the use of the corrected forward model ν and the regularization respectively.

11.5 Conclusion

We introduced a new framework, coined as *Hessian splines*, for reconstructing STXM images. This framework takes into account the sampling location errors (*i.e.*, the fact that the measurements are not acquired on a uniform pixel grid), which are measured using an interferometer. We formulated the reconstruction task as a continuous-domain inverse problem with Hessian nuclear-norm regularization. We then discretized the problem in a cubic spline basis, and solved it using ADMM. On the experimental side, we evaluated our method on simulated data, and showed that it outperforms linear interpolation that is used in state-of-the-art high-resolution STXM. We also applied our algorithm to real STXM data.

12 Conclusion

In this thesis, we studied various continuous-domain inverse problems with sparsity-promoting regularization, and we designed efficient discretization methods and algorithms to solve them computationally. Our particular focus was on devising methods that are *exact* in the continuous domain (in the sense that there is no discretization error), and that reach a desired sparse solution—in some cases, the *sparsest* solution.

Our continuous-domain frameworks have the advantage of being extremely versatile, as they can seamlessly adapt to a wide range of linear forward models. This is a major asset compared to purely discrete methods, where it may not be obvious how to discretize the forward model *a priori*, for example in the case of nonuniform sampling (either in the spatial or Fourier domains). Moreover, a discretized forward model often incurs discretization errors that are absent from continuous methods, which can lead to spurious reconstructions. Finally, our continuous methods have the benefit of mathematical elegance, where the discretization is *a posteriori* and based on theoretical tools such as representer theorems.

In Part I, we laid the mathematical foundations of this thesis by providing some background on generalized total-variation (gTV)-based inverse problems. We then defined the native spaces of higher-order derivative operators L (Chapter 2), whose construction relied on distribution theory [98]. These native spaces then served as search spaces for our gTV-based inverse problems. Existing representer theorems [14, 22] have proved that these problems admit spline solutions, both in nonperiodic and periodic settings. We thus presented some background on polynomial splines and stated these representer theorems in Chapter 3.

12.1 Contributions

Our contributions started in Part II, in which we studied various generic gTV-based inverse problems. We proposed exact discretization methods for each of these settings

using spline bases, and devised efficient multiresolution algorithms using B-splines and the simplex algorithm [116] that were guaranteed to reach a desired sparse solution. Our choices of bases were informed by representer theorems which guaranteed that our problems of interest admit spline solutions. Although exact discretization in spline bases was not a new idea [45], the use of B-splines was a major improvement which made the computational task feasible. Our first work was focused on single-component signal models (Chapter 4), which lead to problems as presented in Section 3.2 and for which a representer theorem was already known [14]. We then considered multicomponent signal models $s = s_1 + s_2$, where each component s_i was assumed to have different characteristics. We formulated inverse problems and proved representer theorems showing that these problems had spline solutions that conformed with the signal model. Firstly, we proposed a *hybrid spline* framework (Chapter 5), where both components were assumed to be sparse in different dictionary bases. Finally, we proposed a framework for multicomponent signals where s_1 and s_2 are assumed to be sparse and smooth, respectively.

In Part III, we studied various specific gTV-based problems and we precisely described their solution sets using tools from duality theory [26]. Specifically, we identified cases of uniqueness, and when uniqueness did not hold, we identified the *sparsest* solution. We then designed efficient algorithms informed by our theoretical study that are guaranteed to reach these sparsest solutions. The first problem we considered consisted in interpolating data points with second-order TV regularization (Chapter 7). We showed that depending on the observation vector $\mathbf{y} \in \mathbb{R}^M$, such problems may or may not have a unique solution. In the exact interpolation scenario—which we call the (g-BPC)—, we identified the sparsest solution and designed a simple algorithm that reaches it in linear time $\mathcal{O}(M)$. We then showed that the noisy interpolation scenario—called the (g-BLASSO)—could be recast as a (g-BPC) problem at the expense of solving a standard finite-dimensional ℓ_1 -regularized problem. Hence, the sparsest solution of the (g-BLASSO) could also be reached efficiently by combining the two steps. Next, we studied the same (g-BLASSO) problem but with an added Lipschitz constraint to favor stable solutions (Chapter 8). We showed that the same conclusions applied to this problem, in the sense that uniqueness depends on the observation vector \mathbf{y} , and that it could be recast as a (g-BPC) by solving an ℓ_1 -regularized problem. Hence, the sparsest solution of this Lipschitz-constrained problem could also be reached via a simple two-step algorithm. Finally, we studied the problem of reconstructing periodic functions based on low-pass Fourier series coefficients with generalized TV regularization. We showed that such problems always admit a unique solution, and we designed a B-spline-based algorithm similar to that of Chapter 4. We proved that the latter converges in uniform norm towards this unique solution as the grid size goes to zero.

Finally, in Part IV, we applied our continuous-domain sparsity-based frameworks to real-world applications. The first was the problem of fitting a two-dimensional sparse curve to contour points (Chapter 10). We formulated this task as a continuous-domain inverse problem with a novel regularization functional which we called *rotation-invariant TV*. We

proved a representer theorem that stated that our problem had sparse spline solutions, and we adapted our B-spline-based algorithm from Chapter 4 to reach a desired sparse-curve solution. Inspired by our hybrid spline framework from Chapter 5, we then extended our method to hybrid curve models with varying smoothness properties. Finally, we proposed an image-reconstruction method for scanning transmission X-ray microscopy (STXM) with nonuniform sampling patterns. We formulated the reconstruction task as a continuous-domain inverse problem with sparsity-promoting Hessian-Schatten regularization [27]. Using our B-spline-based exact discretization techniques from Part II, we then proposed an efficient algorithm to solve these problems computationally and demonstrated the feasibility of our approach on both simulated and real data.

12.2 Outlook

As demonstrated in our literature review in Section 1.2, the study of continuous-domain inverse problems with sparsity-promoting regularization is still a very active field of research, both on theoretical and algorithmic aspects. However, the field is sufficiently established to have some hindsight. This is true of low-dimensional settings, and particularly one-dimensional ones—as in the majority of this thesis. These low-dimensional settings are particularly interesting to the signal-processing community, for example in the context of (spatio-temporal) image reconstruction. Concerning one-dimensional problems, although specific settings can be studied *ad infinitum* and give rise to new ideas such as in Part III of this thesis, existing theories (such as dual certificates [47], geometrical properties of convex sets [61], or duality mappings [112]) provide a wide range of tools to tackle these questions. Low-dimensional settings with dimension higher than one are still being actively studied, and will most likely continue to be in the years to come. Indeed, recent theoretical advances have been achieved even for very classical models such as Rudin *et al.*'s [28] total-variation regularization for functions, *e.g.*, [61, 295].

On the algorithmic side, the sparse spikes deconvolution problem has reached a certain state of maturity with some well-established grid-free methods such as the sliding Frank-Wolfe algorithm [51]. However, new ideas, improvements, and fields of applications are constantly being proposed [52, 296, 297]. Algorithms to recover smoother functions with gTV-type regularization are more recent and are still under active development, as shown by our literature review in Section 1.2.4 and by this thesis. These methods are interesting not only for their mathematical elegance, but for their applicability to a vast number of signal-processing applications (see for example Part IV and [51, 95, 96]). The techniques developed in this thesis, particularly the grid-based methods from Part II, can hopefully be useful in that respect, as exemplified by our works in Part IV.

Another major future direction for gTV-based optimization—which deviates from the more traditional signal-processing applications mentioned above—will also likely be the study of high-dimensional problems. This is partly due to the recent connections that

have been uncovered between variational formulations and the training of neural networks (see for example Proposition 8.1 and [195, 298]), which deal with very high-dimensional problems. Hence, a deeper understanding of these variational approaches could provide useful insights into the training of neural networks, which is currently a major topic of interest in data science. A lot is yet to be done in these variational high-dimensional settings, both on theoretical and algorithmic aspects. However, this line of research can rely on the solid foundations from the low-dimensional settings described above, and perhaps take inspiration from them. The curse of dimensionality prohibits the use of grid-based approaches such as those from Part II in high dimensions; however, the grid-free approaches mentioned above could be an interesting starting point. Finally, as demonstrated by our works in Part III, theoretical insights can lead to algorithmic advances, which could help tackle the issue of scalability to high dimensions.

Appendices

A ADMM for Discrete ℓ_1 -based Problems

In this appendix, we describe how we apply the alternating direction method of multipliers (ADMM) described in [19] to discrete optimization problems of the form

$$\min_{\mathbf{x} \in \mathbb{R}^N} \left(F(\mathbf{x}) + \sum_{r=1}^R G_r(\mathbf{L}_r \mathbf{x}) \right) \quad (\text{A.1})$$

where $F : \mathbb{R}^N \rightarrow \mathbb{R}$ is a differentiable functional, the $\mathbf{L}_r \in \mathbb{R}^{N_r \times N}$ are matrices with no particular structure, and $G_r : \mathbb{R}^{N_r} \rightarrow \mathbb{R}$ are *proximable* functionals.

A.1 Proximity Operators

Being proximable relates to *proximity operators*, which is defined for a functional $G : \mathbb{R}^N \rightarrow \mathbb{R}$ as

$$\text{prox}_G(\mathbf{v}) \triangleq \arg \min_{\mathbf{x} \in \mathbb{R}^N} \frac{1}{2} \|\mathbf{x} - \mathbf{v}\|_2^2 + G(\mathbf{x}). \quad (\text{A.2})$$

We say that a functional G is proximable if for any $\lambda > 0$, the proximity operator $\text{prox}_{\lambda G}$ is *simple*, *i.e.*, it either has a closed-form expression, or it can be computed efficiently. For an intuitive interpretation of the proximity operator, we refer to [17, Chapter 3]. Proximity operators have become ubiquitous in convex optimization methods, due to the fact that many popular cost functionals such as the ℓ_1 norm are nondifferentiable, but are proximable. Being proximable can be seen as the next best thing to being differentiable; gradient-based algorithms are then replaced by *proximal algorithms*, of which ADMM is an example among many others [18, 38–41].

Example A.1. *Examples of common proximity operators that are used in this thesis are:*

- The ℓ_1 norm $\|\cdot\|_1$; its proximity operator is the componentwise soft-thresholding

function

$$\left[\text{prox}_{\lambda \|\cdot\|_1}(\mathbf{v}) \right]_n = \begin{cases} v_n - 1 & v_n > \lambda \\ 0 & |v_n| \leq \lambda, \\ v_n + 1 & v_n < -\lambda \end{cases}, \quad 1 \leq n \leq N. \quad (\text{A.3})$$

- The ℓ_∞ norm $\|\cdot\|_\infty$; its proximity operator has computationally cheap implementations [17, Section 6.5.2].
- The nuclear norm or 1-Schatten norm for matrices $\|\cdot\|_*$, defined as $\|\mathbf{M}\|_* \triangleq \sum_{p=1}^P \sigma_p(\mathbf{M})$, where $\sigma_p(\mathbf{M})$ is the p th singular value of the matrix $\mathbf{M} \in \mathbb{R}^{P \times P}$. Its proximity operator has computationally cheap implementations [27].
- Indicator functions ι_V (defined as $\iota_V(\mathbf{x}) = 0$ if $\mathbf{x} \in V$ and $\iota_V(\mathbf{x}) = +\infty$ otherwise), where $V \subset \mathbb{R}^N$ is a convex set. The latter provide a way of rewriting constrained optimization problems as unconstrained ones. By definition, the proximity operator of ι_V is then the least-squares projector onto the set V . Many such projectors have simple expressions, such as:

- The nonnegativity constraint $\iota_{\geq 0}$ (i.e., $V = \mathbb{R}^+$); its proximity operator is the componentwise ReLU $\left[\text{prox}_{\iota_{\geq 0}}(\mathbf{v}) \right]_n = \max(0, v_n)$.
- The ℓ_∞ box constraint (i.e., $V = \{\mathbf{x} \in \mathbb{R}^N : \|\mathbf{x}\|_\infty \leq C\}$ for some $C \geq 0$); its proximity operator is the componentwise thresholding operator

$$\left[\text{prox}_{\|\cdot\|_\infty \leq C}(\mathbf{v}) \right]_n = \begin{cases} C & v_n > C \\ v_n & |v_n| \leq C \\ -C & v_n < -C \end{cases}.$$

A.2 ADMM Iterates

The idea behind ADMM is to rewrite Problem (A.1) in primal-dual form via the *augmented Lagrangian* with penalty parameters $\boldsymbol{\rho} \in (0, +\infty)^R$, which is given by

$$\begin{aligned} \mathcal{L}_{\boldsymbol{\rho}}(\mathbf{x}, \mathbf{z}_1, \dots, \mathbf{z}_R, \mathbf{y}_1, \dots, \mathbf{y}_R) &\triangleq F(\mathbf{x}) + \sum_{r=1}^R \left(G_r(\mathbf{z}_r) + \mathbf{y}_r^T (\mathbf{L}_r \mathbf{x} - \mathbf{z}_r) + \frac{\rho_r}{2} \|\mathbf{L}_r \mathbf{x} - \mathbf{z}_r\|_2^2 \right) \\ &= F(\mathbf{x}) + \sum_{r=1}^R \left(G_r(\mathbf{z}_r) + \frac{\rho_r}{2} \left\| \mathbf{L}_r \mathbf{x} - \mathbf{z}_r + \frac{\mathbf{y}_r}{\rho_r} \right\|_2^2 \right) + C, \end{aligned} \quad (\text{A.4})$$

where the $\mathbf{z}_r \in \mathbb{R}^{N_r}$ are intermediate primal variables, the $\mathbf{y}_r \in \mathbb{R}^{N_r}$ are the dual variables, and $C \in \mathbb{R}$ is a constant. Clearly, for any $\boldsymbol{\rho} \in (0, +\infty)^R$, Problem (A.1) is equivalent to

the min-max formulation

$$\min_{\substack{\mathbf{x} \in \mathbb{R}^N, \mathbf{z}_r \in \mathbb{R}^{N_r} \\ 1 \leq r \leq R}} \left(\sup_{\substack{\mathbf{y}_r \in \mathbb{R}^{N_r} \\ 1 \leq r \leq R}} \mathcal{L}_\rho(\mathbf{x}, \mathbf{z}_1, \dots, \mathbf{z}_R, \mathbf{y}_1, \dots, \mathbf{y}_R) \right), \quad (\text{A.5})$$

and we must have $\mathbf{z}_r = \mathbf{L}_r \mathbf{x}_r$ for all $1 \leq r \leq R$ at the optimum. ADMM then solves Problem (A.5) by iteratively minimizing \mathcal{L}_ρ over the primal variables \mathbf{x} and \mathbf{z}_r , $1 \leq r \leq R$, and performing a gradient-ascent step over the dual variables \mathbf{y}_r , $1 \leq r \leq R$. More precisely, the k th iteration of ADMM is given by

$$\begin{aligned} \mathbf{x}^{(k+1)} &= \arg \min_{\mathbf{x} \in \mathbb{R}^N} \mathcal{L}_\rho \left(\mathbf{x}, \mathbf{z}_1^{(k)}, \dots, \mathbf{z}_R^{(k)}, \mathbf{y}_1^{(k)}, \dots, \mathbf{y}_R^{(k)} \right) \\ &= \arg \min_{\mathbf{x} \in \mathbb{R}^N} \left(F(\mathbf{x}) + \sum_{r=1}^R \frac{\rho_r}{2} \left\| \mathbf{L}_r \mathbf{x} - \mathbf{z}_r^{(k)} + \frac{\mathbf{y}_r^{(k)}}{\rho_r} \right\|_2^2 \right) \end{aligned} \quad (\text{A.6})$$

$$\begin{aligned} \mathbf{z}_r^{(k+1)} &= \arg \min_{\mathbf{z}_r \in \mathbb{R}^{N_r}} \mathcal{L}_\rho \left(\mathbf{x}^{(k+1)}, \mathbf{z}_1^{(k)}, \dots, \mathbf{z}_r, \dots, \mathbf{z}_R^{(k)}, \mathbf{y}_1^{(k)}, \dots, \mathbf{y}_R^{(k)} \right) \\ &= \text{prox}_{\frac{1}{\rho_r} G_r} \left(\mathbf{L}_r \mathbf{x}^{(k+1)} + \frac{\mathbf{y}_r^{(k)}}{\rho_r} \right) \quad 1 \leq r \leq R \end{aligned} \quad (\text{A.7})$$

$$\begin{aligned} \mathbf{y}_r^{(k+1)} &= \mathbf{y}_r^{(k)} + \rho_r \nabla_{\mathbf{y}_r} \mathcal{L}_\rho \left(\mathbf{x}^{(k+1)}, \mathbf{z}_1^{(k+1)}, \dots, \mathbf{z}_R^{(k+1)}, \mathbf{y}_1^{(k)}, \dots, \mathbf{y}_R^{(k)} \right) \\ &= \mathbf{y}_r^{(k)} + \rho_r \left(\mathbf{L}_r \mathbf{x}_r^{(k+1)} - \mathbf{z}_r^{(k+1)} \right) \quad 1 \leq r \leq R. \end{aligned} \quad (\text{A.8})$$

The algorithm then terminates when a suitable stopping criterion is reached (see [118] for examples). Hence, in ADMM, the difficult Problem (A.5) is decomposed into smaller, computationally cheaper subproblems, which are solved alternately. Indeed, as shown in (A.7), the \mathbf{z}_r -update step amounts to an application of the proximal operator of $\frac{1}{\rho_r} G_r$ which is simple when G_r is proximable. The gradient-ascent updates of \mathbf{y}_r are also computationally inexpensive. The most challenging and computationally demanding subproblem is typically the \mathbf{x} -update step (A.6). However, contrary to the original nondifferentiable Problem (A.1), Problem (A.6) has a differentiable cost functional for which gradient-based methods can be applied. Moreover, in many particular but common cases of F , Problem (A.6) is computationally cheap to solve, as exemplified in the following remark.

Remark A.1. *In the case of the classical quadratic data-fidelity loss*

$$F(\mathbf{x}) = \frac{1}{2} \|\mathbf{H}\mathbf{x} - \mathbf{y}\|_2^2 \quad (\text{A.9})$$

with $\mathbf{H} \in \mathbb{R}^{M \times N}$ and $\mathbf{y} \in \mathbb{R}^M$, Problem (A.6), amounts to solving the linear system

$$(\mathbf{H}^T \mathbf{H} + \sum_{r=1}^R \rho \mathbf{L}_r^T \mathbf{L}_r) \mathbf{x} = \mathbf{H}^T \mathbf{y} + \sum_{r=1}^R \rho_r \mathbf{L}_r^T \left(\rho_r \mathbf{z}_r^{(k+1)} - \mathbf{y}_r^{(k+1)} \right) \quad (\text{A.10})$$

over $\mathbf{x} \in \mathbb{R}^N$. This can be achieved directly when the matrix $\tilde{\mathbf{H}} \triangleq \left(\mathbf{H}^T \mathbf{H} + \sum_{r=1}^R \rho \mathbf{L}_r^T \mathbf{L}_r \right)$ is small enough to be inverted numerically, which is mostly the case in the one-dimensional settings of this thesis. However, for larger problems, $\tilde{\mathbf{H}}$ and its inverse can be too large to store in memory, which calls for matrix-free methods. In some cases (e.g., in Chapter 7), the matrix $\tilde{\mathbf{H}}$ is banded and can thus be inverted efficiently without needing to store $N \times N$ matrices in memory. When $\tilde{\mathbf{H}}$ is large and has no useful structure, one typically uses matrix-free inner-loop iterative solvers such as the conjugate gradient method [294], as we do in Chapter 11.

A.3 Convergence and Computational Complexity

It is well known that ADMM has a $\mathcal{O}(1/k)$ convergence rate in general, where k is the number of iterations [299]. The cost per iteration of ADMM depends on how the \mathbf{x} -minimization step is performed, which may depend on the choice of F . In the standard quadratic case (A.9), as discussed in Remark A.1, the computational bottleneck is to solve the linear system (A.10). When the inverse matrix $\tilde{\mathbf{H}}^{-1}$ has no structure but is small enough to be computed and stored in memory, the computational complexity per iteration of ADMM is $\mathcal{O}(N^2)$ both in time and in space.

B Proofs of Representer Theorems

In this appendix, we compile the proofs of all the representer theorems that are stated in this thesis.

B.1 Proof of Theorem 5.1

The proof is divided into two parts. We first prove the existence of a solution, and then that some solutions are of the form (5.20).

B.1.1 Existence of a Solution

Consider a sequence $(f_k)_{k \in \mathbb{N}} \in \mathcal{M}_{L_1, \phi_0}(\mathbb{R}) + \mathcal{M}_{L_2}(\mathbb{R})$ that monotonically decreases to the infimum value \mathcal{J}^0 of the cost functional. For every k , consider the decomposition $f_k = f_{1,k} + f_{2,k}$ given by (5.7), such that $\mathcal{R}_{\text{hyb}}(f_k) = (1 - \alpha)\|L_1\{f_{1,k}\}\|_{\mathcal{M}} + \alpha\|L_2\{f_{2,k}\}\|_{\mathcal{M}}$.

Expanding $f_{1,k}$ and $f_{2,k}$ using (5.8) yields

$$f_{1,k} = L_{\tilde{\phi}_1}^{-1}\{w_{1,k}\} + \mathbf{c}_{1,k}^T \mathbf{p}_1, \quad (\text{B.1})$$

$$f_{2,k} = L_{\tilde{\phi}_2}^{-1}\{w_{2,k}\} + \mathbf{c}_{0,k}^T \mathbf{p}_0 + \mathbf{c}_{2,k}^T \mathbf{p}_2. \quad (\text{B.2})$$

The proof will now consist in extracting a converging subsequence of $(f_k)_{k \in \mathbb{N}}$. To achieve this, the first step is to prove that the sequences $(w_{i,k})_{k \in \mathbb{N}}$ ($i \in \{1, 2\}$) and $(\mathbf{c}_{i,k})_{k \in \mathbb{N}}$ ($i \in \{0, 1, 2\}$) are bounded.

Due to the assumption of monotonic decrease, we have that

$$\|w_{1,k}\|_{\mathcal{M}}, \|w_{2,k}\|_{\mathcal{M}} \leq C_1 \quad \forall k \in \mathbb{N}, \quad (\text{B.3})$$

where $C_1 \triangleq \frac{\mathcal{J}(f_1)}{\lambda \min(\alpha, 1 - \alpha)}$. Next, we bound the $(\mathbf{c}_{i,k})_{k \in \mathbb{N}}$ sequences. Using the triangle

inequality, we get that

$$\|\nu(f_k)\|_2 \geq \left\| \nu \left(\sum_{i=0}^2 \mathbf{c}_{i,k}^T \mathbf{p}_i \right) \right\|_2 - \left\| \nu \left(L_{\tilde{\phi}_1}^{-1}\{w_{1,k}\} + L_{\tilde{\phi}_2}^{-1}\{w_{2,k}\} \right) \right\|_2. \quad (\text{B.4})$$

Using Proposition 8 in [14], the well-posedness assumption in Theorem 5.1 is equivalent to the existence of a constant $B > 0$ such that

$$\left\| \nu \left(\sum_{i=0}^2 \mathbf{c}_{i,k}^T \mathbf{p}_i \right) \right\|_2 \geq B \sum_{i=0}^2 \|\mathbf{c}_{i,k}\|_2. \quad (\text{B.5})$$

Next, to handle the second term in (B.4), we prove Lemma B.1.

Lemma B.1. *The operators $\nu \left(L_{\tilde{\phi}_i}^{-1}\{\cdot\} \right) : \mathcal{M}(\mathbb{R}) \rightarrow \mathbb{R}^M$ with $i \in \{1, 2\}$ are weak*-continuous.*

Proof. Let $(w_n)_{n \in \mathbb{N}} \in \mathcal{M}(\mathbb{R})$ be a sequence that converges to $w \in \mathcal{M}(\mathbb{R})$ for the weak*-topology. Since ν is weak*-continuous, it is sufficient to prove that $L_{\tilde{\phi}_i}^{-1}\{w_n - w\} \rightarrow 0$ for the weak*-topology in $\mathcal{M}_{L_i}(\mathbb{R})$. By Theorem 2.3, we have that $(C_{L_i}(\mathbb{R}))' = \mathcal{M}_{L_i}(\mathbb{R})$, where $C_{L_i}(\mathbb{R}) = L_i^*\{C_0(\mathbb{R})\} \oplus \text{span}\{\phi_n\}_{n=1}^{N_0}$. Let $f \in C_{L_i}(\mathbb{R})$; there exists functions $(f_1, f_2) \in C_0(\mathbb{R}) \times \text{span}\{\phi_n\}_{n=1}^{N_0}$ such that $f = L_i^*\{f_1\} + f_2$. By the boundary conditions in Theorem 2.1, we have that $\tilde{\phi}_i \left(L_{\tilde{\phi}_i}^{-1}\{w_n - w\} \right) = \mathbf{0}$ which implies that $\langle L_{\tilde{\phi}_i}^{-1}\{w_n - w\}, f_2 \rangle = 0$. We thus have that

$$\left\langle L_{\tilde{\phi}_i}^{-1}\{w_n - w\}, f \right\rangle = \left\langle L_{\tilde{\phi}_i}^{-1}\{w_n - w\}, L_i^*\{f_1\} \right\rangle = \langle w_n - w, f_1 \rangle \rightarrow 0 \quad (\text{B.6})$$

since $w_n \rightarrow w$ for the weak*-topology. This proves that $L_{\tilde{\phi}_i}^{-1}\{w_n - w\} \rightarrow 0$ for the weak*-topology and, thus, the desired result. \square

Next, Lemma B.1 implies that the operators $\nu \left(L_{\tilde{\phi}_i}^{-1}\{\cdot\} \right)$ are continuous for the strong topology $\|\cdot\|_{\mathcal{M}}$ (see Section 2.1.3). This implies the existence of a constant $A > 0$ such that

$$\left\| \nu \left(L_{\tilde{\phi}_1}^{-1}\{w_{1,k}\} + L_{\tilde{\phi}_2}^{-1}\{w_{2,k}\} \right) \right\|_2 \leq A(\|w_{1,k}\|_{\mathcal{M}} + \|w_{2,k}\|_{\mathcal{M}}). \quad (\text{B.7})$$

Combining (B.3), (B.5), and (B.7) yields that

$$\|\nu(f_k)\|_2 \geq B \left(\sum_{i=0}^2 \|\mathbf{c}_{i,k}\|_2 \right) - A(\|w_{1,k}\|_{\mathcal{M}} + \|w_{2,k}\|_{\mathcal{M}}) \geq B \left(\sum_{i=0}^2 \|\mathbf{c}_{i,k}\|_2 \right) - 2AC_1. \quad (\text{B.8})$$

Using the assumption of monotonic decrease and the triangular inequality, we get that

$$\forall k \in \mathbb{N} : \sqrt{2\mathcal{J}(f_0)} + \|\mathbf{y}\|_2 \geq \|\boldsymbol{\nu}(f_k)\|_2. \quad (\text{B.9})$$

This shows that the norms of $\mathbf{c}_{0,k}$, $\mathbf{c}_{1,k}$, and $\mathbf{c}_{2,k}$ are bounded by a constant, which was our initial goal. Together with (B.3), this implies the existence of a subsequence $(f_{k_n})_{n \in \mathbb{N}}$ such that

- the sequences w_{i,k_n} are converging to a limit $w_{i,\text{lim}} \in \mathcal{M}(\mathbb{R})$ for the weak*-topology for $i \in \{1, 2\}$ (Banach-Alaoglu theorem);
- for $i \in \{0, 1, 2\}$, the sequences \mathbf{c}_{i,k_n} converge to a limit $\mathbf{c}_{i,\text{lim}}$.

Using Lemma B.1, we thus have that $\lim_{n \rightarrow +\infty} \boldsymbol{\nu}(f_{k_n}) = \boldsymbol{\nu}(f_{\text{lim}})$ where

$$f_{\text{lim}} = \mathbb{L}_{\tilde{\phi}_1}^{-1}\{w_{1,\text{lim}}\} + \mathbb{L}_{\tilde{\phi}_2}^{-1}\{w_{2,\text{lim}}\} + \sum_{i=0}^2 \mathbf{c}_{i,\text{lim}}^T \mathbf{p}_i. \quad (\text{B.10})$$

This yields that

$$\begin{aligned} \mathcal{J}^0 &= \lim_{n \rightarrow +\infty} \mathcal{J}(f_{k_n}) \\ &= \frac{1}{2} \|\boldsymbol{\nu}(f_{\text{lim}}) - \mathbf{y}\|_2^2 + \lambda((1 - \alpha)\|w_{1,\text{lim}}\|_{\mathcal{M}} + \alpha\|w_{2,\text{lim}}\|_{\mathcal{M}}) \end{aligned} \quad (\text{B.11})$$

$$= \mathcal{J}(f_{\text{lim}}), \quad (\text{B.12})$$

using the weak*-continuity of the $\|\cdot\|_{\mathcal{M}}$ norm in (B.11). This proves the desired result $f_{\text{lim}} \in \mathcal{V}$.

B.1.2 Form of the Solutions

We first prove the following lemma.

Lemma B.2. *All solutions of Problem (5.19) yield the same measurement vector $\mathbf{z}_\lambda \in \mathbb{R}^M$.*

Proof. Let f, g be two solutions of Problem (5.19) and define $h = \frac{f+g}{2}$. Since \mathcal{R}_{hyb} (Proposition 5.1) and $\|\cdot\|_2^2$ are convex functionals, we have that

$$\mathcal{R}_{\text{hyb}}(h) \leq \frac{\mathcal{R}_{\text{hyb}}(f) + \mathcal{R}_{\text{hyb}}(g)}{2} \quad \text{and} \quad (\text{B.13})$$

$$\|\boldsymbol{\nu}(h) - \mathbf{y}\|_2^2 \leq \frac{\|\boldsymbol{\nu}(f) - \mathbf{y}\|_2^2 + \|\boldsymbol{\nu}(g) - \mathbf{y}\|_2^2}{2}. \quad (\text{B.14})$$

Summing these inequalities yields $\mathcal{J}(h) \leq \mathcal{J}^0$. Yet $\mathcal{J}(h) = \mathcal{J}^0$, which implies that the cases of equality are met. Since $\|\cdot\|_2^2$ is strictly convex, we necessarily have that $\nu(f) = \nu(g) = \mathbf{z}_\lambda$. \square

Using Lemma B.2, Problem (5.19) can be reformulated as

$$\mathcal{V} = \arg \min_{f \in \mathcal{M}_{L_1, \phi_0}(\mathbb{R}) + \mathcal{M}_{L_2}(\mathbb{R})} \mathcal{R}_{\text{hyb}}(f) \quad \text{s.t.} \quad \nu(f) = \mathbf{z}_\lambda. \quad (\text{B.15})$$

Now, consider the problem

$$\min_{\substack{w_1, w_2 \in \mathcal{M}(\mathbb{R}) \\ q \in \mathcal{N}_{L_1} + \mathcal{N}_{L_2} \\ F(w_1, w_2, q) = \mathbf{z}_\lambda}} ((1 - \alpha)\|w_1\|_{\mathcal{M}} + \alpha\|w_2\|_{\mathcal{M}}), \quad \text{where} \quad (\text{B.16})$$

$$F(w_1, w_2, q) = \nu \left(L_{\tilde{\phi}_1}^{-1}\{w_1\} + L_{\tilde{\phi}_2}^{-1}\{w_2\} + q \right) \quad (\text{B.17})$$

is a continuous linear functional. Due to the well-posedness assumption,

$$\forall q \in \mathcal{N}_{L_1} + \mathcal{N}_{L_2} : \quad F(0, 0, q) \geq B\|q\|_{\mathcal{N}_{L_1} + \mathcal{N}_{L_2}}. \quad (\text{B.18})$$

Using the vector-valued Fisher-Jerome theorem [94, Theorem C.1], Problem (B.16) has a solution (w_1, w_2, q) such that

$$w_i = \sum_{k=1}^{K_i} a_{i,k} \delta(\cdot - x_{i,k}), \quad i \in \{1, 2\}, \quad (\text{B.19})$$

where $a_{i,n}, x_{i,k} \in \mathbb{R}$ and $K_1 + K_2 \leq M$. Next, let

$$s = L_{\tilde{\phi}_1}^{-1}\{w_1\} + L_{\tilde{\phi}_2}^{-1}\{w_2\} + q. \quad (\text{B.20})$$

Clearly, $\nu(s) = \mathbf{z}_\lambda$ and $s = s_1 + s_2$ where the components s_i are of the form (5.20). Now, assume by contradiction that s is not a solution of Problem (5.19). According to Lemma B.2, this implies the existence of a function $g \in \mathcal{M}_{L_1, \phi_0}(\mathbb{R}) + \mathcal{M}_{L_2}(\mathbb{R})$ such that $\nu(g) = \mathbf{z}_\lambda$ and $\mathcal{R}_{\text{hyb}}(g) < \mathcal{R}_{\text{hyb}}(s)$. As stated in (5.8), g can be rewritten as

$$g = L_{\tilde{\phi}_1}^{-1}\{u_1\} + L_{\tilde{\phi}_2}^{-1}\{u_2\} + r, \quad (\text{B.21})$$

where $u_1, u_2 \in \mathcal{M}(\mathbb{R})$ and $r \in \mathcal{N}_{L_1} + \mathcal{N}_{L_2}$. Thus, $F(u_1, u_2, r) = \mathbf{z}_0$ and

$$(1 - \alpha)\|u_1\|_{\mathcal{M}} + \alpha\|u_2\|_{\mathcal{M}} = \mathcal{R}_{\text{hyb}}(g) < \mathcal{R}_{\text{hyb}}(s) = (1 - \alpha)\|w_1\|_{\mathcal{M}} + \alpha\|w_2\|_{\mathcal{M}}, \quad (\text{B.22})$$

which is in contradiction with (w_1, w_2, q) being a solution of (B.16). This proves that s is indeed the solution of (5.19), and thus the desired result.

B.2 Proof of Theorem 6.1

The main technical part of our proof is to show the existence of a minimizer; once this is ensured, the optimization problem can be decoupled into two separate problems. Then, we can apply representer theorems proven in [45] for these problems to obtain a parametric form of the solution of the original problem. Finally, the uniqueness of the smooth component follows from the strict convexity of the associated regularization penalty. We now recall some relevant notions that will be used throughout the proof.

B.2.1 Preliminaries

We extend the biorthogonal system (ϕ_0, \mathbf{p}_0) for \mathcal{N}_0 to the biorthogonal systems $(\tilde{\phi}_1, \tilde{\mathbf{p}}_1)$ and $(\tilde{\phi}_2, \tilde{\mathbf{p}}_2)$ for \mathcal{N}_{L_1} and \mathcal{N}_{L_2} , respectively, where $\tilde{\phi}_i = \begin{bmatrix} \phi_0 & \phi_i \end{bmatrix}$ and $\tilde{\mathbf{p}}_i = \begin{bmatrix} \mathbf{p}_0 & \mathbf{p}_i \end{bmatrix}$ for $i \in \{1, 2\}$. By (2.43), any function $s_1 \in \mathcal{M}_{L_1}(\mathbb{R})$ has the unique decomposition

$$s_1 = L_{1, \tilde{\phi}_1}^{-1} \{w\} + \tilde{\mathbf{c}}_0^T \mathbf{p}_0 + \mathbf{c}_1^T \mathbf{p}_1, \quad (\text{B.23})$$

where $w \in \mathcal{M}(\mathbb{R})$, $\tilde{\mathbf{c}}_0 = \phi_0(s_1) \in \mathbb{R}^{N_0}$, $\mathbf{c}_1 = \phi_1(s_1) \in \mathbb{R}^{N_{0,1} - N_0}$, and $L_{1, \tilde{\phi}_1}^{-1}$ is the pseudoinverse operator of L_1 for the biorthogonal system $(\tilde{\phi}_1, \tilde{\mathbf{p}}_1)$ defined in Theorem 2.1. Using this decomposition, by (2.44), we can equip the space $\mathcal{M}_{L_1}(\mathbb{R})$ with the norm

$$\|s_1\|_{\mathcal{M}_{L_1, \tilde{\phi}_1}} \triangleq \|w\|_{\mathcal{M}} + \left\| \tilde{\phi}_1(s_1) \right\|_2. \quad (\text{B.24})$$

Finally, an element $s_1 \in \mathcal{M}_{L_1}(\mathbb{R})$ is in the restricted search space $\mathcal{M}_{L_1, \phi_0}(\mathbb{R})$ if and only if $\tilde{\mathbf{c}}_0 = \mathbf{0}$.

Similarly, for any $s_2 \in \mathcal{H}_{L_2}(\mathbb{R})$, there is a unique decomposition

$$s_2 = L_{2, \tilde{\phi}_2}^{-1} \{h\} + \mathbf{c}_0^T \mathbf{p}_0 + \mathbf{c}_2^T \mathbf{p}_2, \quad (\text{B.25})$$

where $\mathbf{c}_0 = \phi_0(s_2) \in \mathbb{R}^{N_0}$, $\mathbf{c}_2 = \phi_2(s_2) \in \mathbb{R}^{N_{0,2} - N_0}$, and $h \in L_2(\mathbb{R})$. Consequently, the associated norm for the space $\mathcal{H}_{L_2}(\mathbb{R})$ is defined as

$$\|s_2\|_{\mathcal{H}_{L_2, \tilde{\phi}_2}} \triangleq \|h\|_{L_2} + \left\| \tilde{\phi}_2(s_2) \right\|_2. \quad (\text{B.26})$$

B.2.2 Existence of a Solution

The first step is to prove that (6.10) has a minimizer. We do so by reformulating the problem as the minimization of a weak*-lower semicontinuous functional over a weak*-compact domain. We then prove the existence by relying on the generalized Weierstrass theorem.

We denote the cost at the trivial point $(0, 0)$ as $\mathcal{J}_0 = \mathcal{J}(0, 0) = E(\mathbf{0}, \mathbf{y})$. Adding the constraint $\mathcal{J}(s_1, s_2) \leq \mathcal{J}_0$ does not change the solution set of the original problem, as it must hold for any minimizer of (6.10). Hence, from now on, we assume that the cost functional is upper-bounded by \mathcal{J}_0 . This readily implies that

$$E(\boldsymbol{\nu}(s_1 + s_2), \mathbf{y}) \leq \mathcal{J}_0, \tag{B.27}$$

$$\|\mathbf{L}_1\{s_1\}\|_{\mathcal{M}} \leq \frac{\mathcal{J}_0}{\lambda_1}, \text{ and} \tag{B.28}$$

$$\|\mathbf{L}_2\{s_2\}\|_{L_2} \leq \sqrt{\frac{\mathcal{J}_0}{\lambda_2}}. \tag{B.29}$$

The coercivity of $E(\cdot, \mathbf{y})$ implies the existence of a constant $C_1 > 0$ such that $E(\mathbf{z}, \mathbf{y}) \leq \mathcal{J}_0 \Rightarrow \|\mathbf{z}\|_2 \leq C_1$. Together with (B.27), this yields

$$\|\boldsymbol{\nu}(s_1 + s_2)\|_2 \leq C_1. \tag{B.30}$$

Moreover, as discussed in Section 2.1.3, since $\boldsymbol{\nu}$ is weak*-continuous over $\mathcal{M}_{L_1}(\mathbb{R})$, it is also continuous for the strong topology. Moreover, by assumption, $\boldsymbol{\nu}$ is continuous over $\mathcal{H}_{L_2}(\mathbb{R})$. Hence, there exists a second constant $C_2 > 0$ such that

$$\|f_1\|_{\mathcal{M}_{L_1, \tilde{\phi}_1}} + \|f_2\|_{\mathcal{H}_{L_2, \tilde{\phi}}} \leq \frac{\mathcal{J}_0}{\lambda_1} + \sqrt{\frac{\mathcal{J}_0}{\lambda_2}} \Rightarrow \|\boldsymbol{\nu}(f_1 + f_2)\|_2 \leq C_2. \tag{B.31}$$

Now, by taking

$$\begin{aligned} f_1 &\triangleq s_1 - \boldsymbol{\phi}_1(s_1)^T \mathbf{p}_1, \\ f_2 &\triangleq s_2 - \boldsymbol{\phi}_0(s_2)^T \mathbf{p}_0 - \boldsymbol{\phi}_2(s_2)^T \mathbf{p}_2, \end{aligned} \tag{B.32}$$

and, together with (B.28) and (B.29), we deduce that

$$\|\boldsymbol{\nu}(s_1 - \boldsymbol{\phi}_1(s_1)^T \mathbf{p}_1 + s_2 - \boldsymbol{\phi}_0(s_2)^T \mathbf{p}_0 - \boldsymbol{\phi}_2(s_2)^T \mathbf{p}_2)\|_2 \leq C_2. \tag{B.33}$$

By using the triangle inequality and the two bounds (B.30) and (B.33), we have

$$\|\boldsymbol{\nu}(\boldsymbol{\phi}_1(s_1)^T \mathbf{p}_1 + \boldsymbol{\phi}_0(s_2)^T \mathbf{p}_0 + \boldsymbol{\phi}_2(s_2)^T \mathbf{p}_2)\|_2 \leq C_1 + C_2. \tag{B.34}$$

Finally, the well-posedness assumption in Theorem 6.1 ensures the existence of a constant

$B > 0$ such that

$$\forall q \in \mathcal{N}_{L_1} + \mathcal{N}_{L_2} : B \|\phi_i(q)\|_2 \leq \|\nu(q)\|_2, \quad i \in \{0, 1, 2\}. \quad (\text{B.35})$$

Hence, by taking

$$q = \phi_1(s_1)^T \mathbf{p}_1 + \phi_0(s_2)^T \mathbf{p}_0 + \phi_2(s_2)^T \mathbf{p}_2 \quad (\text{B.36})$$

and by applying the inequality (B.35), we have that

$$\|\phi_1(s_1)\|_2, \|\phi_0(s_2)\|_2, \|\phi_2(s_2)\|_2 \leq \frac{C_1 + C_2}{B}. \quad (\text{B.37})$$

Therefore, the original Problem (6.10) is equivalent to the constrained minimization problem

$$\min_{\substack{s_1 \in \mathcal{M}_{L_1, \phi_0}(\mathbb{R}) \\ s_2 \in \mathcal{H}_{L_2}(\mathbb{R})}} \mathcal{J}(s_1, s_2) \quad \text{s.t.} \quad \|s_1\|_{\mathcal{M}_{L_1, \tilde{\phi}_1}} \leq A_1, \quad \|s_2\|_{\mathcal{H}_{L_2, \tilde{\phi}}} \leq A_2, \quad (\text{B.38})$$

where $A_1 \triangleq \frac{\mathcal{J}_0}{\lambda_1} + 2\frac{C_1+C_2}{B}$ and $A_2 \triangleq \sqrt{\frac{\mathcal{J}_0}{\lambda_2}} + 2\frac{C_1+C_2}{B}$. The cost functional of Problem (B.38), which is the same as in (6.10), is weak* lower-semicontinuous. Moreover, the constraint cube is weak*-compact in the product topology due to the Banach-Alaoglu theorem [15, Theorem 3.15]. Hence, (B.38) reaches its infimum, and so does (6.10).

B.2.3 Form of the Solution

Let $(\tilde{s}_1, \tilde{s}_2)$ be a solution of (6.10) and consider the minimization problem

$$\min_{s_1 \in \mathcal{M}_{L_1, \phi_0}(\mathbb{R})} \|\mathbf{L}_1\{s_1\}\|_{\mathcal{M}} \quad \text{s.t.} \quad \nu(s_1) = \nu(\tilde{s}_1). \quad (\text{B.39})$$

By Theorem 3.3, (B.39) has a minimizer s_1^* of the form (6.11). One can also readily verify that (s_1^*, \tilde{s}_2) is a minimizer of the original problem. Similarly, one can consider the minimization problem

$$\min_{s_2 \in \mathcal{H}_{L_2}(\mathbb{R})} \|\mathbf{L}_2\{s_2\}\|_{L_2} \quad \text{s.t.} \quad \nu(s_2) = \nu(\tilde{s}_2). \quad (\text{B.40})$$

It is known from [45, Theorem 3] that (B.40) has a minimizer s_2^* of the form (6.12). Again, (s_1^*, s_2^*) is a solution of the original problem, which matches the form specified by Theorem 6.1.

B.2.4 Uniqueness of the Second Component

To prove the final statement of Theorem 6.1, we consider two arbitrary pairs of solutions (\bar{f}_1, \bar{f}_2) and $(\tilde{f}_1, \tilde{f}_2)$ of Problem (6.10) and we denote by \mathcal{J}^0 their minimal cost value. The convexity of the cost functional yields that, for any $\alpha \in (0, 1)$ and $(f_{\alpha,1}, f_{\alpha,2}) = \alpha(\bar{f}_1, \bar{f}_2) + (1 - \alpha)(\tilde{f}_1, \tilde{f}_2)$, we have

$$\mathcal{J}(f_{\alpha,1}, f_{\alpha,2}) \leq \alpha\mathcal{J}(\bar{f}_1, \bar{f}_2) + (1 - \alpha)\mathcal{J}(\tilde{f}_1, \tilde{f}_2) = \mathcal{J}^0. \quad (\text{B.41})$$

The optimality of (\bar{f}_1, \bar{f}_2) and $(\tilde{f}_1, \tilde{f}_2)$ implies that (B.41) must be an equality. In particular, we must have that

$$\|L_2\{f_{\alpha,2}\}\|_{L_2}^2 = \alpha\|L_2\{\bar{f}_2\}\|_{L_2}^2 + (1 - \alpha)\|L_2\{\tilde{f}_2\}\|_{L_2}^2. \quad (\text{B.42})$$

Now, due to the strict convexity of $\|L_2\{\cdot\}\|_{L_2}^2$, we deduce that $L_2\{\bar{f}_2 - \tilde{f}_2\} = 0$, and hence that $(\bar{f}_2 - \tilde{f}_2) \in \mathcal{N}_{L_2}$. This implies that all solutions have the same second component up to a term in the null space of L_2 .

B.3 Proof of Theorem 8.2

B.3.1 Items 1 and 2

The first step is to show that the sampling functional $\delta(\cdot - x_0) : f \mapsto f(x_0)$ is weak*-continuous in $\text{Lip}(\mathbb{R})$. To that end, we identify the predual Banach space \mathcal{X} such that $\text{Lip}(\mathbb{R}) = \mathcal{X}'$ and then show that shifted Dirac impulses are included in \mathcal{X} , which is equivalent to $\delta(\cdot - x_0)$ being weak*-continuous [113, Theorem IV.20]. We recall that following (8.10), we can view $\text{Lip}(\mathbb{R})$ as the native Banach space associated to the pair $(L_\infty(\mathbb{R}), D)$. This allows us to deploy the machinery of [91] to identify its predual space. In short, it follows from [91] that the predual space has the direct-sum structure $\mathcal{X} = D(L_1(\mathbb{R})) \oplus \text{span}\left(e^{-(\cdot)^2}\right)$. In other words, any function $f \in \mathcal{X}$ can be decomposed as $f = D\{g\} + ce^{-(\cdot)^2}$, where $g \in L_1(\mathbb{R})$ and $c \in \mathbb{R}$. One can formally verify that $\delta = D\{\text{sgn} - \text{erf}\} + \frac{2}{\sqrt{\pi}}e^{-(\cdot)^2}$, where sgn is the sign function and erf is the Gauss error function. Due to the rapid decay of the erf function at $t = -\infty$ and the symmetry of $(\text{sgn} - \text{erf})$, we deduce that $\text{sgn} - \text{erf} \in L_1(\mathbb{R})$ and, hence, that $\delta \in \mathcal{X}$. Finally, due to the shift-invariant structure of \mathcal{X} , we deduce the weak*-continuity of the sampling functional $\delta(\cdot - x_0)$ for any $x_0 \in \mathbb{R}$.

Next, we apply the general representer theorem for Banach semi-norms [300, Theorem 3] to deduce that the solution set \mathcal{V}_{Lip} of (8.16) is a nonempty, convex, weak*-compact set whose elements all pass through a fixed set of points. Put differently, the vector $\mathbf{z} = (z_m)$ with $z_m = f(x_m)$ is invariant to the choice of $f \in \mathcal{V}_{\text{Lip}}$. This means that adding the

constraints $z_m = f(x_m)$ for $m = 1, \dots, M$ does not change the solution set, *i.e.*,

$$\mathcal{V}_{\text{Lip}} = \arg \min_{f \in \text{Lip}(\mathbb{R})} \left(\sum_{m=1}^M E(f(x_m), y_m) + \lambda L(f) \right) \quad \text{s.t.} \quad f(x_m) = z_m, \quad 1 \leq m \leq M \quad (\text{B.43})$$

$$= \arg \min_{f \in \text{Lip}(\mathbb{R})} L(f) \quad \text{s.t.} \quad f(x_m) = z_m, \quad (\text{B.44})$$

where the last equality is obtained by observing that $\sum_{m=1}^M E(f(x_m), y_m)$ is constant within the solution set \mathcal{V}_{Lip} . Consequently, we can represent \mathcal{V}_{Lip} as a solution set of a constrained problem of the form (8.17).

B.3.2 Item 3

We first recall the definition of the canonical CPWL interpolant of a collection of 1D data points introduced in Chapter 7.

Definition B.1 (Definition 7.4). *For a series of data points (x_m, z_m) with $m = 1, \dots, M$, the canonical interpolant $f_{\text{cano}} : \mathbb{R} \rightarrow \mathbb{R}$ is the unique CPWL function that passes through these points and is differentiable over $\mathbb{R} \setminus \{x_2, \dots, x_{M-1}\}$.*

We first prove that f_{cano} is a solution of (8.17). Clearly, the Lipschitz constant of f_{cano} is equal to $L(f_{\text{cano}}) = L_{\min}$, where L_{\min} is given in (8.18). Moreover, any function f that passes through the data points (x_m, z_m) necessarily has a Lipschitz constant greater than or equal to L_{\min} . This implies that f_{cano} is a solution of (8.17) and L_{\min} is the minimal value of the Lipschitz constant. Consequently, any function that satisfies the interpolation constraints and is L_{\min} -Lipschitz is a solution of (8.17).

B.3.3 Item 4

Consider a generic point $(x, y) \in \mathcal{E}$, and let m be such that $x \in (x_{m-1}, x_m)$. By definition of \mathcal{E} , there exists a function $f \in \mathcal{V}_{\text{Lip}}$ such that $y = f(x)$. From Item 3, we deduce that $L(f) = L_{\min}$. Hence, we have the inequalities

$$\left| \frac{y - z_{m-1}}{x - x_{m-1}} \right|, \left| \frac{y - z_m}{x - x_m} \right| \leq L_{\min}. \quad (\text{B.45})$$

These inequalities can readily be translated into the inclusion $(x, y) \in \mathcal{R}_{m-1} \cap \mathcal{L}_m$, which implies that $\mathcal{E} \subseteq \bigcup_{m=1}^M (\mathcal{R}_{m-1} \cap \mathcal{L}_m)$. To show the reverse inclusion, consider a point in $(x, y) \in \mathcal{R}_{m-1} \cap \mathcal{L}_m$ for some $m \in \{1, \dots, M+1\}$ and denote by \tilde{f}_{cano} the canonical interpolant of $\{(x_m, z_m)\}_{m=1}^M \cup \{(x, y)\}$. Following Item 3, the Lipschitz constant of \tilde{f}_{cano}

is given by

$$L(\tilde{f}_{\text{cano}}) = \max \left(L_{\min}, \left| \frac{y - z_{m-1}}{x - x_{m-1}} \right|, \left| \frac{y - z_m}{x - x_m} \right| \right) = L_{\min}, \quad (\text{B.46})$$

where we establish the last equality by translating the inclusion $(x, y) \in \mathcal{R}_{m-1} \cap \mathcal{L}_m$ into the inequalities in (B.45). This implies that \tilde{f}_{cano} is a solution of (8.17) and so, by definition, we have that $(x, y) \in \mathcal{E}$.

B.3.4 Item 5

By Proposition 7.5, f_{cano} is also a solution of (8.23). We therefore need to prove that any solution f_{opt} of (8.23) has the same Lipschitz constant $L(f_{\text{opt}}) = L(f_{\text{cano}}) = L_{\min}$. Due to the interpolation constraints, we necessarily have that $L(f_{\text{opt}}) \geq L(f_{\text{cano}})$; we must now prove the reverse inequality $L(f_{\text{opt}}) \leq L(f_{\text{cano}})$. By Theorem 7.1, f_{opt} must follow f_{cano} in $\mathbb{R} \setminus [x_2, x_{M-1}]$. Moreover, in each interval $[x_m, x_{m+1}]$ for $m \in \{2, \dots, M-2\}$, f_{opt} either follows f_{cano} or is concave or convex over the interval $[x_{m-1}, x_{m+2}]$. Hence, it suffices to prove that, for any $m \in \{2, \dots, M-2\}$, we have that $L_m(f_{\text{opt}}) \leq L(f_{\text{cano}})$, where $L_m(f)$ denotes the Lipschitz constant of f restricted to the interval $[x_m, x_{m+1}]$.

Let m be an index for which f_{opt} need not follow f_{cano} in $[x_m, x_{m+1}]$. (If no such index exists, then the result is trivially true.) Assume that f_{opt} is convex in the interval $[x_{m-1}, x_{m+2}]$; the concave scenario is derived in a similar fashion. This implies that, in this interval, the function $(\tilde{x}_1, \tilde{x}_2) \mapsto \frac{f_{\text{opt}}(\tilde{x}_2) - f_{\text{opt}}(\tilde{x}_1)}{\tilde{x}_2 - \tilde{x}_1}$ is increasing in both its variables.

Hence, for any $\tilde{x}_1, \tilde{x}_2 \in [x_m, x_{m+1}]$ with $\tilde{x}_1 \neq \tilde{x}_2$, we have that $\frac{z_m - z_{m-1}}{x_m - x_{m-1}} \leq \frac{f_{\text{opt}}(\tilde{x}_2) - f_{\text{opt}}(\tilde{x}_1)}{\tilde{x}_2 - \tilde{x}_1} \leq \frac{z_{m+2} - z_{m+1}}{x_{m+2} - x_{m+1}}$. This directly implies the desired result $L_m(f_{\text{opt}}) \leq L(f_{\text{cano}})$.

B.4 Proof of Theorem 8.3

B.4.1 Existence

We rewrite the problem in (8.24) as an unconstrained minimization problem

$$\mathcal{V}_{\text{hyb}} = \arg \min_{f \in \mathcal{M}_{\mathbb{D}^2}(\mathbb{R})} \left(\sum_{m=1}^M \mathbb{E}(f(x_m), y_m) + \lambda \|D^2 f\|_{\mathcal{M}} + i_{L(f) \leq \bar{L}} \right), \quad (\text{B.47})$$

where i_E denotes the characteristic function of the set E and is defined as

$$i_E(f) \triangleq \begin{cases} 0 & f \in E \\ +\infty & \text{otherwise} \end{cases}. \quad (\text{B.48})$$

To prove the existence of a minimizer, we use a standard technique in convex analysis which involves the generalized Weierstrass theorem [301] to show that the cost functional of (B.47) is coercive and lower semicontinuous (in the weak*-topology), which is a sufficient condition for the existence of a solution.

The cost functional in (8.24) consists of three terms: (i) an empirical loss term $H(f) \triangleq \sum_{m=1}^M \mathbb{E}(f(x_m), y_m)$; (ii) a second-order TV regularization term $R(f) \triangleq \lambda \|D^2 f\|_{\mathcal{M}}$; and (iii) a Lipschitz constraint i_E , where $E \triangleq \{L(f) \leq \bar{L}\}$. It is known (see [45] for a more general statement) that the functional $H(f) + R(f)$ is coercive and weak*-lower semicontinuous. This, together with the non-negativity of i_E , yields the coercivity of the total cost. The only missing item is the weak*-lower semicontinuity of i_E , for which it is sufficient to prove that E is a closed set for the weak*-topology.

Let $f_n \in \mathcal{M}_{D^2}(\mathbb{R})$ be a sequence of functions with $L(f_n) \leq \bar{L}$ converging in the weak*-topology to $f_{\text{lim}} \in \mathcal{M}_{D^2}(\mathbb{R})$. To prove the weak*-closedness of E , we need to show that $L(f_{\text{lim}}) \leq \bar{L}$, which is equivalent to $|f_{\text{lim}}(a) - f_{\text{lim}}(b)| \leq \bar{L}|a - b|$ for any $a, b \in \mathbb{R}$.

For any $n \in \mathbb{N}$, we have that

$$|f_{\text{lim}}(a) - f_{\text{lim}}(b)| \leq |f_{\text{lim}}(a) - f_n(a)| + |f_n(a) - f_n(b)| + |f_n(b) - f_{\text{lim}}(b)|. \quad (\text{B.49})$$

Using the weak*-continuity of the sampling functionals $\delta(\cdot - a)$ and $\delta(\cdot - b)$ in $\mathcal{M}_{D^2}(\mathbb{R})$ (see Example 3.2), we deduce that $f_n(a) \rightarrow f_{\text{lim}}(a)$ and $f_n(b) \rightarrow f_{\text{lim}}(b)$. Moreover, we have the upper bound $|f_n(a) - f_n(b)| \leq \bar{L}|a - b|$ for any $n \in \mathbb{N}$. We get the desired bound by letting $n \rightarrow +\infty$ in (B.49).

B.4.2 Form of the Solution Set

Now that we have proved the existence of a solution $f_0^* \in \mathcal{V}_{\text{hyb}}$, we can apply a standard argument based on the strict convexity of $E(\cdot, \cdot)$ (see Proposition 7.7) to deduce that for any $f^* \in \mathcal{V}_{\text{hyb}}$, we have that $f^*(x_m) = f_0^*(x_m)$ for $m = 1, \dots, M$. Hence, the original Problem (8.24) is equivalent to

$$\mathcal{V}_{\text{hyb}} = \arg \min_{f \in \mathcal{M}_{D^2}(\mathbb{R})} \|D^2 f\|_{\mathcal{M}} \quad \text{s.t.} \quad \begin{cases} L(f) \leq \bar{L} \\ f(x_m) = f_0^*(x_m) \quad m = 1, \dots, M \end{cases}. \quad (\text{B.50})$$

Since $f_0^* \in \mathcal{V}_{\text{hyb}}$, we deduce that

$$L_0 \triangleq \max_{2 \leq m \leq M} \left| \frac{f_0^*(x_m) - f_0^*(x_{m-1})}{x_m - x_{m-1}} \right| \leq L(f_0^*) \leq \bar{L}.$$

Yet, Item 5 in Theorem 8.2 implies that any solution f^* of the problem

$$\arg \min_{f \in \mathcal{M}_{\mathbb{D}^2}(\mathbb{R})} \|D^2 f\|_{\mathcal{M}} \quad \text{s.t.} \quad f(x_m) = f_0^*(x_m) \quad m = 1, \dots, M \quad (\text{B.51})$$

is a solution of (8.17) with $z_m = f_0^*(x_m)$. Hence, by Item 3 of Theorem 8.2, we have that $L(f^*) = L_0 \leq \bar{L}$. This means that adding the Lipschitz constraint $L(f) \leq \bar{L}$ does not change the solution set of Problem (B.51). Hence, we have that

$$\mathcal{V}_{\text{hyb}} = \arg \min_{f \in \mathcal{M}_{\mathbb{D}^2}(\mathbb{R})} \|D^2 f\|_{\mathcal{M}} \quad \text{s.t.} \quad f(x_m) = f_0^*(x_m), \quad 1 \leq m \leq M. \quad (\text{B.52})$$

The solution set of (B.52) has been fully described in Theorem 7.1 in the previous chapter, which yields the announced characterization.

B.5 Proof of Theorem 10.2

We start by providing the necessary tools before going into the proof of Theorem 10.2; for more details on these tools, we refer to Chapter 2. Specifically, we first describe the topological structure of the search space $\mathcal{X}_L(\mathbb{T}_M)$. We then identify the set of extreme points of the RI-TV unit ball. Finally, we provide a full characterization of the solution set \mathcal{V} , which concludes the proof of Theorem 10.2.

B.5.1 Search Space

We first recall some concepts and results from Chapter 2. The only difference is that $\mathbb{T} = [0, 2\pi]$ is replaced with $\mathbb{T}_M = [0, M]$ in our current setting, which does not affect the validity of the elements presented here.

The space of periodic finite Radon measures is denoted by $\mathcal{M}(\mathbb{T}_M)$. It is a Banach space equipped with the TV norm $\|\cdot\|_{\mathcal{M}}$ defined in (2.18).

The native space of L is defined as $\mathcal{M}_L(\mathbb{T}_M) \triangleq \{f \in \mathcal{S}'(\mathbb{T}_M) : \|L\{f\}\|_{\mathcal{M}} < +\infty\}$. By Theorem 2.4, $\mathcal{M}_L(\mathbb{T}_M)$ is isometrically isomorphic to $\mathcal{M}_0(\mathbb{T}_M) \times \mathbb{R}$, where $\mathcal{M}_0(\mathbb{T}_M) \triangleq \{w \in \mathcal{M}(\mathbb{T}_M) : \langle w, 1 \rangle = 0\}$ is the space of Radon measures with zero mean, via the inverse mappings

$$\begin{aligned} \mathcal{M}_L(\mathbb{T}_M) &\rightarrow \mathcal{M}_0(\mathbb{T}_M) \times \mathbb{R} : f \mapsto (L\{f\}, \langle f, 1 \rangle), \\ \mathcal{M}_0(\mathbb{T}_M) \times \mathbb{R} &\rightarrow \mathcal{M}_L(\mathbb{T}_M) : (w, a) \mapsto L^\dagger\{w\} + a, \end{aligned} \quad (\text{B.53})$$

where L^\dagger is the pseudoinverse of L defined in 2.6. Finally, we recall that the Green's function of $L = D^{N_d}$, defined as $g_L \triangleq L^\dagger\{\text{III}\}$, is a continuous periodic function for all integers $N_d \geq 2$ (Proposition 2.7)

We are now ready to characterize the topological structure of the search space $\mathcal{X}_L(\mathbb{T}_M)$ defined in (10.24).

Proposition B.1. *The search space $\mathcal{X}_L(\mathbb{T}_M)$ can be expressed as*

$$\mathcal{X}_L(\mathbb{T}_M) = \mathcal{M}_L(\mathbb{T}_M) \times \mathcal{M}_L(\mathbb{T}_M). \quad (\text{B.54})$$

Moreover, the mapping

$$\begin{aligned} T : \mathcal{X}_L(\mathbb{T}_M) &\rightarrow \mathcal{M}_0(\mathbb{T}_M)^2 \times \mathbb{R}^2 \\ T(\mathbf{r}) &= (\mathbf{L}\{r_1\}, \mathbf{L}\{r_2\}, \langle r_1, 1 \rangle, \langle r_2, 1 \rangle) \end{aligned} \quad (\text{B.55})$$

is an isomorphism between $\mathcal{X}_L(\mathbb{T}_M)$ and $\mathcal{M}_0(\mathbb{T}_M)^2 \times \mathbb{R}^2$ whose inverse is

$$\begin{aligned} T^{-1} : \mathcal{M}_0(\mathbb{T}_M)^2 \times \mathbb{R}^2 &\rightarrow \mathcal{M}_L(\mathbb{T}_M) \\ T^{-1}(\mathbf{w}, \mathbf{a}) &= (\mathbf{L}^\dagger\{w_1\} + a_1, \mathbf{L}^\dagger\{w_2\} + a_2). \end{aligned} \quad (\text{B.56})$$

Proof. Let $\mathbf{r} = (r_1, r_2) \in \mathcal{X}_L(\mathbb{T}_M)$. We have that

$$\begin{aligned} \mathcal{R}(\mathbf{L}\{\mathbf{r}\}) &= \sup_{\substack{\varphi \in \mathcal{S}(\mathbb{T}_M)^2 \\ \|\varphi\|_{2, L_\infty} = 1}} (\langle \mathbf{L}\{r_1\}, \varphi_1 \rangle + \langle \mathbf{L}\{r_2\}, \varphi_2 \rangle) \\ &\geq \sup_{\substack{\varphi_1 \in \mathcal{S}(\mathbb{T}_M) \\ \|(\varphi_1, 0)\|_{2, L_\infty} = 1}} \langle \mathbf{L}\{r_1\}, \varphi_1 \rangle = \sup_{\substack{\varphi_1 \in \mathcal{S}(\mathbb{T}_M) \\ \|\varphi_1\|_\infty = 1}} \langle \mathbf{L}\{r_1\}, \varphi_1 \rangle = \|\mathbf{L}\{r_1\}\|_{\mathcal{M}}, \end{aligned} \quad (\text{B.57})$$

from which we deduce that $r_1 \in \mathcal{M}_L(\mathbb{T}_M)$. Similarly, we get that $r_2 \in \mathcal{M}_L(\mathbb{T}_M)$ and, hence, we have that $\mathcal{X}_L(\mathbb{T}_M) \subseteq (\mathcal{M}_L(\mathbb{T}_M))^2$. For the reverse inclusion, let $r_1, r_2 \in \mathcal{M}_L(\mathbb{T}_M)$. Using the inequalities $\|\varphi\|_{2, L_\infty} \geq \|\varphi_i\|_\infty$ for $i = 1, 2$, we deduce that

$$|\langle \mathbf{L}\{r_i\}, \varphi_i \rangle| \leq \|\mathbf{L}\{r_i\}\|_{\mathcal{M}} \|\varphi_i\|_{L_\infty} \leq \|\mathbf{L}\{r_i\}\|_{\mathcal{M}} \|\varphi\|_{2, L_\infty}. \quad (\text{B.58})$$

Hence, we have that

$$\langle \mathbf{L}\{r_1\}, \varphi_1 \rangle + \langle \mathbf{L}\{r_2\}, \varphi_2 \rangle \leq (\|\mathbf{L}\{r_1\}\|_{\mathcal{M}} + \|\mathbf{L}\{r_2\}\|_{\mathcal{M}}) \|\varphi\|_{2, L_\infty}, \quad (\text{B.59})$$

which implies that

$$\mathcal{R}(\mathbf{L}\{\mathbf{r}\}) \leq \|\mathbf{L}\{r_1\}\|_{\mathcal{M}} + \|\mathbf{L}\{r_2\}\|_{\mathcal{M}} < +\infty. \quad (\text{B.60})$$

Hence, we have the inclusion $\mathbf{r} \in \mathcal{X}_L(\mathbb{T}_M)$.

Following (B.60) and (B.57), we deduce that the norm topology of $\mathcal{X}_L(\mathbb{T}_M)$ is equivalent to the product topology induced from $\mathcal{M}_L(\mathbb{T}_M) \times \mathcal{M}_L(\mathbb{T}_M)$. This, together with the fact that $\mathcal{M}_L(\mathbb{T}_M)$ is isometrically isomorphic to $\mathcal{M}_0(\mathbb{T}_M) \times \mathbb{R}$, implies that T is an isomorphism. Its inverse is readily deduced from (B.53). \square

B.5.2 Extreme Points of the RI-TV Unit Ball

Our strategy to characterize the solution set \mathcal{V} defined in (10.23) consists in applying the main result of Boyer *et al.* [61], which requires the knowledge of the form of extreme points of the unit ball of the regularization functional. To that end, we prove that the extreme points of the RI-TV unit ball are vector-valued Dirac combs.

Proposition B.2. *An element $\mathbf{w}^* \in \mathcal{M}(\mathbb{T}_M)^2$ is an extreme point of the RI-TV unit ball $B = \{\mathbf{w} \in \mathcal{M}(\mathbb{T}_M)^2 : \mathcal{R}(\mathbf{w}) = 1\}$ if and only if it is a vector-valued Dirac comb of the form $\mathbf{w}^* = \mathbf{a}\text{III}_M(\cdot - t_0)$ for some $t_0 \in \mathbb{T}_M$ and $\mathbf{a} \in \mathbb{R}^2$ with $\|\mathbf{a}\|_2 = 1$.*

Proof. Assume by contradiction that there exists an extreme point \mathbf{w}^* of B that is not a Dirac comb. This implies that there exists an interval $I \subseteq \mathbb{T}_M$ such that $\mathbf{w}_1 = \mathbf{w}^* \mathbb{1}_I$ and $\mathbf{w}_2 = \mathbf{w}^* \mathbb{1}_{I^c}$ are both nonzero Radon measures that satisfy $\mathbf{w}^* = \mathbf{w}_1 + \mathbf{w}_2$. Due to their disjoint support, we have that $\mathcal{R}(\mathbf{w}^*) = \mathcal{R}(\mathbf{w}_1) + \mathcal{R}(\mathbf{w}_2)$. We now define the measures $\mathbf{w}_+ = (1 + \epsilon)\mathbf{w}_1 + (1 - \delta)\mathbf{w}_2$ and $\mathbf{w}_- = (1 - \epsilon)\mathbf{w}_1 + (1 + \delta)\mathbf{w}_2$, where $\epsilon, \delta > 0$ are small constants such that $\epsilon\mathcal{R}(\mathbf{w}_1) = \delta\mathcal{R}(\mathbf{w}_2)$. By observing that $\mathcal{R}(\mathbf{w}_+) = \mathcal{R}(\mathbf{w}_-) = 1$ and $\mathbf{w}^* = \frac{\mathbf{w}_+ + \mathbf{w}_-}{2}$, we conclude that \mathbf{w}^* is not an extreme point of B , which yields a contradiction. Hence, the extreme points of B can only be vector-valued Dirac combs.

To prove the reverse inclusion, let $\mathbf{w}^* = \mathbf{a}\text{III}_M(\cdot - t_0)$ with $\|\mathbf{a}\|_2 = 1$. We now prove that \mathbf{w}^* is an extreme point of B . Assume that there exist $\mathbf{w}_1, \mathbf{w}_2 \in B$ such that $\mathbf{w}^* = \frac{1}{2}(\mathbf{w}_1 + \mathbf{w}_2)$. We define the measure $\mathbf{w}_0 = \mathbf{w}_1 \mathbb{1}_{t \neq t_0} \in \mathcal{M}(\mathbb{T}_M)^2$ so that $\mathbf{w}_1 = \mathbf{w}_0 + \mathbf{a}_1 \text{III}_M(\cdot - t_0)$ for some $\mathbf{a}_1 \in \mathbb{R}^2$. We then must have $\mathbf{w}_2 = (-\mathbf{w}_0) + \mathbf{a}_2 \text{III}_M(\cdot - t_0)$ with $\mathbf{a} = \frac{1}{2}(\mathbf{a}_1 + \mathbf{a}_2)$. The construction implies that

$$1 = \mathcal{R}(\mathbf{w}_i) = \mathcal{R}(\mathbf{w}_0) + \|\mathbf{a}_i\|_2 \geq \|\mathbf{a}_i\|_2, \quad i = 1, 2. \tag{B.61}$$

This, together with the triangle inequality, yields

$$2 = \|\mathbf{a}\|_2 \leq \|\mathbf{a}_1\|_2 + \|\mathbf{a}_2\|_2 \leq 1 + 1 = 2. \tag{B.62}$$

Hence, all inequalities must be saturated. In particular, we must have that $\mathbf{w}_0 = \mathbf{0}$ and $\|\mathbf{a}\|_2 = \frac{1}{2}(\|\mathbf{a}_1\|_2 + \|\mathbf{a}_2\|_2)$. Finally, using the strict convexity of the ℓ_2 norm, we conclude that $\mathbf{a} = \mathbf{a}_1 = \mathbf{a}_2$ and, thus, that $\mathbf{w}_1 = \mathbf{w}_2$, which in turn implies that \mathbf{w}^* is an extreme point of B . \square

B.5.3 Representer Theorem

We now provide a complete characterization of the solution set \mathcal{V} in (10.23) from which we readily deduce Theorem 10.2 as a corollary.

Theorem B.1. *The solution set (10.23) is nonempty, convex, and weak*-compact. Moreover, any extreme point \mathbf{r}^* of \mathcal{V} is a periodic L-spline that satisfies (10.25).*

Proof. We define the cost functional $E : \mathcal{M}(\mathbb{T}_M)^2 \times \mathbb{R}^2 \rightarrow \mathbb{R} \cup \{+\infty\}$ as

$$E(\mathbf{w}, \mathbf{a}) \triangleq \sum_{m=0}^{M-1} \|\boldsymbol{\nu}_m(\mathbf{w}) + \mathbf{a} - \mathbf{p}[m]\|_2^2 + \sum_{i=1}^2 \iota_{\mathcal{M}_0(\mathbb{T}_M)}(w_i), \quad (\text{B.63})$$

where ι_A denotes the characteristic function of the set A , and $\boldsymbol{\nu}_m = (\nu_{m,1}, \nu_{m,2})$ with

$$\begin{aligned} \nu_{m,i}(\mathbf{w}) &\triangleq \left(\mathbf{L}^\dagger \{w_i\}(t) \right) \Big|_{t=m} \\ &= \langle \mathbf{L}^\dagger \{w_i\}, \mathbf{III}(\cdot - m) \rangle = \langle w_i, \mathbf{L}^{\dagger*} \{ \mathbf{III}(\cdot - m) \} \rangle = \langle w_i, g_L(m - \cdot) \rangle \end{aligned} \quad (\text{B.64})$$

for $i = 1, 2$. Due to the inclusion $g_L \in \mathcal{C}(\mathbb{T}_M)$ (Proposition 2.7), we have that $\boldsymbol{\nu}_m$ is weak*-continuous over $\mathcal{M}_L(\mathbb{T}_M)$ (see Remark 3.4).

Then, we formulate the optimization problem

$$\tilde{\mathcal{V}} \triangleq \arg \min_{\substack{\mathbf{w} \in \mathcal{M}(\mathbb{T}_M)^2 \\ \mathbf{a} \in \mathbb{R}^2}} \left(\underbrace{E(\mathbf{w}, \mathbf{a}) + \lambda \mathcal{R}(\mathbf{w})}_{\mathcal{J}(\mathbf{w}, \mathbf{a})} \right), \quad (\text{B.65})$$

whose solution set $\tilde{\mathcal{V}}$ we characterize using standard convex-optimization techniques (see for example the proof of [300, Theorem 3]). One readily verifies that the cost functional \mathcal{J} is weak*-lower semicontinuous (due to the fact that $g_L \in \mathcal{C}(\mathbb{T}_M)$), coercive, and convex. This implies that $\tilde{\mathcal{V}}$ is nonempty. Next, we use the strict convexity of $\|\cdot - \mathbf{p}[m]\|_2^2$ for $m = 0, \dots, M-1$ to deduce the existence of $\mathbf{p}_0[m] \in \mathbb{R}^2$, $m = 0, \dots, M-1$, and $\mathbf{a}_0 \in \mathbb{R}^2$ such that

$$\tilde{\mathcal{V}} = \arg \min_{\substack{\mathbf{w} \in \mathcal{M}(\mathbb{T}_M)^2 \\ \mathbf{a} \in \mathbb{R}^2}} \mathcal{R}(\mathbf{w}) \quad \text{s.t.} \quad \begin{cases} \mathbf{a} = \mathbf{a}_0 \\ \boldsymbol{\nu}_m(\mathbf{w}) = \mathbf{p}_0[m] \quad \forall m \\ \langle w_i, 1 \rangle = 0 \quad i = 1, 2 \end{cases} \quad (\text{B.66})$$

(see for example Proposition 7.7). We can go one step further by denoting the minimal regularization cost by $\mathcal{R}^0 > 0$ and rewriting $\tilde{\mathcal{V}}$ as $\tilde{\mathcal{V}} = \mathcal{W} \times \{\mathbf{a}_0\}$, where

$$\mathcal{W} \triangleq \left\{ \mathbf{w} \in \mathcal{M}(\mathbb{T}_M)^2 : \mathcal{R}(\mathbf{w}) \leq \mathcal{R}^0, \boldsymbol{\nu}_m(\mathbf{w}) = \mathbf{p}_0[m] \quad \forall m, \langle w_i, 1 \rangle = 0 \quad \forall i \in \{1, 2\} \right\}. \quad (\text{B.67})$$

The set \mathcal{W} is the intersection of an RI-TV ball (which is weak*-compact due to the Banach-Alaoglu theorem [15, Theorem 3.15]) with an affine subspace of codimension $(2M + 2)$. Hence, \mathcal{W} is weak*-compact, and since $\tilde{\mathcal{V}} = \mathcal{W} \times \{\mathbf{a}_0\}$, so is $\tilde{\mathcal{V}}$. Finally, we use the main result of [61] together with Proposition B.2 to deduce that any extreme point of $\tilde{\mathcal{V}}$ can be written as $(\mathbf{w}^*, \mathbf{a}_0)$, where $\mathbf{w}^* = \sum_{k=1}^K \mathbf{a}_k \mathbf{III}(\cdot - t_k)$ with $K \leq 2M + 2$ (the

total number of linear constraints in the definition of \mathcal{W}) for some $\mathbf{a}_k \in \mathbb{R}^2$ and $t_k \in \mathbb{T}_M$.

The final step is to observe that the isomorphism T defined in Proposition B.1 allows us to write

$$E(T(\mathbf{r})) = \sum_{m=0}^{M-1} \|\mathbf{r}(t)|_{t=m} - \mathbf{p}[m]\|_2^2 \quad (\text{B.68})$$

for any $\mathbf{r} \in \mathcal{X}_L(\mathbb{T}_M)$, from which we conclude that $\tilde{\mathcal{V}} = T(\mathcal{V})$. Hence, the solution set $\mathcal{V} = T^{-1}(\tilde{\mathcal{V}})$ is nonempty, convex, and weak*-compact, and any extreme point \mathbf{r}^* of \mathcal{V} induces an extreme point $(\mathbf{w}^*, \mathbf{a}^*) = T(\mathbf{r}^*)$ of $\tilde{\mathcal{V}}$. In particular, we have that

$$\mathbf{L}\{\mathbf{r}^*\} = \mathbf{w}^* = \sum_{k=0}^{K-1} \mathbf{a}_k \text{III}(\cdot - t_k) \quad (\text{B.69})$$

where $a_k \in \mathbb{R}^2$, $t_k \in \mathbb{T}_M$, and $K \leq 2M + 2$, which concludes the proof. \square

B.6 Proof of Theorem 10.3

Using the isomorphism between $\mathcal{X}_{L_i}(\mathbb{T}_M)$ and $\mathcal{M}_0(\mathbb{T}_M)^2 \times \mathbb{R}^2$ for $i = 1, 2$, which we denote by T_{L_i} , from Proposition B.1, we deduce the existence a bijection between \mathcal{V}_{hyb} and the solution set

$$\tilde{\mathcal{V}}_{\text{hyb}} \triangleq \arg \min_{\substack{\mathbf{w}_1, \mathbf{w}_2 \in \mathcal{M}_0(\mathbb{T})^2 \\ \mathbf{a}_1, \mathbf{a}_2 \in \mathbb{R}^2}} \left(E(\mathbf{w}_1, \mathbf{a}_1, \mathbf{w}_2, \mathbf{a}_2) + \lambda_1 \mathcal{R}(\mathbf{w}_1) + \lambda_2 \mathcal{R}(\mathbf{w}_2) \right), \quad (\text{B.70})$$

where the data fidelity cost $E : (\mathcal{M}(\mathbb{T})^2 \times \mathbb{R}^2)^2 \rightarrow \mathbb{R}_{\geq 0}$ satisfies

$$E(T_{L_1}(\mathbf{r}_1), T_{L_1}(\mathbf{r}_2)) = \sum_{m=0}^{M-1} \|\mathbf{r}_1(t)|_{t=m} + \mathbf{r}_2(t)|_{t=m} - \mathbf{p}[m]\|_2^2. \quad (\text{B.71})$$

This implies that there is a bijection between \mathcal{V}_{hyb} and $\tilde{\mathcal{V}}_{\text{hyb}}$. The last step is to note that for any extreme point $(\mathbf{w}_1^*, \mathbf{w}_2^*)$ of the unit ball $\{(\mathbf{w}_1, \mathbf{w}_2) \in \mathcal{M}(\mathbb{T})^4 : \lambda_1 \mathcal{R}(\mathbf{w}_1) + \lambda_2 \mathcal{R}(\mathbf{w}_2) \leq 1\}$, we have that $\mathbf{w}_1^* = \mathbf{0}$ or $\mathbf{w}_2^* = \mathbf{0}$. This together with [300, Theorem 3] concludes the proof.

Bibliography

- [1] Y. Chi, L. Scharf, A. Pezeshki, and A. Calderbank, “Sensitivity to basis mismatch in compressed sensing”, *IEEE Transactions on Signal Processing*, vol. 59, no. 5, pp. 2182–2195, May 2011.
- [2] M. Guerquin-Kern, L. Lejeune, K. Pruessmann, and M. Unser, “Realistic analytical phantoms for parallel magnetic resonance imaging”, *IEEE Transactions on Medical Imaging*, vol. 31, no. 3, pp. 626–636, Mar. 2012.
- [3] B. Adcock and A. Hansen, “Generalized sampling and infinite-dimensional compressed sensing”, *Foundations of Computational Mathematics*, vol. 16, no. 5, pp. 1263–1323, Aug. 2015.
- [4] A. Tikhonov, “Solution of incorrectly formulated problems and the regularization method”, *Soviet Mathematics*, vol. 4, pp. 1035–1038, 1963.
- [5] S. Mallat, *A Wavelet Tour of Signal Processing*. Elsevier, 2009.
- [6] S. Chen, D. Donoho, and M. Saunders, “Atomic decomposition by basis pursuit”, *SIAM Review*, vol. 43, no. 1, pp. 129–159, Jan. 2001.
- [7] R. Tibshirani, “Regression shrinkage and selection via the lasso”, *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, Jan. 1996.
- [8] D. Donoho, “Compressed sensing”, *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [9] E. Candès, “Compressive sampling”, in *Proceedings of the International Congress of Mathematicians*, vol. 3, Madrid, Spain: European Mathematical Society Publishing House, 2006, pp. 1433–1452.
- [10] Y. Eldar and G. Kutyniok, Eds., *Compressed Sensing*. Cambridge University Press, 2009.
- [11] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*. Springer Basel AG, Aug. 20, 2013.
- [12] K. Bredies and H. Pikkarainen, “Inverse problems in spaces of measures”, *ESAIM: Control, Optimisation and Calculus of Variations*, vol. 19, no. 1, pp. 190–218, Mar. 2012.

-
- [13] E. Candès and C. Fernandez-Granda, “Super-resolution from noisy data”, *Journal of Fourier Analysis and Applications*, vol. 19, no. 6, pp. 1229–1254, Aug. 2013.
- [14] M. Unser, J. Fageot, and J. Ward, “Splines are universal solutions of linear inverse problems with generalized TV regularization”, *SIAM Review*, vol. 59, no. 4, pp. 769–793, Dec. 2017.
- [15] W. Rudin, *Functional Analysis*. New York: McGraw-Hill, 1991.
- [16] T. Hastie, R. Tibshirani, and M. Wainwright, *Statistical Learning with Sparsity*. Chapman and Hall/CRC, May 2015.
- [17] N. Parikh, “Proximal algorithms”, *Foundations and Trends® in Optimization*, vol. 1, no. 3, pp. 127–239, 2014.
- [18] A. Beck and M. Teboulle, “A fast iterative shrinkage-thresholding algorithm for linear inverse problems”, *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, Jan. 2009.
- [19] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, “Distributed optimization and statistical learning via the alternating direction method of multipliers”, *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2010.
- [20] M. Unser, J. Fageot, and H. Gupta, “Representer theorems for sparsity-promoting ℓ_1 regularization”, *IEEE Transactions on Information Theory*, vol. 62, no. 9, pp. 5167–5180, Sep. 2016.
- [21] R. Tibshirani, “The LASSO problem and uniqueness”, *Electronic Journal of Statistics*, vol. 7, no. 0, pp. 1456–1490, 2013.
- [22] J. Fageot and M. Simeoni, “TV-based reconstruction of periodic functions”, *Inverse Problems*, vol. 36, no. 11, p. 115 015, Nov. 2020.
- [23] I. Schoenberg, “Contributions to the problem of approximation of equidistant data by analytic functions”, *Quarterly of Applied Mathematics*, vol. 4, no. 1, pp. 45–99, 1946.
- [24] C. de Boor, *A Practical Guide to Splines*. Springer-Verlag GmbH, Nov. 29, 2001.
- [25] M. Unser, “Splines: A perfect fit for signal and image processing”, *IEEE Signal Processing Magazine*, vol. 16, no. 6, pp. 22–38, Nov. 1999, IEEE-SPS best paper award.
- [26] I. Ekeland and R. Temam, *Convex Analysis and Variational Problems*. Philadelphia: Society for Industrial and Applied Mathematics, 1976.
- [27] S. Lefkimmiatis, J. Ward, and M. Unser, “Hessian Schatten-norm regularization for linear inverse problems”, *IEEE Transactions on Image Processing*, vol. 22, no. 5, pp. 1873–1888, May 2013.
- [28] L. Rudin, S. Osher, and E. Fatemi, “Nonlinear total variation based noise removal algorithms”, *Physica D: Nonlinear Phenomena*, vol. 60, no. 1-4, pp. 259–268, Nov. 1992.

BIBLIOGRAPHY

- [29] P. Bühlmann and S. van de Geer, *Statistics for High-Dimensional Data*. Springer Berlin Heidelberg, 2011.
- [30] I. Rish and G. Grabarnik, *Sparse Modeling: Theory, Algorithms, and Applications*. Boca Raton, FL: CRC Press, 2015.
- [31] E. Candès, J. Romberg, and T. Tao, “Stable signal recovery from incomplete and inaccurate measurements”, *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006.
- [32] E. Candès, “The restricted isometry property and its implications for compressed sensing”, *Comptes Rendus Mathématique*, vol. 346, no. 9-10, pp. 589–592, May 2008.
- [33] G. Tang, B. Bhaskar, P. Shah, and B. Recht, “Compressed sensing off the grid”, *IEEE Transactions on Information Theory*, vol. 59, no. 11, pp. 7465–7490, Nov. 2013.
- [34] B. Adcock, A. Hansen, C. Poon, and B. Roman, “Breaking the coherence barrier: A new theory for compressed sensing”, *Forum of Mathematics, Sigma*, vol. 5, 2017.
- [35] Y. Traonmilin, G. Puy, R. Gribonval, and M. Davies, “Compressed sensing in Hilbert spaces”, in *Compressed Sensing and its Applications*, Springer International Publishing, 2017, pp. 359–384.
- [36] A. Bhandari and Y. Eldar, “Sampling and super resolution of sparse signals beyond the Fourier domain”, *IEEE Transactions on Signal Processing*, vol. 67, no. 6, pp. 1508–1521, Mar. 2019.
- [37] M. Lustig, D. Donoho, and J. Pauly, “Sparse MRI: The application of compressed sensing for rapid MR imaging”, *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [38] P. Combettes and V. Wajs, “Signal recovery by proximal forward-backward splitting”, *Multiscale Modeling & Simulation*, vol. 4, no. 4, pp. 1168–1200, Jan. 2005.
- [39] A. Beck and M. Teboulle, “Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems”, *IEEE Transactions on Image Processing*, vol. 18, no. 11, pp. 2419–2434, Nov. 2009.
- [40] A. Chambolle and T. Pock, “A first-order primal-dual algorithm for convex problems with applications to imaging”, *Journal of Mathematical Imaging and Vision*, vol. 40, no. 1, pp. 120–145, Dec. 2010.
- [41] L. Condat, “A primal–dual splitting method for convex optimization involving lipschitzian, proximable and linear composite terms”, *Journal of Optimization Theory and Applications*, vol. 158, no. 2, pp. 460–479, Dec. 2012.
- [42] B. Schölkopf, R. Herbrich, and A. Smola, “A generalized representer theorem”, in *Lecture Notes in Computer Science*, D. Helmbold and R. Williamson, Eds., ser. LNCS, Max-Planck-Gesellschaft, vol. 2111, Berlin, Germany: Springer, 2001, pp. 416–426.

-
- [43] G. Wahba, *Spline Models for Observational Data*. Philadelphia, USA: Society for Industrial and Applied Mathematics, Jan. 1990.
- [44] A. Berlinet and C. Thomas-Agnan, *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Springer US, Jun. 2011, 355 pp.
- [45] H. Gupta, J. Fageot, and M. Unser, “Continuous-domain solutions of linear inverse problems with Tikhonov *versus* generalized TV regularization”, *IEEE Transactions on Signal Processing*, vol. 66, no. 17, pp. 4670–4684, Sep. 2018.
- [46] A. Badoual, J. Fageot, and M. Unser, “Periodic splines and Gaussian processes for the resolution of linear inverse problems”, *IEEE Transactions on Signal Processing*, vol. 66, no. 22, pp. 6047–6061, Nov. 2018.
- [47] Y. de Castro and F. Gamboa, “Exact reconstruction using Beurling minimal extrapolation”, *Journal of Mathematical Analysis and Applications*, vol. 395, no. 1, pp. 336–354, Nov. 2012.
- [48] B. Bhaskar, G. Tang, and B. Recht, “Atomic norm denoising with applications to line spectral estimation”, *IEEE Transactions on Signal Processing*, vol. 61, no. 23, pp. 5987–5999, Dec. 2013.
- [49] E. Candès and C. Fernandez-Granda, “Towards a mathematical theory of super-resolution”, *Communications on Pure and Applied Mathematics*, vol. 67, no. 6, pp. 906–956, Apr. 2013.
- [50] H. Pan, T. Blu, and M. Vetterli, “Towards generalized FRI sampling with an application to source resolution in radioastronomy”, *IEEE Transactions on Signal Processing*, vol. 65, no. 4, pp. 821–835, Feb. 2017.
- [51] Q. Denoyelle, V. Duval, G. Peyré, and E. Soubies, “The sliding Frank-Wolfe algorithm for the BLASSO”, in *Proceedings of the Workshop on Signal Processing with Adaptive Sparse Structured Representations (SPARS’19)*, paper no. 172, Toulouse, French Republic, Jul. 2019.
- [52] J. Courbot and B. Colicchio, “A fast homotopy algorithm for gridless sparse recovery”, *Inverse Problems*, vol. 37, no. 2, p. 025 002, Jan. 2021.
- [53] A. Eftekhari, J. Tanner, A. Thompson, B. Toader, and H. Tyagi, “Sparse non-negative super-resolution — simplified and stabilised”, *Applied and Computational Harmonic Analysis*, vol. 50, pp. 216–280, Jan. 2021.
- [54] M. Vetterli, P. Marziliano, and T. Blu, “Sampling signals with finite rate of innovation”, *IEEE Transactions on Signal Processing*, vol. 50, no. 6, pp. 1417–1428, Jun. 2002, IEEE-SPS best paper award.
- [55] S. Simeoni, A. Adrien Besson, P. Hurley, and M. Vetterli, “CPGD: Cadzow plug-and-play gradient descent for generalised FRI”, *IEEE Transactions on Signal Processing*, vol. 69, pp. 42–57, 2021.
- [56] R. Schmidt, “Multiple emitter location and signal parameter estimation”, *IEEE Transactions on Antennas and Propagation*, vol. 34, no. 3, pp. 276–280, Mar. 1986.

BIBLIOGRAPHY

- [57] G. Plonka and M. Tasche, “Prony methods for recovery of structured functions”, *GAMM-Mitteilungen*, vol. 37, no. 2, pp. 239–258, Nov. 2014.
- [58] W. Liao and A. Fannjiang, “MUSIC for single-snapshot spectral estimation: Stability and super-resolution”, *Applied and Computational Harmonic Analysis*, vol. 40, no. 1, pp. 33–67, Jan. 2016.
- [59] S. Zuhovickii, “On approximation of real functions in the sense of P.L. Chebyshev”, *AMS Translations of Mathematical Monographs*, vol. 19, no. 2, pp. 221–252, 1962.
- [60] M. Krein and A. Nudelman, *The Markov Moment Problem and Extremal Problems*. American Mathematical Society, Dec. 1977.
- [61] C. Boyer, A. Chambolle, Y. De Castro, V. Duval, F. de Gournay, and P. Weiss, “On representer theorems and convex regularization”, *SIAM Journal on Optimization*, vol. 29, no. 2, pp. 1260–1281, Jan. 2019.
- [62] C. Fernandez-Granda, “Super-resolution of point sources via convex programming”, *Information and Inference: A Journal of the IMA*, vol. 5, no. 3, pp. 251–303, Apr. 2016.
- [63] J.-M. Azaïs, Y. de Castro, and F. Gamboa, “Spike detection from inaccurate samplings”, *Applied and Computational Harmonic Analysis*, vol. 38, no. 2, pp. 177–195, Mar. 2015.
- [64] G. Tang, B. Bhaskar, and B. Recht, “Near minimax line spectral estimation”, *IEEE Transactions on Information Theory*, vol. 61, no. 1, pp. 499–512, Jan. 2015.
- [65] V. Duval and G. Peyré, “Exact support recovery for sparse spikes deconvolution”, *Foundations of Computational Mathematics*, vol. 15, no. 5, pp. 1315–1355, Oct. 2014.
- [66] —, “Sparse regularization on thin grids I: The LASSO”, *Inverse Problems*, vol. 33, no. 5, p. 055 008, Mar. 2017.
- [67] —, “Sparse spikes super-resolution on thin grids II: The continuous basis pursuit”, *Inverse Problems*, vol. 33, no. 9, p. 095 008, Aug. 2017.
- [68] C. Poon, N. Keriven, and G. Peyré, “Support localization and the Fisher metric for off-the-grid sparse regularization”, in *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, K. Chaudhuri and M. Sugiyama, Eds., ser. Proceedings of Machine Learning Research, vol. 89, PMLR, Apr. 2019, pp. 1341–1350.
- [69] Q. Denoyelle, V. Duval, and G. Peyré, “Support recovery for sparse super-resolution of positive measures”, *Journal of Fourier Analysis and Applications*, vol. 23, no. 5, pp. 1153–1194, Oct. 2017.
- [70] C. Poon and G. Peyré, “Multidimensional sparse super-resolution”, *SIAM Journal on Mathematical Analysis*, vol. 51, no. 1, pp. 1–44, Jan. 2019.

-
- [71] G. Schiebinger, E. Robeva, and B. Recht, “Superresolution without separation”, *Information and Inference: A Journal of the IMA*, vol. 7, no. 1, pp. 1–30, May 2017.
- [72] V. Duval, “A characterization of the non-degenerate source condition in super-resolution”, *Information and Inference: A Journal of the IMA*, Mar. 2019.
- [73] H. García, C. Hernández, M. Junca, and M. Velasco, “Approximate super-resolution of positive measures in all dimensions”, *Applied and Computational Harmonic Analysis*, vol. 52, pp. 251–278, May 2021.
- [74] Y. Chi and M. Ferreira Da Costa, “Harnessing sparsity over the continuum: Atomic norm minimization for superresolution”, *IEEE Signal Processing Magazine*, vol. 37, no. 2, pp. 39–57, Mar. 2020.
- [75] A. Eftekhari and M. Wakin, “Greed is super: A new iterative method for super-resolution”, in *2013 IEEE Global Conference on Signal and Information Processing*, IEEE, Dec. 2013.
- [76] C. Elvira, R. Gribonval, C. Soussen, and C. Herzet, “OMP and continuous dictionaries: Is k-step recovery possible?”, in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, May 2019.
- [77] —, “When does OMP achieve exact recovery with continuous dictionaries?”, *Applied and Computational Harmonic Analysis*, vol. 51, pp. 374–413, Mar. 2021.
- [78] C. Aubel, D. Stotz, and H. Bölcskei, “A theory of super-resolution from short-time Fourier transform measurements”, *Journal of Fourier Analysis and Applications*, vol. 24, no. 1, pp. 45–107, Apr. 2017.
- [79] Y. de Castro, F. Gamboa, D. Henrion, and J.-B. Lasserre, “Exact solutions to super resolution on semi-algebraic domains in higher dimensions”, *IEEE Transactions on Information Theory*, vol. 63, no. 1, pp. 621–630, Jan. 2017.
- [80] P. Catala, V. Duval, and G. Peyré, “A low-rank approach to off-the-grid sparse superresolution”, *SIAM Journal on Imaging Sciences*, vol. 12, no. 3, pp. 1464–1500, Jan. 2019.
- [81] Y. Traonmilin, J. Aujol, and A. Leclaire, “Projected gradient descent for non-convex sparse spike estimation”, *IEEE Signal Processing Letters*, vol. 27, pp. 1110–1114, 2020.
- [82] M. Frank and P. Wolfe, “An algorithm for quadratic programming”, *Naval Research Logistics Quarterly*, vol. 3, no. 1-2, pp. 95–110, Mar. 1956.
- [83] N. Boyd, G. Schiebinger, and B. Recht, “The alternating descent conditional gradient method for sparse inverse problems”, *SIAM Journal on Optimization*, vol. 27, no. 2, pp. 616–639, Jan. 2017.

BIBLIOGRAPHY

- [84] J.-B. Courbot, V. Duval, and B. Legras, “Sparse analysis for mesoscale convective systems tracking”, *Signal Processing: Image Communication*, vol. 85, p. 115–154, Jul. 2020.
- [85] A. Flinth, F. de Gournay, and P. Weiss, “On the linear convergence rates of exchange and continuous methods for total variation minimization”, *Mathematical Programming*, vol. 190, no. 1-2, pp. 221–257, Jun. 2020.
- [86] L. Chizat, “Sparse optimization on measures with over-parameterized gradient descent”, *Mathematical Programming*, Mar. 2021.
- [87] S. Fisher and J. Jerome, “Spline solutions to L^1 extremal problems in one and several variables”, *Journal of Approximation Theory*, vol. 13, no. 1, pp. 73–83, Jan. 1975.
- [88] C. de Boor, “On “best” interpolation”, *Journal of Approximation Theory*, vol. 16, no. 1, pp. 28–42, Jan. 1976.
- [89] K. Bredies and M. Carioni, “Sparsity of solutions for variational inverse problems with finite-dimensional data”, *Calculus of Variations and Partial Differential Equations*, vol. 59, no. 1, pp. 1–26, Dec. 2019.
- [90] A. Flinth and P. Weiss, “Exact solutions of infinite dimensional total-variation regularized problems”, *Information and Inference: A Journal of the IMA*, vol. 8, no. 3, pp. 407–443, Sep. 2019.
- [91] M. Unser and J. Fageot, “Native Banach spaces for splines and variational inverse problems”, *arXiv preprint arXiv:1904.10818*, 2019.
- [92] M. Simeoni, “Functional inverse problems on spheres: Theory, algorithms and applications”, EPFL Thesis no. 7174 (2020), 326 p. Swiss Federal Institute of Technology Lausanne (EPFL), Feb. 2020.
- [93] F. Filbir, K. Schröder, and A. Veselovska, “Super-resolution on the two-dimensional unit sphere”, *arXiv preprint arXiv:2107.10762*, Jul. 2021.
- [94] S. Aziznejad and M. Unser, “Multikernel regression with sparsity constraint”, *SIAM Journal on Mathematics of Data Science*, vol. 3, no. 1, pp. 201–224, 2021.
- [95] M. Simeoni, “Functional penalised basis pursuit on spheres”, *Applied and Computational Harmonic Analysis*, vol. 53, pp. 1–53, Jul. 2021.
- [96] Y. de Castro, V. Duval, and R. Petit, “Towards off-the-grid algorithms for total variation regularized inverse problems”, *arXiv preprint arXiv:2104.06706*, Apr. 2021.
- [97] A. Jarret, J. Fageot, and M. Simeoni, “A fast and scalable polyatomic frank-wolfe algorithm for the LASSO”, *IEEE Signal Processing Letters*, vol. 29, pp. 637–641, 2022.
- [98] L. Schwartz, *Théorie des distributions*. Paris, France: Hermann Paris, 1951, vol. 2.

-
- [99] W. Rudin, *Real and Complex Analysis*. McGraw-Hill Education, May 1, 1986, 432 pp.
- [100] M. Unser and T. Blu, “Fractional splines and wavelets”, *SIAM Review*, vol. 42, no. 1, pp. 43–67, Mar. 2000.
- [101] M. Unser, “Cardinal exponential splines: Part II—Think analog, act digital”, *IEEE Transactions on Signal Processing*, vol. 53, no. 4, pp. 1439–1449, Apr. 2005.
- [102] M. Unser and T. Blu, “Cardinal exponential splines: Part I—Theory and filtering algorithms”, *IEEE Transactions on Signal Processing*, vol. 53, no. 4, pp. 1425–1438, Apr. 2005.
- [103] M. Unser, “A representer theorem for deep neural networks”, *Journal of Machine Learning Research*, vol. 20, no. 110, pp. 1–30, Jan. 2019.
- [104] T. Blu, P. Thévenaz, and M. Unser, “MOMS: Maximal-order interpolation of minimal support”, *IEEE Transactions on Image Processing*, vol. 10, no. 7, pp. 1069–1080, Jul. 2001.
- [105] I. J. Schoenberg, *Cardinal Spline Interpolation*. Society for Industrial and Applied Mathematics, Jan. 1973.
- [106] M. Unser, A. Aldroubi, and M. Eden, “B-Spline signal processing: Part II—Efficient design and applications”, *IEEE Transactions on Signal Processing*, vol. 41, no. 2, pp. 834–848, Feb. 1993.
- [107] —, “B-Spline signal processing: Part I—Theory”, *IEEE Transactions on Signal Processing*, vol. 41, no. 2, pp. 821–833, Feb. 1993, IEEE-SPS best paper award.
- [108] D. Van De Ville, T. Blu, B. Forster, and M. Unser, “Isotropic-polyharmonic B-Splines and wavelets”, in *Proceedings of the 2004 IEEE International Conference on Image Processing (ICIP’04)*, Singapore, Singapore, Oct. 2004, pp. 661–664.
- [109] M. Unser and I. Daubechies, “On the approximation power of convolution-based least squares *versus* interpolation”, *IEEE Transactions on Signal Processing*, vol. 45, no. 7, pp. 1697–1711, Jul. 1997.
- [110] G. Strang and G. Fix, “A Fourier analysis of the finite element variational method”, in *Constructive Aspects of Functional Analysis*, Springer Berlin Heidelberg, 2011, pp. 793–840.
- [111] M. McCann and M. Unser, “Biomedical image reconstruction: From the foundations to deep neural networks”, *Foundations and Trends[®] in Signal Processing*, vol. 13, no. 3, pp. 280–359, Dec. 2019.
- [112] M. Unser, “A unifying representer theorem for inverse problems and machine learning”, *Foundations of Computational Mathematics*, vol. 21, no. 4, pp. 941–960, Sep. 2020.
- [113] M. Reed and B. Simon, *Methods of Modern Mathematical Physics*. Elsevier, 1972.

BIBLIOGRAPHY

- [114] T. Debarre, J. Fageot, H. Gupta, and M. Unser, “B-Spline-based exact discretization of continuous-domain inverse problems with generalized TV regularization”, *IEEE Transactions on Information Theory*, vol. 65, no. 7, pp. 4457–4470, Jul. 2019.
- [115] M. Arigovindan, M. Sühling, P. Hunziker, and M. Unser, “Variational image reconstruction from arbitrarily spaced samples: A fast multiresolution spline solution”, *IEEE Transactions on Image Processing*, vol. 14, no. 4, pp. 450–460, Apr. 2005.
- [116] G. Dantzig, A. Orden, and P. Wolfe, “The generalized simplex method for minimizing a linear form under linear inequality restraints”, *Pacific Journal of Mathematics*, vol. 5, no. 2, pp. 183–195, Jun. 1955.
- [117] D. Luenberger and Y. Ye, *Linear and Nonlinear Programming*. Springer International Publishing, 2021.
- [118] E. Soubies, F. Soulez, M. McCann, T.-a. Pham, L. Donati, T. Debarre, D. Sage, and M. Unser, “Pocket guide to solve inverse problems with GlobalBioIm”, *Inverse Problems*, vol. 35, no. 10, pp. 1–20, Oct. 2019, paper no. 104006.
- [119] E. Candès and T. Tao, “Near-optimal signal recovery from random projections: Universal encoding strategies?”, *IEEE Transactions on Information Theory*, vol. 52, no. 12, pp. 5406–5425, Dec. 2006.
- [120] A. Chambolle, V. Caselles, D. Cremers, M. Novaga, and T. Pock, “An introduction to total variation for image analysis”, in *Theoretical Foundations and Numerical Methods for Sparse Recovery*, M. Fornasier, Ed., DE GRUYTER, Jan. 2010.
- [121] T. Debarre, S. Aziznejad, and M. Unser, “Hybrid-spline dictionaries for continuous-domain inverse problems”, *IEEE Transactions on Signal Processing*, vol. 67, no. 22, pp. 5824–5836, Nov. 2019.
- [122] D. Donoho and P. Stark, “Uncertainty principles and signal recovery”, *SIAM Journal on Applied Mathematics*, vol. 49, no. 3, pp. 906–931, Jun. 1989.
- [123] D. Donoho and X. Huo, “Uncertainty principles and ideal atomic decomposition”, *IEEE Transactions on Information Theory*, vol. 47, no. 7, pp. 2845–2862, 2001.
- [124] M. Elad and A. Bruckstein, “A generalized uncertainty principle and sparse representation in pairs of bases”, *IEEE Transactions on Information Theory*, vol. 48, no. 9, pp. 2558–2567, Sep. 2002.
- [125] R. Gribonval and M. Nielsen, “Sparse representations in unions of bases”, *IEEE Transactions on Information Theory*, vol. 49, no. 12, pp. 3320–3325, Dec. 2003.
- [126] D. Donoho and M. Elad, “Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization”, *Proceedings of the National Academy of Sciences*, vol. 100, no. 5, pp. 2197–2202, Feb. 2003.
- [127] J.-J. Fuchs, “On sparse representations in arbitrary redundant bases”, *IEEE Transactions on Information Theory*, vol. 50, no. 6, pp. 1341–1344, Jun. 2004.

-
- [128] M. Elad, M. Figueiredo, and Y. Ma, “On the role of sparse and redundant representations in image processing”, *Proceedings of the IEEE*, vol. 98, no. 6, pp. 972–982, Jun. 2010.
- [129] D. Donoho and G. Kutyniok, “Microlocal analysis of the geometric separation problem”, *Communications on Pure and Applied Mathematics*, vol. 66, no. 1, pp. 1–47, Aug. 2012.
- [130] H. Rauhut, K. Schnass, and P. Vandergheynst, “Compressed sensing and redundant dictionaries”, *IEEE Transactions on Information Theory*, vol. 54, no. 5, pp. 2210–2219, May 2008.
- [131] J.-L. Starck, M. Elad, and D. Donoho, “Redundant multiscale transforms and their application for morphological component separation”, in *Advances in Imaging and Electron Physics*, Elsevier, 2004, pp. 287–348.
- [132] —, “Image decomposition via the combination of sparse representations and a variational approach”, *IEEE Transactions on Image Processing*, vol. 14, no. 10, pp. 1570–1582, Oct. 2005.
- [133] J.-L. Starck, Y. Moudden, J. Bobin, M. Elad, and D. Donoho, “Morphological component analysis”, in *Wavelets XI*, M. Papadakis, A. F. Laine, and M. A. Unser, Eds., SPIE, Aug. 2005.
- [134] M. Elad, J.-L. Starck, P. Querre, and D. Donoho, “Simultaneous cartoon and texture image inpainting using morphological component analysis (MCA)”, *Applied and Computational Harmonic Analysis*, vol. 19, no. 3, pp. 340–358, Nov. 2005.
- [135] J. Bobin, J.-L. Starck, J. Fadili, Y. Moudden, and D. Donoho, “Morphological component analysis: An adaptive thresholding strategy”, *IEEE Transactions on Image Processing*, vol. 16, no. 11, pp. 2675–2681, Nov. 2007.
- [136] T. Goldstein and S. Osher, “The split Bregman method for L1-regularized problems”, *SIAM Journal on Imaging Sciences*, vol. 2, no. 2, pp. 323–343, Jan. 2009.
- [137] M. Elad, P. Milanfar, and R. Rubinstein, “Analysis *versus* synthesis in signal priors”, *Inverse Problems*, vol. 23, no. 3, pp. 947–968, Apr. 2007.
- [138] J.-F. Cai, S. Osher, and Z. Shen, “Split Bregman methods and frame based image restoration”, *Multiscale Modeling & Simulation*, vol. 8, no. 2, pp. 337–369, Jan. 2010.
- [139] R. Otazo, E. Candès, and D. Sodickson, “Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components”, *Magnetic Resonance in Medicine*, vol. 73, no. 3, pp. 1125–1136, Apr. 2014.
- [140] E. Candès, Y. Eldar, D. Needell, and P. Randall, “Compressed sensing with coherent and redundant dictionaries”, *Applied and Computational Harmonic Analysis*, vol. 31, no. 1, pp. 59–73, Jul. 2011.

BIBLIOGRAPHY

- [141] J. Lin, S. Li, and Y. Shen, “Compressed data separation with redundant dictionaries”, *IEEE Transactions on Information Theory*, vol. 59, no. 7, pp. 4309–4315, Jul. 2013.
- [142] T. Debarre, S. Shayan Aziznejad, and M. Unser, “Continuous-domain formulation of inverse problems for composite sparse-plus-smooth signals”, *IEEE Open Journal of Signal Processing*, vol. 2, pp. 545–558, Sep. 2021.
- [143] C. De Mol and M. Defrise, “Inverse imaging with mixed penalties”, in *Proceedings URSI EMTS*, Pisa, Italy, 2004, pp. 798–800.
- [144] A. Gholami and S. Hosseini, “A balanced combination of Tikhonov and total variation regularizations for reconstruction of piecewise-smooth signals”, *Signal Processing*, vol. 93, no. 7, pp. 1945–1960, Jul. 2013.
- [145] V. Naumova and S. Peter, “Minimization of multi-penalty functionals by alternating iterative thresholding and optimal parameter choices”, *Inverse Problems*, vol. 30, no. 12, p. 125 003, Oct. 2014.
- [146] M. Grasmair, T. Klock, and V. Naumova, “Adaptive multi-penalty regularization based on a generalized Lasso path”, *Applied and Computational Harmonic Analysis*, vol. 49, no. 1, pp. 30–55, Nov. 2018.
- [147] V. Chernozhukov, C. Hansen, and Y. Liao, “A lava attack on the recovery of sums of dense and sparse signals”, *The Annals of Statistics*, vol. 45, no. 1, pp. 39–76, 2017.
- [148] V. Debarnot, P. Escande, T. Mangeat, and P. Weiss, “Learning low-dimensional models of microscopes”, *IEEE Transactions on Computational Imaging*, vol. 7, pp. 178–190, 2021.
- [149] I. Daubechies, M. Defrise, and C. De Mol, “Sparsity-enforcing regularisation and ISTA revisited”, *Inverse Problems*, vol. 32, no. 10, p. 104 001, Aug. 2016.
- [150] M. Belge, M. Kilmer, and E. Miller, “Efficient determination of multiple regularization parameters in a generalized l-curve framework”, *Inverse Problems*, vol. 18, no. 4, pp. 1161–1183, Jul. 2002.
- [151] S. Roth and M. Black, “Fields of experts: A framework for learning image priors”, in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, vol. 2, San Diego, California: IEEE, Jun. 2005, pp. 860–867.
- [152] Z. Chen, Y. Lu, Y. Xu, and H. Yang, “Multi-parameter Tikhonov regularization for linear ill-posed operator equations”, *Journal of Computational Mathematics*, vol. 26, no. 1, pp. 37–55, 2008.
- [153] S. Lu and S. Pereverzev, “Multi-parameter regularization and its numerical realization”, *Numerische Mathematik*, vol. 118, no. 1, pp. 1–31, Jul. 2010.

-
- [154] Z. Wang, “Multi-parameter tikhonov regularization and model function approach to the damped Morozov principle for choosing regularization parameters”, *Journal of Computational and Applied Mathematics*, vol. 236, no. 7, pp. 1815–1832, Jan. 2012.
- [155] R. Abhishake and S. Sivananthan, “Multi-penalty regularization in learning theory”, *Journal of Complexity*, vol. 36, pp. 141–165, Oct. 2016.
- [156] H. Zou and T. Hastie, “Regularization and variable selection via the elastic net”, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 67, no. 2, pp. 301–320, Apr. 2005.
- [157] Y. Meyer, *Oscillating Patterns in Image Processing and Nonlinear Evolution Equations: The Fifteenth Dean Jacqueline B. Lewis Memorial Lectures*. Providence RI, USA: American Mathematical Society, Sep. 2001, vol. 22.
- [158] L. Vese and S. Osher, “Modeling textures with total variation minimization and oscillating patterns in image processing”, *Journal of Scientific Computing*, vol. 19, no. 1/3, pp. 553–572, Dec. 2003.
- [159] —, “Image denoising and decomposition with total variation minimization and oscillatory functions”, *Journal of Mathematical Imaging and Vision*, vol. 20, no. 1/2, pp. 7–18, Jan. 2004.
- [160] D. Mumford and J. Shah, “Optimal approximations by piecewise smooth functions and associated variational problems”, *Communications on Pure and Applied Mathematics*, vol. 42, no. 5, pp. 577–685, Jul. 1989.
- [161] I. Daubechies, *Ten Lectures on Wavelets*. Philadelphia PA, USA: Society for Industrial and Applied Mathematics, Jan. 1992.
- [162] P. Bohra and M. Unser, “Computation of “best” interpolants in the L_p sense”, in *Proceedings of the Forty-Fifth IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP’20)*, Barcelona, Kingdom of Spain, May 2020, pp. 5505–5509.
- [163] L. Dadi, S. Aziznejad, and M. Unser, “Generating sparse stochastic processes using matched splines”, *IEEE Transactions on Signal Processing*, vol. 68, pp. 4397–4406, Jul. 2020.
- [164] M. Unser and T. Blu, “Generalized smoothing splines and the optimal discretization of the Wiener filter”, *IEEE Transactions on Signal Processing*, vol. 53, no. 6, pp. 2146–2159, Jun. 2005.
- [165] A. Ali and R. Tibshirani, “The generalized lasso problem and uniqueness”, *Electronic Journal of Statistics*, vol. 13, no. 2, Jan. 2019.
- [166] C. Pouchol and O. Verdier, “The ML–EM algorithm in continuum: Sparse measure solutions”, *Inverse Problems*, vol. 36, no. 3, p. 035 013, Feb. 2020.
- [167] —, “Linear inverse problems with nonnegativity constraints through divergences: Sparsity of optimisers”, *arXiv preprint arXiv:2006.15845*, Jun. 2020.

BIBLIOGRAPHY

- [168] A. Eftekhari, T. Bendory, and G. Tang, “Stable super-resolution of images: Theoretical study”, *Information and Inference: A Journal of the IMA*, vol. 10, no. 1, pp. 161–193, Nov. 2020.
- [169] E. Candes and Y. Plan, “A probabilistic and RIPless theory of compressed sensing”, *IEEE Transactions on Information Theory*, vol. 57, no. 11, pp. 7235–7254, Nov. 2011.
- [170] T. Debarre, Q. Denoyelle, M. Unser, and J. Fageot, “Sparsest piecewise-linear regression of one-dimensional data”, *Journal of Computational and Applied Mathematics*, vol. 406, p. 114044, May 2022.
- [171] G. Cybenko, “Approximation by superpositions of a sigmoidal function”, *Mathematics of Control, Signals, and Systems*, vol. 2, no. 4, pp. 303–314, Dec. 1989.
- [172] K. Hornik, “Approximation capabilities of multilayer feedforward networks”, *Neural Networks*, vol. 4, no. 2, pp. 251–257, 1991.
- [173] M. Leshno, V. Lin, A. Pinkus, and S. Schocken, “Multilayer feedforward networks with a nonpolynomial activation function can approximate any function”, *Neural Networks*, vol. 6, no. 6, pp. 861–867, Jan. 1993.
- [174] T. Mitchell, *Machine Learning*. McGraw-Hill Education - Europe, Mar. 1, 1997, 432 pp.
- [175] C. Ekanadham, D. Tranchina, and E. Simoncelli, “Recovery of sparse translation-invariant signals with continuous basis pursuit”, *IEEE Transactions on Signal Processing*, vol. 59, no. 10, pp. 4735–4744, Oct. 2011.
- [176] R. Koenker, P. Ng, and S. Portnoy, “Quantile smoothing splines”, *Biometrika*, vol. 81, no. 4, pp. 673–680, 1994.
- [177] E. Mammen and S. van de Geer, “Locally adaptive regression splines”, *The Annals of Statistics*, vol. 25, no. 1, pp. 387–413, Feb. 1997.
- [178] A. Pinkus, “On smoothest interpolants”, *SIAM Journal on Mathematical Analysis*, vol. 19, no. 6, pp. 1431–1441, Nov. 1988.
- [179] M. Unser, “Representer theorems for the design of deep neural networks and the resolution of continuous-domain inverse problems”, in *First International Workshop on Machine Learning for Medical Image Reconstruction (MLMIR’18)*, Keynote address, Granada, Kingdom of Spain, Sep. 2018.
- [180] V. Duval, “An epigraphical approach to the representer theorem”, *arXiv preprint arXiv:1912.13224*, Dec. 2019.
- [181] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, <http://www.deeplearningbook.org>.
- [182] G. Montufar, R. Pascanu, K. Cho, and Y. Bengio, “On the number of linear regions of deep neural networks”, in *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, Eds., Curran Associates, Inc., 2014, pp. 2924–2932.

-
- [183] R. Balestriero and R. Baraniuk, “Mad max: Affine spline insights into deep learning”, *Proceedings of the IEEE*, vol. 109, no. 5, pp. 704–727, May 2021.
- [184] F. Agostinelli, M. Hoffman, P. Sadowski, and P. Baldi, “Learning activation functions to improve deep neural networks”, in *Third International Conference on Learning Representations (ICLR) Workshop*, San Diego, USA: ICLR, 2015.
- [185] S. Aziznejad, H. Gupta, J. Campos, and M. Unser, “Deep neural networks with trainable activations and controlled Lipschitz constant”, *IEEE Transactions on Signal Processing*, vol. 68, pp. 4688–4699, Aug. 2020.
- [186] I. Daubechies, R. DeVore, S. Foucart, B. Hanin, and G. Petrova, “Nonlinear approximation and (deep) ReLU networks”, *Constructive Approximation*, Apr. 2021.
- [187] T. Poggio, L. Rosasco, A. Shashua, N. Cohen, and F. Anselmi, “Notes on hierarchical splines, DCLNs and i-theory”, Center for Brains, Minds and Machines (CBMM), CBMM Memo, 2015.
- [188] H. Bölcskei, P. Grohs, G. Kutyniok, and P. Petersen, “Optimal approximation with sparsely connected deep neural networks”, *SIAM Journal on Mathematics of Data Science*, vol. 1, no. 1, pp. 8–45, Jan. 2019.
- [189] D. Yarotsky, “Error bounds for approximations with deep ReLU networks”, *Neural Networks*, vol. 94, pp. 103–114, Oct. 2017.
- [190] P. Petersen and F. Voigtlaender, “Optimal approximation of piecewise smooth functions using deep ReLU neural networks”, *Neural Networks*, vol. 108, pp. 296–330, Dec. 2018.
- [191] R. Gribonval, G. Kutyniok, M. Nielsen, and F. Voigtlaender, “Approximation spaces of deep neural networks”, *arXiv preprint arXiv:1905.01208*, May 3, 2019.
- [192] J. de Dios and J. Bruna, “On sparsity in overparametrised shallow ReLU networks”, *arXiv preprint arXiv:2006.10225*, Jun. 18, 2020.
- [193] P. Savarese, I. Evron, D. Soudry, and N. Srebro, “How do infinite width bounded norm networks look in function space?”, in *Proceedings of the Thirty-Second Conference on Learning Theory*, A. Beygelzimer and D. Hsu, Eds., ser. Proceedings of Machine Learning Research, vol. 99, Phoenix, USA: PMLR, Jun. 2019, pp. 2667–2690.
- [194] R. Parhi and R. Nowak, “The role of neural network activation functions”, *IEEE Signal Processing Letters*, vol. 27, pp. 1779–1783, 2020.
- [195] —, “Banach space representer theorems for neural networks and ridge splines”, *Journal of Machine Learning Research*, vol. 22, no. 43, pp. 1–40, 2021.
- [196] G. Ongie, R. Willett, D. Soudry, and N. Srebro, “A function space view of bounded norm infinite width ReLU nets: The multivariate case”, in *International Conference on Learning Representations*, 2020.

BIBLIOGRAPHY

- [197] R. Rockafellar, *Convex Analysis*. Princeton University Press, Dec. 1970.
- [198] A. Ben-Israel and T. Greville, *Generalized Inverses*. Springer-Verlag, 2003.
- [199] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski, “Optimization with sparsity-inducing penalties”, *Foundations and Trends® in Machine Learning*, vol. 4, no. 1, pp. 1–106, 2011.
- [200] M. Osborne, B. Presnell, and B. Turlach, “On the LASSO and its dual”, *Journal of Computational and Graphical Statistics*, vol. 9, no. 2, p. 319, Jun. 2000.
- [201] S. Aziznejad, T. Debarre, and M. Unser, “Sparsest univariate learning models under Lipschitz constraint”, *IEEE Open Journal of Signal Processing*, vol. 3, pp. 140–154, Mar. 2022.
- [202] L. Györfi, M. Kohler, A. Krzyżak, and H. Walk, *A Distribution-Free Theory of Nonparametric Regression*. Springer New York, 2002.
- [203] T. Hastie, R. Tibshirani, and J. Friedman, “Overview of supervised learning”, in *The Elements of Statistical Learning*, Springer New York, Dec. 2008, pp. 9–41.
- [204] V. Katkovnik, K. Egiazarian, and J. Astola, “A spatially adaptive nonparametric regression image deblurring”, *IEEE transactions on Image Processing*, vol. 14, no. 10, pp. 1469–1478, 2005.
- [205] V. Katkovnik, “A multiresolution nonparametric regression for spatially adaptive image de-noising”, *IEEE Signal Processing Letters*, vol. 11, no. 10, pp. 798–801, 2004.
- [206] G. Mateos and G. Giannakis, “Robust nonparametric regression via sparsity control with application to load curve data cleansing”, *IEEE Transactions on Signal Processing*, vol. 60, no. 4, pp. 1571–1584, 2011.
- [207] T. Poggio and F. Girosi, “Networks for approximation and learning”, *Proceedings of the IEEE*, vol. 78, no. 9, pp. 1481–1497, 1990.
- [208] —, “Regularization algorithms for learning that are equivalent to multilayer networks”, *Science*, vol. 247, no. 4945, pp. 978–982, 1990.
- [209] G. Kimeldorf and G. Wahba, “Some results on Tchebycheffian spline functions”, *Journal of Mathematical Analysis and Applications*, vol. 33, no. 1, pp. 82–95, 1971.
- [210] B. Schölkopf and A. J. Smola, *Learning with Kernels*. The MIT Press, 2018.
- [211] T. Evgeniou, M. Pontil, and T. Poggio, “Regularization networks and support vector machines”, *Advances in Computational Mathematics*, vol. 13, no. 1, p. 1, 2000.
- [212] A. Christmann and I. Steinwart, *Support Vector Machines*. Springer New York, Oct. 2014, 620 pp.
- [213] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning”, *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.

-
- [214] A. Krizhevsky, I. Sutskever, and G. Hinton, “ImageNet classification with deep convolutional neural networks”, in *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, Eds., Curran Associates, Inc., 2012, pp. 1097–1105.
- [215] K. Jin, M. McCann, E. Froustey, and M. Unser, “Deep convolutional neural network for inverse problems in imaging”, *IEEE Transactions on Image Processing*, vol. 26, no. 9, pp. 4509–4522, Sep. 2017.
- [216] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation”, in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds., Cham: Springer International Publishing, 2015, pp. 234–241.
- [217] S. Moosavi-Dezfooli, A. Fawzi, and P. Frossard, “Deepfool: A simple and accurate method to fool deep neural networks”, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 2574–2582.
- [218] A. Fawzi, S. Moosavi-Dezfooli, and P. Frossard, “The robustness of deep networks: A geometrical perspective”, *IEEE Signal Processing Magazine*, vol. 34, no. 6, pp. 50–62, 2017.
- [219] V. Antun, F. Renna, C. Poon, B. Adcock, and A. Hansen, “On instabilities of deep learning in image reconstruction and the potential costs of AI”, *Proceedings of the National Academy of Sciences*, vol. 117, no. 48, pp. 30 088–30 095, 2020.
- [220] J. MacDonald, S. Wäldchen, S. Hauch, and G. Kutyniok, “A rate-distortion framework for explaining neural network decisions”, *arXiv preprint arXiv:1905.11092*, 2019.
- [221] C. Heiß, R. Levie, C. Resnick, G. Kutyniok, and J. Bruna, “In-distribution interpretability for challenging modalities”, *arXiv preprint arXiv:2007.00758*, 2020.
- [222] M. Arjovsky, S. Chintala, and L. Bottou, “Wasserstein generative adversarial networks”, in *Proceedings of the 34th International Conference on Machine Learning*, ser. ICML’17, vol. 70, Sydney, NSW, Australia: JMLR.org, 2017, pp. 214–223.
- [223] A. Bora, A. Jalal, E. Price, and A. Dimakis, “Compressed sensing using generative models”, in *Proceedings of the 34th International Conference on Machine Learning*, vol. 70, 2017, pp. 537–546.
- [224] B. Neyshabur, S. Bhojanapalli, D. McAllester, and N. Srebro, “Exploring generalization in deep learning”, in *Advances in Neural Information Processing Systems*, 2017, pp. 5947–5956.
- [225] W. Zhang K. and Zuo, Y. Chen, D. Meng, and L. Zhang, “Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising”, *IEEE Transactions on Image Processing*, vol. 26, no. 7, pp. 3142–3155, 2017.

BIBLIOGRAPHY

- [226] H. Gupta, J. Fageot, and M. Unser, “Optimality of splines for the resolution of linear inverse problems with Tikhonov or total-variation regularization”, in *SIAM Conference on Imaging Science (SIS’18)*, session MS47-1, Bologna, Italian Republic, Jun. 2018.
- [227] Y. Sun, J. Liu, and U. Kamilov, “Block coordinate regularization by denoising”, *IEEE Transactions on Computational Imaging*, vol. 6, pp. 908–921, 2020.
- [228] J. Liu, Y. Sun, C. Eldeniz, W. Gan, H. An, and U. Kamilov, “RARE: Image reconstruction using deep priors learned without ground truth”, *IEEE Journal of Selected Topics in Signal Processing*, vol. 14, no. 6, pp. 1088–1099, 2020.
- [229] Z. Wu, Y. Sun, A. Matlock, J. Liu, L. Tian, and U. Kamilov, “SIMBA: Scalable inversion in optical tomography using deep denoising priors”, *IEEE Journal of Selected Topics in Signal Processing*, vol. 14, no. 6, pp. 1163–1175, 2020.
- [230] P. Bohra, D. Perdios, A. Goujon, S. Emery, and M. Unser, “Learning Lipschitz-controlled activation functions in neural networks for plug-and-play image reconstruction methods”, in *NeurIPS 2021 Workshop on Deep Learning and Inverse Problems*, 2021.
- [231] D. McClure, “Perfect spline solutions of L_∞ extremal problems by control methods”, *Journal of Approximation Theory*, vol. 15, no. 3, pp. 226–242, 1975.
- [232] S. Karlin, “Interpolation properties of generalized perfect splines and the solutions of certain extremal problems. i”, *Transactions of the American Mathematical Society*, vol. 206, pp. 25–25, 1975.
- [233] C. de Boor, “How small can one make the derivatives of an interpolating function?”, *Journal of Approximation Theory*, vol. 13, no. 2, pp. 105–116, 1975.
- [234] C. Micchelli, T. Rivlin, and S. Winograd, “The optimal recovery of smooth functions”, *Numerische Mathematik*, vol. 26, no. 2, pp. 191–200, Jun. 1976.
- [235] U. von Luxburg and O. Bousquet, “Distance-based classification with Lipschitz functions”, *Journal of Machine Learning Research*, vol. 5, no. Jun, pp. 669–695, 2004.
- [236] R. Arora, A. Basu, P. Mianjy, and A. Mukherjee, “Understanding deep neural networks with rectified linear units”, in *International Conference on Learning Representations*, 2018.
- [237] R. Parhi and R. Nowak, “What kinds of functions do deep neural networks learn? insights from variational spline theory”, *arXiv preprint arXiv:2105.03361*, May 2021.
- [238] N. Weaver, *Lipschitz Algebras*. World Scientific, Jul. 1999.
- [239] P. Bohra, J. Campos, H. Gupta, S. Aziznejad, and M. Unser, “Learning activation functions in deep (spline) neural networks”, *IEEE Open Journal of Signal Processing*, vol. 1, pp. 295–309, Nov. 2020.

-
- [240] T. Debarre, Q. Denoyelle, and J. Fageot, “On the uniqueness of solutions for the basis pursuit in the continuum”, *arXiv preprint arXiv:2009.11855*, Feb. 2022.
- [241] ———, “TV-based spline reconstruction with Fourier measurements: Uniqueness and convergence of grid-based methods”, *arXiv preprint arXiv:2202.05059*, Feb. 2022.
- [242] M. Powell, *Approximation Theory and Methods*. Cambridge University Press, Mar. 1981.
- [243] J. Kelley, *General Topology*. Mineola, New York: Dover Publications, 2017.
- [244] G. Dal Maso, *An Introduction to Γ -Convergence*. Boston, MA: Birkhäuser Boston, 1993.
- [245] P. Heins, “A novel regularization technique and an asymptotic analysis of spatial sparsity priors”, PhD thesis, Westfälische Wilhelms Universität Münster (WWU Münster), 2015.
- [246] I. Lloréns Jover, T. Debarre, S. Aziznejad, and M. Unser, “Coupled splines for sparse curve fitting”, *arXiv preprint arXiv:2202.01641*, Feb. 2022.
- [247] F. Chang and C.-J. Chen, “A component-labeling algorithm using contour tracing technique”, in *Seventh International Conference on Document Analysis and Recognition, 2003. Proceedings.*, Citeseer, vol. 3, 2003, pp. 741–741.
- [248] F. Chang, C.-J. Chen, and C.-J. Lu, “A linear-time component-labeling algorithm using contour tracing technique”, *Computer Vision and Image Understanding*, vol. 93, no. 2, pp. 206–220, 2004.
- [249] S. Suzuki and A. Keiichi, “Topological structural analysis of digitized binary images by border following”, *Computer Vision, Graphics, and Image Processing*, vol. 30, no. 1, pp. 32–46, 1985.
- [250] O. Hori and S. Tanigawa, “Raster-to-vector conversion by line fitting based on contours and skeletons”, in *Proceedings of 2nd International Conference on Document Analysis and Recognition (ICDAR '93)*, IEEE Comput. Soc. Press, Oct. 1993.
- [251] M. Awadallah, L. Abbott, and S. Ghannam, “Segmentation of sparse noisy point clouds using active contour models”, in *2014 IEEE International Conference on Image Processing (ICIP)*, IEEE, 2014, pp. 6061–6065.
- [252] M. Kass, A. Witkin, and D. Terzopoulos, “Snakes: Active contour models”, *International Journal of Computer Vision*, vol. 1, no. 4, pp. 321–331, Jan. 1988.
- [253] R. Delgado-Gonzalo, V. Uhlmann, D. Schmitter, and M. Unser, “Snakes on a plane: A perfect snap for bioimage analysis”, *IEEE Signal Processing Magazine*, vol. 32, no. 1, pp. 41–48, Jan. 2015.
- [254] A. Jain, Y. Zhong, and M.-P. Dubuisson-Jolly, “Deformable template models: A review”, *Signal Processing*, vol. 71, no. 2, pp. 109–129, Dec. 1998.

BIBLIOGRAPHY

- [255] M. Jacob, T. Blu, and M. Unser, “Efficient energies and algorithms for parametric snakes”, *IEEE Transactions on Image Processing*, vol. 13, no. 9, pp. 1231–1244, Sep. 2004.
- [256] V. Uhlmann, R. Delgado-Gonzalo, C. Conti, L. Romani, and M. Unser, “Exponential Hermite splines for the analysis of biomedical images”, in *Proceedings of the Thirty-Ninth IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP’14)*, Firenze, Italian Republic, May 2014, pp. 1650–1653.
- [257] R. Delgado-Gonzalo and M. Unser, “Spline-based framework for interactive segmentation in biomedical imaging”, *IRBM—Ingénierie et Recherche Biomédicale / BioMedical Engineering and Research*, vol. 34, no. 3, pp. 235–243, Jun. 2013.
- [258] P. Thévenaz and M. Unser, “Snakuscles”, *IEEE Transactions on Image Processing*, vol. 17, no. 4, pp. 585–593, Apr. 2008.
- [259] X. Chen, B. Williams, S. Vallabhaneni, G. Czanner, R. Williams, and Y. Zheng, “Learning active contour models for medical image segmentation”, in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 11 632–11 640.
- [260] M. Zhang, B. Dong, and Q. Li, “Deep active contour network for medical image segmentation”, in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, 2020, pp. 321–331.
- [261] V. Uhlmann, J. Fageot, and M. Unser, “Hermite snakes with control of tangents”, *IEEE Transactions on Image Processing*, vol. 25, no. 6, pp. 2803–2816, Jun. 2016.
- [262] T. Pavlidis, “Curve fitting with conic splines”, *ACM Transactions on Graphics (TOG)*, vol. 2, no. 1, pp. 1–31, 1983.
- [263] H. Pottmann, S. Leopoldseder, and M. Hofer, “Approximation with active b-spline curves and surfaces”, in *10th Pacific Conference on Computer Graphics and Applications, 2002. Proceedings.*, IEEE, 2002, pp. 8–25.
- [264] W. Wang, H. Pottmann, and Y. Liu, “Fitting B-spline curves to point clouds by curvature-based squared distance minimization”, *ACM Transactions on Graphics (ToG)*, vol. 25, no. 2, pp. 214–238, 2006.
- [265] W. Zheng, P. Bo, Y. Liu, and W. Wang, “Fast B-spline curve fitting by L-BFGS”, *Computer Aided Geometric Design*, vol. 29, no. 7, pp. 448–462, 2012.
- [266] M. Grossman, “Parametric curve fitting”, *The Computer Journal*, vol. 14, no. 2, pp. 169–172, Feb. 1971.
- [267] M. Plass and M. Stone, “Curve-fitting with piecewise parametric cubics”, in *Proceedings of the 10th annual conference on Computer graphics and interactive techniques - SIGGRAPH ’83*, ACM Press, Jul. 1983.
- [268] T. Lyche and K. Mørken, “Knot removal for parametric b-spline curves and surfaces”, *Computer Aided Geometric Design*, vol. 4, no. 3, pp. 217–230, 1987.

-
- [269] ———, “A discrete approach to knot removal and degree reduction algorithms for splines”, in *Algorithms for Approximation*. USA: Clarendon Press, 1987, pp. 67–82.
- [270] ———, “A Data-Reduction Strategy for Splines with Applications to the Approximation of Functions and Data”, *IMA Journal of Numerical Analysis*, vol. 8, no. 2, pp. 185–208, Apr. 1988.
- [271] M. Eck and J. Hadenfeld, “Knot removal for B-spline curves”, *Computer Aided Geometric Design*, vol. 12, no. 3, pp. 259–282, 1995.
- [272] O. Bingol and A. Krishnamurthy, “NURBS-Python: An open-source object-oriented NURBS modeling framework in Python”, *SoftwareX*, vol. 9, pp. 85–94, 2019.
- [273] L. Piegl and W. Tiller, “Fundamental geometric algorithms”, in *The NURBS Book*, Springer, 1995, pp. 141–228.
- [274] A. Goshtasby, “Grouping and parameterizing irregularly spaced points for curve fitting”, *ACM Transactions on Graphics (TOG)*, vol. 19, no. 3, pp. 185–203, 2000.
- [275] J. Seo, S. Chae, J. Shim, D. Kim, C. Cheong, and T.-D. Han, “Fast contour-tracing algorithm based on a pixel-following method for image sensors”, *Sensors*, vol. 16, no. 3, 2016.
- [276] X. Wang, Y. Xu, K. Xu, A. Tagliasacchi, B. Zhou, A. Mahdavi-Amiri, and H. Zhang, “Pie-net: Parametric inference of point cloud edges”, *Advances in Neural Information Processing Systems*, vol. 33, pp. 20 167–20 178, 2020.
- [277] F. Scholz and B. Jüttler, “Parameterization for polynomial curve approximation via residual deep neural networks”, *Computer Aided Geometric Design*, vol. 85, p. 101 977, 2021.
- [278] Y. Guo, H. Wang, Q. Hu, H. Liu, L. Liu, and M. Bennamoun, “Deep learning for 3d point clouds: A survey”, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 12, pp. 4338–4364, 2020.
- [279] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski, “Structured sparsity through convex optimization”, *Statistical Science*, vol. 27, no. 4, Nov. 2012.
- [280] Y. Hu and M. Jacob, “Higher degree total variation (HDTV) regularization for image recovery”, *IEEE Transactions on Image Processing*, vol. 21, no. 5, pp. 2559–2571, May 2012.
- [281] W. Dong, G. Shi, X. Li, Y. Ma, and F. Huang, “Compressive sensing via nonlocal low-rank regularization”, *IEEE Transactions on Image Processing*, vol. 23, no. 8, pp. 3618–3632, Aug. 2014.
- [282] D. Ulyanov, A. Vedaldi, and V. Lempitsky, “Deep image prior”, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2018.
- [283] L. Ambrosio, N. Fusco, and D. Pallara, *Functions of Bounded Variation and Free Discontinuity Problems*. Oxford University Press, May 2000, 456 pp.

BIBLIOGRAPHY

- [284] G. Folland, *Real Analysis: Modern Techniques and Their Applications*. John Wiley & Sons, 1999, vol. 40.
- [285] G. Nardi, G. Peyré, and F.-X. Vialard, “Geodesics on shape spaces with bounded variation and Sobolev metrics”, *arXiv preprint arXiv:1402.6504*, Feb. 2014.
- [286] T. Debarre, B. Watts, B. Rösner, and M. Unser, “Hessian splines for scanning transmission x-ray microscopy”, in *Proceedings of the Seventeenth IEEE International Symposium on Biomedical Imaging (ISBI’20)*, Iowa City IA, USA, Apr. 2020, pp. 199–202.
- [287] B. Watts and H. Ade, “Nexafs imaging of synthetic organic materials”, *Materials Today*, vol. 15, no. 4, pp. 148–157, 2012.
- [288] B. Rösner, F. Koch, F. Döring, J. Bosgra, V. Guzenko, E. Kirk, M. Meyer, J. Ornelas, R. Fink, S. Stanescu, *et al.*, “Exploiting atomic layer deposition for fabricating sub-10 nm x-ray lenses”, *Microelectronic Engineering*, vol. 191, pp. 91–96, 2018.
- [289] B. Rösner, F. Koch, F. Döring, V. Guzenko, M. Meyer, J. Ornelas, A. Späth, R. Fink, S. Stanescu, S. Swaraj, *et al.*, “7 nm spatial resolution in soft x-ray microscopy”, *Microscopy and Microanalysis*, vol. 24, no. S2, pp. 272–273, 2018.
- [290] M. Unser, “Sampling—50 Years after Shannon”, *Proceedings of the IEEE*, vol. 88, no. 4, pp. 569–587, Apr. 2000.
- [291] S. Lefkimmiatis and M. Unser, “Poisson image reconstruction with Hessian Schatten-norm regularization”, *IEEE Transactions on Image Processing*, vol. 22, no. 11, pp. 4314–4327, Nov. 2013.
- [292] L. Liu, X. Li, K. Xiang, J. Wang, and S. Tan, “Low-dose CBCT reconstruction using Hessian Schatten penalties”, *IEEE Transactions on Medical Imaging*, vol. 36, no. 12, pp. 2588–2599, 2017.
- [293] P. Thévenaz, T. Blu, and M. Unser, “Interpolation revisited”, *IEEE Transactions on Medical Imaging*, vol. 19, no. 7, pp. 739–758, Jul. 2000.
- [294] M. Hestenes and E. Stiefel, “Methods of conjugate gradients for solving linear systems”, *Journal of Research of the National Bureau of Standards*, vol. 49, no. 6, p. 409, Dec. 1952.
- [295] K. Bredies, J. Iglesias, and G. Mercier, “Boundedness and unboundedness in total variation regularization”, *arXiv preprint arXiv:2203.03264*, Mar. 2022.
- [296] J.-B. Courbot and B. Colicchio, “Boosting the sliding Frank-Wolfe solver for 3D deconvolution”, *arXiv preprint arXiv:2009.05473*, Sep. 2020.
- [297] G. Chardon and U. Boureau, “Gridless three-dimensional compressive beamforming with the sliding Frank-Wolfe algorithm”, *The Journal of the Acoustical Society of America*, vol. 150, no. 4, pp. 3139–3148, Oct. 2021.
- [298] M. Unser, “Ridges, neural networks, and the Radon transform”, *arXiv preprint arXiv:2203.02543*, Mar. 2022.

- [299] B. He and X. Yuan, “On the $O(1/n)$ convergence rate of the Douglas-Rachford alternating direction method”, *SIAM Journal on Numerical Analysis*, vol. 50, no. 2, pp. 700–709, Jan. 2012.
- [300] M. Unser and S. Aziznejad, “Convex optimization in sums of Banach spaces”, *Applied and Computational Harmonic Analysis*, vol. 56, no. 1, pp. 1–25, Jan. 2022.
- [301] A. Kurdila and M. Zabrankin, *Convex Functional Analysis*. Birkhäuser-Verlag, 2005.

Thomas Debarre
Email: thomas.debarre@epfl.ch
French/USA citizen



EDUCATION

- 2017 - 2022
(expected) Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland
PhD student at the Biomedical Imaging Group (BIG)
Project: Variational Methods For Continuous-Domain Inverse Problems:
the Quest for the Sparsest Solution
Advisors: Prof. Michael Unser and Dr. Julien Fageot
- 2016 - 2017 Ecole Normale Supérieure de Cachan, France
MVA Master's degree in applied mathematics (mathematics, computer
vision, and machine learning).
- 2013 - 2016 Mines ParisTech, France
MSc in engineering, applied mathematics major
- 2011 - 2013 Louis-le-Grand preparatory school, Paris, France
Preparatory classes for the "Grandes Ecoles" national competitive exams,
math and physics section (MPSI-MP*)

INTERNSHIPS AND RESEARCH PROJECTS

- 2017
(5 months) Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland
Master project at the Biomedical Imaging Group (BIG)
Project: B-spline implementation to find the solution of continuous-
domain total-variation minimization problems
Advisors: Dr. Julien Fageot, Harshit Gupta, and Prof. Michael Unser
- 2016
(5 months) French-German Research Institute of Saint-Louis (ISL), France
Aerodynamics and exterior ballistic intern
Project: Position and attitude determination of a vehicle in free flight
based on 3D high-speed videos
Advisors: Dr. Simona Dobre and Dr. Claude Berner
- 2015
(3 months) Morpho Detection (Safran group), Newark (CA), USA
Machine vision intern
Project: Investigations of ramp filter in filtered back-projection reconstruction
Advisors: Dr. Robin Palit and Dr. Todd Gable
- 2014
(5 months) University of New South Wales, Sydney, Australia.
Research exchange semester
Project: An introduction to open-loop quantum control
Advisor: Prof. Hendra Nurdin

RESEARCH INTERESTS

I am interested in theoretical and algorithmic aspects of inverse problems in signal processing with sparsity priors. More specifically, I study inverse problems that are formulated in the continuous domain, and develop spline-based discretization methods to solve them algorithmically. Keywords:

- continuous-domain inverse problems
- sparsity
- splines
- convex optimization
- proximal algorithms

PUBLICATIONS

Preprints:

1. I. Lloréns Jover, **T. Debarre**, S. Aziznejad, and M. Unser, “Coupled splines for sparse curve fitting,” *ArXiv preprint*, February 2022, link: <https://arxiv.org/abs/2202.01641>
2. **T. Debarre**, Q. Denoyelle, and J. Fageot, “TV-based spline reconstruction with Fourier measurements: Uniqueness and convergence of grid-based methods,” *ArXiv preprint*, February 2022, link: <https://arxiv.org/abs/2202.05059>
3. —, “On the uniqueness of solutions for the basis pursuit in the continuum,” *ArXiv preprint*, February 2022, link: <https://arxiv.org/abs/2009.11855>

Journal papers:

1. S. Aziznejad*, **T. Debarre***, and M. Unser, “Sparsest univariate learning models under Lipschitz constraint,” *IEEE Open Journal of Signal Processing*, vol. 3, pp. 140–154, March 7, 2022
2. **T. Debarre**, Q. Denoyelle, M. Unser, and J. Fageot, “Sparsest continuous piecewise-linear representation of one-dimensional data,” *Journal of Computational and Applied Mathematics*, vol. 406, pp. 1–30, May 1, 2022, paper no. 114044
3. **T. Debarre**, S. Aziznejad, and M. Unser, “Continuous-domain formulation of inverse problems for composite sparse-plus-smooth signals,” *IEEE Open Journal of Signal Processing*, vol. 2, no. 22, pp. 545–558, September 29, 2021
4. —, “Hybrid-spline dictionaries for continuous-domain inverse problems,” *IEEE Transactions on Signal Processing*, vol. 67, no. 22, pp. 5824–5836, November 15, 2019
5. E. Soubies, F. Soulez, M. McCann, T.-a. Pham, L. Donati, **T. Debarre**, D. Sage, and M. Unser, “Pocket guide to solve inverse problems with GlobalBioIm,” *Inverse Problems*, vol. 35, no. 10, pp. 1–20, October 2019
6. **T. Debarre**, J. Fageot, H. Gupta, and M. Unser, “B-spline-based exact discretization of continuous-domain inverse problems with generalized TV regularization,” *IEEE Transactions on Information Theory*, vol. 65, no. 7, pp. 4457–4470, July 2019

*: equal contribution

Conferences:

1. **T. Debarre**, B. Watts, B. Rösner, and M. Unser, “Hessian splines for scanning transmission x-ray microscopy,” in *Proceedings of the Seventeenth IEEE International Symposium on Biomedical Imaging (ISBI’20)*, Iowa City IA, USA, April 5-7, 2020, pp. 199–202

LANGUAGES

French : Native
English : Native
German : Basics

PERSONAL

I have been playing go (a strategy board game originating from China) competitively since my childhood. Notable results:

- 6 times French national champion (2010-2013, 2018, and 2021)
- 4 times European team champion playing for France (2015, 2019-2021)
- 3 times semi-finalist at the European championships (2012, 2013, 2015)
- 4th place at the World amateur go championships 2011 (5 participations representing France)

I enjoy playing all kinds of strategy games.

I am an avid badminton player, I have played competitively for many years. I am involved in my badminton club as a board member and team captain.

I love doing sports in general: racket sports, biking, and hiking.