

# Geometric Variational Models for Inverse Problems in Imaging

THÈSE N° 5785 (2013)

PRÉSENTÉE LE 11 JUILLET 2013

À LA FACULTÉ DES SCIENCES ET TECHNIQUES DE L'INGÉNIEUR  
LABORATOIRE DE TRAITEMENT DES SIGNAUX 5  
PROGRAMME DOCTORAL EN GÉNIE ÉLECTRIQUE

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Virginia ESTELLERS CASAS

acceptée sur proposition du jury:

Prof. P. Frossard, président du jury  
Prof. J.-Ph. Thiran, directeur de thèse  
Dr T. Pock, rapporteur  
Prof. P. Vandergheynst, rapporteur  
Prof. L. Vese, rapporteur



ÉCOLE POLYTECHNIQUE  
FÉDÉRALE DE LAUSANNE

Suisse  
2013



# Acknowledgements

A PhD resembles an iceberg, and this manuscript is only its tip. The invisible part, the support, has been developed through four years of work, collaborations, and interactions with colleagues. I am glad to have the opportunity to acknowledge the contribution of several people, and thank many others.

First of all, I would like to thank my advisor, Prof. Jean-Philippe Thiran, for all the opportunities and freedom he gave me. Jean-Philippe always let me explore new ideas, jump from problem to problem without reserve, and even change topics at the middle of my PhD. His generosity with his time, my independence, and his friendliness have been a constant throughout the years.

At the same time, I want to thank Dr. Xavier Bresson for all his technical guidance and discussions. Xavier acted as non-official advisor for several projects, providing the experienced view but also being ready to discuss the smallest detail of any equation. Working with Xavier was always constructive, stimulating, and easy.

I am also extremely grateful to my co-authors for their share of work and technical discussions. They provided the questioning and critic eye that forced my ideas to mature, be more contrasted, and hopefully far-sighted. I would like to both thank and acknowledge them.

Finally, on the technical side, I would like to thank the members of my PhD jury for providing useful feedback and for challenging me with new questions.

A PhD might be an academic achievement, but it is only possible because of the personal support given by many, from the colleagues that make lab life a joyful experience to friends and family. Thanks are due to all of them in different forms.

I would like to thank my colleagues at EPFL for making the signal processing labs a fun and agreeable place to work and make friends. I enjoyed corridor chats, social coffees, and, of course, ski-days and fondues. Two colleagues, Ale and Gorthi, deserve special credit. I would like to thank my office-mate Gorthi for all our discussions on indian culture, religion, and ethics. Gorthi answered all my questions and heard my opinions with an open mind; knowing what a curious person I am, he had to explain many details before I would be satisfied. I would also like to thank my coffee-mate Alessandro for sharing his postdoc views on academia and for his friendship. Both of them helped me to “see a little further” as a person and researcher.

It is also a great pleasure to thank the colleagues of my internship at IBM Research–Zurich. I would like to thank my manager Ronald for his enthusiasm, Maria for her technical lead, and my office-mate Nathalie for her introduction to IBM ways and friendly atmosphere. During

## Acknowledgements

---

my internship, Rosetta welcomed me at her house, making me feel at home and sharing her life experiences and adventures with me. I am extremely grateful to her for her unreserved hospitality and generosity.

Finally, I want to thank the friends that make me feel at home wherever I go because they are always there in person, Skype, or mail. Any explicit list will fall short, but the constancy of a *petit comité* deserves special mention. My brother, who loves fiery arguments as much as I do, because his crude honesty is a healthy mirror to keep close-by. My good friend Maria Ines, who is always ready to help, listen, and see the funny side of things, because she is now a sister. Pip, who is driven by the same curiosity, views, and passions as myself, because life is a shared experience.

Special thanks to my parents for always encouraging me to try, and being there if I failed.

*Lausanne, June 2013*

Virginia Estellers



# Preface

The opening line of *Anna Karenina*, “All happy families resemble one another, but each unhappy family is unhappy in its own way”, is a useful analogy for the analysis of inverse problems in comparison with their respective forward problems in imaging<sup>1</sup>. Forward problems are a happy family of physical models of the data acquisition system or the object being imaged; perhaps they do not represent the same physical entity, but we can always write them in a well-posed form and analyse them with a small number of techniques. Inverse problems, on the other hand, are a varied collection of problems, each unhappy in its own way; they require specific analysis of their individual ill-posed natures and a different set of techniques to solve them.

Inverse problems are thus studied case by case, or chapter by chapter in this dissertation. This results here in four self-contained chapters that investigate different inverse problems in imaging and can be read in any order. The rest of the manuscript, the introduction and the conclusions, simply present a coherent frame for this dissertation, they are light in references by design to avoid distraction, and leave the actual work and literature review to the appropriate chapters.

*Lausanne, March 2013*

Virginia Estellers

---

<sup>1</sup>This analogy was inspired by the comparison between ordinary and partial differential equations of Chapter 8 in Iserles [2009]. The original quote, of course, is from Lev Tolstoy.



# Abstract

This dissertation develops geometric variational models for different inverse problems in imaging that are ill-posed, designing at the same time efficient numerical algorithms to compute their solutions.

Variational methods solve inverse problems by the following two steps: formulation of a variational model as a minimization problem, and design of a minimization algorithm to solve it. This dissertation is organized in the same manner. It first formulates minimization problems associated with geometric models for different inverse problems in imaging, and it then designs efficient minimization algorithms to compute their solutions. The minimization problem summarizes both the data available from the measurements and the prior knowledge about the solution in its objective functional; this naturally leads to the combination of a measurement or data term and a prior term. Geometry can play a role in any of these terms, depending on the properties of the data acquisition system or the object being imaged. In this context, each chapter of this dissertation formulates a variational model that includes geometry in a different manner in the objective functional, depending on the inverse problem at hand.

In the context of compressed sensing, the first chapter exploits the geometric properties of images to include an alignment term in the sparsity prior of compressed sensing; this additional prior term aligns the normal vectors of the level curves of the image with the reconstructed signal, and it improves the quality of reconstruction. A two-step recovery method is designed for that purpose: first, it estimates the normal vectors to the level curves of the image; second, it reconstructs an image matching the compressed sensing measurements, the geometric alignment of normals, and the sparsity constraint of compressed sensing. The proposed method is extended to non-local operators in graphs for the recovery of textures.

The harmonic active contours of Chapter 2 make use of differential geometry to interpret the segmentation of an image as a minimal surface manifold. In this case, geometry is exploited in both the measurement term, by coupling the different image channels in a robust edge detector, and in the prior term, by imposing smoothness in the segmentation. The proposed technique generalizes existing active contours to higher dimensional spaces and non-flat images; in the plane, it improves the segmentation of images with inhomogeneities and weak edges.

## Abstract

---

Shape-from-shading is investigated in Chapter 3 for the reconstruction of a silicon wafer from images of printed circuits taken with a scanning electron microscope. In this case, geometry plays a role in the image acquisition system, that is, in the measurement term of the objective functional. The prior term involves a smoothness constraint on the surface and a shape prior on the expected pattern in the circuit. The proposed reconstruction method also estimates a deformation field between the ideal pattern design and the reconstructed surface, substituting the model of shape variability necessary in shape priors with an elastic deformation field that quantifies deviations in the manufacturing process.

Finally, the techniques used for the design of efficient numerical algorithms are explained with an example problem based on the level set method. To this purpose, Chapter 4 develops an efficient algorithm for the level set method when the level set function is constrained to remain a signed distance function. The distance function is preserved by the introduction of an explicit constraint in the minimization problem, the minimization algorithm is efficient by the adequate use of variable-splitting and augmented Lagrangian techniques. These techniques introduce additional variables, constraints, and Lagrange multipliers in the original minimization problem, and they decompose it into sub-optimization problems that are simple and can be efficiently solved. As a result, the proposed algorithm is five to six times faster than the original algorithm for the level set method.

**Key words:** Variational methods, image segmentation, image reconstruction, shape-from-shading, minimization.

# Résumé

Cette thèse développe des modèles variationnels géométriques pour différents problèmes inverses en imagerie qui sont mal posés, en développant en même temps des algorithmes numériques efficaces pour calculer leurs solutions.

Les méthodes variationnelles résolvent des problèmes inverses en deux étapes : formulation d'un modèle variationnel comme un problème de minimisation, et développement d'un algorithme de minimisation pour le résoudre. Cette thèse est organisée de la même manière. D'abord, des problèmes de minimisation associés aux modèles géométriques sont formulés pour différents problèmes inverses en imagerie, puis des algorithmes efficaces de minimisation sont développés pour calculer leurs solutions. Le problème de minimisation tient compte à la fois des données mesurées et des connaissances *a priori* sur la solution dans sa fonction objectif, ce qui conduit naturellement à la combinaison d'un terme de données et d'un terme *a priori*. La géométrie peut jouer un rôle dans l'un ou l'autre de ces termes, en fonction des propriétés du système d'acquisition des images ou de l'objet dont les données sont mesurées. Dans ce contexte, chaque chapitre de cette thèse élabore un modèle variationnel qui introduit la géométrie d'une manière différente dans la fonction objective, en fonction du problème inverse étudié.

Dans le cadre de *compressed sensing*, le premier chapitre exploite les propriétés géométriques des images et introduit un terme d'alignement dans le terme *a priori* de parcimonie de *compressed sensing*, ce terme supplémentaire aligne les vecteurs normaux des courbes de niveau de l'image avec le signal reconstruit et améliore la qualité de la reconstruction. A cet effet, une méthode en deux étapes a été conçue qui d'abord estime les vecteurs normaux aux courbes de niveau de l'image, et en deuxième lieu reconstruit une image qui s'ajuste à l'alignement géométrique des normales, et aux données et au terme *a priori* de parcimonie de *compressed sensing*. Cette méthode en deux étapes est directement étendue à des opérateurs non-locaux pour la reconstruction d'images avec textures.

Le contours actifs harmoniques proposés dans le deuxième chapitre interprètent la segmentation d'une image en tant que variété différentielle de surface minimale. Dans ce cas, la géométrie est exploitée à la fois dans le terme de données, par un couplage des différentes composantes d'une image vectorielle qui résulte en un détecteur de contours robuste, et dans le terme *a priori*, en imposant de la régularité à la segmentation. La technique proposée

utilise la géométrie différentielle pour généraliser plusieurs contours actifs à des espaces de dimension supérieure et à des images non planes, et elle améliore la segmentation d'images avec des inhomogénéités et des contours faibles.

*Shape-from-shading* est étudiée pour la reconstruction des surfaces des wafers de silicium à partir des images des circuits imprimés acquises avec un microscope électronique. Dans ce cas, la géométrie joue un rôle dans le système d'acquisition de l'image et le terme de données de la fonction objectif. Le terme *a priori* est composé d'une contrainte de régularité sur la surface et d'un *a priori* de forme sur le dessin prévu dans le circuit. La méthode de reconstruction proposée dans le Chapitre 3 estime aussi un champ de déformation entre un modèle du dessin du circuit et la surface reconstruite, en remplaçant ainsi le modèle de variabilité de forme nécessaire sur l'*a priori* de forme pour un champ de déformation élastique qui permet de quantifier les déviations dans le processus de fabrication.

Enfin, les techniques utilisées pour le développement d'algorithmes numériques efficaces sont expliquées en analysant un exemple de problème utilisant la méthode des surfaces de niveau (*level set method*). À cet effet, le Chapitre 4 développe un algorithme efficace pour la méthode des surfaces de niveau lorsque la fonction de niveau est contrainte de rester une fonction de distance signée. La fonction de distance est conservée par l'introduction d'une contrainte explicite dans le problème de minimisation, tandis que l'accélération de la procédure de minimisation est due aux techniques de séparation de variables et du Lagrangien. Ces techniques introduisent des variables supplémentaires, des contraintes, et des multiplicateurs de Lagrange dans le problème de minimisation d'origine pour le décomposer en sous-tâches d'optimisation qui sont plus faciles à résoudre. En conséquence, l'algorithme proposé est cinq fois plus rapide que l'algorithme original de la méthode des surfaces de niveau.

**Mots-clés :** Méthodes variationnelles, segmentation d'images, reconstruction de l'image, shape-from-shading, minimisation.

# Contents

<b>Acknowledgements</b>	<b>iii</b>
<b>Preface</b>	<b>v</b>
<b>Abstract (English/Français)</b>	<b>vii</b>
<b>Contents</b>	<b>xi</b>
<b>List of Figures</b>	<b>xv</b>
<b>List of Tables</b>	<b>xvii</b>
<b>Introduction</b>	<b>1</b>
0.1 Brief Context . . . . .	1
0.1.1 Motivation: Imaging . . . . .	1
0.1.2 Goals: Inverse Problems . . . . .	1
0.1.3 Techniques: Geometric Variational Models . . . . .	2
0.2 Variational methods . . . . .	3
0.2.1 Minimization Variables . . . . .	3
0.2.2 Objective Functional . . . . .	4
0.2.3 Minimization Techniques . . . . .	7
0.3 Outline of the Dissertation . . . . .	7
0.4 Contributions . . . . .	10
<b>1 Enhanced Compressed Sensing Recovery with Level Set Normals</b>	<b>13</b>
1.1 Introduction to compressed sensing . . . . .	13
1.2 Compressed sensing with recovered normals . . . . .	15
1.2.1 Estimation of level set normals . . . . .	16
1.2.2 Matching normals and compressed sensing measurements . . . . .	18
1.3 Related approaches in compressed sensing . . . . .	19
1.3.1 Edge-guided compressed sensing . . . . .	20
1.3.2 Bregman methods . . . . .	21
1.4 Extension to non-local methods . . . . .	22
1.4.1 Non-local operators . . . . .	22
1.4.2 Compressed sensing with non-local normals . . . . .	25

1.4.3	Estimation of non-local normals . . . . .	26
1.5	Numerical minimization . . . . .	27
1.6	Experimental results . . . . .	27
1.7	Conclusions . . . . .	39
<b>2</b>	<b>Harmonic Active Contours</b>	<b>41</b>
2.1	Introduction . . . . .	41
2.2	Differential geometry in image processing . . . . .	44
2.3	Harmonic Active Contours . . . . .	48
2.3.1	Image segmentation as a Riemannian manifold . . . . .	48
2.3.2	Minimal surface manifold as segmentation criterion . . . . .	50
2.3.3	Feature definition for edge- and region-based segmentation . . . . .	51
2.3.4	Relation to <i>subjective surfaces</i> and anisotropic diffusion . . . . .	52
2.3.5	Relation to robust edge integration . . . . .	53
2.4	Numerical minimization . . . . .	54
2.5	Experimental results . . . . .	55
2.6	Conclusions . . . . .	57
<b>3</b>	<b>Surface Reconstruction From Microscopic Images of Printed Circuits</b>	<b>59</b>
3.1	Introduction: a few words on lithography . . . . .	60
3.2	Shape-from-shading with a shape prior . . . . .	61
3.2.1	Shape-from-shading from SEM images . . . . .	62
3.2.2	Review of surface reconstruction from SEM images of printed circuits . .	63
3.2.3	Variational formulation of shape-from-shading . . . . .	64
3.2.4	Shape prior in shape-from-shading . . . . .	65
3.3	Numerical minimization . . . . .	68
3.4	Estimation of the reflectance map . . . . .	69
3.5	Experimental results . . . . .	69
3.6	Conclusions . . . . .	72
<b>4</b>	<b>An efficient algorithm for the level set method preserving distance functions</b>	<b>79</b>
4.1	Algorithms for the level set method and distance functions . . . . .	80
4.2	Image segmentation and surface reconstruction . . . . .	82
4.3	Efficient level set algorithm preserving signed distance functions . . . . .	84
4.4	Efficient solution to sub-minimization problems . . . . .	85
4.4.1	Sub-minimization with respect to $\phi$ and $u$ . . . . .	86
4.4.2	Sub-minimization with respect to $\varphi$ . . . . .	87
4.4.3	Sub-minimization with respect to $\mathbf{q}$ . . . . .	88
4.4.4	Sub-minimization with respect to $\mathbf{p}$ . . . . .	89
4.5	Experiments and discussions . . . . .	89
4.5.1	Image segmentation and surface reconstruction . . . . .	89
4.5.2	Cortex segmentation with coupled surfaces . . . . .	90
4.5.3	Comparison with other level set methods preserving distance functions	96



4.6	Conclusions . . . . .	101
<b>5</b>	<b>Conclusions and Future Work</b>	<b>103</b>
5.1	Conclusions . . . . .	103
5.1.1	On Geometric Models . . . . .	104
5.1.2	On Variational Methods and Minimization . . . . .	105
5.2	Future Work . . . . .	105
<b>A</b>	<b>Minimization Algorithms of Chapter 1: Local Recovery Method</b>	<b>107</b>
A.0.1	Notation . . . . .	107
A.1	Estimate $u$ matching measurements, sparsity prior, and normals . . . . .	108
A.1.1	Minimization problem with respect to $u$ . . . . .	109
A.1.2	Minimization problem with respect to $\mathbf{d}$ . . . . .	110
A.2	Regularization of normal vectors . . . . .	110
A.2.1	Minimization problem with respect to $\mathbf{n}$ . . . . .	110
A.2.2	Minimization problem with respect to $\mathbf{d}$ and $\mathbf{e}$ . . . . .	111
A.2.3	Minimization problem with respect to $\mathbf{m}$ . . . . .	112
<b>B</b>	<b>Minimization Algorithms of Chapter 1: Non-local Recovery Method</b>	<b>113</b>
B.0.4	Notation . . . . .	113
B.1	Estimate $u$ matching measurements, non-local sparsity prior, and non-local normals . . . . .	113
B.1.1	Minimization problem with respect to $s$ . . . . .	115
B.1.2	Minimization problem with respect to $\mathbf{d}$ . . . . .	115
B.2	Regularization of non-local divergence of normal vectors . . . . .	116
B.2.1	Minimization problem with respect to $\nu$ . . . . .	117
B.2.2	Minimization problem with respect to $\mathbf{d}$ . . . . .	117
<b>C</b>	<b>Minimization Algorithm of Chapter 2</b>	<b>119</b>
C.0.3	Notation . . . . .	120
C.0.4	Minimization associated with $\phi$ and $f^i$ . . . . .	120
C.0.5	Minimization associated with vector fields $\mathbf{p}$ and $\mathbf{q}^1, \dots, \mathbf{q}^k$ . . . . .	122
C.0.6	Minimization associated with $\varphi$ . . . . .	124
<b>D</b>	<b>Minimization Algorithm of Chapter 3</b>	<b>127</b>
D.0.7	Notation . . . . .	127
D.1	Numerical minimization algorithm . . . . .	128
D.1.1	Minimization with respect to $z$ . . . . .	129
D.1.2	Minimization with respect to $\mathbf{d}$ . . . . .	130
D.1.3	Minimization with respect to $\mathbf{u}$ . . . . .	132
D.1.4	Minimization with respect to $\mathbf{V}$ . . . . .	133
	<b>Bibliography</b>	<b>134</b>

## Contents

---

Curriculum Vitae	149
------------------	-----

# List of Figures

1	Illustration of an inverse ill-posed problem. . . . .	2
1.1	Definition of non-local interactions in an image in terms of a weighted graph. .	24
1.2	Compressed sensing reconstruction of Shepp-Logan phantom from 8% of measurements in Fourier domain. . . . .	32
1.3	Compressed sensing reconstruction of MRI brain image from 12% of measurements in Fourier domain. . . . .	33
1.4	Compressed sensing reconstruction of MRI images and the estimated normals to their level curves from 12% of measurements in Fourier domain. . . . .	34
1.5	Compressed sensing reconstruction of two synthetic images characterized by ridges from 10% and 14% of measurements in Fourier domain. . . . .	34
1.6	Compressed sensing reconstruction of Lena from 20% of measurements in Fourier domain. . . . .	35
1.7	Compressed sensing reconstruction of Lena and Barbara from 12% of measurements in Fourier domain. . . . .	36
1.8	Compressed sensing reconstruction of Barbara and baboon images from 20% of measurements in Fourier domain. . . . .	37
1.9	Compressed sensing reconstruction of fingerprint from 20% of measurements in Fourier domain. . . . .	38
2.1	Comparison of segmentation methods for a gray-scale image with a strong inhomogeneity. . . . .	43
2.2	Comparison of segmentation methods for images with ridge structures, inhomogeneities, and noise. . . . .	56
2.3	Comparison of segmentation methods for medical images. . . . .	56
2.4	Comparison of segmentation methods for color images. . . . .	57
2.5	Effects of initialization on the segmentation of images subject to inhomogeneities. . . . .	58
3.1	Binary masks associated with different patterns. . . . .	70
3.2	Image of a line-end structure acquired with a scanning electron microscope. .	71
3.3	Reconstructed surface and estimated deformation field of a line-end structure at multiple scales. . . . .	73
3.4	Reconstructed surfaces and estimated deformation fields of complex patterns. .	74
3.5	Reconstructed surfaces from noisy images. . . . .	75

## List of Figures

---

3.6	Reconstructed surfaces of complex patterns where the lithographic printing failed.	76
3.7	Reconstructed surface of a line-end structure ignoring the shape prior. . . . .	77
3.8	Comparison of the critical distances obtained with our reconstruction method and with lithographic simulations. . . . .	77
4.1	Graphical explanation of shrinkage operator . . . . .	88
4.2	Segmentation results with the proposed level set algorithm applied to natural images. . . . .	91
4.3	Reconstruction of the Stanford bunny from scattered points with a multi-resolution approach. . . . .	92
4.4	Reconstruction of the Stanford bunny, dragoon, and hand surfaces from scattered data points at different resolutions. . . . .	93
4.5	Reconstruction of the Stanford Budha and horse surfaces from scattered data points at different resolutions. . . . .	94
4.6	Segmentation of grey-white matter interface on MRI images of human cortex. .	95
4.7	Evolution of the level set method when the level set function is not constrained to be a signed distance function. . . . .	96
4.8	Qualitative comparison of state-of-the-art algorithms for the level set method with a synthetic image. . . . .	98
4.9	Summary of quantitative comparison of state-of-the-art algorithms for the level set method with 72 images. . . . .	99
4.10	Comparison of state-of-the-art algorithms for the level set method for the segmentation of grey-white matter interface on MRI images of human cortex. . . .	100



## List of Tables

1	Definition of objective functionals. . . . .	10
2	Summary of contributions by chapter. . . . .	11
1.1	Comparison of compressed sensing reconstruction of MRI images. . . . .	29
1.2	Comparison of compressed sensing reconstruction of noisy MRI images. . . . .	29
1.3	Comparison of compressed sensing reconstruction of textured images. . . . .	30



# Introduction

## 0.1 Brief Context

The title *Geometric variational models for inverse problems in imaging* summarizes the motivation, goals, and techniques of this dissertation. However, a brief description is still necessary to put the different terms in context. This is the goal of this first section.

### 0.1.1 Motivation: Imaging

Imaging is a multidisciplinary field that studies the generation, analysis, and visualization of images. The core of imaging as a formal science lies in the definition of the “image chain”, i.e., the process that results in the creation of an image from a collection of physical measurements. The wide spectrum of these measurements leads to a wide range of imaging modalities; from digital photography to specialized measurements in astronomy, microscopy, or medicine. The range of applications, tasks, and sub-disciplines is equally broad. Image processing, the core of this dissertation, is concerned with the analysis and characterization of images in relation to the physical world. Imaging, in a nutshell, is an extended eyeball that helps us understand the universe from the macro- to the nano-scale.

The importance of research in such a field is unquestionable, and yet the number of open questions is too large to be presented in a single dissertation. Consequently, this dissertation focuses on the study of four specific problems, but develops techniques and principles that can be applied to other imaging tasks.

### 0.1.2 Goals: Inverse Problems

Inverse problems can be informally defined as the determination of the “cause”, from observation of the “effects”. Each inverse problem is thus coupled with a forward problem, a model that explains the “effects” given the “cause”. Inverse problems arise naturally in imaging by following backward the image chain, and constitute an extensive source of research problems. Some examples include the reconstruction of an image from reduced measurements, its segmentation into salient structures, or the reconstruction of a three-dimensional scene from two-dimensional images.




Figure 1: Inverse ill-posed problem. Non-uniqueness refers to the existence of multiple signals  $u$  leading to the same measurements  $f$ , while ill-conditioning refers to an inversion process where small errors in the data lead to large errors in the solution.

Formally, inverse problems can be defined as the recovery of an unknown function  $u$  from the observed measurements  $f = K(u) + n$ . The operator  $K$  is determined by the forward model, while the noise variable  $n$  is introduced to account for the limitations of this model.

Interesting inverse problems are ill-posed, that is, they suffer from non-uniqueness of the solution and ill-conditioning of the inversion process. To overcome these obstacles and find a solution that is physically meaningful, it is necessary to consider information about the solution that goes beyond the measurements. This additional information will guide the inversion process and rule out solutions physically irrelevant or due to noise. In other words, inverse problems are solved by making assumptions about the solution based on prior knowledge of the physical world. See Figure 1.

### 0.1.3 Techniques: Geometric Variational Models

The prior knowledge used to “invert the problem” can take many forms, but two common assumptions in imaging are that the solution  $u$  is **smooth** and has some **geometric properties**.

*Natura non facit saltus*<sup>2</sup> has been a principle in natural sciences since Aristotle, it is a common assumption in physics, and has naturally lead to the same premise in imaging. As a consequence, the unknown signal  $u = u(x)$ ,  $x \in \Omega \subset \mathbb{R}^n$ , is considered a function in space and constrained to some degree of smoothness<sup>3</sup>, which is usually measured in terms of its first-order derivatives.

Geometric models have also been explored in imaging as a source of prior knowledge to solve inverse problems. In this case, the assumptions on the solution  $u$  are formulated in terms of the normal vectors to its level curves, their mean curvature, or the surface of a manifold defined by the unknown signal and the image. The smoothness constraints on  $u$  are then substituted by geometric terms that involve second-order derivatives and produce more faithful and complex models of the physical world. As a common rule, more detailed models lead to more accurate solutions, but also result in problems that are more difficult to solve.

---

<sup>2</sup>Latin for “nature does not make jumps”

<sup>3</sup>Images have jumps, but they are smooth along lower-dimensional hypersurfaces, e.g., edges.



Variational methods, in particular, solve inverse problems by formulating minimization problems that summarize both the forward model and the necessary assumptions about the solution. The minimizer of the resulting variational problem is considered a meaningful solution of the original inverse problem, it thus defines a variational model.

The formulation and solution of accurate variational models is the subject of this thesis and requires, therefore, a more detailed introduction.

## 0.2 Variational methods

This dissertation makes use of geometry to formulate variational models that provide physically meaningful solutions to different inverse problems in imaging, developing also efficient numerical algorithms to compute them.

The variational formulation that we adopt reads

$$\min_u E(u) = \min_u J(u) + \alpha R(u) \quad u \in BV(\Omega), \quad (1)$$

where  $\alpha$  is a weighing parameter, and  $BV(\Omega)$  is the Banach space of all functions with bounded variation in  $\Omega \subset \mathbb{R}^n$ . With this formulation, the inversion techniques that we propose are described by the three key ingredients of the optimization problem (1): the minimization variables  $u$ , the objective functional  $E(u)$ , and the minimization algorithm.

### 0.2.1 Minimization Variables

The physical meaning of the minimization variables depends on the inverse problem at hand and the way we describe its solution. A correct description of the the solution should reflect the prior information we have about it, easily translate the assumptions introduced in the problem into the objective functional, and simplify the minimization problem as far as possible. This is better understood with an example.

Image segmentation is interested in the decomposition of an image into homogeneous regions, and the minimization variable is a parametrization of the contours that describe the boundary between these regions. If the topology of the different regions is not known *a priori*, an implicit parametrization of the contours (as the zero-level set of a two-dimensional function) simplifies the design of algorithms that naturally handle changes of topology. In this context, the *active contours without edges* of Chan and Vese [2001] decomposes the image into piecewise constant regions according to the model of Mumford and Shah [1989], and results in an objective functional that is widely used but not convex. If the original level set function is further constrained to take values in the range  $[0, 1]$ , then the previous functional can be relaxed to the convex model of Chan et al. [2006], Strandmark et al. [2009], and fast minimization algorithms can be adopted (Goldstein et al. [2009]). In this case, the choice of the correct minimization variables allows the reformulation of the piecewise constant model

of Mumford and Shah into a convex minimization problem that is easy to solve<sup>4</sup>.

It is equally important to define the space of admissible solutions in a coherent manner with the previous assumptions and the domain requirements of the objective functional. To that purpose, variational models consider the unknown signals as  $L^2$  functions constrained to some degree of regularity in the spatial domain  $\Omega \subset \mathbb{R}^n$ . In particular, the space of bounded variation  $BV(\Omega)$  is widely used in image processing because it models a key feature of images, the edges, and admits the analysis of basic cartoon-like images.

### 0.2.2 Objective Functional

The objective functional defines the criterion used to recover the unknown signal  $u$  from both the observed measurements and the prior knowledge available about the solution. It can thus be decomposed into two terms: a data or measurement term  $R(u)$ , and a prior term  $J(u)$ . The data term  $R(u)$  describes the relation between the measurements  $f$  and the unknown variable  $u$ ; the prior term  $J(u)$  describes the assumptions on  $u$  that lead to physically meaningful solutions.

Some inverse problems consider one of these terms, usually  $R(u)$ , a minimization constraint instead of an additive term in the objective functional. In that case, the minimization problem of Equation (1) can be reformulated into the constrained minimization problem

$$\min_u J(u) \quad \text{subject to} \quad R(u) < \sigma, \quad (2)$$

where  $\sigma$  is related to the level of noise present in the measurements. Indeed, optimization theory tells us that, under some considerations, problems (1) and (2) are equivalent, that is, they have the same minimizer and solving any of the two determines the parameter  $(\sigma, \alpha)$  of the other (Boyd and Vandenberghe [2004]). In this dissertation we only consider the Lagrangian formulation of Equation (1).

The data term  $R(u)$  is formulated in terms of the forward operator  $K$  associated with the inverse problem. In particular,  $R(u) = \rho(K(u) - f) = \rho(n)$  measures the deviation between the forward model ( $f = K(u) + n$ ) and the observed measurements with the penalty function  $\rho(\cdot)$ . The choice of this penalty function depends on the limitations of the forward model or on the statistical model of the noise present in the measurements. In a Bayesian framework, in fact,  $R(u)$  corresponds to the log-likelihood of the measurements given the unknown signal, and optimal penalty functionals can be defined for each noise model by maximum likelihood estimates. The definition of maximum a posteriori estimates requires an additional statistical model of the unknown signal  $u$ , which is given by the prior term  $J(u)$ .

---

<sup>4</sup> Chan et al. [2006] assumes that the mean intensity values associated with each region of the image are known; while the piecewise constant model of Mumford and Shah [1989] includes them in the minimization problem. For this reason, the convex model of Chan et al. [2006] is in fact a relaxed equivalent to the piecewise constant model of Mumford and Shah. The formulation of Strandmark et al. [2009] solves this issue.

The prior knowledge about the solution is summarized in the functional  $J(u)$ , which describes the physical properties of the object being imaged and can take a multitude of forms. Most assumptions, nevertheless, fall in one of the following three categories: smoothness constraints, geometric properties of the solution, or prior knowledge of the shape or topology of the object being imaged. Although this dissertation focuses on geometric models, a brief description of each category is still presented in order to understand their limitations and use.

### Smoothness

Regularization techniques impose smoothness on the solution  $u$  to overcome the ill-conditioning of the problem, that is, they provide stable solutions by controlling those parts of the inversion process that are particularly sensitive to noise. In imaging, the smoothness of  $u$  is usually measured in terms of its first-order derivatives with a functional of the form

$$J(u) = \int_{\Omega} |\nabla u|^p, \quad (3)$$

where  $p$  is a parameter that chooses the type of regularization applied. Dirichlet or  $H^1$  regularization is obtained with  $p = 2$  and leads to solutions overly smooth that do not suit natural images. Total variation (TV) is obtained with  $p = 1$ , and it has become the default regularizer for image reconstruction since its first appearance in Rudin et al. [1992] for image denoising. The success of TV is due to the following two properties: it preserves edges in the form of jumps because it is the natural regularizer for cartoon-like images, and it is a convex functional that offers good minimization properties.

Natural images, however, do not agree with a cartoon-like model, and they require regularizers with higher-order derivatives to avoid the staircase effect of TV reconstructions. To that purpose, a generalized TV regularizer has been proposed recently in Bredies et al. [2010]. This generalization of TV is still a convex functional, but it prefers piecewise smooth images over staircase images because it finds a balance between first- and second-order derivatives. In some applications, however, the staircase effect is not the only image artifact that must be overcome and more complex models with higher-order derivatives are necessary; clearly, information that goes beyond smoothness is necessary to reconstruct the geometry and topology of the physical world.

### Geometry

Geometric models interpret an image  $I$  as a surface in the image plane  $\Omega \subset \mathbb{R}^2$ , where  $I(x)$  represents the height of the surface at point  $x$ . This interpretation of images endows the level sets of the image with certain geometric properties and measures, and geometric variational models of minimal length curves, minimal mean curvature, or minimal area surface can then be defined to segment or reconstruct an image. The two key geometric concepts used in imaging are the normal vector to a curve or surface and its curvature; they are both defined in

terms of differential operators, but present more complex forms than standard regularizers.

The interpretation of images as surfaces in the image plane  $\mathbb{R}^2$  has been extended to multi-channel images and non-Euclidean spaces by means of differential geometry in Sochen et al. [1998]. In this case, the image is considered a space-feature manifold embedded in a higher dimensional space, and the usual norms, curvatures, and areas are measured in terms of the metric of the manifold. The resulting methods rely therefore on differential geometry to define functionals that naturally take into account the coupling of the different image channels, they can be extended to higher dimensional spaces and general parametric manifolds.

In general, geometric models lead to variational minimization problems with non-convex functionals and second-order derivatives; this results in additional theoretical and computational issues. In terms of theory, it is necessary to identify admissible function spaces associated with regularized normal vectors or curvature functionals, and to study the existence and uniqueness of solutions of non-convex minimization problems. In terms of computation, it is necessary to develop specific numerical algorithms that overcome the speed limitations imposed by second-order derivatives, and avoid the local minima of non-convex functionals. Indeed, variational minimization techniques are based on the numerical evolution of partial differential equations (PDEs), where the time and spatial steps of the flow are limited by the order of the PDE derivatives, and result in extremely slow algorithms for geometric models. At the same time, to overcome the non-convexity of the objective functional and to avoid local minima these algorithms require good initialization strategies or multi-resolution approaches. Consequently, geometric models call for additional optimization efforts to compute a solution of the original inverse problem.

### Prior Knowledge of the Shape or Topology of the Solution

In some applications, prior knowledge about the shape or topology of the solution is available and can be exploited to improve variational techniques. A clear example of this kind of methods corresponds to atlas-based segmentation of medical images, where an atlas is introduced to summarize the constraints on the topology and shapes of the segmented regions imposed by human anatomy. The simplest case is the segmentation of a single structure, where the atlas reduces to a shape prior. In this case, the standard variational approach defines a measure of similarity between the evolving shape and the shape prior, and defines  $J(u)$  to penalize deviations with respect to the expected shape that do not fit a model of shape deformation. These methods obtain accurate segmentations robust to noise and image inhomogeneities, but they require specific models of the object being imaged and are limited to very specific applications.

### 0.2.3 Minimization Techniques

Eventually, to solve the inverse problem and find the unknown signal  $u$ , it is necessary to solve a variational minimization problem. In the context of the calculus of variations, the objective functional is traditionally minimized with an iterative procedure equivalent to gradient descent. As the minimization variable is here a function, the usual gradient operator is substituted by the first variation  $\delta$  and results in a descent flow  $\frac{\partial u}{\partial t} = -\delta E(u)$  that is given in terms of partial differential equations. This PDE flow is then discretized both in time and space, and leads to descent algorithms where time and space steps are limited by the order of the derivatives by Courant-Friedrichs-Lewy (CFL) stability condition (Courant et al. [1967]). These conditions assure convergence and accuracy in the solution, but make explicit PDE evolution extremely slow for geometric models with second-order derivatives.

In the last few years, PDE methods have been combined and replaced by more efficient optimization techniques to provide faster and more stable algorithms. These optimization techniques are based on splitting variables and on the introduction of equality constraints that are solved by quadratic-penalties, Bregman iterations, or the equivalent augmented Lagrangian method (see e.g. Wang et al. [2008], Goldstein and Osher [2009], Wu and Tai [2010]). The key idea of variable-splitting is to decompose the original problem into independent sub-optimizations for each variable; by design, these sub-optimizations are easier to solve because they reduce to convex problems, involve only first-order derivatives, or are decoupled for each pixel. At the same time, the constraints associated with the split variables are easily introduced in the minimization algorithm by additional penalty terms in the objective functional or by dual variables that are periodically updated. The initial minimization is thus divided into small and easier subtasks that are efficiently solved, and the resulting minimization algorithm is faster than the original PDE flow of problem.

## 0.3 Outline of the Dissertation

As we have seen, variational methods solve inverse problems in two steps: formulation of a minimization problem, and design of a numerical minimization algorithm to solve it. It is then natural to structure this dissertation in the same manner.

Chapters 1-3 investigate different inverse problems in imaging and propose geometric variational models to solve them. Naturally, each application exploits geometry in a different way, and results in a variational method that incorporates geometry in a different term of the objective functional. For a quick view, Table 1 summarizes the role of geometry in each chapter. Chapter 4 explains the techniques used to develop efficient minimization algorithms to compute these solutions.

The rest of this section describes explicitly how each chapter fits in the previous variational context and provides a “big picture” of the dissertation.

### Chapter 1: Enhanced Compressed Sensing Recovery with Level Set Normals

In the context of compressed sampling, Chapter 1 presents a method to reconstruct an image  $u$  from reduced measurements  $f$  in the Fourier domain. The forward operator  $K$  is a sampling matrix that combines Fourier's sensing basis with the binary matrix that extracts the corresponding measurements in  $f$ . The noise in the measurements is assumed Gaussian, and the data term simply measures the variance of the noise in the measurements. The prior functional  $J(u)$  is composed of two terms, associated with two different assumptions about the reconstructed image  $u$ . First, a TV regularizer that promotes images with sparse edges and acts as the sparsity term characteristic of compressive sensing. Second, a geometric term that aligns the estimated level set normals with the reconstructed image gradients. The resulting variational model, therefore, exploits the geometric properties of images to complement the TV regularizer in the prior term  $J(u)$ .

The first part of Chapter 1 explains in detail the introduction of this geometric term in compressed sensing, justifying why the alignment of the level curve normals with the gradients of the reconstructed signal improves image recovery. The second part of Chapter 1 extends this method to non-local operators and images defined on graphs. In this case the proposed technique exploits the geometry of the graph defined by the non-local operators, and it is able to recover finer details and structures of textured images. Both local and non-local methods reconstruct the image by iterating the following two steps: first, estimation and regularization of the normal vectors to the level curves of the image; second, reconstruction of an image fitting the normal vectors, the compressed sensing measurements, and the sparsity constraint. The resulting reconstruction methods are thus reduced to a series of convex minimization problems that can be efficiently solved.

The experimental results of Chapter 1 show that the proposed method improves image recovery in two ways. First, the introduction of the geometric term enables the recovery sharp edges as well as smoothly varying image regions, avoiding the staircase effect of total variation. Second, the regularization of the estimated normals makes the proposed reconstruction method robust to noise and a reduced number of measurements.

### Chapter 2: Harmonic Active Contours

In the context of image segmentation, Chapter 2 partitions an image into different regions by parametrizing the boundaries of these regions with a level set function  $\phi$ , which becomes the minimization variable. The observed measurements  $f$  correspond here to the different channels of the image, and the objective functional couples the data and prior smoothness terms in a single geometric functional that measures the surface of the manifold associated with the segmentation.

To that purpose, Chapter 2 makes use of differential geometry, and interprets each possible segmentation as a 2-dimensional manifold embedded in a higher dimensional space, from

which it inherits a metric. The proposed metric allows the definition of distances between points in the manifold that consider simultaneously the spatial distance between the points, the values of the image at these points, and the labels assigned to them in the segmentation. Consequently, the surface of this manifold naturally measures the smoothness of the segmentation and the agreement between the values of the image at each point and the label assigned to them in the segmentation; this couples both data and regularity measures into a single geometric functional, and assigns a minimal surface manifold to an optimal segmentation of the image. The proposed geometric model offers an additional advantage: it aligns the gradients of the level set function with the gradients of the image, and exploits this directional information to overcome fragmented edges and image inhomogeneities.

The potential of this geometric approach lies in the general definition of Riemannian manifolds, which naturally generalizes the existing segmentation methods of Kichenassamy et al. [1995], Caselles et al. [1997], Chan and Vese [2001], and Kimmel and Bruckstein [2002] to higher dimensional spaces, non-flat images, and feature spaces. As a consequence, the proposed technique improves the segmentation of multi-channel images, images with inhomogeneities, and images characterized by geometric structures like ridges or valleys.

## Chapter 3: Surface Reconstruction from Microscopic Images of Printed Circuits

In the context of shape-from-shading, Chapter 3 reconstructs 3D surfaces of silicon wafers from 2D images of printed circuits taken with a scanning electron microscope. In this case, the surface of the wafer is parametrized by the height  $z(x)$  associated to the surface at each point  $x$  of the image plane, while the measurements are given by the shading observed in the image at that point. This shading depends on the orientation between the light source and the surface, and results in a geometric forward operator that is directly written in terms of the alignment between the normal vector to the surface and the light source. Geometry, therefore, is intrinsic to the definition of shape-from-shading.

The method developed in Chapter 3 combines the physical model of the optical acquisition system with prior knowledge about the shapes of the patterns in the circuit. To that purpose, the objective functional combines a measurement term based on the irradiance equation with a shape prior that constrains the shape of the surface to agree with the expected shape of the pattern. To account for the variability of the manufacturing process, the model allows a non-linear elastic deformation between the expected pattern and the reconstructed surface. As a consequence, the resulting minimization problem involves two unknowns, and the reconstruction method provides two outputs: a reconstructed surface and a deformation field. The reconstructed surface is derived from the shading observed in the image and the prior knowledge about the circuit pattern, which results in a shape-from-shading technique robust to noise. The deformation field produces a mapping between the expected shape and the reconstructed surface, which provides a quantitative measure of deviation in the manufacturing process.

## List of Tables

Table 1: Each applications makes use of geometry in a different term of the objective functional. While compressed sensing incorporates it in the prior term and shape-from-shading in the measurements, the proposed segmentation technique couples the regularity of the image and the level set function in a geometric functional that measures the surface of a manifold.

Chapter	Inverse problem	$R(u)$	$J(u)$
1	Image reconstruction	Fourier measurements	Sparsity, Geometry
2	Image segmentation	Geometry – coupling of data and smoothness term	
3	Shape-from-shading	Geometry	Shape prior

### Chapter 4: an efficient minimization algorithm for the level set method

Solving inverse problems requires the formulation of a minimization problem, but also the design of a minimization algorithm to eventually compute a solution. For this reason, a considerable amount of time has been invested in the development of efficient numerical algorithms for the previous inverse problems. The resulting algorithms make use of variable-splitting and constrained optimization to overcome the limitations of functionals that are non-differentiable, involve second-order derivatives, or are non-convex. It is unnecessary, however, to present the algorithms developed for all the previous problems to understand the principles that guide these techniques; instead, it seems more interesting to analyse in detail the algorithm designed for the level set method in Chapter 4, and leave the algorithms of Chapters 1-3 for the appendices.

In the context of image segmentation and surface reconstruction, Chapter 4 designs an efficient algorithm for the level set method when the level set function is required to remain a signed distance function. In this case, the original level set method of Osher and Sethian [1988] is limited by two important numerical issues. First, the minimizing flow does not preserve the level set function as a distance function, and it requires periodic reinitialization procedures. Second, the time step of this flow is limited by the CFL condition, and the resulting algorithm is slow. To avoid these limitations, Chapter 4 makes use of variable-splitting, augmented Lagrangians, and constrained minimization techniques to develop a fast algorithm for the level set method that naturally preserves the distance function. The proposed algorithm is fast because it is not limited by the CFL condition, and it does not require periodic reinitialization procedures because the level set function is explicitly constrained to be a distance function in the minimization problem. As a result, the algorithm is five to six times faster than related state-of-the-art methods.

## 0.4 Contributions

This dissertation investigates different inverse problems in imaging in a chapter-by-chapter manner, therefore its contributions organize naturally in the same manner in Table 2.



In terms of publications, the material of Chapters 1 and 4 has been published in two journal papers (Estellers et al. [2013a, 2012]), while the material of Chapters 2 and 3 has been presented in two journal papers (Estellers et al. [2013b,c]) that are in peer review.

Table 2: Summary of the contributions of this dissertation. The main contributions of Chapters 1-3 are the formulation of new geometric variational models, however, efficient algorithms are also developed for the respective minimization problems. On the other hand, Chapter 4 focuses on the design of an efficient minimization algorithm for the level set method.

Chapter	Contributions
1	Introduction of normal alignment in reconstruction criterion of compressed sensing. Extension to the non-local geometry of the graph associated with the image.
2	Interpretation of the segmentation of an image as a minimal surface manifold. Generalization of existing segmentation methods (Caselles et al. [1997] Chan and Vese [2001], Kimmel and Bruckstein [2003]) to Beltrami framework.
3	Introduction of a shape prior in shape-from-shading. Introduction of an elastic deformation field instead of a statistical shape model.
4	Introduction of a signed distance function constraint in level set minimization problem. Efficient minimization of the resulting constrained optimization problem.

Due to the diversity of topics, nevertheless, the full understanding of these contributions calls for a clear context of each inverse problem, that is, for the lecture of each chapter.



# 1 Enhanced Compressed Sensing Recovery with Level Set Normals

This chapter presents a recovery method for compressed sensing that exploits the geometric properties of images to recover signals of high quality from few measurements. The data term of this inverse problem corresponds to reduced measurements in Fourier domain, while the prior term combines the sparsity constraint characteristic of compressed sensing with a geometric alignment term. This additional term aligns the gradients of the reconstructed image with an estimate of the normal vectors to the level curves of the image. As a result, the proposed method first estimates the geometry of the image and then incorporates it in the reconstruction criterion of compressed sensing.

The reconstruction of an image is thus done by iterating the following two steps: first, the normal vectors to the level curves of the image are estimated by regularizing a point-wise estimate from the previous iteration; second, an image is reconstructed with a recovery model that fits the normal vectors, the compressed sensing measurements, and the sparsity constraint. This results in a series of convex optimization problems that are efficiently solved.

The proposed technique is extended in a natural manner to non-local operators in graphs, where it exploits the geometry of the graph associated with the image to recover fine detail structures and textured images. The key point, in this case, is the regularization of non-local gradients defined in the graph, which is done in terms of their graph divergence.

Both in the local and non-local case, experiments show a clear improvement in the quality of the reconstructed images and the robustness of the proposed method to noise and a reduced number of measurements.

## 1.1 Introduction to compressed sensing

Compressed sensing (CS) is founded on the principle that, through optimization, the sparsity of a signal can be exploited to recover it from a reduced number of measurements. This

---

The material of this chapter has been partially published in Estellers et al. [2013a].

simple and yet powerful idea is intriguing because it brings Shannon's sampling theorem into question. Compressed sensing is in fact the equivalent of Shannon's theorem from the point of view of sparsity: while Shannon states that to recover a band limited signal the sampling rate must be at least twice the maximum frequency present in the signal; CS relates the sparsity of a signal in certain basis with the number of measurements in another basis necessary to recover it from a minimization problem. A few definitions are necessary to understand the formulation of CS as an inverse recovery problem.

We say that a signal  $u \in \mathbb{R}^n$  is  $s$ -sparse in the basis or dictionary  $\Psi$  if it can be expressed by  $s$  non-zero coefficients in that basis, i.e.  $\|\Psi u\|_0 = s$ ; while  $u$  is compressible if most of the energy in  $\Psi u$  is contained in its largest  $s$  coefficients.

Given  $\Phi$  and  $\Psi$  two orthobasis or dictionaries of  $\mathbb{R}^n$  and a signal  $u$  that is sparse in basis  $\Psi$ , the CS recovery problem is formulated as the reconstruction of  $u$  from  $m$  linear measurements  $f$  in the sensing basis  $\Phi$ , with  $m < n$ . Ideally we should measure the  $n$  projections of  $u$  in basis  $\Phi$ , that is  $\Phi u$ , but we only observe a small subset  $f = Au$  of size  $m$ . The sampling matrix  $A = R\Phi$  results from the combination of the sensing basis  $\Phi$  and the matrix  $R$  that extracts the corresponding measurements in  $f$ . Consequently, the system  $f = Au$  is undetermined, and the sparsity of the signal  $u$  must be exploited to "invert" the problem and obtain a correct reconstruction.

The obvious strategy would be to recover the sparsest  $u$  agreeing with the measurements, that is, to solve the following non convex problem

$$\min_u \|\Psi u\|_0 \text{ subject to } Au = f. \quad (1.1)$$

Problem (1.1) is NP-hard due to the  $\ell_0$  penalty and only approximate solutions can be used in real applications. Relaxing the  $\ell_0$  penalty to an  $\ell_1$  norm, problem (1.1) becomes the convex problem

$$\min_u \|\Psi u\|_1 \text{ subject to } Au = f. \quad (1.2)$$

Recent results in CS, see Candès et al. [2011] and Needell and Ward [2012], prove that the minimization problem (1.2) exactly recovers  $s$ -sparse signals with an overwhelming probability when the number of measurements is  $\mathcal{O}(s \log n)$ . In addition, if the sampling matrix  $A$  verifies certain restricted isometry condition, then problem (1.2) actually recovers the signal associated with the  $s$  largest coefficients of  $u$  in basis  $\Psi$ , i.e., exact recovery for  $s$ -sparse signals and recovery of the  $s$ -sparse  $\ell_2$  approximation for compressible signals. In other words, in CS the NP-hard problem (1.1) can be safely substituted by the convex problem (1.2) if matrix  $A$  is chosen appropriately.

When the measurements are contaminated with noise, the constraint  $Au = f$  on the

measurements is relaxed. In particular, under Gaussian noise the recovery is given by

$$\min_u \|\Psi u\|_1 \text{ subject to } \|Au - f\|_2 \leq \sigma, \quad (1.3)$$

where  $\sigma$  is related to the noise level. From optimization theory, we know that problem (1.3) is equivalent to

$$\min_u \|\Psi u\|_1 + \frac{\alpha}{2} \|Au - f\|_2^2, \quad (1.4)$$

in the sense that solving any of the two determines the parameter  $(\sigma, \alpha)$  in the other and both have the same minimizer, see e.g., Boyd and Vandenberghe [2004].

Designing the sparsifying basis  $\Psi$  depends on the signal at hand. For images a common choice are orthogonal wavelets or the discretized total variation regularizer. TV assumes that the edges of an image are sparse, and it is extensively used in image reconstruction because it recovers sharp edges and preserves the convexity of problem (1.2). The discrete sparsity prior of CS  $\|\Psi u\|_1$  is thus substituted in imaging by a discretization of the TV regularizer  $J(u) = \int_{\Omega} |\nabla u|$  in an abuse of notation.

The basic CS recovery model considered in this chapter is then

$$\min_u J(u) + \frac{\alpha}{2} \|Au - f\|_2^2, \quad (1.5)$$

where  $A$  samples random Fourier measurements<sup>1</sup>, and the prior term  $J(u)$  of total variation acts as sparsity criterion. The rest of this chapter shows how we can improve this recovery model by exploiting the geometry of images, introducing an additional prior term in the objective functional.

In particular, Section 1.2 presents the proposed geometric model, and Section 1.3 explains its relation to similar techniques. The proposed method is then extended to non-local operators in Section 1.4, and the minimization algorithms are briefly sketched in Section 1.5. Finally, experiments are presented in Section 1.6 and conclusions drawn in Section 1.7.

## 1.2 Compressed sensing with recovered normals

The main idea behind the proposed method is that the recovered normals of an image can significantly improve CS reconstruction. This observation raises two main questions: How can we recover normals robustly and accurately from CS measurements? How can we introduce the estimated normals in the CS recovery model (1.5)? The answer proposed in this chapter is a two-step iterative method that we initialise with the baseline solution of (1.5).

---

<sup>1</sup>The proposed matrix  $A$  satisfies the restricted isometry condition with high probability and is therefore a common choice in MRI, see Rudelson and Vershynin [2006]. Nevertheless, other basis or dictionaries could be equally adopted in our formulation, e.g., Lustig et al. [2005], Candès and Romberg [2005], Ma et al. [2008].

In the first step of each iteration, we estimate the normal vectors to the level curves of the image by regularizing a point-wise approximation with a vectorial extension of the ROF model (Rudin et al. [1992]). To that purpose, we extend the weighted TV regularizer  $J_w(\cdot)$  to vector fields, with a weighting function  $w$  designed to verify  $w \approx 0$  near the edges and  $w \approx 1$  on flat regions of the estimated image.

Once the normals are estimated, we can find an image that fits the measurements, the estimated normals, and the sparsity criterion. The process is then iterated and can be summarized as

$$\begin{cases} \mathbf{n}_k = \underset{|\mathbf{n}| \leq 1}{\operatorname{argmin}} J_w(\mathbf{n}) + \frac{\mu}{2} \|\mathbf{n} - \hat{\mathbf{n}}\|_2^2 \\ u_k = \underset{u}{\operatorname{argmin}} J(u) + \gamma \langle \operatorname{div} \mathbf{n}_k, u \rangle + \frac{\alpha}{2} \|Au - f\|_2^2, \end{cases} \quad (1.6)$$

where  $\mathbf{n}_k$  and  $u_k$  are the normal field and the image estimated at iteration  $k$ ,  $\hat{\mathbf{n}}$  is a point-wise estimate of the normals obtained from  $u_{k-1}$ , and  $w = g(|\nabla u_{k-1}|)$  is an inverse edge detector.

This sections details each of these two steps, which both reduce to convex optimizations that can be efficiently solved. Combining the two stages into one would lead to a non-convex model of higher order, and the resulting minimization would be slower and suffer from local minima. A two-step method is computationally more efficient in the same way than splitting variables in Section 1.5 simplifies the minimization problems and leads to closed form solutions. The drawback of this two-step procedure is the lack of rigorous theory and proof of convergence of the resulting algorithm. Nevertheless, experimental results show that a single iteration of the proposed method already improves the standard TV recovery model (1.5), while the optimal performance in terms of SNR is attained after a few iterations. A similar limitation affects the iterative edge-guided CS algorithm of Guo and Yin [2010].

### 1.2.1 Estimation of level set normals

At each iteration  $k$ , the normals of the image are estimated in two steps. First, a noisy point-wise estimate  $\hat{\mathbf{n}}$  is obtained from the previous solution  $u_{k-1}$ ; this estimate is afterwards regularized solving a minimization problem to obtain  $\mathbf{n}_k$ .

Given an estimate of the underlying image  $u_{k-1}$ , the normal vectors associated with its level set curves are defined as

$$\hat{\mathbf{n}} = \begin{cases} \frac{\nabla u_{k-1}}{|\nabla u_{k-1}|} & \text{if } |\nabla u_{k-1}| > 0 \\ 0 & \text{if } |\nabla u_{k-1}| = 0. \end{cases} \quad (1.7)$$

This first estimate of the normals  $\hat{\mathbf{n}}$  is then denoised with a combination of the vectorial ROF model of Blomgren and Chan [1998] with the constraint  $|\mathbf{n}| \leq 1$ . In particular the vector

field  $\mathbf{n}_k = (n_x, n_y)_k$  is defined as the solution of the following variational problem

$$\min_{|\mathbf{n}| \leq 1} J_w(n_x, n_y) + \frac{\mu}{2} \|n_x - \hat{n}_x\|_2^2 + \frac{\mu}{2} \|n_y - \hat{n}_y\|_2^2, \quad (1.8)$$

where  $J_w(n_x, n_y)$  is the extension of the weighted TV regularizer to vector fields, and  $w = g(|\nabla u_{k-1}|)$  is an edge detector designed to verify  $w \approx 0$  near the edges and  $w \approx 1$  on flat regions of  $u_{k-1}$ .

Weighting the TV regularizer with an edge detector  $w = g(|\nabla u_{k-1}|)$  encourages the edges of the regularized solution to coincide with the main edges of the noisy signal  $u_{k-1}$ . To be robust against false edges, we can use the robust edge detector of Black et al. [1998], which provides a statistical interpretation of the edge-stopping functions of the anisotropic diffusion of Perona and Malik [1990]. In this statistical interpretation, edges are considered outliers in the normal distribution of  $|\nabla u_{k-1}|$  associated with noisy piecewise constant regions, and the edge-stopping functions  $g(\cdot)$  are derived from error norms robust to outliers. The edge detectors have thus a parameter  $\sigma_e$  that acts as a soft-threshold in the detection of outliers, and can be estimated a priori from the values of  $|\nabla u_{k-1}|$  in the image. Based on the results of Black et al. [1998], the edge detector is defined as

$$g(x) = \begin{cases} \frac{1}{2} \left(1 - \frac{x^2}{\sigma_e^2}\right)^2 & |x| \leq \sigma_e \\ 0 & |x| > \sigma_e \end{cases}, \quad (1.9)$$

with  $\sigma_e = 1.4826 \text{median}(|\nabla u_{k-1}| - \text{median}(|\nabla u_{k-1}|))$ . In fact, using other robust edge detectors in the estimation of normals leads to similar results in the final CS recovery method.

The constraint  $|\mathbf{n}| \leq 1$  in (1.8) corresponds to a relaxation of the condition  $|\mathbf{n}| = 1$  inherent to the definition of normals. It is numerically necessary in flat regions, where  $\nabla u = \mathbf{0}$  and we cannot numerically normalize the gradient vector.

In the context of image denoising and inpainting, Hahn et al. [2011] shows that a combination of vectorial TV with an  $\ell_1$  data term provides better estimates of the normal fields of images contaminated with salt-and-pepper noise. In compressed sensing, however, the first estimate of the normals  $\hat{\mathbf{n}}$  is obtained from a CS recovery algorithm, and we cannot assume any particular noise model but certain kinds of artifacts. Experimentally, both the introduction of a weight in the TV regularizer and the relaxation  $|\mathbf{n}| \leq 1$  lead to more robust estimates of the normal field for CS recovery. Indeed, the use of an edge detector as the weighting function in  $J_w(\cdot)$  gives more weight to the data term in the estimation of the normals near the edges, where  $\nabla u$  is large and  $\hat{\mathbf{n}}$  is clearly defined; in regions with no clear orientation, where  $\hat{\mathbf{n}}$  is affected by noise, the regularization averages neighbouring vectors and produces normal fields close to zero. Similarly the relaxation  $|\mathbf{n}| \leq 1$  experimentally leads to smaller norm values in noisier areas. In non-flat areas, in fact, the value of  $|\mathbf{n}_k|$  can be considered an experimental measure of quality in the estimation of the normal field.

### 1.2.2 Matching normals and compressed sensing measurements

Once the normal field  $\mathbf{n}_k$  is computed, we can find an image that matches this field by including the term  $-\langle \mathbf{n}_k, \nabla u \rangle$  in the standard CS recovery model (1.5). This term tries to maximize the alignment of the estimated normals of the signal  $\mathbf{n}_k$  with the normals of the reconstruction  $\frac{\nabla u}{|\nabla u|}$ , and acts as an additional prior to  $J(u)$  when the unknown signal is an image. The resulting recovery model is

$$u_{k+1} = \arg \min_u J(u) - \gamma \langle \mathbf{n}_k, \nabla u \rangle + \frac{\alpha}{2} \|Au - f\|_2^2. \quad (1.10)$$

Taking into account that the divergence is the adjoint operator of the gradient, the previous minimization can be rewritten as

$$u_{k+1} = \arg \min_u J(u) + \gamma \langle \operatorname{div} \mathbf{n}_k, u \rangle + \frac{\alpha}{2} \|Au - f\|_2^2. \quad (1.11)$$

The proposed method exploits the geometry of the image in the recovery process and obtains better regularization properties than the standard TV recovery (1.5). In particular, the proposed geometric model preserves edges like TV by encouraging the gradients to be sparse with the initial prior  $J(u)$ , and it recovers smooth regions by aligning the gradients of the reconstructed image with the estimated normals with the term  $\langle \mathbf{n}_k, \nabla u \rangle$ .

The reformulation of (1.10) into model (1.11) is necessary to overcome the following limitation of (1.10) in smoothly varying regions. In these regions, the minimization term  $-\langle \mathbf{n}_k, \nabla u \rangle$  is negligible because the reconstructed gradient  $\nabla u$  is small, and the data and regularity terms dominate the minimization. In the model (1.11), however, this term is rewritten as  $\langle \operatorname{div} \mathbf{n}_k, u \rangle$ , it is not affected by the value of  $\nabla u$ , and the smoothness and orientation of the region are only summarized by the divergence estimated from the previous step. The divergence of the normals accumulates the orientation of the normal field for the different neighbours around each pixel; it has constant non-zero values in smoothly varying regions, large magnitude close to the edges of the image, and is close to zero in flat areas contaminated with noise. The magnitude of the divergence, that is, the weight given to the alignment term, is proportional to the coherence in local orientation of the estimated normals, and the recovery model is more robust to noise.

In principle we could also use a smooth estimate of the gradients  $\mathbf{v} = \nabla u_{k-1}$ , instead of the normals, to align the gradients of the reconstructed signal. However, discontinuities of the image would have a contribution to  $\mathbf{v}$  proportional to their jump, and the resulting term  $\langle \mathbf{v}, \nabla u \rangle$  would give different weights to discontinuities of different sizes. From a geometric perspective, if we want to recover the shapes of the image independently of their contrast, we need to consider the normal vectors derived from its level sets. By the use of level sets, we treat all shapes equally and the term  $\gamma \langle \mathbf{n}_{k+1}, \nabla u \rangle$  only accounts for geometric quantities. Indeed,



with the co-area formula we can rewrite the alignment term as

$$\int_{\Omega} \mathbf{n}_k \cdot \nabla u \, d\mathbf{x} = \int_{\Omega} \mathbf{n}_k \cdot \frac{\nabla u}{|\nabla u|} |\nabla u| \, d\mathbf{x} = \int_{-\infty}^{\infty} \int_{u^{-1}(c)} |\mathbf{n}_k| \cos \theta(s) \, ds \, dc, \quad (1.12)$$

where  $\theta(s)$  is the angle between the normal vector to the level curve  $u^{-1}(c) = \{\mathbf{x} \mid u(\mathbf{x}) = c\}$  and the regularized normal vector  $\mathbf{n}_k$ . The alignment term it thus is purely sensitive to angles, and its contribution is weighted by  $|\mathbf{n}_k|$ . In flat regions, where  $|\mathbf{n}_k| = 0$ , the alignment term vanishes and the proposed reconstruction method simplifies to TV recovery. In smoothly varying regions and close to the edges,  $|\mathbf{n}_k|$  is an experimental measure of confidence of the estimated  $\mathbf{n}_k$ , and the alignment term of Equation (1.12) is less sensitive to noise because it is weighted proportionally to this confidence.

### 1.3 Related approaches in compressed sensing

The method proposed in this chapter is inspired by image denoising and inpainting methods that align an estimate of the normals with the reconstructed image. In the context of image denoising, Lysaker et al. [2004] first regularizes the unit gradient of the noisy image and then improves reconstruction by fitting this gradient into the regularized vector. The resulting method outperforms the ROF model of Rudin et al. [1992] and similar higher-order PDE methods of Lysaker et al. [2003]. Dong et al. [2009] improves this model by regularizing the angles instead of the vectors and introducing an edge indicator as an extra weight. In image inpainting, an equivalent two-step method is proposed by Ballester et al. [2001], later improved with the divergence-free constraint by Tai et al. [2009, 2007] and adapted to image decomposition and denoising by Hahn et al. [2010, 2011]. This alignment between normal estimates and image gradients, however, has not been exploited before for CS recovery of images.

In the field of CS, several methods have been proposed to improve the quality of the  $\ell_1$  recovery model (1.5). For general signals, greedy algorithms and  $\ell_p$  minimizations with  $0 < p < 1$  approximate the solution of the  $\ell_0$  model (1.1) and improve the sparsity of  $\ell_1$  penalty in (1.5), but the resulting minimization problems are not convex, suffer from local minima, and have slow numerical algorithms (see e.g., Rao and Kreutz-Delgado [1999], Chartrand [2007], Tropp and Gilbert [2007], Needell and Tropp [2009], Blumensath and Davies [2009]). To improve the sparsity of  $\ell_1$  recovery without increasing its complexity, Candès et al. [2008] proposes an iterative process solving a weighted  $\ell_1$  problem at each iteration. The weights are defined inversely proportional to the value of the recovered signal in the previous iterate, approximating the behaviour of the  $\ell_0$  penalty and promoting the recovery of sparser signals than  $\ell_1$ . The resulting method efficiently solves a convex problem at each iteration and experimentally improves signal recovery; for this reason, it has been adapted to image processing with TV regularization by the edge-guided CS of Guo and Yin [2010]. Edge-guided CS incorporates information about the magnitude of the gradient in the recovery process, and it is therefore related to the geometric model proposed in this chapter. Our geometric recovery

model, however, is additive and more robust to noise, and it exploits both the magnitude and directional information of the gradients.

CS recovery of images has also been improved modifying the data term  $\|Au - f\|_2^2$  inspired by image denoising techniques. In particular, the Bregman iterations proposed by Osher et al. [2005] for image denoising and deblurring have been applied to CS by He et al. [2006]. In this case, Bregman iterations improve CS recovery for phantom MRI data, but fail in the recovery of real images because reconstructing a signal from partial measurements is a more ill-posed problem than denoising. For the particular case of TV regularisation, the first Bregman iteration has a geometric interpretation similar to the second step of the proposed recovery method. However, Bregman iterations do not include a regularization step for the normals and therefore fail for noisy and real MRI signals.

The following section summarizes each of these methods and clarifies their relationship with our geometric technique.

### 1.3.1 Edge-guided compressed sensing

Edge-guided CS of Guo and Yin [2010] improves recovery of MRI images by exploiting edge information with an iterative process that weights the TV regularizer with an edge detector associated with the image recovered in the previous iteration. The key idea is that edges correspond to locations where  $|\nabla u|$  is large, and TV corresponds to the  $\ell_1$  norm of the gradient; therefore an inverse edge detector can be used to re-weight TV and approximate the  $\ell_0$  penalty in a similar fashion to the re-weighted  $\ell_1$  of Candès et al. [2008] for general signals.

The method starts with the standard CS solution of model (1.5) to obtain a first estimate of the image  $u_1$ , it then defines the weights  $w_1 = g(|\nabla u_1|)$  as inversely proportional to  $|\nabla u_1|$ , and recovers an image with sparser edges by solving the re-weighted TV problem. The process is iterated, leading to the following two-step algorithm:

$$\begin{cases} u_{k+1} = \arg \min_u J_{w_k}(u) + \frac{\alpha}{2} \|Au - f\|_2^2 \\ w_{k+1} = g(|\nabla u_{k+1}|) \end{cases} \quad (1.13)$$

There is no stopping criterion or guarantee of convergence for this iterative process, and after a few iterations the reconstruction usually does not improve or even degrades. Moreover, the multiplicative model of edge-guided CS is very sensitive to false edge detection. Indeed, if an edge is detected in a wrong location, the weight associated with it on the next iteration will encourage the creation of an edge in this location, and CS recovery will degrade with any new iteration.

The iterative re-weighting process is designed to improve sparsity of the signal and the recovery of piecewise constant functions, but it fails in the recovery of smooth image regions. Compared with the proposed geometric model, edge-guided CS incorporates only information

about the magnitude of  $\nabla u$ , while our method also uses its directional information; edge-guided CS simply re-estimates the edges at each iteration, while our geometric model includes a regularization step of the estimated normals.

### 1.3.2 Bregman methods

The proposed geometric model also share similarities with Bregman methods, whose original idea was to restore normals and image intensities simultaneously. However, Bregman methods cannot recover normals as accurately and robustly as the proposed method because they do not regularize the estimated normals. This improvement comes at the price of loosing global convexity.

Bregman iterations substitute the minimization problem (1.5) for a sequence of convex optimizations substituting  $J(u)$  for its Bregman distance to the previous iterate. In particular, the first Bregman iteration has a geometric interpretation closely related to the proposed method. Starting with  $u = 0, v = 0$ , the Bregman iterative process can be summarized as

$$\begin{cases} u_{k+1} = \underset{u}{\operatorname{argmin}} J(u) + \frac{\alpha}{2} \|f + v_k - Au\|_2^2 \\ v_{k+1} = v_k + f - Au_{k+1}. \end{cases} \quad (1.14)$$

The first Bregman iteration corresponds to the standard TV model (1.5), while the second iteration implicitly exploits the normals of the image recovered in the first iteration to improve CS recovery.

For simplicity, the connection to proposed method is explained with a continuous formulation, where  $A(\cdot)$  is the continuous functional operator of CS and  $A^*$  its adjoint. For the first iteration  $u = 0, v = 0$  and the method solves

$$u_1 = \underset{u \in \mathbb{R}^n}{\operatorname{argmin}} \int_{\Omega} |\nabla u| + \frac{\alpha}{2} \|f - A(u)\|_2^2, \quad (1.15)$$

The optimality condition associated with problem (1.15) is derived from its the Euler-Lagrange equation, and reads

$$\operatorname{div} \frac{\nabla u_1}{|\nabla u_1|} = -\alpha A^*(u_1) (f - A(u_1)), \quad (1.16)$$

where  $\mathbf{n}_1 = \frac{\nabla u_1}{|\nabla u_1|}$  corresponds to the normals of  $u_1$ . At the next iteration we can introduce a term  $\langle \mathbf{n}_1, \nabla u \rangle$  aligning the normals of the reconstructed signal with the estimate of the normals from the previous iteration, that is

$$u_2 = \underset{u}{\operatorname{argmin}} \int_{\Omega} |\nabla u| - \langle \mathbf{n}_1, \nabla u \rangle + \frac{\alpha}{2} \|f - A(u)\|_2^2. \quad (1.17)$$

Integrating by parts and substituting  $\text{div } \mathbf{n}_1$  in Equation (1.16) we have

$$\begin{aligned} - \langle \mathbf{n}_1, \nabla u \rangle &= \langle \text{div } \mathbf{n}_1, u \rangle = - \langle \alpha A^*(f - A(u_1)), u \rangle = \\ &= - \alpha \langle f - A(u_1), A(u) \rangle = - \alpha \langle v_1, A(u) \rangle, \end{aligned} \quad (1.18)$$

with  $v_1 = f - A(u_1)$  as defined in (1.14). If we substitute (1.18) in the minimization (1.17) and group together the terms with  $A(u)$ , we end up with the Bregman update rule

$$u_2 = \int_{\Omega} |\nabla u| + \frac{\alpha}{2} \|f + v_1 - A(u)\|_2^2. \quad (1.19)$$

For the rest of iterations the geometric interpretation of the update is lost.

Compared with Bregman iterations, the proposed geometric method explicitly uses the normals in the recovery model for all iterations, not only the second one, and it is not restricted to TV regularization. In fact, Bregman iterations can also be used with other convex functionals  $J(u)$ , but the geometric interpretation of (1.17) is only possible for TV regularization. Moreover, the proposed method is more robust to noise thanks to the regularization of the normals and, unlike the Bregman iteration, experimentally improves CS reconstruction for both phantom and real MRI data. In addition, the proposed recovery method extends to non-local operators to exploit graph geometry and recover textured images.

### 1.4 Extension to non-local methods

Total variation regularization is designed to recover images with sharp edges and, as other methods based on local gradients, it is not suited for textured images with fine structures. In this section we extend the geometric model of Section 1.2 to textured images using both non-local TV regularization and a term aligning the estimated non-local normals with the non-local gradients of the reconstructed image.

#### 1.4.1 Non-local operators

Non-local TV is a variational extension of the non-local means filter proposed for image denoising in Buades et al. [2005]. Non-local means exploits the repetition of patterns in natural and textured images to reconstruct sharp edges as well as the fine structures that appear repeatedly in an image. This principle is the basis of the different non-local regularization methods used in imaging, which outperform classic regularization techniques by incorporating global information in the regularization process.

In a variational framework, Gilboa and Osher [2008] uses graph theory to extended the classical TV to a non-local functional. In the discrete setting, Zhou and Schölkopf [2005] and Elmoataz et al. [2008] use graph Laplacians to define similar non-local regularization operators. The resulting non-local methods have been used to solve several inverse problems in imaging,

showing in every case an improvement over their equivalent local methods based on classic differential operators. In CS reconstruction, non-local TV has been recently proposed by Peyré et al. [2008] and Zhang et al. [2010].

In order to generalize the geometric concept of normals, it is necessary to adopt the discrete formulation of the continuous model presented in Gilboa and Osher [2008]. This non-local framework considers the image domain as a graph  $G = (\Omega, E)$ , where  $\Omega$  is the set of nodes of the graph, one for each pixel in the image, and  $E$  is the set of edges connecting the nodes. The edge connecting nodes  $i$  and  $j$  is weighted with a positive symmetric function  $w(i, j)$  that is inversely proportional to the squared distance between the two nodes in graph terms. Consequently, two pixels  $i$  and  $j$  spatially far away in the image can be considered neighbours in the graph and interact if  $w(i, j) > 0$  (we write then  $i \sim j$ ). For that reason, the resulting approach is considered non-local.

Given an image  $u$  defined on the graph, the non-local gradient  $\nabla_G u$  at node  $i$  is defined as the vector of all directional derivatives associated with the neighbours of  $i$ , that is

$$\nabla_G u(i, j) = (u(j) - u(i))\sqrt{w(i, j)} \quad \forall j \in \Omega. \quad (1.20)$$

In the graph, vectors  $\mathbf{d} = d(i, j)$  are therefore functions defined in the domain  $\Omega \times \Omega$ . In this setting we can define the standard  $L_2$  inner product between functions as

$$\langle u, v \rangle_G = \sum_{i \in \Omega} u(i) v(i). \quad (1.21)$$

For vectors, the dot product is defined pixel-wise as

$$(\mathbf{d} \cdot \mathbf{e})_G(i) = \sum_{j \sim i} d(i, j) e(i, j), \quad (1.22)$$

and the inner product on the graph is given by

$$\langle \mathbf{d}, \mathbf{e} \rangle_G = \sum_i (\mathbf{d} \cdot \mathbf{e})_G(i) = \sum_i \sum_{j \sim i} d(i, j) e(i, j). \quad (1.23)$$

In order to have an equivalent to the TV regularizer, it is necessary to define a norm function on the graph  $|\cdot|_G$ . With the previous definitions, the magnitude of a vector at node  $i$  is given by

$$|\mathbf{d}|_G(i) = \sqrt{(\mathbf{d} \cdot \mathbf{d})_G(i)} = \sqrt{\sum_{j \sim i} d(i, j)^2}. \quad (1.24)$$

The standard TV is then naturally extended to a non-local version as the  $\ell_1$  norm of the graph

norm  $|\cdot|_G$  associated with the non-local gradient, that is,

$$TV_G(u) = J_G(u) = \sum_i |\nabla_G u|_G(i). \quad (1.25)$$

With the above inner products, the non-local divergence of a vector  $\mathbf{d}$  is defined as the adjoint of the non-local gradient, that is

$$\operatorname{div}_G \mathbf{d}(i) = \sum_{j \sim i} (d(i, j) - d(j, i)) \sqrt{w(i, j)}. \quad (1.26)$$

With these definitions, if we consider only the immediate pixels as neighbours and fix their weights to  $w(i, j) = 1$ , the non-local TV reduces to the standard TV definition. If we define the weighting function proposed by the non-local means filter of Buades et al. [2005], the non-local operators incorporate global information, and the standard regularization process is greatly improved. The weighting function, therefore, has an important impact in the performance of the non-local regularizers and should be chosen carefully.

For this reason, we adopt the successful weighting of non-local means and, given a reference image  $u_0$ , we compute the weight  $w_0(i, j)$  measuring the difference between the patches around each pixel as

$$w_0(i, j) = \exp^{-\frac{\|P_0(i) - P_0(j)\|^2}{2h^2}}, \quad (1.27)$$

where  $h$  is a scaling factor, and  $P_0(i)$  is a patch of  $u_0$  centred at pixel  $i$ . The effect of this weighting function in the definition of non-local neighbours is illustrated in Figure 1.1.




Figure 1.1: Definition of non-local interactions in a microscopic image in terms of a weighted graph. The nodes in the resulting graph connect only pixels that belong to the same structures, that is, either background, one cell (red), groups of cells (green).

This weighting function is designed to reduce Gaussian noise while preserving the textures of the image. To that purpose, the reference image should be as close as possible to the true image in order to incorporate valid information in the non-local operators. Accordingly, the weighting function is initialized with the solution  $u_0$  to the standard CS reconstruction of

model (1.5), a better estimate of the solution is then obtained with the non-local recovery method, and the weights are updated with this new solution. Consequently, the basic non-local CS recovery is given by the following two-step iterative method

$$\begin{cases} \nabla_{G_k} \leftarrow \text{estimate non-local operators from } u_{k-1} \\ u_k = \argmin_u J_{G_k}(u) + \frac{\alpha}{2} \|Au - f\|_2^2. \end{cases} \quad (1.28)$$

#### 1.4.2 Compressed sensing with non-local normals

Symmetrizing the proposed local technique, we can develop a non-local method that exploits the geometry of images defined in non-local graphs to improve the CS recovery method (1.28). The first step of each iteration estimates the non-local normals  $\mathbf{n}_G$  associated with the level set curves of the image in the graph. Once the non-local normals are estimated, the second step finds an image that fits the non-local normals, the CS measurements, and the non-local TV regularizer. As in the local case, these two steps are iterated alternatively.

In the local setting, the normal vectors of the level curves of an image  $u$  are defined as  $\mathbf{n} = \frac{\nabla u}{|\nabla u|}$ . We can extend this definition to the non-local framework by normalizing the components of the non-local gradient  $\nabla_G u$  node-wise, i.e., all the components associated with node  $i$  are normalized by  $|\nabla_G u|_G(i)$ .

Given an estimate of the non-local normals  $\mathbf{n}_G$ , we can then include a term in the CS reconstruction (1.28) that maximizes the alignment of the reconstructed signal with the non-local normals. The resulting minimization problem is

$$u = \argmin_u J_G(u) - \gamma \langle \mathbf{n}_G, \nabla_G u \rangle_G + \frac{\alpha}{2} \|Au - f\|_2^2. \quad (1.29)$$

Exploiting the adjoint relation of the non-local divergence and gradient, we can re-write the previous expression as

$$u = \argmin_u J_G(u) + \gamma \langle \text{div}_G \mathbf{n}_G, u \rangle_G + \frac{\alpha}{2} \|Au - f\|_2^2. \quad (1.30)$$

As before, the process is iterated and we obtain the following non-local CS recovery method

$$\begin{cases} \nabla_{G_k} \leftarrow \text{estimate non-local operators from } u_{k-1} \\ \text{div}_{G_k} \mathbf{n}_{G_k} = \argmin_v J_{G_k}(v) + \frac{\mu}{2} \|v - \hat{v}\|^2 \quad \text{with } \hat{v} = (1 - g(|\nabla_G u_{k-1}|_G)) \text{div}_G \frac{\nabla_G u_{k-1}}{|\nabla_G u_{k-1}|} \\ u_k = \argmin_u J_{G_k}(u) + \gamma \langle \text{div}_{G_k} \mathbf{n}_{G_k}, u \rangle_G + \frac{\alpha}{2} \|Au - f\|_2^2 \end{cases}$$

The third step of the proposed non-local method is naturally derived from its local version, but the second step, the regularization of non-local normals, requires a more detailed analysis.

### 1.4.3 Estimation of non-local normals

The non-local gradient operator, and consequently the non-local normals, do not correspond to the discretization of standard vector fields in a grid. Indeed,  $\nabla_G u$  has a different number of components for each pixel, and the direction associated with  $\nabla_G u(i, j)$  depends on the relative position of the node  $i$  and its neighbour  $j$ . Therefore, we cannot use standard techniques to regularize these vector fields, and we must regularize the term  $\text{div}_G \mathbf{n}_G$  posteriorly used in the recovery algorithm. Compared with the regularization of the non-local normals, regularizing their divergence loses directional information, but the resulting method is simpler and faster.

Given an estimate of the reconstructed signal  $u_{k-1}$ , we can first compute a noisy estimate of the non-local normals and their divergence pixel-wise, and then denoise it with standard regularization methods. To that purpose, the non-local normals are estimated as

$$\hat{\mathbf{n}}_G = \frac{\nabla_G u_{k-1}}{|\nabla_G u_{k-1}|}, \quad (1.31)$$

and a rough estimate of the non-local divergence is computed as

$$\hat{v} = (1 - g(|\nabla_G u_{k-1}|_G)) \text{div}_G \hat{\mathbf{n}}_G, \quad (1.32)$$

where  $g(x)$  is a function designed to verify  $g \approx 0$  when  $x$  is large and  $g \approx 1$  when  $x$  is small. In fact,  $g(|\nabla_G u_{k-1}|_G)$  acts as the equivalent edge detector presented in Section 1.2.1 and is defined with the same Equation (1.9). As in the local case, we can adopt the statistical interpretation of the edge detector  $g(|\nabla_G u|_G)$  presented in Black et al. [1998], where the edges are considered as outliers in the normal distribution of  $|\nabla_G u|_G$  associated with homogeneous regions. Since the edge detector  $g$  is derived from error norms robust to outliers, weighting the estimate of the divergence with the function  $1 - g(|\nabla_G u_{k-1}|_G)$  in Equation (1.32) is equivalent to soft-thresholding the non-local divergence when we suspect that the variations in  $u_{k-1}$  are caused by noise inside homogeneous regions.

Finally, we can regularize  $\hat{v}$  to obtain a smoother estimate of the non-local divergence, which will be used in the last step of the proposed iterative method. There are two natural approaches for this regularization: we can either ignore the non-local nature of the divergence and gradient operators and use any local model to regularize  $\hat{v}$ , or use the non-local neighbours to denoise  $\hat{v}$  with Equation (1.34), that is, use the natural distance and neighbouring relations inherent to the definition of  $\hat{v}$  to denoise it.

$$\text{div}_{G_k} \mathbf{n}_{G_k} = \arg \min_v J(v) + \frac{\mu}{2} \|v - \hat{v}\|^2 \quad (1.33)$$

$$\text{div}_{G_k} \mathbf{n}_{G_k} = \arg \min_v J_G(v) + \frac{\mu}{2} \|v - \hat{v}\|^2 \quad (1.34)$$

In particular, we choose to regularize the non-local divergence with the ROF model of Equation (1.33) based on experimental results.



## 1.5 Numerical minimization

As explained in the introduction, the subject of research of this chapter is the inclusion of geometric terms in the CS recovery of images, not the minimization techniques. For this reason, the details of the minimizations are presented in Appendices A and B; here only the relation between the proposed algorithms and the state-of-the-art is briefly explained.

The minimizations associated with each of the local steps involve both a TV and a quadratic term similar to the ROF model of Rudin et al. [1992]. Consequently, the resulting algorithms apply a similar strategy to overcome the non linearity and non differentiability of TV than the multitude of algorithms proposed for the ROF model. In the original ROF model, Rudin et al. propose a time marching method to solve the Euler-Lagrange PDE of the model, but the resulting algorithm is slow due to the constraint on the time step by the CFL condition. In the last years more efficient methods have been proposed for the ROF model due to its extensive use in imaging. A popular class of methods is based on the dual formulations of the ROF model, e.g., the projection method of Chambolle [2004] or the primal-dual approaches of Chan et al. [1999], Zhu et al. [2010], Chambolle and Pock [2010]. Other options are based on variable-splitting and equality constrained optimization; which is solved by quadratic penalties (Wang et al. [2008]), Bregman iterations (Goldstein and Osher [2009], Yin et al. [2008]) or the equivalent augmented Lagrangian method (Wu and Tai [2010]). In the case of CS, dual solvers are not usually adopted because they suffer from matrices  $A$  that are large-scale and dense. In particular, for matrices corresponding to transforms with fast implementations, like the Fourier transform of this chapter, splitting methods are a good choice because they can easily exploit fast transforms to compute  $Au$  and  $A^T u$ . The algorithms proposed here fall in this last category. To that purpose, the different problems are re-written as constrained minimizations, and augmented Lagrangians are used to solve them. The resulting Lagrangians are minimized with respect to each variable independently, and the multipliers are then updated in a cyclic way. Since all the minimizations can be analytically solved, the resulting algorithms are extremely fast and easy to implement.

Similarly, the minimization algorithm for the non-local problem (1.30) is closely related to the minimization of the non-local ROF model proposed in Gilboa and Osher [2008]. In particular, the non-local CS problem has been solved with a combination of forward-backward splitting and Bregman iteration in Zhang et al. [2010], but for uniformity this chapter uses the same combination of splitting and augmented Lagrangians than the rest of the dissertation.

## 1.6 Experimental results

This section presents some of the numerical results obtained with our geometric models and compares them with other techniques. In particular, the local version of our method is compared with the standard CS recovery algorithm of model (1.5), with the edge-guided CS of Guo and Yin [2010], and with the Bregman method of (1.14). The non-local version of our

method is compared to the non-local CS recovery of Equation (1.28), which does not take into account the geometric information of the non-local gradients in the recovery process.

The experiments use partial Fourier measurements and perform radial sampling on  $R$ ; the size of the measurements  $m$  in relation to the size of the signal  $n$  is specified by the ratio or percentage  $\frac{m}{n}$ . For a fair comparison, the same robust edge detector of Equation (1.9) is used in both the edge-guided CS and the proposed method, and all the minimization problems are solved with the same splitting and augmented Lagrangian techniques. The parameter  $\alpha$ , which is related to the noise present in the CS measurements, has been manually tuned to obtain best reconstruction with the standard CS recovery models (1.5) and (1.28), and posteriorly used with edge-guided CS and the proposed normal methods. This parameter was reduced by an order of magnitude for the Bregman iterations, as suggested in Osher et al. [2005]. The other parameters of our model,  $\gamma$  and  $\mu$ , have also been chosen manually to obtain good CS recovery in terms of SNR. The parameter  $\gamma$ , which controls the weight given to the alignment of the normals, takes similar values for the same kind of images (textured or brain MRI images) and remains stable for different sparsity and noise levels. The parameter  $\mu$ , which controls the smoothness of the estimated normals, decreases when the number of measurements decrease or the noise level increases because the point-wise estimated normals are less accurate and require more regularization.

The first set of experiments tests our method with MRI images, first with the Shepp-Logan phantom and then with a real MRI brain image. Table 1.1 shows the quantitative results of the different CS reconstruction methods for MRI images. The proposed geometric model always outperforms the standard TV reconstruction and the edge-guided CS technique. In the experiments, both the edge-guided CS and our method are initialized with the TV solution and, therefore, always improve TV reconstruction. In comparison with TV reconstruction, however, our method more than doubles the gain of edge-guided CS, and outperforms Bregman iterations for real images or a reduced number of measurements (Bregman is only able to exploit normal information when the TV reconstruction is accurate and the normal vectors directly estimated from it do not require regularization). In general, performance improves with non-local regularization, and our non-local geometric method outperforms the non-local TV reconstruction in all the experiments. As expected, the gain in relation to non-local TV is lower than in the local approach because the non-local geometric model loses part of the directional information of the normals by denoising their divergence instead of the vector fields.

The Figures 1.2-1.3 show qualitatively the improvement of our method over TV reconstruction and edge-guided CS. In the case of the phantom, our geometric models are able to better reconstruct the phantom with fewer measurements both in the local and non-local case, while with a real MRI image our reconstruction is able to capture better non dominant edges of the white-grey matter interface. Figure 1.4 explicitly compares the normals associated with the TV solution and the regularized normals of our local reconstruction for the real brain MRI image. We observe that our local method is able to better reconstruct the normals, and therefore the

Table 1.1: Comparison of CS reconstruction for MRI images. first four columns show the results with the standard TV regularization: TV stands for the model of Equation (1.5), edge CS for (1.13), Bregman iterations (1.14), and normal CS for the proposed geometric method. The last two columns correspond to the definition of NL TV: NL-TV corresponds to standard non-local CS recovery (1.28) and NL normal CS for the proposed non-local method.

Image	$\frac{m}{n}$	local CS				non-local CS	
		TV	edge CS	Bregman	normal CS	TV	normal CS
Phantom	8%	7.33 dB	7.37 dB	7.50 dB	<b>12.78 dB</b>	28.28 dB	<b>33.13 dB</b>
Phantom	12%	38.60 dB	45.33 dB	<b>89.12 dB</b>	56.14 dB	61.84 dB	<b>74.57 dB</b>
Brain	12%	17.14 dB	17.38 dB	17.16 dB	<b>17.71 dB</b>	18.96 dB	<b>20.39 dB</b>
Brain	20%	22.16 dB	22.35 dB	22.17 dB	<b>23.82 dB</b>	23.13 dB	<b>24.12 dB</b>

shapes, of the image.

For each image in the experiments, we can also add different levels of Gaussian noise ( $\sigma_n$ ) to the signal in order to investigate the robustness of the proposed model to noise. Results are shown in table 1.2. We observe that our methods are more robust to noise than edge-guided CS and Bregman iterations thanks to regularization step on the estimation of the normals. As before, non-local regularization improves CS reconstruction, and we observe that the proposed non-local method outperforms again the non-local TV, and is robust to noise.

Table 1.2: Comparison of CS reconstruction for noisy MRI images with 12% of samples and different levels  $\sigma_n$  of Gaussian noise. first four columns show the results with the standard TV regularization: TV stands for the model of Equation (1.5), edge CS for (1.13), Bregman iterations (1.14), and normal CS for the proposed geometric method. The last two columns correspond to the definition of NL TV: NL-TV corresponds the standard non-local CS recovery (1.28) and NL normal CS for the proposed non-local method.

Image	noise $\sigma_n$	local CS				non-local CS	
		TV	edge CS	Bregman	normal CS	TV	normal CS
Phantom	5%	11.90 dB	11.91 dB	12.63 dB	<b>12.90 dB</b>	17.92 dB	<b>18.36 dB</b>
	10%	8.37 dB	8.38 dB	9.41 dB	<b>9.44 dB</b>	12.15 dB	<b>13.03 dB</b>
	15%	6.59 dB	6.59 dB	7.14 dB	<b>7.28 dB</b>	10.09 dB	<b>10.27 dB</b>
Brain	5%	13.37 dB	13.36 dB	13.67 dB	<b>13.78 dB</b>	14.86 dB	<b>15.00 dB</b>
	10%	10.88 dB	10.88 dB	11.48 dB	<b>11.57 dB</b>	12.31 dB	<b>12.50 dB</b>
	15%	9.89 dB	9.89 dB	10.31 dB	<b>10.48 dB</b>	10.94 dB	<b>11.19 dB</b>

The next experiment investigates the properties of the local and non-local versions of our algorithm with two synthetic images. The first step of our iterative method is designed to preserve the discontinuities in the level set normals of the reconstructed image, while the second step introduces this geometric information in the CS reconstruction model. This property is specially interesting for the recovery of images with geometric structures like

ridges or valleys, as shown in Figure 1.5. We observe that the proposed local method is able to recover ridges and valleys to certain extent and outperforms TV reconstruction in this kind of structures, but it is not able to capture the repetitive nature of the image like its non-local version.

The third set of experiments investigates CS recovery with natural images containing textures, where edge detection by itself is a difficult task, and the images cannot be considered piecewise constant. With these images, local regularization loses all texture information, while the non-local approaches can recover repetitive patterns and better exploit the geometric information of the image. The results with the proposed geometric method are presented in Table 1.3, with some of the reconstructed images shown in Figures 1.6-1.9 to qualitatively analyse the performance of the different methods.

A quantitative comparison of the different methods with textured images is presented in Table 1.3. We observe that the inclusion of an edge detector in edge-guided CS does not

Table 1.3: Comparison of CS reconstruction for textured images. The first four columns show the results with the standard TV regularization: TV stands for the model of Equation (1.5), edge CS for (1.13), Bregman iterations (1.14), and normal CS for the proposed geometric method. The last two columns correspond to the definition of NL TV: NL-TV corresponds to the standard non-local CS recovery (1.28) and NL normal CS for the proposed non-local method.

$\frac{m}{n}$	image	local CS				non-local CS	
		TV	edge CS	Bregman	normal CS	TV	normal CS
12%	Lena	14.53 dB	14.47 dB	14.55 dB	<b>14.86 dB</b>	15.82 dB	<b>16.79 dB</b>
	Barbara	13.35 dB	13.31 dB	13.25 dB	<b>13.59 dB</b>	15.00 dB	<b>15.52 dB</b>
	fingerprint	4.13 dB	4.11 dB	4.09 dB	<b>4.13 dB</b>	5.97 dB	<b>5.98 dB</b>
	baboon	7.40 dB	7.25 dB	7.18 dB	<b>7.40 dB</b>	7.65 dB	<b>7.65 dB</b>
20%	Lena	18.44 dB	18.36 dB	18.45 dB	<b>19.27 dB</b>	19.95 dB	<b>21.09 dB</b>
	Barbara	16.71 dB	16.62 dB	16.55 dB	<b>17.13 dB</b>	18.37 dB	<b>18.93 dB</b>
	fingerprint	5.70 dB	5.62 dB	5.59 dB	<b>5.70 dB</b>	9.03 dB	<b>9.07 dB</b>
	baboon	9.13 dB	8.91 dB	8.90 dB	<b>9.14 dB</b>	9.63 dB	<b>9.74 dB</b>
39%	Lena	25.39 dB	25.30 dB	25.44 dB	<b>26.71 dB</b>	26.39 dB	<b>27.51 dB</b>
	Barbara	20.83 dB	20.68 dB	20.59 dB	<b>21.36 dB</b>	24.68 dB	<b>25.33 dB</b>
	fingerprint	12.02 dB	11.84 dB	11.57 dB	<b>12.03 dB</b>	14.52 dB	<b>14.56 dB</b>
	baboon	13.30 dB	13.14 dB	13.28 dB	<b>13.41 dB</b>	13.44 dB	<b>13.82 dB</b>


improve the TV reconstruction because the partially reconstructed images are not accurate enough to detect edges and, as a result, the weighted TV term of edge-guided CS encourages edges in wrong positions. This effect is not observed in our method because it is an additive and not a multiplicative model; moreover, it exploits the directional information of the regularized normals, which can partially capture texture information better than an edge detector. As a consequence, the proposed local method always outperforms the TV recon-

struction and edge-guided CS methods. Similarly, Bregman iterations are not able to exploit the geometric alignment of normals in the CS reconstruction because Bregman methods do not regularize the first estimate of the normal vectors obtained from the TV solution. For the non-local regularizations, our method also outperforms non-local TV, but the gain in some cases is negligible (fingerprint and baboon for  $\frac{m}{n} = 12\%$  or  $20\%$ ). In fact, the non-local methods require a good estimate of the reconstruction to initialize the non-local operators. Since the proposed method requires both gradient and divergence to estimate the non-local normals, it can only improve non-local TV when the initialization (in this case the standard TV solution) has a minimum level of accuracy. More measurements are required for fingerprint and baboon simply because these images are less sparse than Lena or Barbara in terms of TV and CS requires more samples for their reconstruction.


In the reconstruction of Lena in Figures 1.6a-1.6f we can qualitatively observe the advantages of our method in comparison to local and non-local TV reconstructions. In the local case, the use of geometry avoids the staircase effect that is visible in the TV reconstruction of Lena's cheek. In the non-local case, the geometric model captures the slowly varying textures of Lena's skin or hat. In both cases the improvements are due to the regularization of the level set normals of the image.




(a) Shepp-Logan phantom



(b) TV reconstruction, 7.33 dB



(c) edgeCS, 7.37 dB



(d) non-local TV reconstruction, 28.28 dB



(e) proposed local method, 12.78 dB



(f) proposed non-local method, 31.26 dB

Figure 1.2: Compressed sensing reconstruction of Shepp-Logan phantom from 8% of measurements in Fourier domain.




Figure 1.3: Compressed sensing reconstruction of MRI brain image from 12% of measurements in Fourier domain.




Figure 1.4: Zoom on reconstructed brain MRI image from 12% of measurements in Fourier domain. We superpose the reconstructed signals with the normals associated with their level sets for the standard TV solution (left) and for the local version of our method (right). The proposed geometric method is able to better reconstruct the normals and shapes of the image.

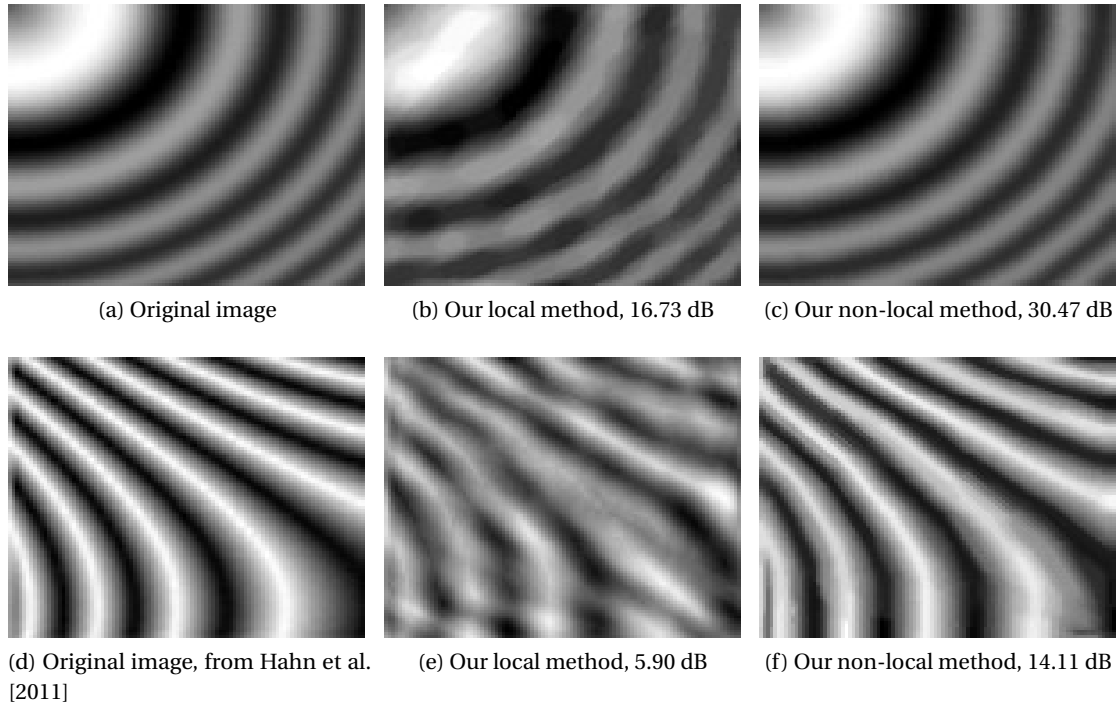


Figure 1.5: Reconstruction of two synthetic images characterized by ridges from 10% and 14% of measurements in Fourier domain. Local and non-local TV results in SNR of 16.30 and 30.15 dB for the top image, and 5.86 and 12.99 dB for the bottom one.





Figure 1.6: Compressed sensing reconstruction of Lena from 20% of measurements in Fourier domain.



(a) Lena



(b) Barbara



(c) proposed local method, 14.86 dB



(d) proposed local method, 13.59 dB



(e) proposed non-local method, 16.79 dB




(f) proposed non-local method, 15.52 dB

Figure 1.7: Left column: reconstruction of Lena phantom from 12% of measurements in Fourier domain. Right column: reconstruction of Barbara from 12% of measurements in Fourier domain.



(a) Barbara



(b) Baboon



(c) proposed local method, 17.13 dB



(d) proposed local method, 9.14 dB




(e) proposed non-local method, 18.92 dB




(f) proposed non-local method, 9.74 dB


Figure 1.8: Left column: reconstruction of Barbara from 20% of measurements in Fourier domain. Right column: reconstruction of baboon from 20% of measurements in Fourier domain.



(a) Fingerprint



(b) proposed local method, 5.70 dB



(c) proposed non-local method, 9.07 dB

Figure 1.9: Compressed sensing reconstruction of fingerprint from 20% of measurements in Fourier domain.

## 1.7 Conclusions

This chapter has proposed a recovery method for compressed sensing that iteratively aligns the normal vectors to the level curves of the reconstructed image. The alignment of these normal vectors has already been exploited in image denoising and inpainting methods, but existing compressed sensing techniques have just neglected it because this information is buried in the compressed sensing measurements and cannot be directly exploited.

To overcome this issue, the proposed technique alternates between the estimation of the normals to the level curves of the image, and a compressed sensing reconstruction that matches the estimated normals, the compressed sensing measurements, and the sparsity constraint of total variation. The proposed method is also extended to non-local operators in graphs to recover textured images, where it exploits the repetitive nature of patterns to recover the relevant structures of the image. The proposed non-local method can also be used to improve existing non-local denoising and deblurring methods.

Although there is no proof of convergence for the proposed two-step procedure, our experiments show a clear improvement over state-of-the-art algorithms. The introduction of geometry experimentally improves the reconstruction of images in two ways: the reconstruction is more robust to noise and a reduced number of measurements, and the recovery model is able to recover sharp edges as well as smoothly varying image regions, overcoming the staircase effect of total variation.



## 2 Harmonic Active Contours

This chapter presents a segmentation method based on the geometric representation of images as 2-dimensional manifolds embedded in a higher dimensional space. The segmentation is formulated as a variational problem, where the contours are described by a level set function, and the objective functional measures the surface of a manifold defined by the segmentation of the image. The minimizer of this functional, i.e., the proposed segmentation, is therefore a geometric model of a minimal surface manifold.

In this geometric framework, both the data and prior terms of the objective functional are coupled at pixel level and lead to a variational model that intrinsically aligns the gradients of the level set function with the gradients of the image and exploits this directional information to overcome image inhomogeneities and fragmented contours. The proposed formulation combines this robust alignment of gradients with attractive properties of methods previously developed in the same framework: the natural coupling of image channels developed in Sochen et al. [1998] for anisotropic diffusion, and the ability of the *subjective surfaces* of Sarti et al. [2002] to detect weak edges and close fragmented boundaries.

The potential of such a geometric approach lies in the general definition of Riemannian manifolds, which naturally generalizes the *Geodesic Active Contours* of Caselles et al. [1997], the *Active Contours without Edges* of Chan and Vese [2001], and the robust edge integrator of Kimmel and Bruckstein [2003] to higher-dimensional spaces, non-flat images and feature spaces. As a consequence, the proposed technique improves the segmentation of multi-channel images, images subject to inhomogeneities, and images characterized by geometric structures like ridges or valleys.

### 2.1 Introduction

Image segmentation is a first fundamental step in many applications of computer vision and machine learning because it simplifies the understanding of an image from thousands of

---

The material of this chapter has been published in Estellers et al. [2013c].

pixels to a few regions. The goal of image segmentation is thus to partition the image domain  $\Omega$  into homogeneous regions corresponding to individual objects or, equivalently, to find the contours  $\mathcal{C}$  that define the boundaries of these objects. To that purpose, the segmentation of an image is formulated as a minimization problem, where the objective functional specifies the segmentation criterion, and the unknown variables describe the contours of the different regions. The level set method of Osher and Sethian [1988], for instance, adopts an implicit parametrization of the contours as the zero-level set of a function  $\phi$ , which becomes then the minimization variable, and writes the objective functional in terms of this level set function.

In this context, two representative segmentation methods are the *Geodesic Active Contours* (GAC) of Caselles et al. [1997] and the *Active Contours without Edges* (ACWE) of Chan and Vese [2001]. The GAC model, or the equivalent model of Kichenassamy et al. [1995], proposes an edge-based segmentation criterion and defines an objective functional that weights the length of the contour with an inverse edge detector on the image. On the other hand, ACWE proposes a segmentation criterion based on the homogeneity of the different regions defined by the segmentation, and its objective functional measures the variance of the grey-level values of the image within each region. Both segmentation criteria can be combined and result in a segmentation method simple and yet extensively used in imaging, for which fast convex implementations have been proposed in Chan et al. [2006], Bresson et al. [2007], Goldstein et al. [2009]. The resulting segmentation methods are extremely fast and reliable for cartoon-like images, but fail in the segmentation of complex images with low contrast, inhomogeneities, or ridge structures. To tackle these cases, more complex segmentation criteria are necessary, and objective functionals based on second-order derivatives of the image along the contours have been proposed in Kimmel and Bruckstein [2003] and Vasilevskiy and Siddiqi [2002]. The robust edge integrator of Kimmel and Bruckstein [2003] is based on the key observation that the direction of the image gradient is a good estimator of the orientation of an edge contour, and the segmentation criterion can be improved by introducing a geometric term in the objective functional that aligns the normal vectors to the contours with the gradients of the image. With a level set parametrization, the objective functional of the robust edge integrator of Kimmel and Bruckstein is formulated as

$$- \int_{\mathcal{C}} |\nabla I \cdot \frac{\nabla \phi}{|\nabla \phi|}| \quad (2.1)$$

and tries to align the normal vectors to the contours  $\frac{\nabla \phi}{|\nabla \phi|}$  with the gradient of the image  $\nabla I$  along the contour  $\mathcal{C}$ . The inclusion of this alignment term in the objective functionals of GAC or ACWE leads to more accurate segmentations, but the resulting objective functionals are not convex and suffer therefore from local minima and slow numerical minimizations. Figure 2.1a illustrates the main idea behind this alignment term, and Figures 2.1b-2.1c show the improvement it brings to GAC and ACWE segmentation methods.

The method that we propose goes a step further and aligns the gradients of the image  $\nabla I$  with the gradients of the level set function  $\nabla \phi$  for all the level sets of the image, not only for the evolving active contour  $\mathcal{C}$ . This results in a functional that is able to also exploit the



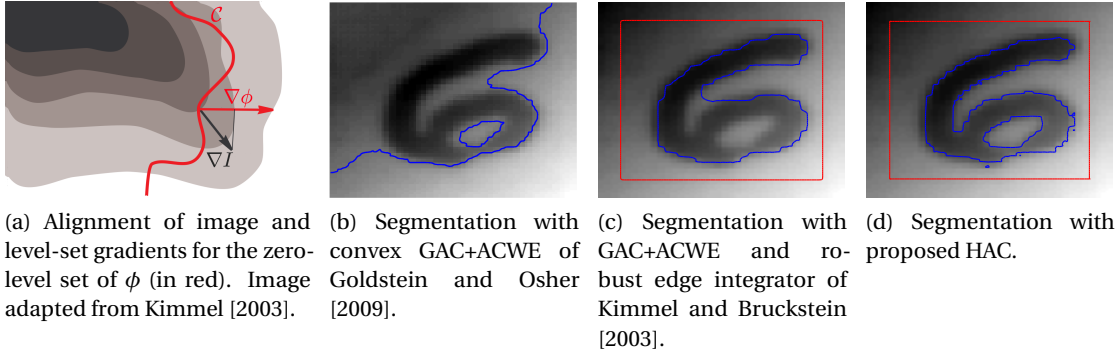


Figure 2.1: Figure 2.1a illustrates the mechanism of the proposed alignment term for the active contour  $\mathcal{C}$ . Figures 2.1b-2.1d show the segmentation of a gray-scale image subject to inhomogeneity with different methods, where the initial level set set appears in red and the final segmentation in blue.

alignment of the neighbouring level sets to pull the contours to the right position. Compared with the robust edge integrator of Kimmel and Bruckstein [2003], we are able to detect new contours from the alignment of the neighbouring level sets, which makes us more robust to local minima and less sensitive to initialization. A clear example is shown in Figure 2.1d.

Our method also answers the following question: how can we align gradients and normal vectors in images with  $k$  channels? Generalizing the alignment term proposed by Kimmel and Bruckstein to multi-channel images is not trivial because if we consider images as vector fields in the image domain  $\Omega$ , the coupling of the different channels must be defined heuristically. The natural way to treat multi-channel images is to interpret them as 2-dimensional manifolds or surfaces embedded in  $\mathbb{R}^{k+2}$  and make use of differential geometry to define equivalent alignment terms. To that purpose, we develop our method in the Beltrami framework, which was originally proposed for color denoising in Sochen et al. [1998] and later extended to image segmentation by the *subjective surfaces* of Sarti et al. [2002]. In that sense, our method introduces an alignment term in the *subjective surface* model, in the same way that Kimmel and Bruckstein introduced an alignment term in the GAC model.

The proposed harmonic active contours (HAC) also include region- and edge-based segmentation criteria. Our objective functional, however, is not a sum of terms with region- and edge-based functionals, but is defined through the embedding of different image features in the Beltrami framework. In this framework, the objective functional of GAC has been generalized to 2-dimensional manifolds in the *subjective surfaces* of Sarti et al. [2002] by embedding the grey-level values of the image and by defining the objective functional to be the weighted surface of the embedded manifold. Our method presents a generalization of GAC, ACWE, and the robust edge integrator of Kimmel and Bruckstein to higher dimensional spaces by embedding the image channels, the region features, and the level set function and by defining the objective functional to be the surface of the embedded manifold. This functional offers

three main advantages. First, the alignment term between the normal vectors to the contours and the gradients of the image is naturally extended to all the level sets of  $\phi$ , and the resulting method exploits the alignment of neighbouring level sets to overcome local minima. Second, the definition of images as manifolds can naturally include and handle multi-channel images, and it takes into account the coupling of the different channels. Third, the proposed formulation is easily generalized to non-flat parametric manifolds and feature spaces, where the usual segmentation methods cannot be directly applied, see e.g., Kimmel and Sochen [2002], Bogdanova et al. [2007].

The rest of this chapter is organized as follows. Section 2.2 reviews basic concepts of differential geometry and introduces the Beltrami framework. Section 2.3 defines the HAC model as a minimal surface manifold, specifying the proposed embedding, features, and metric. Section 2.4 sketches the minimization algorithm, which is detailed in Appendix C, and Section 2.5 presents the experimental results. Finally, conclusions are drawn in Section 2.6.

### 2.2 Differential geometry in image processing

In order to understand the Beltrami framework, it is necessary to formalize two basic concepts of differential geometry: smooth manifolds and the metric tensor.

We review here the basic definitions of differential geometry to provide both an intuitive and a formal understanding of the geometric ideas behind the equations of the Beltrami framework. We introduce this framework by particularizing all the definitions to images, which will allow us later to draw the connections between our method and previous literature. For a more detailed description on differential geometry, please refer to do Carmo [1976].

The simplest manifold that we can consider is a curve in the plane, which corresponds to a 1-dimensional manifold embedded in  $\mathbb{R}^2$ . Intuitively, a curve in the plane is a flexible cord that can be straightened to locally look like the real line  $\mathbb{R}$ . Formally, we describe it by means of a smooth parametrization between  $\mathbb{R}$  and the embedding space  $\mathbb{R}^2$ , ensuring that there are no critical self-intersections in the curve by constraining the parametrization to be locally invertible. In order to measure the length of the curve or the angle between its different directions, we consider its trajectory in  $\mathbb{R}^2$  and measure length and angles in terms of the usual coordinates and scalar product of  $\mathbb{R}^2$ . In this section, we generalize these concepts to  $n$ -dimensional manifolds embedded in  $\mathbb{R}^m$ ,  $m > n$ . The definitions, however, just formalize the previous intuitive explanation of a curve in  $\mathbb{R}^2$  (the manifold) and how we measure its length (the metric).

We say that  $\mathcal{M} \subset \mathbb{R}^m$  is an  **$n$ -dimensional smooth manifold** in  $\mathbb{R}^m$  if for every point  $p$  in  $\mathcal{M}$  there is a local chart  $(\sigma, U)$  satisfying

- $U \subset \mathbb{R}^n$  and  $V \subset \mathbb{R}^m$  are open sets with  $p \in V$
- $\sigma$  is a parametrization of  $\mathcal{M}$  around  $p$  such that

$$\begin{aligned}\sigma: U \subset \mathbb{R}^n &\longrightarrow V \subset \mathbb{R}^m \\ (u_1, \dots, u_n) &\mapsto \sigma(u_1, \dots, u_n).\end{aligned}\tag{2.2}$$

- $\sigma$  is a differential function with rank  $J\sigma = n$ , where  $J\sigma$  is the Jacobian of  $\sigma$ .
- $\sigma: U \longrightarrow V \cap \mathcal{M}$  is an homeomorphism, i.e.,  $\sigma^{-1}$  exists and is continuous.

The collection of local charts  $\{\sigma, U\}$  of  $\mathcal{M}$  is called an **atlas**.

The inherent idea of the definition is the same than for the 1-dimensional curve in  $\mathbb{R}^2$ . The first point is only a technical condition to ensure differentiability. The second condition formalizes the idea that an  $n$ -dimensional smooth manifold is something that locally looks like  $\mathbb{R}^n$  and can be parametrized with a set of differential functions, the atlas, that are locally invertible, i.e., the manifold has no critical points associated with self-intersections.

For instance, the plane  $\mathbb{R}^2$  is a 2-dimensional manifold that can be parametrized with local coordinates  $(x, y)$  and a single atlas  $\{\text{id}, \mathbb{R}^2\}$ . A more interesting case is the sphere  $S^2 \subset \mathbb{R}^3$ , a 2-dimensional manifold in  $\mathbb{R}^3$  which cannot be covered by one single atlas and has curvilinear coordinates. In this case we can construct an atlas with the stereographic projections from the north and south poles, which cover all the sphere except the projecting point. The stereographic projection from the north pole, for instance, reads  $\sigma(x, y) = \left( \frac{2x}{1+x^2+y^2}, \frac{2y}{1+x^2+y^2}, \frac{x^2+y^2-1}{1+x^2+y^2} \right)$ . This will be useful later.

In the Beltrami framework introduced in Sochen et al. [1998], images are considered as 2-dimensional Riemannian manifolds embedded in the so-called space-feature manifold. For a grey-scale image defined in a rectangle  $\Omega \subset \mathbb{R}^2$ , for instance, the manifold is parametrized with the following single chart:

$$\begin{aligned}\sigma: \Omega \subset \mathbb{R}^2 &\longrightarrow \Omega \times \mathbb{R} \subset \mathbb{R}^3 \\ (x, y) &\mapsto (x, y, f(x, y)),\end{aligned}\tag{2.3}$$

where  $f(x, y)$  is the grey level associated with point  $(x, y)$ . Assuming  $f$  to be differentiable, it is trivial to prove that  $\sigma$  verifies the properties of a local chart, and the image can be considered a 2-dimensional manifold  $\mathcal{M}$  embedded in a higher dimensional space  $\mathcal{M} \subset \Omega \times \mathbb{R}$ . The potential of this geometric framework lies in the general definition of the space-feature manifold and the choice of its metric. The features are not restricted to scalar values, but include vector features encountered in color, texture, or multi-spectral image analysis (Kimmel et al. [1997], Sagiv et al. [2006]). Similarly, the embedding is not limited to 2-dimensional images and generalizes naturally to  $n$ -dimensional manifolds associated with volumetric data or time varying images. Moreover, the choice of the metric enables the study of complex geometries inherent to scale-space methods and non-flat images generated by catadioptric or omnidirectional cameras (Bresson et al. [2006a], Bogdanova et al. [2007]).

To measure distances, angles, and areas in the manifold  $\mathcal{M}$  we require the concept of a metric, that is, a scalar product and a vectorial space. The definition of these concepts in a manifold is naturally inherited from  $\mathbb{R}^m$  by means of the parametrizations.

Given an  $n$ -dimensional smooth manifold  $\mathcal{M} \subset \mathbb{R}^m$  and  $p \in \mathcal{M}$ , we say that  $\mathbf{v} \in \mathbb{R}^m$  is a **tangent vector** to  $\mathcal{M}$  in  $p$  if there is  $\epsilon > 0$  and a parametric curve  $\gamma: (-\epsilon, \epsilon) \rightarrow \mathcal{M}$  such that  $\gamma(0) = p$  and  $\gamma'(0) = \mathbf{v}$ .

The collection of all tangent vectors to  $\mathcal{M}$  at point  $p$  is called the **tangent space**  $T_p\mathcal{M}$ .

It is easy to prove that  $T_p\mathcal{M}$  is a vectorial space of dimension  $n$ , with basis  $\{\frac{\partial \sigma}{\partial u_1}, \dots, \frac{\partial \sigma}{\partial u_n}\}_{|u=\sigma^{-1}(p)}$ .

As in the case of the curve in  $\mathbb{R}^2$ , distances and angles in the manifold  $\mathcal{M}$  are inherited from the standard scalar product in the embedding space  $\mathbb{R}^m$ . This operation, denoted  $\sigma^*$ , is known as **pull-back** of the metric of  $\mathbb{R}^m$  to the manifold.

To that purpose, for every point  $p \in \mathcal{M}$  we consider the vectorial space  $T_p\mathcal{M} \subset \mathbb{R}^m$  with the inner product  $\cdot$  induced by  $\mathbb{R}^m$ . The pull-back operation  $\sigma^*$  gives us a linear symmetric definite positive form  $\mathbb{I}_p$  for every  $p \in \mathcal{M}$ , that we call the **metric or first fundamental form** of the manifold  $\mathcal{M}$

$$\begin{aligned} \mathbb{I}_p: T_p\mathcal{M} \times T_p\mathcal{M} &\longrightarrow \mathbb{R} \\ (\mathbf{u}, \mathbf{v}) &\longmapsto \mathbf{u} \cdot \mathbf{v}. \end{aligned} \quad (2.4)$$

Given a local chart  $(\sigma, U)$  and the basis  $\{\frac{\partial \sigma}{\partial u_1}, \dots, \frac{\partial \sigma}{\partial u_n}\} = \{\sigma_1, \dots, \sigma_n\}$  of  $T_p\mathcal{M}$ , we can compute the expression of  $\mathbb{I}_p$  in that basis. The resulting  $n \times n$  matrix  $G^*$  has entries  $g_{ij}^* = \sigma_i \cdot \sigma_j$ .

As a bilinear form,  $\mathbb{I}_p$  does not depend on the local coordinates, but its matrix expression  $G^*$  depends on the basis used for  $T_p\mathcal{M}$  and, consequently, on the local chart.

In our previous examples, the plane can be described by one single atlas  $\{\text{id}, \mathbb{R}^2\}$  and we have  $T_p\mathbb{R}^2 = \mathbb{R}^2$ , with metric given by the identity matrix  $I_2$ . The sphere  $S^2$ , on the contrary, requires multiple atlases, and the induced metric with the stereographic projection reads  $G^* = \frac{4}{(1+x^2+y^2)^2} I_2$ .

In the case of images with the parametrization defined by Equation (2.3), the basis of the tangent space is given by  $\{\sigma_1 = (1, 0, f_x), \sigma_2 = (0, 1, f_y)\}$ , and the induced metric in this basis reads

$$G^* = \begin{pmatrix} 1 + f_x^2 & f_x f_y \\ f_x f_y & 1 + f_y^2 \end{pmatrix}. \quad (2.5)$$

In this image manifold, the distance between two points  $p_1 = (x, y)$  and  $p_2 = (x + dx, y + dy)$  in the image plane is measured by the length of vector  $\mathbf{v} = \overrightarrow{p_1 p_2}$  in the metric of the manifold. Therefore, the squared distance between these two points is given by  $\mathbf{v}^T G^* \mathbf{v} = (dx)^2 + (dy)^2 + (D_{\mathbf{v}} f)^2$ , where  $D_{\mathbf{v}} f = \nabla f \cdot \mathbf{v}$  is the directional derivative of  $f$  in the direction  $\mathbf{v}$ . As a consequence, two points in the manifold are close if they are physically close in the coordinate space and their grey level values are similar. In other words, the notion of distance between two points in the image refers not only to the spatial distance, but also to the information available in the feature space. As a consequence, a scaling factor  $\alpha$  is usually introduced in order to bring feature and space variables to the same scale.

To measure areas in the manifold, we recall first that two non-parallel vectors  $\mathbf{u}, \mathbf{v} \in \mathbb{R}^3$  define a parallelogram with an area given by  $\sqrt{|\mathbf{u} \wedge \mathbf{v}|} = \sqrt{\|\mathbf{u}\|^2 \|\mathbf{v}\|^2 - (\mathbf{u} \cdot \mathbf{v})^2} = \sqrt{\det(\mathbf{u}, \mathbf{v})^T (\mathbf{u}, \mathbf{v})}$ , where  $(\mathbf{u}, \mathbf{v})$  is the matrix of columns  $\mathbf{u}$  and  $\mathbf{v}$ .

For a 2-dimensional manifold in general, the area of a bounded region  $R \subset \mathcal{M}$  is measured by the “curvilinear parallelogram” formed by the tangent vectors associated to its coordinate parametrization of this region, that is,

$$\int_{\sigma^{-1}(R)} \sqrt{\|\sigma_1\|^2 \|\sigma_2\|^2 - (\sigma_1 \cdot \sigma_2)^2} = \int_{\sigma^{-1}(R)} \sqrt{\det G^*}.$$

For higher dimensional manifolds  $\int_{\Omega} \sqrt{\det G^*}$  measures the hyper-surface of the manifold in its metric, and  $g^* = \det G^*$  is defined as the squared hyper-surface element on the manifold.

In the case of grey-scale images in the plane, the surface of the manifold is given by  $\int_{\Omega} \sqrt{1 + \alpha^2 |\nabla f|^2}$  and reduces to a regularization term on the intensity values of the pixels. The role of the scale parameter  $\alpha$  allows us to consider different norms. If  $\alpha \rightarrow \infty$ , the 1 inside the square root becomes negligible and the surface of the manifold approaches a total variation regularizer; the minimization of this functional corresponds to an anisotropic smoothing where the diffusion coefficient is defined by the grey-level values of the image. On the other hand, if  $\alpha \rightarrow 0$ , the minimizing flow approaches the heat diffusion and results in an isotropic diffusion where the smoothing is completely determined by the spatial coordinates in the image plane. In other words, the surface of the manifold defined by an image couples in a single functional the smoothness of the image in both spatial and feature coordinates, and the weight given to each term is controlled by the metric parameter  $\alpha$ .

For images with  $k$  feature channels  $f^1, \dots, f^k$ , we make use of Einstein’s convention<sup>1</sup> to write the surface element as  $g^* = 1 + \alpha_i^2 |\nabla f^i|^2 + \alpha_i \alpha_j [\nabla f^i, \nabla f^j]^2$ , which takes into account the coupling of the different channels in the image. Indeed, the terms  $[\nabla f^i, \nabla f^j] = f_x^i f_y^j - f_y^i f_x^j$  correspond to the magnitude of the cross product of the vectors  $\nabla f^i$  and  $\nabla f^j$  and measure the coupling of feature channels  $f^i$  and  $f^j$  in terms of their gradients. As a consequence, Kimmel et al. [1997] observes that the determinant of the induced metric works as a generalized edge indicator. When the determinant of the metric has a value larger than unity, it indicates the presence of a strong gradient on the manifold. A value which is close to unity indicates a region where the manifold is almost flat. In the case of vectorial images, it does not only exploit  $\nabla f^i$  in terms of its norm, but also in terms of its orientation, i.e., it makes use of the directional information to measure the coupling of the image channels in the definition of edges. In fact, this is the main advantage of this geometric framework for image segmentation.

Finally, the notion of distances in  $\mathcal{M}$  allows the definition of differential operators on functions defined on the manifold. The differential operators are defined taking into account the metric of the manifold  $G = (g_{ij})$ . For instance, the gradient is defined imposing that the directional derivative of function  $f$  in direction  $\mathbf{v}$ , denoted as  $D_{\mathbf{v}} f$ , verifies  $D_{\mathbf{v}} f = \nabla f \cdot \mathbf{v}$  with the scalar product defined with the metric of the manifold. The gradient operator corresponds

<sup>1</sup>Summation is assumed for variables with the same sub- and super-indexes

then to  $\nabla_{\mathcal{M}} f = g^{ij} f_j \mathbf{d} u^i$ , and the squared norm of the gradient of a function is given by  $\|\nabla_{\mathcal{M}} f\|_{\mathcal{M}}^2 = g^{ij} f_i f_j$ , where  $(g^{ij})$  is the inverse of the metric matrix  $G$  and  $f_i = \frac{\partial f}{\partial u^i}$ .

## 2.3 Harmonic Active Contours

We formulate our segmentation method as a minimization problem. Our minimization variable is the level set function that describes the active contour of the segmentation, and the objective functional measures the surface of the manifold defined by the joint embedding of the image and the level set function. Consequently, the proposed segmentation corresponds to a minimal surface manifold, and the minimizing function is an harmonic map; this identifies our technique as an harmonic active contour.

For simplicity, our experiments consider only images in the plane, but the proposed harmonic active contours extend easily to higher-dimensional and non-Euclidean spaces. To illustrate this flexibility, this section presents the general formulation of the embedding, manifold, and metric that characterize the proposed segmentation technique; the expressions derived for images in the plane are only a particularization of the general HAC model.

### 2.3.1 Image segmentation as a Riemannian manifold

In our formulation, the segmentation is defined by the zero-level set of a function  $\phi$  and the segmentation criterion is given by the contour of the features  $f^1, \dots, f^k$ , which might depend on  $\phi$ . Both the features and level set function are defined in the same space manifold  $\Sigma$  than the images, take values in  $\mathbb{R}$  and are considered as differentiable functions in  $\Sigma$ .

The space manifold  $\Sigma$  has coordinates  $u = (u^1, \dots, u^n)$ , with metric  $G = (g_{\mu\nu})$  on these coordinates<sup>2</sup>. For images in the plane, for instance, we have  $u = (x, y)$ ,  $\Sigma = \Omega \subset \mathbb{R}^2$ , and  $G = I_2$ , while for images in the sphere parametrized with the stereographic projection we have  $G = \frac{4}{(1+x^2+y^2)^2} I_2$ .

We now make use of the Beltrami framework and consider an  $n$ -dimensional manifold  $\mathcal{M}$  associated with each possible segmentation of the image. This manifold  $\mathcal{M}$  is defined by the following embedding into the space-feature  $\Sigma \times \mathbb{R}^{k+1}$

$$\begin{aligned} \sigma: \Sigma &\longrightarrow \Sigma \times \mathbb{R}^{k+1} \\ u &\mapsto (u, f^1, \dots, f^k, \phi). \end{aligned} \tag{2.6}$$

Compared to previous approaches, our space-feature manifold includes both image features and level set function in the embedding. We keep the name space-feature manifold for simplicity, but we point out that the inclusion of both the level set function and the image

---

<sup>2</sup>To simplify notation, we use Greek indexes  $\mu, \nu$  to refer to space coordinates in  $\Sigma$  and latin ones  $i, j$  for the features. To that purpose also, the usual gradient and norms are assumed in  $\mathbb{R}^n$  and we specify with a sub-index  $\nabla_{\Sigma}$ ,  $\|\cdot\|_{\Sigma}$  any other case.

features is the key point that allows the definition of an alignment term between the gradients of the level set function and the gradients of the image.

We now pull-back the metric defined in  $\mathcal{M}$  by space and features together. To this purpose, in  $\Sigma \times \mathbb{R}^{k+1}$  we consider the metric

$$\begin{pmatrix} G & 0 & \dots & 0 & 0 \\ 0 & \alpha_1 & \dots & 0 & 0 \\ 0 & 0 & \ddots & 0 & 0 \\ 0 & 0 & \dots & \alpha_k & 0 \\ 0 & 0 & \dots & 0 & \beta \end{pmatrix}, \quad (2.7)$$

which offers the following interpretation: the distances on the coordinates  $u$  associated with the space manifold are measured in the corresponding metric  $G$  of  $\Sigma$ , which is considered orthogonal to the feature space, while both the feature channels  $f^1, \dots, f^k$  and the level set function  $\phi$  live in the Euclidean space  $\mathbb{R}^{k+1}$  and are considered orthogonal to each other. The pull-back operator on (2.7) induces the following space-feature metric ( $g_{\mu\nu}^*$ ) on the manifold  $\mathcal{M}$

$$g_{\mu\nu}^* = g_{\mu\nu} + \alpha_i f_\mu^i f_\nu^i + \beta \phi_\mu \phi_\nu. \quad (2.8)$$

With the proposed embedding both the image features and segmenting functions are included in the induced metric of Equation (2.8), leading to terms weighted by  $\alpha_i$  and  $\beta$ . The *subjective surface* model of Sarti et al. [2002] includes only the level set function in the embedding and weights the surface element with an edge detector to drive the active contours to the edges of the image; this results in a generalization of GAC to 2-dimensional manifolds that ignores region-based segmentation criteria. We avoid this issue by introducing additional dimensions on the embedding. The additional dimensions in the feature space, however, do not affect the intrinsic dimension of the image manifold  $\mathcal{M}$ , which does not depend on the embedding but on the space manifold  $\Sigma$ .

In this work we limit ourselves to 2-dimensional images, i.e.,  $u = (x, y)$ , the determinant of the induced metric tensor is computed as  $g^* = g_{11}^* g_{22}^* - g_{12}^* g_{21}^*$  and the squared surface element given by

$$g^* = g + \alpha_i g g^{\mu\nu} f_\mu^i f_\nu^i + \beta g g^{\mu\nu} \phi_\mu \phi_\nu + \frac{1}{2} \alpha_i \alpha_j [\nabla f^i, \nabla f^j]^2 + \alpha_i \beta [\nabla f^i, \nabla \phi]^2. \quad (2.9)$$

In Equation (2.9),  $\nabla f$  is the gradient computed in the coordinates  $u$ ,  $g$  is the surface element associated with the original metric tensor ( $g_{\mu\nu}$ ) on  $\Sigma$  and ( $g^{\mu\nu}$ ) its inverse. If we take into account the definition of differential operators in the manifold  $\Sigma$ , Equation (2.9) simplifies to

$$g^* = g + \alpha_i g \|\nabla_\Sigma f^i\|_\Sigma^2 + \beta g \|\nabla_\Sigma \phi\|_\Sigma^2 + \frac{1}{2} \alpha_i \alpha_j [\nabla f^i, \nabla f^j]^2 + \alpha_i \beta [\nabla f^i, \nabla \phi]^2. \quad (2.10)$$

This allows us to interpret the terms  $\|\nabla_{\Sigma} f^i\|_{\Sigma}^2$  and  $\|\nabla_{\Sigma} \phi\|_{\Sigma}^2$  as a simple measure of smoothness of the features and the level set function in the original metric of the space  $\Sigma$ . On the other hand, the terms  $[\nabla f^i, \nabla f^j] = f_x^i f_y^j - f_y^i f_x^j$  correspond to the magnitude of the cross-product of the vectors  $\nabla f^i$  and  $\nabla f^j$  and measure the coupling of the image features  $f^i$  and  $f^j$ . An equivalent term for the coupling of the level set function and the different feature channels is given by the terms  $[\nabla f^i, \nabla \phi]$ . Naturally, as the image features and the level set function are orthogonal to each other in (2.7), their associated terms do not involve the metric of the space manifold ( $g_{\mu\nu}$ ).

For images on manifolds conformally equivalent to the Euclidean plane, i.e., whose metric matrix is diagonal  $G = \sqrt{g}I_2$ , the squared surface element simplifies to

$$g^* = g + \alpha_i \sqrt{g} \|\nabla f^i\|^2 + \beta \sqrt{g} \|\nabla \phi\|^2 + \frac{1}{2} \alpha_i \alpha_j [\nabla f^i, \nabla f^j]^2 + \alpha_i \beta [\nabla f^i, \nabla \phi]^2.$$

This type of manifolds are commonly encountered in image processing, from the usual images defined on the plane ( $\sqrt{g} = 1$ ) to non flat images on the sphere ( $\sqrt{g} = \frac{2}{1+x^2+y^2}$  with the stereographic projection).

On the following, we will consider images on the plane, where we have

$$g^* = 1 + \alpha_i \|\nabla f^i\|^2 + \beta \|\nabla \phi\|^2 + \frac{1}{2} \alpha_i \alpha_j [\nabla f^i, \nabla f^j]^2 + \alpha_i \beta [\nabla f^i, \nabla \phi]^2. \quad (2.11)$$

### 2.3.2 Minimal surface manifold as segmentation criterion

We formulate our segmentation technique as the following minimization problem

$$\min_{\phi} \int_{\Omega} \sqrt{1 + \alpha_i \|\nabla f^i\|^2 + \beta \|\nabla \phi\|^2 + \frac{1}{2} \alpha_i \alpha_j [\nabla f^i, \nabla f^j]^2 + \alpha_i \beta [\nabla f^i, \nabla \phi]^2}. \quad (2.12)$$

We claim that the optimal segmentation is given by the level set function that minimizes the surface of the manifold defined by the joint embedding of the image and the level set function. Indeed, as can be seen in Equation (2.12), minimization of the surface element aligns the gradients of the level set function  $\phi$  with the gradients of the image features, that is, with the edges of the image. Let us analyse it in more detail.

Inside homogeneous regions, the minimization is driven by the smoothness penalty  $\|\nabla \phi\|$  because  $\sqrt{g} \approx \sqrt{1 + \beta \|\nabla \phi\|^2}$ , the objective functional approximates TV or isotropic diffusion depending on the value of parameter  $\beta$ , and the level set function tends to be flat.

In regions of the image with well-defined edges, the gradients of the image features are large and minimization of the cross-terms  $[\nabla f^i, \nabla \phi]$  attracts the level sets of  $\phi$  to the edges of the image, the spatial gradient of  $\phi$  increases, and the level set function develops a discontinuity. In regions with fragmented edges, the information about size and orientation of neighbouring well-defined edges is smoothly extended by the terms  $|\nabla \phi|$  and  $[\nabla f^i, \nabla \phi]$ , and the level set



function is able to complete the missing boundaries.

In the proposed harmonic active contours, therefore, the segmentation criterion is defined by the choice of the image features, which we particularize next to obtain contour- and region-based segmentation criteria.

### 2.3.3 Feature definition for edge- and region-based segmentation

Edge-based segmentation is obtained by choosing features that are local image descriptors. In the easiest case, we simply embed the grey level or color intensities of the image, but more elaborate features such as semi-local texture descriptors, Wavelet or Gabor coefficients could be equally used. For simplicity we define  $f^i(x, y) = I^i(x, y)$  for each channel  $I^i$  in the image.

At the same time, in order to detect objects that are not defined by local gradients, but by homogeneous regions, we introduce region-based features. We adopt the piecewise constant instance of the segmentation model proposed by Mumford and Shah [1989]. We define two image regions,  $R_+ = \{(x, y) \mid \phi(x, y) > 0\}$  and  $R_- = \{(x, y) \mid \phi(x, y) < 0\}$ , and characterize them by the mean values of the local features inside the region,  $\mu_+^i$  and  $\mu_-^i$ . For each channel, the associated region features are inspired by ACWE and introduce the following region descriptor associated with  $R_+$

$$f^i = (I^i - \mu_+^i)^2 H(\phi) = s^i H(\phi), \quad (2.13)$$

where  $H$  is the Heaviside function, and the resulting feature  $f^i$  measures the square error associated with the piecewise approximation of  $R_+$ . The region descriptor for  $R_-$  can be included in a similar manner, but we omit it here to shorten the equations. Its treatment and minimization are analogous to the region descriptor  $R_+$ . The proposed region features segment the image into two regions in terms of the zero-level set of the function  $\phi$ , that is, they fix the level set that parametrizes the contours of the segmentation.

The proposed segmentation method combines contour- and region-based criteria by including both contour- and region-based features in the embedding. The resulting squared surface element reads

$$g = 1 + \alpha_i \|\nabla I^i\|^2 + \gamma_i \|\nabla f^i\|^2 + \beta \|\nabla \phi\|^2 + \frac{1}{2} \alpha_i \alpha_j [\nabla I^i, \nabla I^j]^2 + \alpha_i \gamma_j [\nabla I^i, \nabla f^j]^2 + \frac{1}{2} \gamma_i \gamma_j [\nabla f^i, \nabla f^j]^2 + \alpha_i \beta [\nabla I^i, \nabla \phi]^2 + \gamma_i \beta [\nabla f^i, \nabla \phi]^2. \quad (2.14)$$

The term  $\|\nabla \phi\|^2$  controls the smoothness of the level set function and  $\|\nabla f^i\|^2$  the smoothness of the region decomposition of the image, the cross-terms  $[\nabla I^i, \nabla f^j]^2$  measure the coherence between region and contour criteria, and the terms  $[\nabla I^i, \nabla \phi]^2$  and  $[\nabla f^i, \nabla \phi]^2$  measure the alignment between the contours of level set function and the edges of the original image and its region decomposition.

The parameter of the metric associated with the image channel  $I^i$  is  $\alpha_i$ , while the parameter  $\gamma_i$  is associated with the region term  $f^i$ . In order to simplify the metric, we normalize the feature range of each channel by  $\alpha_i = \frac{\alpha}{m_\alpha^i}$  and  $\gamma_i = \frac{\gamma}{m_\gamma^i}$ . The resulting segmentation is defined by three parameters:  $\alpha$  controls the weight given to the edge segmentation criterion,  $\gamma$  controls the region criterion associated with a piecewise constant decomposition of the image, and  $\beta$  controls the regularity of the contours and the level set function  $\phi$ .

### 2.3.4 Relation to *subjective surfaces* and anisotropic diffusion

We analyse now the relation of the proposed technique with existing geometric methods proposed in the Beltrami framework. For simplicity we limit the analysis to images in the plane.

We start by reviewing the embedding and resulting metric proposed for anisotropic diffusion and texture denoising in Sochen et al. [1998], Kimmel et al. [1997], Sagiv et al. [2006]. Given a multi-valued image in the plane with features  $f^1, \dots, f^k$ , the natural embedding for denoising is given by

$$\begin{aligned} \sigma^f: \Omega &\longrightarrow \Omega \times \mathbb{R}^k \\ (x, y) &\mapsto (x, y, f^1, \dots, f^k). \end{aligned} \quad (2.15)$$

The metric induced in the image manifold is then

$$G^f = \begin{pmatrix} 1 + \alpha_i f_x^{i2} & \alpha_i f_x^i f_y^i \\ \alpha_i f_x^i f_y^i & 1 + \alpha_i f_y^{i2} \end{pmatrix}, \quad (2.16)$$

and the resulting surface element  $g^f = 1 + \alpha_i^2 |\nabla f^i|^2 + \alpha_i \alpha_j [\nabla f^i, \nabla f^j]^2$  works as a robust edge detector. It defines the edges of the image considering both the gradients of each channel with the terms  $|\nabla f^i|^2$ , and the coupling of the different channels with the cross-products  $[\nabla f^i, \nabla f^j]$ .

A similar geometric framework has been applied to segment images in Sarti et al. [2002], where the level set function is interpreted as a *subjective surface*. To that purpose, Sarti et al. substitute the image features  $f^1, \dots, f^k$  by the level set function  $\phi$  and define the following embedding

$$\begin{aligned} \sigma^\phi: \Omega &\longrightarrow \Omega \times \mathbb{R}^k \\ (x, y) &\mapsto (x, y, \phi). \end{aligned} \quad (2.17)$$

With the usual scalar product in the space-feature  $\Omega \times \mathbb{R} \subset \mathbb{R}^3$ , the induced surface element is given by  $g^\phi = 1 + \beta \|\nabla \phi\|^2$  and measures the area of the level set function seen as a surface. In order to detect contours, the surface element is weighted with an inverse edge detector, and the level set function is evolved to minimize its weighted surface measured by  $\int_\Omega w \sqrt{1 + \beta \|\nabla \phi\|^2}$ .

Consequently, this method is a generalization of GAC to 2-dimensional manifolds, as shown in Bresson et al. [2006a]. An equivalent formulation of this functional is obtained weighting directly the elements of the metric of  $\Omega \times \mathbb{R}$  by the inverse edge detector  $w$ , that is, considering the scalar product in  $\Omega \times \mathbb{R}$  given by the following symmetric, positive semi-definite matrix

$$\begin{pmatrix} w & 0 & 0 \\ 0 & w & 0 \\ 0 & 0 & \beta w \end{pmatrix}. \quad (2.18)$$

The inverse edge detector  $w$  is usually a simple function of the image gradient, like  $w = \frac{1}{1+a|\nabla I|^2}$ , and exhibits none of the properties of the robust edge detector  $g^f$  used in the Beltrami framework for diffusion. Embedding of the level set function  $\phi$  in the definition of the manifold results, however, in a segmentation technique able to cope with weak edges and fragmented boundaries, see Sarti et al. [2002].

The segmentation model of the *subjective surface*, moreover, does not include a region term in the objective functional, and it does not explicitly define a contour because the particular level sets of  $\phi$  that parametrizes the contour is not specified in the objective functional. The inclusion of the region features of Equation (2.13) in our embedding overcomes these two limitations, i.e., the ability to segment images based on the homogeneity of the regions and the explicit parametrization of the contour as the zero-level set of function  $\phi$ .

In fact, it is easy to prove that our embedding corresponds to a generalization of the *subjective surfaces* with the following metric in the space-feature  $\Omega \times \mathbb{R}$

$$\begin{pmatrix} G^f & 0 \\ 0 & \beta \end{pmatrix} = \begin{pmatrix} 1 + \alpha_i f_x^{i2} & \alpha_i f_x^i f_y^i & 0 \\ \alpha_i f_x^i f_y^i & 1 + \alpha_i f_y^{i2} & 0 \\ 0 & 0 & \beta \end{pmatrix}. \quad (2.19)$$

Compared to *subjective surfaces*, our metric in the space-feature  $\Omega \times \mathbb{R}$  substitutes the edge detector  $w$  with the induced metric of anisotropic diffusion by Sochen et al. [1998], Kimmel et al. [1997], Sagiv et al. [2006]. Interpreting the metric  $G^f$  as a robust edge detector, we can easily see that the proposed technique combines the advantages of both geometric approaches: diffusion methods for the robust detection of edges and *subjective surfaces* for its ability to close fragmented and weak edges.

### 2.3.5 Relation to robust edge integration

The proposed HAC shares a strong connection with the segmentation method proposed in Kimmel and Bruckstein [2003], where a geometric alignment term between the normals to the contour and the gradient of the image is combined with the objective functionals of GAC and ACWE to improve image segmentation. This alignment term defines a robust edge integrator that acts in a similar manner to the cross-product terms of the proposed HAC. To understand the differences and similarities of both methods, a more detailed description of the robust

edge integrator of Kimmel and Bruckstein [2003] is necessary.

Kimmel and Bruckstein [2003] observes that the direction associated with the gradient vector of an image  $I$  is a good estimator of the orientation of the edge contour, and introduces a robust edge integrator that encourages the alignment between the normal vector to the contour and the gradient of the image along the contour. With an implicit parametrization of the contours, this alignment term reads

$$-\int_{\mathcal{C}} \left| \nabla I \cdot \frac{\nabla \phi}{|\nabla \phi|} \right|. \quad (2.20)$$

From Equation (2.20), it is clear that this alignment term is only active in  $\mathcal{C}$  for the zero-level set of  $\phi$ . On the other hand, the proposed HAC aligns the gradients of the level set function, not the normals, with the gradients of the image for all the level sets of  $\phi$  by integrating the cross-terms  $[\nabla f^i, \nabla \phi]^2$  throughout the whole image domain  $\Omega$ . This gives us two theoretic differences with respect to the robust edge integrator of Kimmel and Bruckstein for grey-scale images. First, the weight given to the alignment term in HAC is proportional to  $|\nabla \phi|$ , which gives more weight to the alignment of gradients when the level set function has larger gradients, i.e., stronger edges. Second, the HAC aligns the gradients of all the level sets of the level set function, not only along the active contour, and is therefore less likely to get trapped in local minima, see Figures 2.1c and 2.1d and the experiments of Section 2.5.

Moreover, HAC naturally generalizes to multi-channel images, where both the coupling of the different image channels and the coupling of their alignment with the level set function are included in the objective functional of Equation (2.14). Finally, the proposed HAC can also be applied to non-Euclidean parametric manifolds and feature spaces.

In terms of implementation, Kimmel [2003] proposes an efficient algorithm to minimize the resulting objective functional with a PDE descent flow in the level set function. The algorithm combines additive operator-splitting and a narrow-band implementation, but requires periodic redistancing of  $\phi$  as a signed distance function. As a result, and despite the level set formulation, the alignment term is only active in a narrow band close to the zero-level set of  $\phi$ . On the other hand, HAC uses modern optimization techniques to minimize the objective functional and is therefore faster. Our objective functional, in fact, is not compatible with a narrow-band implementation or periodic redistancing of the level set function because it exploits the alignment of all the level sets of  $\phi$  with the image features, not only close to the active contour.

## 2.4 Numerical minimization

The optimization problem associated with the proposed HAC model is given by

$$\min_{\phi} \int_{\Omega} \sqrt{g} \quad \text{with } g \text{ given in Equation (2.14)}. \quad (2.21)$$

In Estellers et al. [2011], we adopted the traditional PDE descent flow used in the Beltrami framework to minimize the objective functional. In particular, we used Euler-Lagrange to obtain the optimality conditions for the minimization problem and evolved  $\phi$  with the following flow

$$\phi_t = -\frac{1}{2\sqrt{g^*}} \left( \frac{\partial g^*}{\partial \phi} + \frac{1}{2g^*} \frac{\partial g^*}{\partial \sigma^\mu} \frac{\partial g^*}{\partial \phi_\mu} - \frac{\partial}{\partial \sigma^\mu} \frac{\partial g^*}{\partial \phi_\mu} \right) \quad (2.22)$$

until a fixed-point was encountered, and the optimality conditions were met. To guarantee stability of the iterative scheme, the time step of the flow was limited by the CFL condition, and the minimization technique was extremely slow.

Several techniques have been proposed recently to accelerate the Beltrami flow for image reconstruction, like the vector extrapolation of Dascal et al. [2007], or the semi-implicit PDE schemes of Dascal et al. [2009]. The experiments of Rosman [2010], however, show that minimization techniques based on variable-splitting and augmented Lagrangians result in faster algorithms.

Based on this observation, we develop here a minimization algorithm that relies on the same variable-splitting and augmented Lagrangian techniques than Chapter 4 to obtain an algorithm that, in average, is 4-5 times faster than the original descent flow (2.22). To avoid repetition, the details of the proposed algorithm are given in Appendix C.

## 2.5 Experimental results

This section evaluates the properties of the proposed HAC by segmenting different kinds of images and compares its performance with related segmentation methods. For a fair comparison, all the code is implemented in Matlab and runs in a laptop with an Intel Core i5 at 2.3GHz, the timings are given in seconds (s).

The first set of experiments is designed to characterize the properties of the proposed HAC model. We start in Figures 2.2a and 2.2b by showing the ability of HAC to exploit the directional information of the image gradients in order to overcome the inhomogeneities present in the images and obtain correct segmentations. We then show the ability of the proposed method to correctly segment ridge structures, first with a noisy image of a car-plate image in Figure 2.2d and then with the segmentation of blood vessels in a fundus image with low-contrast in Figure 2.2c. Next, Figures 2.3a-2.3d show the ability of the proposed method to close fragmented contours and segment medical images with weak edges. Figure 2.4 presents the results of the segmentation of color images, where the proposed method exploits the coupling of the different channels to detect meaningful edges. In this case, we provide also the segmentations obtained with the standard convex GAC+ACWE segmentation method, where neither the coupling of the different channels nor the alignment of image gradients with the contours are exploited to overcome the inhomogeneities of the background (Figure 2.4c) or textured areas




Figure 2.2: Segmentation of images with ridge structures, inhomogeneities and noise with the proposed HAC (top row) and with Kimmel and Bruckstein [2003]'s model (bottom row). Initial level set set in red, final segmentation in blue.

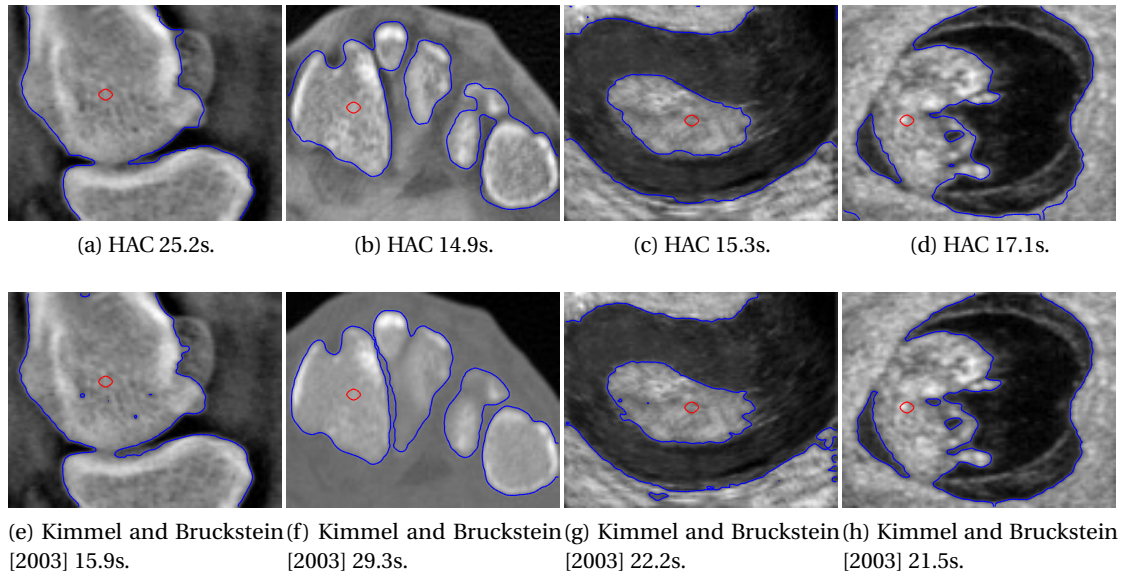


Figure 2.3: Segmentation of medical images with the proposed HAC (top row) and with Kimmel and Bruckstein [2003] model (bottom row). Initial level set set in red, final segmentation in blue.

(Figure 2.4d).

In a second set of experiments we compare the proposed HAC with the robust edge

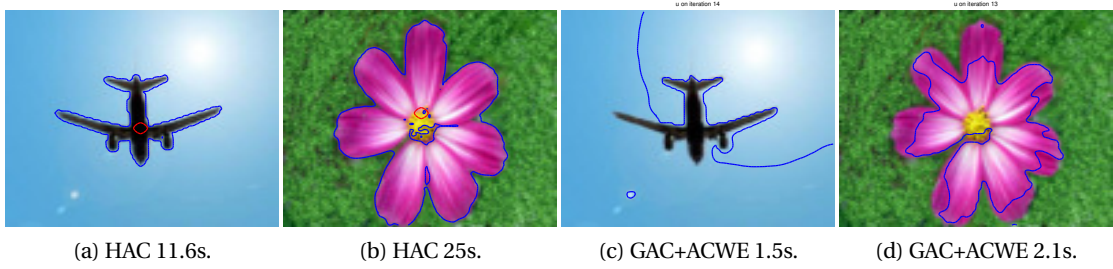


Figure 2.4: Segmentation of color images with HAC model and a the convex formulation of GAC+ACWE models. Initial level set set in red, final segmentation in blue. Note that initialization does not affect the convex GAC+ACWE segmentation model<sup>3</sup>.

integrator of Kimmel and Bruckstein [2003] combined with the GAC and ACWE models. In particular, we adopt the narrow-band implementation proposed in Kimmel [2003], with fast redistancing of the level set function following Sethian [1996]. Figures 2.2e-2.2h and 2.3e-2.3h show the results of the segmentations for Kimmel and Bruckstein's method. For images not subject to inhomogeneities, Figures 2.3e-2.3h, both methods perform similarly in terms of accuracy, but HAC is faster because we make use of variable-splitting and augmented Lagrangian to design an efficient minimization technique. The images of Figures 2.2e-2.2g are subject to inhomogeneities and cannot be segmented with a simple combination of GAC+ACWE models. In these cases, the alignment of image gradients and level set function is necessary to produce correct segmentations, and our experiments show that the proposed HAC produces slightly better segmentations in terms of speed and accuracy, see for instance the fundus image in Figures 2.2c and 2.2g, or the polyhedra of Figures 2.2a and 2.2e.

The key point, however, is the robustness of HAC to different initializations, as shown in Figure 2.5. In these synthetic images, Kimmel and Bruckstein's method cannot exploit the region term due to the image inhomogeneity, and a balloon force is introduced to either shrink or inflate the original contour. With the wrong initialization, therefore, Kimmel and Bruckstein's method is not able to detect new outer contours of the second of object of Figure 2.5a or the inner contours of the number of Figure 2.1c because they are too far from the zero-level set of  $\phi$ . On the other hand, the proposed HAC aligns the gradients of all the level sets of  $\phi$  with the gradients of the image features and is able to propagate the alignment of neighbouring level sets to detect new contours. As a result, the segmentations obtained from different initializations in Figures 2.2a and 2.5c, and Figures 2.1d and 2.5d are not qualitatively different with our method, while they differ with Kimmel and Bruckstein's method.

## 2.6 Conclusions

This chapter has developed a segmentation method in a geometric framework that interprets each possible segmentation of an image as a 2-dimensional manifold embedded in a higher-dimensional space, from which it inherits a metric. This metric defines distances between




Figure 2.5: Comparison of image segmentation results with HAC and the inclusion of the robust edge integrator of Kimmel and Bruckstein [2003] into the GAC+ACWE model. Initial level set set in red, final segmentation in blue.

points in the manifold that consider simultaneously the spatial distance between the points, the values of the image features at these points, and the labels assigned to them in the segmentation. A minimal surface manifold corresponds then to a segmentation that minimizes the usual region, edge, and regularity terms of an active contour model, and additionally takes into account the coupling of the different image channels and the alignment between the gradients of the image and the gradients of the level set function.

Compared with the *Geodesic Active Contours* of Caselles et al. [1997] and the *Active Contour without Edges* of Chan and Vese [2001], the proposed technique overcomes image inhomogeneities because it exploits the directional information of the gradients of the image to align the contours of the segmentation with the edges of the image.

A similar idea is the basis of the robust edge integrator of Kimmel and Bruckstein [2003], where the alignment between the gradients of the image and the normal vectors to the contours drives only the zero-level set of the level set function. The alignment of the proposed model, on the other hand, affects all the level sets and is thus able to exploit the alignment of the neighbouring levels to pull the contours to the right position. As a result, the proposed harmonic active contours are able to discover new edges and are less sensitive to initialization.

The proposed formulation, moreover, is directly extended to higher-dimensional spaces and non-flat images, where usual segmentation methods cannot be applied.



### 3 Surface Reconstruction from Microscopic Images of Printed Circuits

This chapter develops a method to reconstruct 3D surfaces of silicon wafers from 2D images of printed circuits taken with a scanning electron microscope. To solve this inverse problem, the reconstruction method combines the physical model of the optical acquisition system with prior knowledge about the shapes of the patterns in the circuit; the result is a shape-from-shading technique with a shape prior.

In the context of variational methods, the reconstruction of the surface is formulated as an optimization problem with an objective functional that combines a data term on the microscopic image with two prior terms on the surface. The data term models the acquisition system through the irradiance equation characteristic of the microscope; the first prior is a common smoothness penalty, and the second prior constrains the shape of the surface to agree with the expected shape of the pattern in the circuit. Geometry, in this chapter, plays a role in the data term as a result of the irradiance equation, which relates the shading observed in the image with the orientation of the unknown surface.

In order to account for the variability of the manufacturing process, the prior term incorporates a variable that allows a non-linear elastic deformation between the expected pattern and the reconstructed surface; the estimation of this deformation field introduces a registration problem in the proposed variational formulation. As a result, the minimization problem has two unknowns, and the reconstruction method provides two outputs: a reconstructed surface and a deformation field. The reconstructed surface is derived from the shading observed in the image and the prior knowledge about the pattern in the circuit, while the deformation field produces a mapping between the expected shape and the reconstructed surface that provides a measure of deviation between the theoretical models and the real manufacturing process.

---

The material of this chapter has been partially published in Estellers et al. [2013b].

### 3.1 Introduction: a few words on lithography

As circuit designs become more complex due to technology scaling, engineers in the semiconductor industry require new tools to obtain accurate and useful information about the manufacturing process. In particular, the characterization of surface topography is essential in optical lithography in order to optimize the printing process from layout masks to silicon wafers. Scanning electron microscopy (SEM) is the primary tool currently used in industry to measure geometrical features of semiconductor structures, but it provides only limited information about the third dimension of the surfaces. In fact, several parameters associated with the surface of the wafer are crucial to control the lithographic process (the height of the step, the slope of the side walls, or the curvature of edges, see Mack [2007]), and it has now become necessary for the industry to develop image processing techniques to reconstruct accurate 3D surfaces from 2D SEM images. This chapter develops such a technique. It is formulated as a shape-from-shading (SFS) reconstruction that considers prior knowledge about the shape of the patterns in the circuit in order to improve the accuracy and robustness of the reconstruction.

In the past, computer vision has focused on the combination of stereoscopy with SFS for the reconstruction of surfaces from SEM images of printed circuits, but the resulting techniques have not been adopted by the industry. The acquisition of pairs of stereoscopic SEM images at the nanometre scale is not only challenging, but extremely expensive for large wafers, and the throughput requirements of manufacturers call for simpler and more economic methods. Some examples of these techniques are Kayaalp et al. [1990], Beil and Carlsen [1991], Ellison and Taylor [1991], Jones and Taylor [1995], Danzl and Scherer [2002], Reithmeier et al. [2010].

Industrial SEM images are usually analysed with simple image processing techniques that use edge detection and template matching to estimate 1D profiles of the surface along the contours of the SEM image (Morokuma et al. [2005], Tabery et al. [2007, 2006], Shindo et al. [2009], Shishido et al. [2010], Shibahara et al. [2011]). These techniques coarsely detect edges in the SEM image with simple gradient and thresholding methods, locate a preliminary contour and subsequently refine it to match a physically meaningful criterion in terms of the underlying surface. To that purpose, they exploit available information about the patterns in the circuit and analyse 1D profiles of the image along the normals to the contours, which are then modified to match a fixed height level of the underlying surface. These profile measurements are relatively fast to extract, but no real 3D surface is reconstructed. This chapter “fills in the gap” between these two approaches and includes the knowledge of the patterns expected in the circuit into a surface reconstruction method based on shape-from-shading. Compared to the stereoscopic method, the proposed technique does not require SEM image pairs obtained with an expensive dedicated system; instead, it exploits the available information about the shapes of the patterns in the circuit to overcome the ill-posed nature of SFS with a shape prior and obtain a robust reconstruction method.

For each pattern in the printed circuit, the nature and accuracy of the available information

can be considerably different. While reasonable optical proximity correction (OPC) models of the printing process exists for some technology nodes and designs, these models are under calibration in the first cycles of development and only the mask associated with the circuits can be trusted (Mack [2007]). The reconstruction method proposed here considers both situations and defines different shape priors depending on the nature of the available information. At the time, it estimates a deformation field between the shape prior and the reconstructed surface that can be used *a posteriori* to quantify unaccounted errors in the OPC modelling or the lithographic printing process.

The rest of this chapter is organized as follows. Section 3.2 presents the proposed technique and its context, reviewing SFS techniques in optical lithography and shape priors in imaging. The reconstruction of the surface is then formulated as a minimization problem and solved with efficient optimization techniques in Section 3.3, leaving the algorithmic details for the Appendix D. Section 3.5 presents the experimental results, and conclusions are drawn in Section 3.6.

## 3.2 Shape-from-shading with a shape prior

In the context of variational methods, the reconstruction of a surface is formulated here as an optimization problem with a combined matching criterion based on the irradiance equation and a shape prior that constrains the reconstructed surface to agree with the expected shape of the pattern. To account for the variability of the physical printing process, the model allows a non-linear elastic deformation of the expected patterns. We propose two different objective energies to measure this deformation, one based on an OPC model of the expected surface and another on the limited information extracted from the mask of the circuit.

The proposed variational problem has thus two variables, and the reconstruction method provides two outputs: a surface parametrized by its height map  $z$  and a deformation field  $\mathbf{u}$ . The minimization problem that will be described in this section reads

$$\min_{z, \mathbf{u}} \alpha A(\nabla z) + \beta B(z, \mathbf{u}) + \gamma G(\nabla z) + W(\nabla \mathbf{u}). \quad (3.1)$$

The height map of the surface  $z$  is reconstructed from the shading observed in the SEM image and the prior knowledge about circuit patterns, which results in terms  $A$ ,  $B$  and  $G$  in the objective functional. The deformation field  $\mathbf{u}$  produces a mapping between the expected shapes in the circuit and the reconstructed surface, which leads to terms  $B$  and  $W$  in the minimization problem (3.1). The details of each one of these terms are presented next, together with a review of related approaches.

### 3.2.1 Shape-from-shading from SEM images

In computer vision, shape recovery refers to the reconstruction of a 3D scene from 2D images and covers different techniques, called shape-from-x, depending on the nature of the input information. Shape-from-shading, in particular, deals with the recovery of shape from a gradual variation of shading in the image.

To understand the problem of SFS it is necessary to study how the images are formed. In our formulation we attach to the camera a 3D coordinate system  $(x, y, z)$ , where  $(x, y)$  corresponds to the image plane and the axis  $z$  corresponds on the optical axis. If we assume an orthographic projection, the visible part of the scene defines a height map  $z(x, y)$  with the coordinates of the points in the image given by  $(x, y)$ . The SFS problem is then formulated in terms of the “irradiance equation” of the acquisition system:

$$I_0(x, y) = R(\mathbf{n}(x, y)). \quad (3.2)$$

Equation (3.2) relates the grey level  $I_0(x, y)$  measured in the image with the reflectance function  $R(\mathbf{n}(x, y))$  of the surface at that point. In physical terms,  $R(\mathbf{n})$  gives the value of the light re-emitted by the surface as a function of its orientation, i.e., of the vector  $\mathbf{n}$  normal to the surface. This normal can easily be expressed in terms of the gradient of the height map as

$$\mathbf{n} = \frac{(-\nabla_x z, -\nabla_y z, 1)}{\sqrt{|\nabla z|^2 + 1}}. \quad (3.3)$$

The original SFS problem proposed by Horn [1970] consists on the recovery of a height map  $z$ , the unknown of the problem, from the shading observed in the image  $I_0$ , the data available from measurements.

Since the introduction of SFS in the 70s, most efforts have been directed into the recovery of shapes from images taken under normal lighting conditions, and only few investigations have considered the problem associated with SEM images, see e.g. Kayaalp et al. [1990], Beil and Carlsen [1991], Ellison and Taylor [1991], Jones and Taylor [1994, 1995], Danzl and Scherer [2002], Rad et al. [2007], Reithmeier et al. [2010].

In the SEM acquisition system, images are formed scanning the sample with an electron beam and detecting the electrons emitted at the point where the beam strikes the surface. The number of electrons emitted by the surface is related to the angle between the electron beam and the surface normal, resulting in shades in the captured image  $I_0$  proportional to the number of electrons. At high magnifications, the imaging geometry can be modelled as an orthographic projection with a unique light source at infinity. The resulting reflectance function for SEM images is then given by

$$R(\mathbf{n}) = \frac{\nu}{\mathbf{n} \cdot \mathbf{L}}, \quad (3.4)$$

where  $\nu$  is the reflectance coefficient of the surface, and  $\mathbf{L}$  is a unit vector indicating the

direction of the light source. In the case of SEM images of silicon wafers,  $v$  is assumed uniform within an image,  $L \approx (0, 0, 1)$ , and both parameters can be directly estimated from the SEM images as explained in Section 3.4. Such a reflectance map neglects the effects of mutual illumination, but provides a model of the SEM acquisition system that is computationally tractable.

#### 3.2.2 Review of surface reconstruction from SEM images of printed circuits

In the case of silicon wafers of printed circuits, the features derived from SEM images are generally too sparse for standard SFS techniques. For this reason, in research, surface reconstruction from SEM measurements has been formulated as the correspondence problem associated with stereo pairs of images that are obtained tilting the microscope stage or deflecting the electron beam. The combination of stereoscopy and SFS is able to reconstruct accurate surfaces because the two techniques complement each other: SFS works best where the surface is smooth and featureless, while stereo works well where the surface is rough and the features are dense. See Kayaalp et al. [1990], Beil and Carlsen [1991], Ellison and Taylor [1991], Jones and Taylor [1995], Danzl and Scherer [2002], Reithmeier et al. [2010] for the details of such techniques.

In the industry, however, the main approach to 3D surface reconstruction from SEM images is based on 1D measurements and Monte Carlo simulations. The Monte Carlo method models statistically the interaction of the electron beam with the wafer by repeated simulations of electron behaviour, which results in extremely long running times. Consequently, Monte Carlo simulations are only used to obtain a library of SEM signals for simple circuit patterns defined by a few parameters, see Villarrubia et al. [2004], Villarrubia and Ding [2009], Villarrubia et al. [2010]. A measured SEM signal is then compared to the signals in library, and a 1D profile of the surface is reconstructed estimating the library parameters that best match the measurements. In this case it is important to obtain a robust estimate of the SEM profile, and several 1D measurements are usually averaged, or smoothed, to avoid the effects of noise. This technique has proved considerably accurate for simple structures available in the libraries, but it is too restrictive for the complicated 2D patterns currently used in circuit designs. Compared to the SFS approach adopted in computer vision, Monte Carlo simulations provide an accurate model of the reflectance map, but the reconstruction is limited to 1D profiles and does not consider the 3D nature of the surface.

The approach proposed here combines SFS techniques from computer vision with the prior knowledge available in industry about the expected patterns and OPC models. We approximate SEM acquisition system with the irradiance equation of SFS, and we consider the 2D nature of the image measurements, the expected smoothness of the surface, and the prior knowledge about its shape to overcome the limitations of Equation (3.4) to model the electron interaction as accurately as Monte Carlo simulations. Similar to the combination of SFS with stereoscopy, our approach integrates the complementary techniques of SFS and

image registration to merge the best of both worlds: while SFS works well where the image is featureless and the surface is smooth, the registration of the shape prior is guided by edges of the image. Compared with stereoscopy, we do not require special acquisition systems to obtain simultaneous SEM image pairs, instead we make use of the information available from circuit designs about the expected patterns in the image. It is worth noting, however, that our approach could also be modified to include stereoscopic pairs.

#### 3.2.3 Variational formulation of shape-from-shading

A variational formulation of SFS offers here three additional advantages. First, variational SFS techniques are generally more robust than other approaches, and therefore better suited for an application with industrial SEM data (Zhang et al. [1999], Durou et al. [2008]). Second, the objective minimization energy can easily include a term describing the prior information provided by the mask or the OPC model in the shape reconstruction algorithm. Third, the resulting minimization problem can be solved with efficient optimization techniques to speed up the reconstruction. We refer the reader to the review papers of Zhang et al. [1999], Prados et al. [2006], Durou et al. [2008], for an overview of other possible SFS approaches.

Variational SFS methods are defined by the choice of three elements: the unknown  $z$  used to describe the surface, the objective functional  $E$  that is minimized, and the minimization method. Functional  $E$  is defined by two terms or constraints in SFS: the brightness and the smoothness constraint. The brightness term  $A$  ensures that the reconstructed shape produces a similar irradiance as the input image, while the smoothness term  $G$  ensures the reconstruction of a smooth surface, and it is introduced to overcome the ill-posed nature of the problem. As both constraints can be directly written in terms of the gradient  $\nabla z$  or the normal  $\mathbf{n}$  to the surface, the minimization can be solved in terms of  $\nabla z$  or  $\mathbf{n}$ . If the minimization is solved in terms of  $\mathbf{n}$ , the height map of the surface must be recovered *a posteriori* by integration (Ikeuchi and Horn [1981]), and the integrability constraint of Frankot and Chellappa [1988]  $\frac{\partial^2 z}{\partial x \partial y} = \frac{\partial^2 z}{\partial y \partial x}$  must be guaranteed to reconstruct physically meaningful surfaces. The integrability of the surfaces can also be ensured by directly formulating the brightness constraint in terms of the height map  $z$ , which results in variational methods with higher-order derivatives and slower minimization algorithms (see e.g., Horn [1990], Leclerc and Bobick [1991], Szeliski [1991]). We adopt such a formulation, but use recent advances in optimization techniques to obtain a fast minimization algorithm.

#### Proposed data term – a model of the image acquisition system

In the proposed formulation, the brightness term reads

$$A(I) = \frac{1}{2} \int_{\Omega} (R(\mathbf{d}) - I_0)^2 \quad \text{with} \quad \mathbf{d} = (d_x, d_y) = \nabla z, \quad (3.5)$$

which assumes Gaussian noise in the SEM image acquisition process, but other norms instead of  $L_2$  could be used for other noise models.

#### Proposed smoothness term on the surface

We adapt the common smoothness term of SFS,  $\frac{1}{2} \int_{\Omega} (\Delta z + 2\nabla_x \nabla_y z)^2$ , inspired by image segmentation techniques. In particular, our approach mimics the active contours of Kichenassamy et al. [1995] and Caselles et al. [1997] by weighting the common smoothness term with an inverse ridge detector that only penalizes locations where high slopes are unlikely. The resulting penalty is

$$G(z) = \frac{1}{2} \int_{\Omega} w (\Delta z + 2\nabla_x \nabla_y z)^2 = \frac{1}{2} \int_{\Omega} w |\nabla \mathbf{d}|^2 \quad \text{with} \quad \mathbf{d} = \nabla z. \quad (3.6)$$

The weighting function  $w$  is designed to satisfy  $w \approx 0$  on the ridges of the SEM image and  $w \approx 1$  elsewhere because the SEM ridges appear only at locations where the shape of the surface changes abruptly. Weighting the smoothness term acts as the constraint proposed in Zheng and Chellappa [1991], which specifies that the intensity gradients of the reconstructed image and the input image should be close to each other. In particular we define  $w$  as a simple inverse ridge detector adapted from Lindeberg [1996], Frangi et al. [1998] to the SEM image  $I_0$ , but other ridge detectors are equally possible.

#### 3.2.4 Shape prior in shape-from-shading

Humans, and in particular engineers, interpret SEM images of circuits by incorporating high-level knowledge about the expected patterns. Such a knowledge is necessary to disambiguate low-level intensities and artifacts in noisy images and to make any automatic reconstruction method robust. Computer vision has successfully imitated the integration of prior knowledge done by humans into the image segmentation problem, but this information has not yet been exploited in the context of shape-from-shading. Inspired by image segmentation techniques, we incorporate the prior knowledge about the shapes of the patterns in the circuit into the surface reconstruction process.

Shape priors are included in image segmentation techniques to ensure robustness to occlusions and noise, see e.g., Leventon et al. [2000], Tsai et al. [2001], Chen and Tagare [2002], Rousson and Paragios [2002], Cremers et al. [2002], Paragios [2003], Cremers et al. [2004], Chan and Zhu [2005], Foulonneau et al. [2006], Bresson et al. [2006b].

The standard approach in variational methods is to incorporate an additional term in the objective functional that defines the segmentation criterion. This additional term measures the similarity between the evolving shape and a reference one, called shape prior, and penalizes deviations with respect to the expected shape. The resulting method must deal with two important issues: alignment of the shape prior within the image and variability of the shape with respect to the prior. For the question of shape alignment, affine transformations can be

taken into account in the model. This can be done explicitly in the definition of the objective functional, at the price of increasing the complexity of the optimization; or by intrinsic alignment of complex shape descriptors. The second issue is the question of shape variability in the reference template, which is usually handled through statistical models or trained shape models from annotated data. In the case of SFS from SEM images of silicon wafers, these two issues require different answers.

Shape alignment techniques can be directly adapted and included in our SFS model, but they are unnecessary and would only increase the complexity of the minimization. Indeed, SEM images of silicon wafers are usually correctly scaled and rotated because the imaging device takes into account the circuit design in the acquisition system. The question of shape variability is more difficult to handle because obtaining an accurate model of a surface is extremely expensive at the nanometre scale. Atomic force microscopy has been used to measure the surfaces of a few circuit patterns in Seeger [2004], but it is an extremely expensive process that is not viable if several images of each pattern must be acquired. OPC modelling can be used to obtain surface models through simulations, but the resulting models are partly based on heuristics and only predict certain effects of the lithographic printing process (Mack [2007]). Due to these limitations, we substitute the model of shape variability with a dense deformation field in our surface recovery model.

In our approach we simultaneously recover the surface  $z$  and estimate a deformation field  $\mathbf{u}$  between the expected shape and the recovered one. This strategy does not require a shape model and offers two additional advantages. First we can incorporate information of the expected shapes in terms of 3D models or 2D masks of the patterns, depending on the accuracy of the available information. Second, the deformation field directly provides a measure of deviation of the printing process compared to the expected shape (larger deformation fields are associated with larger printing errors), and it can be used in metrology or OPC modelling to detect failing regions, wrong measurements, or models. The problem of shape variability is thus re-formulated as the estimation of a deformation field between two surfaces or, equivalently, the registration of their height maps. The registration technique that we adopt is based on the non-linear elastic model of Guyader and Vese [2009].

In agreement with our variational formulation, we consider the registration problem as the minimization of a functional whose unknown is the displacement field  $\mathbf{u}$  and refer the reader to Modersitzki [2004] for an overview of other registration methods. Denoting as  $z_0$  the height map of the shape prior, image registration is performed by minimizing a functional that combines a measure of similarity between  $z$  and  $z_0$  and a penalty on the displacement field  $\mathbf{u}$ .

#### Proposed shape priors in optical lithography

If we can use OPC simulations to provide a prior model of the surface and its height map  $z_0$ , we can directly measure the distance between the surface and the expected OPC shape with



the  $L_2$ -norm. The term associated with the shape prior is then

$$B(z, \mathbf{u}) = \int_{\Omega} (z(\mathbf{x}) - z_0(\mathbf{x} + \mathbf{u}))^2. \quad (3.7)$$

If a reliable OPC model is not available, we must compare the reconstructed surface to the binary mask of the pattern provided in the circuit design. In this case the minimization of the energy is slower, and the reconstruction method completely neglects the optical and physical proximity effects modelled in OPC. To compare our surface to the mask, we need first to binarize it in order to have a meaningful measure. For simplicity, we threshold the surface at a fixed threshold  $\theta$  and compare it to the mask with the  $L_2$ -norm. The thresholding is written in terms of the Heaviside function  $H(z - \theta)$ , and the mask is described in terms of the signed distance function  $\phi$  to its contour as  $H(\phi)$ . The similarity to the layout is therefore measured with the following shape prior term

$$B(z, \mathbf{u}) = \int_{\Omega} (H(z(\mathbf{x}) - \theta) - H(\phi(\mathbf{x} + \mathbf{u})))^2. \quad (3.8)$$

We choose to describe the mask by means of the distance function to its contour in order to improve the numerical convergence associated with the registration of binary texture-less images.

#### Proposed deformation field on the shape prior

To regularize the displacement field we can either rely on physical models, such as elastic and viscous fluid models, or rely on image regularization techniques (see e.g., Fischer and Modersitzki [2003, 2004] for the regularization approach). In general, physical models are more interesting than image regularizers because they model real materials and follow our intuition. For this reason, we consider that the deformation of the mask or the OPC model follows a physical model, and we adopt energies from material sciences to penalize deformation fields  $\mathbf{u}$  that are non-elastic.

The concept of elasticity is incorporated in image registration by considering the template and reference images as observations of the same elastic body before and after being subjected to a deformation. Deformations of elastic bodies are then measured in terms of the Green-St. Venant strain tensor

$$\boldsymbol{\varepsilon} = \frac{\nabla \mathbf{u} + \nabla \mathbf{u}^T + \nabla \mathbf{u}^T \nabla \mathbf{u}}{2}. \quad (3.9)$$

As penalty for our deformation field  $\mathbf{u}$ , elastic models use the St. Venant-Kirchoff energy stored by the deformed material, which is given by

$$W(\boldsymbol{\varepsilon}(\mathbf{u})) = \int_{\Omega} \frac{\lambda}{2} (\text{tr } \boldsymbol{\varepsilon}(\mathbf{u}))^2 + \mu \text{tr } \boldsymbol{\varepsilon}(\mathbf{u})^2, \quad (3.10)$$

where  $\lambda$  and  $\mu$  are the first and second Lamé parameters characteristic of the material.

The first linear elastic model, proposed by Broit [1981], neglects the quadratic term in the strain tensor and the resulting model is fast to optimize but does not allow large deformations. To overcome this limitation, Christensen et al. [1996] proposes a viscous fluid model where objects evolve as fluids in accordance to Navier-Stokes equations. Compared to the elastic models, fluid equations perform regularization on the velocity field and allow larger deformations, but they are computationally more expensive. Recently, Guyader and Vese [2009] observed that we can obtain an elastic model comparable to the viscous model of Christensen et al. [1996] by simply considering the quadratic term in the strain tensor and defining Equation (3.10) as a penalty on the deformation field. The non linearity in this elastic model allows small strains even when the stress is relatively large and, consequently, large deformation between the shape obtained from the OPC model and the recovered surface or their respective masks.

### 3.3 Numerical minimization

The variational formulation of our surface reconstruction method results in the following non-convex minimization problem

$$\min_{z, \mathbf{u}} \alpha A(\nabla z) + \beta B(z, \mathbf{u}) + \gamma G(\nabla z) + W(\nabla \mathbf{u}). \quad (3.11)$$

To solve this optimization problem, we apply again variable-splitting and augmented Lagrangians to obtain an efficient and easy-to-code algorithm. As these techniques are analysed in detail in Chapter 4, the details of the proposed minimization algorithm are left for Appendix D, and here only the particularities of the SFS problem are analysed.

In this chapter we adopt a multi-resolution approach to speed up the surface reconstruction method and be more robust against local minima. Multi-resolution or scale-space tracking is commonly used in computer vision to solve large non-convex minimization problems. As the name suggests, scale-space tracking finds an approximate solution to the problem at a coarse scale and then tracks it through scale as it solves the problem at higher resolutions. At a large scale, the problem will not suffer from local minima and a first coarse solution is easily found. This solution is used to initialize the algorithm at smaller scales. As the scale is reduced, local minima appear in the minimization, and tracking of the initial solution guarantees that the solution is kept meaningful. In fact, multi-resolution is not certified to find the global minimum of a non-convex minimization problem, only a significant one that appears at large scales.

In terms of efficiency, the algorithm is designed to perform most of its iterations at a large scale, where it finds a coarse solution to the problem fast. As the scale is reduced, the algorithm is initialized closer to a minimum and requires less iterations to converge, which are now more expensive due to the finer resolution of the grid. Multi-resolution makes our method more

fast and robust to local minima.

### 3.4 Estimation of the reflectance map

In order to improve the approximation of the reflectance function by Equation (3.4), we alternatively iterate between the refinement of the reflectance map, in terms of the estimated reflectance coefficient  $\nu$  and the light source direction  $\mathbf{L}$ , and the refinement of the reconstructed surface.

Given a first estimate of the surface and its normal  $\mathbf{n}$ , we can refine the parameters of the reflectance map by solving two least squares minimization problems.

To estimate the value of the reflectance coefficient we assume that the light source direction is known, fixing its value to the previous estimate, and solve

$$\min_{\nu} \int_{\Omega} (\nu - I_0 \mathbf{n} \cdot \mathbf{L})^2. \quad (3.12)$$

This least-squares problem is analytically solved with  $\nu = \int_{\Omega} I_0 \mathbf{n} \cdot \mathbf{L}$ .

Similarly, in the estimation of  $\mathbf{L}$  we fix the value of  $\nu$  to the previous estimate, and solve the following minimization problem

$$\min_{|\mathbf{L}|=1} \int_{\Omega} \left( \frac{\nu}{I_0} - \mathbf{n} \cdot \mathbf{L} \right)^2. \quad (3.13)$$

The constraint  $|\mathbf{L}| = 1$  is not handled explicitly in the minimization by a closed-form solution; instead we adopt the same augmented Lagrangian technique used for the minimization of (3.11).

### 3.5 Experimental results

This section presents the experimental results of the proposed surface reconstruction method. We use large-field-of-view SEM images of circuits in the diffusion layer that cover an area of  $1.25 \times 1.25 \mu\text{m}$ . For simplicity, in all the experiments the parameters of the algorithm are set to  $\alpha = 500$ ,  $\beta = 50$ ,  $\gamma = 100$ ,  $\lambda = \mu = 10^{-4}$ .

The shape prior are based on the mask, that is on Equation (3.8), as reliable OPC models were not accessible. Convergence of the algorithm should be faster with a shape prior based on OPC models, but the reconstructed surface would only be more accurate to the same extent that the OPC models are. In this sense, it is always better to use limited but reliable mask information than unreliable OPC models. Figure 3.1 shows the masks associated with the patterns in the experiments.

A multi-resolution approach in  $z$  has been adopted to speed up convergence, adapting the

shape prior at each resolution with the solution of the previous scale. At each scale, therefore, the absolute deformation field can be computed concatenating the deformations at previous resolutions correctly scaled. Figure 3.3 shows the reconstructed surface for a simple structure at multiple scales, where the estimated deformation field to each scale prior is also included.

In our experiments we can only provide qualitative results of our reconstruction method because no ground truth for the reconstructed surfaces is currently available, but we observe that results agree with our intuitive reconstruction from the SEM images. Moreover, if we superpose the isoline of height  $\theta = 0.5z_{max} + 0.5z_{min}$  of the reconstructed surface in the SEM image in Figure 3.2, we observe that our reconstruction correctly estimates the position of the edges in the image.

To measure the contribution of the shape prior in the reconstruction, we compare the reconstructed surface of Figure 3.3 with the equivalent variational method ignoring the shape-prior, that is, defining  $z$  as the minimizer of  $\min_z \alpha A(\nabla z) + \gamma G(\nabla z)$ . Figure 3.7 presents the results of the reconstruction without the shape prior, where SFS alone fails due to noise, shadows, and SEM artifacts. For instance, the shadow present under the white band in the SEM image results in a deep and narrow valley in the surface reconstructed in Figure 3.7b.

Figure 3.4 presents the results of the reconstruction of more complex patterns from Figures 3.1b-3.1d. In order to test the robustness of our method to noise, we can repeat the experiments adding white Gaussian noise to the previous SEM images, see Figure 3.5. We observe that our SFS method is robust to noise. In particular, the only differences visible in the surfaces in comparison with Figures 3.3 and 3.4 are the gentle and almost negligible waves present in the dark flat areas. This waves are the only residual left in the reconstruction from the noise.

It is interesting to analyse the reconstructions of the surfaces in Figure 3.6, where the printing process partially fails. In the first surface, we observe that the height of the walls degrades in the vertical direction, while for the second surface we can see how the walls of the pattern become thinner and almost break. In both cases, the proposed method produces surfaces that agree with the human interpretation of the corresponding SEM images, and it can be used therefore to detect failures in the lithographic process.




Figure 3.1: Binary masks associated with different patterns.

Due to the lack of reliable ground truth data at the nanometre scale, we can only evaluate indirectly the performance of the proposed reconstruction method. Inspired by the evaluation methods of OPC models in optical lithography, we will compare 1D distance measurements




Figure 3.2: SEM image with associated mask 3.1a. In red we superposed the iso-contours of height  $\theta$  associated with the reconstructed surface of Figure 3.3.

associated with different patterns. In particular, the critical distance of OPC simulations will be compared with the corresponding distances measured from our reconstructed surfaces. To understand the next experiments, therefore, it is first necessary to understand some additional concepts from lithography.

One of the basic tasks of computational lithography is the quantitative characterization of the printing process, from patterns and masks in circuit designs to the shapes carved in silicon wafers. In particular, numerical simulations that model optical and physical properties of the light source and the material of the wafer are performed in order to estimate the shape of the surface with OPC models. For accurate results, these models are iteratively refined by calibration with measurements of the wafer in several cycles of printing-simulation. To that purpose, OPC modellers describe the patterns printed in a wafer in terms of critical distances, that is, 1D measurements of distances associated with the surface of the wafer at a critical position where the engineers expect that the printing process might fail. The next experiment analyses the critical distance associated with the reconstructed surface with its equivalent generated with a calibrated OPC model for 24 patterns that printed correctly. Both distances should coincide with an accurate SFS reconstruction, while wrong SFS reconstructions would lead to a mismatch of the pairs of critical distances. Consequently, we use this process as validation of the proposed reconstruction method.

For the OPC simulations, we use the compact resist model *CM1* of *Mentor Graphics* and run simulations with *Calibre* to estimate the critical distance associated with each one of the 24 patterns. In both simulations and SFS reconstruction, the critical distance is measured at the location of the gauge where the simulated and reconstructed surfaces attain height 0.3 with respect to a normalized surface, as suggested in Mack [2007]. Figure 3.8 shows the scatter plot between the critical distances from OPC simulations and from the proposed method. This experiment provides a quantitative evaluation of the proposed method, but it is only representative to the extent that the OPC model is reliable. We observe that the critical distances are nicely correlated (the difference in scaling is due to the pixel-to-nanometre conversion), except in a few outlier cases. We can conclude that the proposed reconstruction method agrees with the calibrated OPC estimates. Currently we cannot assess if the existing


differences are due to errors in the OPC models, inaccuracy of the SFS reconstruction method, or they simply show effects of the manufacturing process that are not considered in the OPC model. To clarify this point it is necessary to compare the surfaces estimated with the proposed SFS method with real and accurate 3D measurements on the wafer, possibly with atomic force microscopy. At present, this information is not available and the question remains open to future investigations.

## 3.6 Conclusions


This chapter has developed a method to reconstruct surfaces of silicon wafers from images of printed circuits taken with a scanning electron microscope. The proposed method is robust to the noise and the image artifacts characteristic of the electronic microscope because it incorporates knowledge about the shapes of the patterns in the reconstruction process.

To that purpose, the proposed variational formulation defines an intuitive objective functional that includes models of the image acquisition systems, the expected smoothness of the wafer, and the shape of the patterns in the circuit. The proposed technique also estimates a deformation field between the shape prior of the circuit design and the reconstructed surface; this field substitutes the unknown statistical shape models that are necessary to model the variability of the lithographic printing process in the shape prior. Consequently, this deformation field complements the information available from the reconstructed surface and can be directly used to detect failures and critical regions in the wafer.


Both outputs are in fact extremely important in optical lithography, where it is necessary to characterize the surface of silicon wafers and to quantify deviations in the manufacturing process. The proposed technique, therefore, can have an important impact on the chip manufacturing industry and opens a new line of research in shape-from-shading.




(a) Surface at largest scale.




(b) Deformation at largest scale.




(c) Zoom at middle scale.



(d) Deformation at middle scale.



(e) Zoom at smallest scale.



(f) Deformation at smallest scale.

Figure 3.3: Reconstructed surface and deformation fields at three different scales from SEM image of Figure 3.2.






Figure 3.4: First row, SEM images from patterns of Figure 3.1b, 3.1c and 3.1d (right). Second to fifth rows, reconstructed surfaces on the left and the deformation fields on the right.







(a) Noisy image, 1.85 dB SNR




(b) Surface from noisy image 3.5a




(c) Noisy image, 2.36 dB SNR.



(d) Surface from noisy image 3.5c



(e) Noisy image, 1.05 dB SNR



(f) Surface from noisy image 3.5e

Figure 3.5: Reconstructed surfaces from noisy images.




Figure 3.6: SEM images and reconstructed surface for patterns of Figure 3.1e and 3.1f. The lithographic process resulted in pinching in the center of the left's pattern and thin and fading the outer walls on the right pattern.

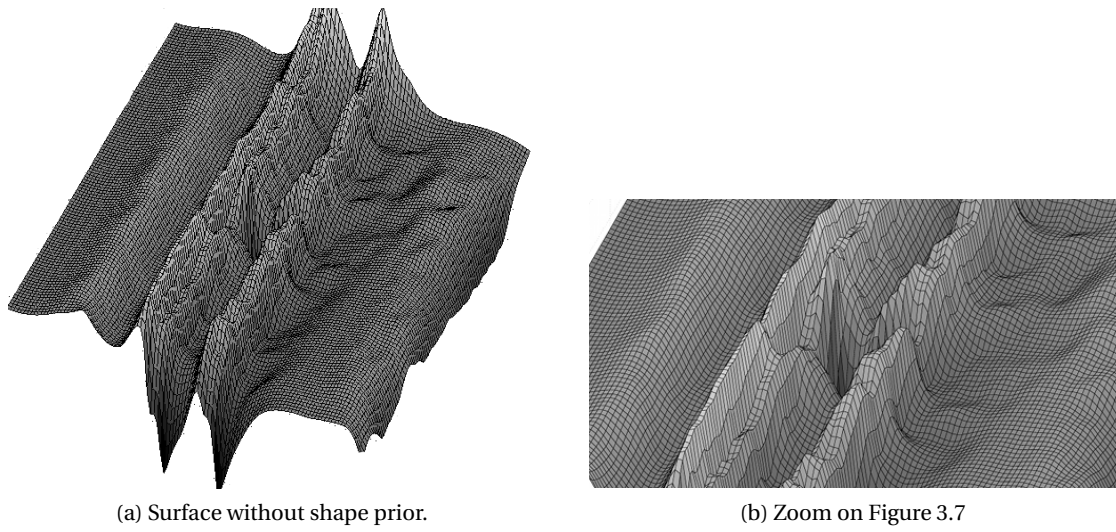


Figure 3.7: Reconstruction results for the SEM image of Figure 3.2 ignoring the shape prior in the SFS algorithm.




Figure 3.8: Scatter plot of the critical distances obtained with our SFS reconstruction method and OPC simulations.



## 4 An efficient algorithm for the level set method preserving distance functions

This chapter presents the techniques used in this dissertation to develop efficient minimization algorithms for different inverse problems. It is unnecessary, however, to analyse the algorithms developed for all the problems in order to understand the principles that guide these techniques, it is more constructive to present in detail a single algorithm. As a single sample, therefore, the minimization problem should be interesting and allows us to explore the issues of a non-convex problem with a non-differentiable functional. For this reason, this chapter presents the algorithm developed for the level set method when the level set function is constrained to be a signed distance function. The rest of the algorithms are left for the Appendices A-D.

The level set method of Osher and Sethian [1988] is a popular technique for tracking moving interfaces in many disciplines because of its high flexibility. The original level set method, however, is still limited by two numerical issues. First, it does not implicitly preserve the level set function as a distance function, which is necessary to accurately estimate geometric features like the curvature or the normal vectors to the contours. Second, the level set algorithm is slow because the time step of its PDE flow is limited by the CFL condition, which is essential to the numerical stability and accuracy of the iterative scheme. Hence, an efficient algorithm for this problem is necessary.

In this context, this chapter presents an algorithm that preserves distance functions in the level set method and avoids the previous two issues. The proposed algorithm makes use of variable-splitting and augmented Lagrangians to both overcome the CFL limitations and to naturally preserve the level set function as a distance function in the minimization; this avoids periodic re-distancing of the level set function and speeds the numerical algorithm. In particular, the algorithm proposed here is applied to image segmentation, where it proves to be 5 to 6 times faster than related state-of-the-art techniques, but it could be equally adapted to any level set method that requires the distance information.

---

The material of this chapter has been partially published in Estellers et al. [2012].

### 4.1 Algorithms for the level set method and distance functions

In the last twenty years, the level set method (LSM) of Osher and Sethian [1988] has become a popular numerical technique for tracking moving interfaces in a wide range of disciplines, from material sciences, to fluid mechanics or computer vision. The main reasons of its success are the high flexibility of the LSM to adapt to different problems, its ability to handle changes of topology like contour breaking and merging seamlessly, and the guarantee of the existence of solutions in the class of viscosity PDEs. Moreover, extensive numerical algorithms based on Hamilton-Jacobi equations have been developed, accurately handling shocks and providing stable numerical schemes for the level set method.

The key idea of the LSM is to implicitly represent a contour or interface as the zero level set of a higher dimensional function, called the level set function (LSF), and formulate the evolution of the contour through the evolution of the level set function. For closed contours, signed distance functions (SDFs) were originally adopted to represent level set functions because they directly provide stability and accuracy to the LSM (Osher and Fedkiw [2002], Osher and Paragios [2006]). In the last years, however, new approaches have proposed to use binary functions, rather than distance functions, to represent the level set function. This change of representation allows the use fast convex optimization techniques, like the graph cut method of Boykov et al. [2001] or the convex relaxation of Chan et al. [2006], but loose the distance information inherent to a SDF. Distance functions are still essential in many applications. For instance, in medical imaging SDFs are used in the segmentation of cortical surfaces of Section 4.5, they provide an efficient estimation of the normal vectors in the reconstruction of surfaces in Section 4.2, and are used in the creation of special effects in computer graphics. For these applications, and many others, it is necessary to develop a fast and accurate algorithm for distance preserving level set methods.

A major issue of distance preserving LSMs is the limited speed of the existing algorithms based on Hamilton-Jacobi equations and upwind schemes (Osher and Sethian [1988], Osher and Paragios [2006]). Two handicaps affect these iterative methods. First, the speed of the algorithms is limited by the CFL condition of Courant et al. [1967], which is necessary to guarantee the stability of the numerical evolution of Hamilton-Jacobi PDEs. Second, the objective functional of the LSM is hard to optimize because it usually involves a total variation penalty that is not differentiable, due to the  $\ell^1$  norm of the gradient, and must be approximated by a smooth function. This approximation does not provide an exact solution to the problem and significantly slows down the minimization process. This chapter will overcome these limitations by the use of efficient techniques in  $\ell^1$  optimization.

Another major issue of the LSM, pointed out by Gomes and Faugeras [1999], is a contradiction between the theory and the implementation when level set functions are represented by a signed distance function. Indeed, the level set method does not preserve the SDF during its evolution because SDFs are not solutions of the Hamilton-Jacobi PDEs of the LSM. Consequently, additional techniques are necessary to preserve the LSF as an SDF during contour

evolution, but no fully satisfying methods have been proposed so far. The most common approaches to fix this problem are periodic re-distancing of the LSF as an SDF, this is called re-distancing or re-initialization, or constraining the LSF to remain an SDF during the evolution of the contour.

Re-distancing is the most common approach. It consists in stopping the evolution of the LSF periodically and re-initializing it as an SDF while preserving the zero level set. This approach introduces the questions of when and how to re-initialize the LSF. It is hard to say when the re-distancing must be applied because there is a trade-off between speed and accuracy: each re-distancing procedure takes time, but the LSM will develop irregularities during the evolution without re-distancing. Similarly, the re-initialization can be performed with different LSM techniques, like the PDE method of Sussman et al. [1994] or the fast marching algorithms of Tsitsiklis [1995] or Sethian [1996]. The main issue with re-distancing, nevertheless, is the exact preservation of the zero level set during the re-distancing process, which might require interpolation to avoid shifting the moving interface to undesired positions.

The second approach aims at constraining the LSF to stay an SDF during the contour evolution, avoiding the previous re-distancing procedure altogether. Gomes and Faugeras [1999] introduces a new level set formulation that restricts the LSF to be an SDF by the use of three coupled PDEs, but the analysis of the existence of solutions and the numerical implementations of the resulting method are more difficult than the standard LSM. More recently, Li et al. [2005, 2010] proposed to add a penalty term in the level set energy to constrain the LSF to be close to an SDF. In this case, the resulting numerical algorithm is simple and efficient, but its time step is still restricted by the CFL condition, and the SDF property is only encouraged but not enforced. We will see that our method overcomes all of these limitations and provides an efficient way to constrain the LSF to be an SDF.

In image processing, the level set method was first introduced to carry out image segmentation in Kichenassamy et al. [1995], Caselles et al. [1997] and Chan and Vese [2001] and then extended to other tasks as stereo reconstruction, object tracking, and object recognition in Faugeras and Keriven [1997], Paragios and Deriche [2000] and Leventon et al. [2000]. In this chapter we focus on image segmentation and surface reconstruction, but the proposed method can be easily extended to other problems.

In image segmentation, the level set method re-formulates the parametric active contour model into a non-parametric minimization problem that is independent of the contour discretization. To this purpose, the active contour is described by the zero level set of the LSF, and the evolution equation of the active contour is applied to the LSF, instead of only to the active contour, in order to drive its zero level set to the edges of the image. This level set formulation easily includes edge-, region-, and shape-based terms in the objective functional that defines the segmentation criterion, which makes the LSM highly flexible to adapt to different tasks and explains its extensive use in imaging. Similar segmentation models, for instance, have been applied to the problem of reconstructing a surface from unorganized data points in Zhao [2000]. In this case, the implicit representation of the surface by an SDF offers the following

two advantages: it avoids complex 3D parametrizations that require prior knowledge of the topology of the surface, and it provides an easy and reliable way to directly estimate surface normals from the level set function. This advantages will be explored in Section 4.2.

In this chapter, we develop a fast and accurate algorithm for the level set method that constrains the LSF to remain an SDF via constrained minimization. This results in a minimization problem that is hard to solve directly, and we therefore propose to split it into sub-optimization problems that are easier to solve, and combine them together using an augmented Lagrangian. This idea is borrowed from the split-Bregman method proposed by Goldstein and Osher [2009] for the  $\ell^1$  minimization of compressed sensing, but the same principles guide the method of alternating direction of multipliers proposed by Lions and Mercier [1979], Glowinski and Le Tallec [1989], Bioucas-dias and Figueiredo [2007] and Setzer [2011]. As a result, the proposed algorithm holds three important properties. First, it is fast because it is not limited by the CFL condition. Indeed, the splitting strategy and its reformulation into a constrained minimization allows us to deal with the non-differentiability of TV and go beyond the CFL time-step restriction. Second, our algorithm preserves the LSF as an SDF, avoiding the classical re-distancing problem and providing desirable properties for applications where the distance information is necessary. In surface reconstruction from unorganized points, for instance, the normal vectors to the surface can be fast and reliably estimated from an SDF, and in medical imaging the distance information can be exploited to include topological restrictions on human anatomy in the segmentation of cortical surfaces. Finally, our algorithm is fast and easy to implement because each step is based on standard minimization problems. The rest of this chapter will analyse in detail these properties and techniques.

In particular, Section 4.2 reviews very briefly the level set method applied to image segmentation and surface reconstruction. Sections 4.3 and 4.4 present the proposed algorithm, and Section 4.5 describes the numerical experiments and results. Finally conclusions are drawn in Section 4.6.

### 4.2 Image segmentation and surface reconstruction with the LSM

The level set method can be applied to perform image segmentation using the active contour method. The segmentation problem is then defined as the following minimization problem with respect to a contour  $\mathcal{C}$ :

$$\min_{\mathcal{C} \subset \Omega} \int_{\mathcal{C}} w_b(s) \, ds + \int_{C_{in}} w_r^{in}(x) \, dx + \int_{C_{out}} w_r^{out}(x) \, dx, \quad (4.1)$$

where  $\Omega$  is the image domain,  $s$  is the arc-length parametrization of the contour  $\mathcal{C}$ , and  $C_{in}, C_{out}$  designate the regions inside and outside the contour. The first term of (4.1) weights the length of  $\mathcal{C}$  by an edge detector  $w_b$ , as proposed in Kichenassamy et al. [1995] and Caselles et al. [1997]. The functions  $w_r^{in}, w_r^{out}$  are region terms that measure the homogeneity of the different regions in the segmentation, like e.g., the model of Chan and Vese [2001]. We adopt



this minimization model for image segmentation because it represents a large class of active contour published in the literature.

Similar models have been used to reconstruct smooth surfaces from unorganized data points in Hoppe et al. [1992] and Zhao [2000]. Given a set of noisy points  $\{x_i\}_{1 \leq i \leq N}$  lying close to the unknown surface, a smooth estimate of the surface can be reconstructed by minimizing the area of the surface weighted by the distance to the set of points, i.e.,  $w_b(x) = \min_{1 \leq i \leq N} |x - x_i|$ . We propose here a model which also includes a region term  $w_r^{in}$  to improve the quality of reconstruction with sparse data sets. In particular, we adapt the model proposed in Lempitsky and Boykov [2007] and align the normals of the reconstructed surface to the precomputed normal vectors at the data points  $\{n_i\}_{1 \leq i \leq N}$ . In our variational formulation, however, first a semi-dense vector field of the surface normal  $\hat{n}$  must be estimated everywhere in space, and then the unknown surface can be reconstructed maximizing the alignment of its normal  $\mathcal{N}$  with the estimated normal field  $\hat{n}$ .

The proposed region term that maximizes the alignment of normals can be identified as  $w_r^{in} = -\operatorname{div} \hat{n}$  by making use of the divergence theorem. Indeed,

$$\arg \max_{\mathcal{C} \subset \Omega} \int_{\mathcal{C}} \hat{n} \cdot \mathcal{N} \, ds = \arg \min_{\mathcal{C} \subset \Omega} \int_{C_{in}} -\operatorname{div} \hat{n} \, dx, \quad (4.2)$$

where  $\mathcal{C}$  represents here the unknown surface. Consequently, it is only necessary to estimate the flux of the semi-dense normal field at the point  $x$ , which can be done with the following formula proposed in Ye et al. [2010]

$$\operatorname{div} \hat{n}(x) = \sum_{1 \leq i \leq N} \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{|x-x_i|^2}{2\sigma^2}} \langle x - x_i, n_i \rangle, \quad (4.3)$$

where  $n_i$  is the surface normal estimated at  $x_i$ . Unlike the method of Lempitsky and Boykov, the method proposed here does not need to know *a priori* the normals to the surface at the data points. Our method, instead, reliably estimates these normals from the signed distance function during the reconstruction process by  $n_i = \nabla \phi(x_i)$ , where  $\phi$  is the signed distance function. To sum up, for the reconstruction of a surface from unorganized data points, the functions in problem (4.1) are re-defined as  $w_b(x) = \min_{1 \leq i \leq N} |x - x_i|$ ,  $w_r^{in} = -\operatorname{div} \hat{n}$ , and  $w_r^{out} = 0$ .

The level set method can be applied to solve problem (4.1) by re-writing the minimization in terms of a level set function  $\phi$  as follows:

$$\min_{\phi} \int_{\Omega} w_b(x) |\nabla H(\phi)| + w_r(x) H(\phi) \quad \text{subject to} \quad |\nabla \phi| = 1, \quad (4.4)$$

where the  $\mathcal{C} \subset \mathbb{R}^n$  is described by the zero level set of  $\phi$ , and  $H$  is the Heaviside function. The dimension of the embedding space and the level set interface depend on the application, with  $n = 2$  for image segmentation and  $n = 3$  for surface reconstruction. The additional constraint  $|\nabla \phi| = 1$  guarantees that the level set function is a signed distance function (Sussman et al. [1994]).

### 4.3 Efficient level set algorithm preserving signed distance functions

In this section, we introduce an efficient algorithm to solve the level set minimization problem (4.4) and preserve a signed distance function. The main idea is to split the original hard problem (4.4) into sub-optimization problems that are well-known and easy to solve, and combine them together using an augmented Lagrangian. This idea is borrowed from the split-Bregman method of Goldstein and Osher [2009], which is an efficient optimization method for  $\ell^1$ -norm problems recently introduced in image processing for compressed sensing recovery.

Let us consider the following constrained minimization problem, which is equivalent to the original LSM problem (4.4):

$$\min_{\phi, \varphi, u, \mathbf{q}, \mathbf{p}} \int_{\Omega} w_b(x)|q(x)| + w_r(x)u(x) \quad \text{subject to} \quad \begin{cases} u = H(\varphi) \\ \varphi = \phi \\ \mathbf{q} = \nabla u \\ \mathbf{p} = \nabla \phi \\ |\mathbf{p}| = 1 \end{cases}, \quad (4.5)$$

where we have introduced additional variables, the functions  $(\varphi, u)$  and the  $n$ -dimensional vector fields  $(\mathbf{q}, \mathbf{p})$ , to split the original objective functional into simpler independent terms, as a result this step is known as variable-splitting. The proposed splitting approach makes the original problem (4.4) easier to solve because (4.5) can better handle the non-differentiability and non-linearity of (4.4). Indeed, it is known from Wang et al. [2008], Goldstein and Osher [2009] that the minimization of the TV term  $|\nabla \phi|$  can be carried out efficiently by decoupling the  $\ell^1$  norm and the gradient operator  $\nabla$ ; the term  $|\nabla \phi|$  in (4.4) is thus replaced by  $|\mathbf{p}|$ , and  $\mathbf{p} = \nabla \phi$  in (4.5). Similarly, the variable  $\varphi = \phi$  is introduced to handle the non-linear term  $u = H(\phi)$ , and the term  $|\nabla H(\phi)|$  is rewritten as  $|q|$ , with  $q = \nabla u$ .

Next, we want to reformulate this constrained minimization problem as an unconstrained optimization task. This can be done with the augmented Lagrangian approach of Glowinski and Le Tallec [1989], which translates the constraints into pairs of a Lagrangian multiplier and a penalty term. Let us define the augmented Lagrangian energy associated with (4.5):

$$\begin{aligned} \mathcal{L}(\phi, \varphi, u, \mathbf{q}, \mathbf{p}, \Lambda) = & \int_{\Omega} w_b|q| + w_r u + \lambda_1(\varphi - \phi) + \frac{r_1}{2}(\varphi - \phi)^2 + \lambda_2(u - H(\varphi)) + \frac{r_2}{2}(u - H(\varphi))^2 \\ & + \lambda_3 \cdot (\mathbf{q} - \nabla u) + \frac{r_3}{2}|\mathbf{q} - \nabla u|^2 + \lambda_4 \cdot (\mathbf{p} - \nabla \phi) + \frac{r_4}{2}|\mathbf{p} - \nabla \phi|^2 \quad \text{subject to} \quad |\mathbf{p}| = 1 \end{aligned} \quad (4.6)$$

where  $\Lambda = (\lambda_1, \lambda_2, \lambda_3, \lambda_4)$  are the Lagrangian functions, and  $r_1, \dots, r_4$  are positive constants. The constrained minimization problem (4.5) reduces then to finding the saddle-point of the augmented Lagrangian energy  $\mathcal{L}$ , and its solution can be approximated by the iterative Algorithm 1, see Glowinski and Le Tallec [1989].

We start by initializing  $\phi^0$  and  $\varphi^0$  with the signed distance function to the initial con-

**Algorithm 1** Augmented Lagrangian method for distance preserving level set methods

---

- 1: Initialize  $\phi, \varphi, u, \mathbf{q}, \mathbf{p}, \Lambda$
- 2: Find a minimizer of  $\mathcal{L}$  with respect to variables  $(\phi, \varphi, u, \mathbf{q}, \mathbf{p})$  with fixed Lagrange multipliers  $\Lambda^{k-1}$ :

$$(\phi^k, \varphi^k, u^k, \mathbf{q}^k, \mathbf{p}^k) = \underset{\substack{\phi, \varphi, u, \mathbf{q}, \mathbf{p} \\ |\mathbf{p}|=1}}{\operatorname{argmin}} \mathcal{L}(\phi, \varphi, u, \mathbf{q}, \mathbf{p}, \Lambda^{k-1}) \quad (4.7)$$

- 3: Update Lagrange multipliers

$$\lambda_1^k = \lambda_1^{k-1} + r_1(\varphi^k - \phi^k) \quad (4.8)$$

$$\lambda_2^k = \lambda_2^{k-1} + r_2(u^k - H(\varphi^k)) \quad (4.9)$$

$$\lambda_3^k = \lambda_3^{k-1} + r_3(\mathbf{q}^k - \nabla u^k) \quad (4.10)$$

$$\lambda_4^k = \lambda_4^{k-1} + r_4(\mathbf{p}^k - \nabla \phi^k) \quad (4.11)$$

- 4: Stop the iterative process when  $\|\phi^k - \phi^{k-1}\|_2 < \epsilon$ .
- 

tour,  $u^0 = H(\phi_0)$ ,  $\mathbf{q}^0 = \nabla u^0$ ,  $\mathbf{p}^0 = \nabla \phi^0$ , and the Lagrange multipliers to zero. At each iteration, an alternating minimization method is used to find an approximate minimizer of  $\mathcal{L}(\phi, \varphi, u, \mathbf{q}, \mathbf{p}, \Lambda^{k-1})$  with respect to the variables  $\phi, \varphi, u, \mathbf{q}$  and  $\mathbf{p}$ . Then the Lagrange multipliers are updated with the residual associated with each constraint, and the process is repeated until the change of the level set function  $\phi$  falls below a certain threshold  $\epsilon$ , which happens at convergence.

In general, it is difficult to find the exact minimizer of the minimization problem (4.7) because the functional (4.6) is not convex with respect to the variables  $\phi, \varphi, u, \mathbf{q}$  and  $\mathbf{p}$ , and there is no guarantee of convergence to a global minimum. However, experiments show that a good approximation can be found by the alternating direction method of multipliers, see Glowinski and Le Tallec [1989]. An approximate solution is thus computed by iteratively alternating the minimization of  $\mathcal{L}(\phi, \varphi, u, \mathbf{q}, \mathbf{p}, \Lambda^{k-1})$  with respect to each variable while considering the others fixed. This leads to Algorithm 2.

The next step is to determine the solutions of the five sub-minimization problems (4.12) - (4.16), which can actually be computed efficiently.

## 4.4 Efficient solution to sub-minimization problems

In this section, we simplify notation by omitting the super-index and the tilde symbol in the sub-minimization problems (4.12)-(4.16).

In our implementation, we discretize the image domain  $\Omega \subset \mathbb{R}^2$  with a regular grid of size  $n = n_x \times n_y$ , and use forward differences to compute the discrete gradients and backward differences for the divergence in order to preserve the adjoint relationship  $\operatorname{div} = -\nabla^*$  in the

---

**Algorithm 2** Alternate minimization for an approximate solution of (4.7)

---

- 1: Initialize  $\tilde{\phi}^0 = \phi^{k-1}, \tilde{\varphi}^0 = \varphi^{k-1}, \tilde{u}^0 = u^{k-1}, \tilde{\mathbf{q}}^0 = \mathbf{q}^{k-1}, \tilde{\mathbf{p}}^0 = \mathbf{p}^{k-1}$ .
- 2: For  $l = 1, \dots, L$  and fixed Lagrange multipliers  $\Lambda^k$ , solve the following sub-problems alternatively:

$$\tilde{\phi}^l = \arg \min_{\phi} \mathcal{L}(\phi, \tilde{\varphi}^{l-1}, \tilde{u}^{l-1}, \tilde{\mathbf{q}}^{l-1}, \tilde{\mathbf{p}}^{l-1}, \Lambda^{l-1}) \quad (4.12)$$

$$\tilde{\varphi}^l = \arg \min_{\varphi} \mathcal{L}(\tilde{\phi}^l, \varphi, \tilde{u}^{l-1}, \tilde{\mathbf{q}}^{l-1}, \tilde{\mathbf{p}}^{l-1}, \Lambda^{l-1}) \quad (4.13)$$

$$\tilde{u}^l = \arg \min_u \mathcal{L}(\tilde{\phi}^l, \tilde{\varphi}^l, u, \tilde{\mathbf{q}}^{l-1}, \tilde{\mathbf{p}}^{l-1}, \Lambda^{l-1}) \quad (4.14)$$

$$\tilde{\mathbf{q}}^l = \arg \min_{\mathbf{q}} \mathcal{L}(\tilde{\phi}^l, \tilde{\varphi}^l, \tilde{u}^l, \mathbf{q}, \tilde{\mathbf{p}}^{l-1}, \Lambda^{l-1}) \quad (4.15)$$

$$\tilde{\mathbf{p}}^l = \arg \min_{\mathbf{p}} \mathcal{L}(\tilde{\phi}^l, \tilde{\varphi}^l, \tilde{u}^l, \tilde{\mathbf{q}}^l, \mathbf{p}, \Lambda^{l-1}) \quad \text{subject to} \quad |\mathbf{p}| = 1 \quad (4.16)$$

- 3: Set  $(\phi^k, \varphi^k, u^k, \mathbf{q}^k, \mathbf{p}^k) = (\tilde{\phi}^L, \tilde{\varphi}^L, \tilde{u}^L, \tilde{\mathbf{q}}^L, \tilde{\mathbf{p}}^L)$ .
- 

discrete setting. In matrix-vector notation, we efficiently compute the spatial derivatives multiplying the discretized functions arranged as a column vector with the sparse finite difference matrices  $\nabla_x \phi = D_x \phi$ ,  $\nabla_y \phi = D_y \phi$ . Similarly, the discretization of the  $L_2$  inner product in  $\Omega$  corresponds to the usual dot product of vectors.

#### 4.4.1 Sub-minimization with respect to $\phi$ and $u$

The sub-minimization problems (4.12) and (4.14) can be written as follows:

$$\min_{\phi} \int_{\Omega} \frac{r_1}{2} \left( \phi - \left( \varphi + \frac{\lambda_1}{r_1} \right) \right)^2 + \frac{r_4}{2} \left| \nabla \phi - \left( \mathbf{p} + \frac{\lambda_4}{r_4} \right) \right|^2 \quad (4.17)$$

$$\min_u \int_{\Omega} w_r u + \frac{r_2}{2} \left( u - \left( H(\varphi) - \frac{\lambda_2}{r_2} \right) \right)^2 + \frac{r_3}{2} \left| \nabla u - \left( \mathbf{q} + \frac{\lambda_3}{r_3} \right) \right|^2 \quad (4.18)$$

The Euler-Lagrange equation of (4.17) is

$$(-r_4 \Delta + r_1) \phi = -r_4 \operatorname{div} \mathbf{p} - \operatorname{div} \lambda_4 + r_1 \left( \varphi + \frac{\lambda_1}{r_1} \right), \quad (4.19)$$

and can be solved efficiently by the fast Fourier transform (FFT) as proposed in Wang et al. [2008]. Indeed, with the vector notation explained above, Equation (4.19) is discretized as a

linear system of equations  $A\phi = b$ , with matrix  $A$  and vector  $b$  given by

$$\begin{aligned} A &= r_1 I_n + r_4 D_x^T D_x + r_4 D_y^T D_y \\ b &= -r_4 D_x^T \mathbf{p}_x - r_4 D_y^T \mathbf{p}_y + r_1 \left( \phi + \frac{\lambda_1}{r_1} \right). \end{aligned}$$

Matrix  $A$  is symmetric, definite positive, and block-circulant, and we can use the Fourier transform  $\mathcal{F}$  to decompose it as  $A = \mathcal{F}^T D \mathcal{F}$ , with  $D$  a diagonal matrix. Consequently, the system  $A\phi = b$  can easily be solved in the Fourier domain. In practice we use the FFT transform instead of doing the matrix multiplications with  $\mathcal{F}$  and  $\mathcal{F}^T$ , which gives us the following solution  $\phi = \mathcal{F}^T (D^{-1} \mathcal{F} b)$  of complexity  $\mathcal{O}(n \log n)$

It is now straightforward to apply the same procedure to compute the minimizer  $u^*$  of (4.18), whose Euler-Lagrange equation is

$$(-r_3 \Delta + r_2)u = -w_r - r_3 \operatorname{div} \mathbf{q} + \operatorname{div} \lambda_3 + r_2 \left( H(\phi) - \frac{\lambda_2}{r_2} \right). \quad (4.20)$$

#### 4.4.2 Sub-minimization with respect to $\phi$

The sub-minimization problem (4.13) can be written as follows:

$$\min_{\phi} \int_{\Omega} \frac{r_1}{2} \left( \phi - \left( \phi - \frac{\lambda_1}{r_1} \right) \right)^2 + \frac{r_2}{2} \left( H(\phi) - \left( u + \frac{\lambda_2}{r_2} \right) \right)^2 \quad (4.21)$$

Let us call  $z = \phi - \frac{\lambda_1}{r_1}$ ,  $v = u + \frac{\lambda_2}{r_2}$ , and observe that the minimization is decoupled for each pixel. The function that we need to minimize pixel-wise is now

$$F(\phi) = \frac{r_1}{2} (\phi - z)^2 + \frac{r_2}{2} (H_{\epsilon}(\phi) - v)^2. \quad (4.22)$$

Observe that, for practical implementations, the minimization problem (4.22) involves a smooth approximation  $H_{\epsilon}$  of the Heaviside function. For this reason, we propose a two-step method to quickly find a minimizer of (4.22).

*Step 1:* Find a solution  $\phi^0$  of (4.22) for  $\epsilon = 0$ , i.e., for the distributional/non-smooth Heaviside function. A closed-form solution exists for this problem and can be computed as follows. The first term of (4.22) is minimized by  $\phi^0 = z$ . As the distributional Heaviside function can take only values 0 or 1, the second term of (4.22) is minimized by either  $\phi^0 < 0$  when  $v < \frac{1}{2}$  or by  $\phi^0 \geq 0$  when  $v \geq \frac{1}{2}$ . This means that both terms are minimized by  $\phi^0 = z$  when  $v < \frac{1}{2}$  and  $z < 0$  or when  $v \geq \frac{1}{2}$  and  $z \geq 0$ . Otherwise we must choose to minimize the greater of these terms, and set  $\phi^0 = 0$  if  $E(0) < E(z)$  and  $\phi^0 = z$  otherwise.

*Step 2:* Find a solution  $\phi^0$  of (4.22) for  $\epsilon > 0$  using the standard Newton's method with  $\phi^0$

as initialization. The iterative Newton's method for finding the minimizer of (4.22) reads

$$\varphi^{m+1} = \varphi^m - \frac{F'(\varphi^m)}{F''(\varphi^m)}, \quad (4.23)$$

with  $\varphi^{m=0} = \varphi_0$ , in our case it gives

$$\varphi^{m+1} = \varphi^m - \frac{r_1(\varphi^m - z) + r_2(H_\varepsilon(\varphi^m) - v)\delta_\varepsilon(\varphi^m)}{r_1 + r_2(H_\varepsilon(\varphi^m) - v)\delta'_\varepsilon(\varphi^m) + r_2\delta_\varepsilon^2(\varphi^m)},$$

where  $\delta_\varepsilon$  and  $\delta'_\varepsilon$  are smooth approximations of the Dirac function and its derivative.

Newton's method presents quadratic convergence close to the minimizer, which intuitively means that the number of correct digits doubles in every step. With the correct initialization of the first step, therefore, usually two iterations are enough to find a good minimizer of (4.21), and the resulting technique is almost as fast as a closed-form solution.

#### 4.4.3 Sub-minimization with respect to $\mathbf{q}$

The sub-minimization problem (4.15) can be written as follows:

$$\min_{\mathbf{q}} \int_{\Omega} w_b |\mathbf{q}| + \frac{r_3}{2} \left| \mathbf{q} - \left( \nabla u - \frac{\lambda_3}{r_3} \right) \right|^2 \quad (4.24)$$

Let us call  $\mathbf{z} = \nabla u - \frac{\lambda_3}{r_3}$ , observe that the minimization (4.24) is decoupled for each pixel, and the problem reduce to pixel-wise minimization of the following function in the plane

$$F_q(\mathbf{q}) = |\mathbf{q}| + \frac{r_3}{2w_b} |\mathbf{q} - \mathbf{z}|^2. \quad (4.25)$$




Figure 4.1: Graphical explanation of shrinkage operator

The first term in  $F_q$  only depends on the norm of  $\mathbf{q}$ , and therefore the orientation of the minimizer is determined by the second term  $|\mathbf{q} - \mathbf{z}|^2$ . It is clear from Figure 4.1 that this minimizer has the form  $\mathbf{q}^* = \beta \mathbf{z}$ , for some  $\beta \geq 0$ . The minimization of  $F_q$  then reduces to the

minimization of the quadratic function  $F_\beta(\beta) = F_q(\beta\mathbf{z}) = \beta|\mathbf{z}| + (1 - \beta)^2 \frac{r_3}{2w_b} |\mathbf{z}|^2$  with respect to the scalar parameter  $\beta$ . By simple differentiation with respect to  $\beta$  we obtain  $\beta^* = 1 - \frac{w_b}{r_3|\mathbf{z}|}$ , and the constraint  $\beta \geq 0$  leads to the following solution:

$$\mathbf{q}^* = \max\left\{|\mathbf{z}| - \frac{w_b}{r_3}, 0\right\} \frac{\mathbf{z}}{|\mathbf{z}|}. \quad (4.26)$$

This minimizer corresponds, in fact, to soft-thresholding with the shrinkage operator of Donoho [1995], for which we simply provide a graphical explanation.

#### 4.4.4 Sub-minimization with respect to $\mathbf{p}$

The sub-minimization problem (4.16) can be written as follows:

$$\min_{\mathbf{p}} \int_{\Omega} \frac{r_4}{2} \left| \mathbf{p} - \left( \nabla\phi - \frac{\lambda_4}{r_4} \right) \right|^2 \quad \text{subject to} \quad |\mathbf{p}| = 1 \quad (4.27)$$

Let us call  $\mathbf{z} = \nabla\phi - \frac{\lambda_4}{r_4}$ , and observe again that the minimization of (4.27) is decoupled for each pixel. The problem is equivalent to the pixel-wise minimization of the following function on the plane  $F(\mathbf{p}) = \frac{r_4}{2} |\mathbf{p} - \mathbf{z}|^2 = \frac{r_4}{2} (|\mathbf{p}|^2 - |\mathbf{z}|^2 - 2\mathbf{p} \cdot \mathbf{z})$ .

We can introduce the constraint in  $|\mathbf{p}| = 1$  in  $F$  by writing it  $F(\mathbf{p}) = \frac{r_4}{2} (1 - |\mathbf{z}|^2 - 2\mathbf{p} \cdot \mathbf{z})$ . It is obvious that the minimizer of  $F$  is of the form  $\mathbf{p}^* = \frac{\mathbf{z}}{|\mathbf{z}|}$ . Indeed, when vectors  $\mathbf{z}$  and  $\mathbf{p}$  have the same orientation, the scalar product  $\mathbf{p} \cdot \mathbf{z}$  reaches a maximum, and the constraint  $|\mathbf{p}| = 1$  is verified. The minimizer  $\mathbf{p}^*$  of (4.27) is then given by

$$\mathbf{p}^* = \frac{r_4}{|r_4 \nabla\phi - \lambda_4|} \nabla\phi - \frac{1}{|r_4 \nabla\phi - \lambda_4|} \lambda_4. \quad (4.28)$$

## 4.5 Experiments and discussions

To analyse the properties of the proposed algorithm, the experimental section is organized in three different parts. Sections 4.5.1 presents the results of the proposed level set algorithm for image segmentation and surface reconstruction, where our method is accurate and fast. Section 4.5.2 shows the importance of preserving the distance function for the correct segmentation of medical images. Finally, Section 4.5.3 compares our algorithm with existing level set methods that preserve the signed distance function, analysing and comparing their performance in terms of speed and accuracy.

### 4.5.1 Image segmentation and surface reconstruction

We apply our algorithm for image segmentation by defining the edge and region terms proposed in Caselles et al. [1997] and Chan and Vese [2001], which have been extended to handle

color images. Figure 4.2 shows the results obtained for different images from the Berkeley<sup>1</sup>, Weizmann<sup>2</sup> and GrabCut<sup>3</sup> databases. The method behaves as expected, providing the same results as re-distancing or the penalty methods in terms of the final segmentation, but with a considerable speed-up in time.

We also use the proposed model to successfully reconstruct several surfaces from the Stanford dataset<sup>4</sup>. We initialize the method with a sphere containing all the data points and recompute the region term  $w_r^{in}$  every 5 iterations as described in Section 4.2. Unlike the models of Lempitsky and Boykov [2007], we do not need to estimate *a priori* the surface normal at the data points, instead the normal to the surface at the points is directly estimated from the current value of the level set function. Finally, the process is accelerated with a standard multi-resolution approach, see Figure 4.3. The final reconstructed surfaces are shown in figures 4.4 and 4.5.

### 4.5.2 Cortex segmentation with coupled surfaces

Inspired by Gomes and Faugeras [1999], we develop a segmentation algorithm for the cortical layer with two active contours coupled by their relative distance, which constrains the thickness of the cerebral cortex to the expected topology of the brain. Graph cut methods or convex relaxation techniques cannot be directly applied to solve this problem because they use binary function to represent the contour and loose the distance information, which is necessary for a successful segmentation of the cortex.

The cerebral cortex is the layer of the brain bounded by the outer and inner cortical surfaces, that is, the outer interface between cerebral spinal fluid and grey matter, and the inner interface between the grey and the white matter. Locating this cortical surface is a first step in many brain imaging tasks, and measuring its thickness is a common procedure in the diagnosis of many neurological diseases. We will see that the use of SDFs in the segmentation of the cortical surface allows the inclusion of information about the cortical structure into the segmentation problem and, at the same time, provides an estimate of the cortical thickness.

In order to extract the cortical layer we need to extract its two bounding surfaces  $C^1$  and  $C^2$ . In theory, therefore, we could segment the regions defined by the white matter and the exterior parts of the brain by independent minimization of the following functionals associated with

---

<sup>1</sup><http://www.eecs.berkeley.edu/vision>

<sup>2</sup><http://www.wisdom.weizmann.ac.il/vision>

<sup>3</sup><http://research.microsoft.com/en-us/um/cambridge/projects/visionimagevideoediting/segmentation/grabcut.htm>

<sup>4</sup><http://graphics.stanford.edu/data/3Dscanrep/>





Figure 4.2: Proposed level set-based segmentation method applied to natural images. For each result, we show the segmentation result and the level set function. We plot the initial zero level set in blue, the final contour in pink and the  $\pm 1, \pm 2$  level sets of the final function in black.




Figure 4.3: Reconstructed bunny from the Stanford dataset at different resolutions. Linear interpolation of the SDF obtained at lower resolutions was used as initial LSF for higher resolutions.

$C^1$  and  $C^2$

$$\min_{C^1} \mathcal{E}(w_b, w_{r1}, C^1) = \int_{C^1} w_b(s) ds + \int_{C_{in}^1} w_{r1}(x) dx, \quad (4.29)$$

$$\min_{C^2} \mathcal{E}(w_b, w_{r2}, C^2) = \int_{C^1} w_b(s) ds + \int_{C_{in}^2} w_{r2}(x) dx, \quad (4.30)$$

where  $w_b$  is a boundary indicator, and  $w_{r1}$ ,  $w_{r2}$  are region descriptors for the white matter and the exterior areas of the brain. In practice, however, the boundaries between grey and white matter are not clear, MRI images suffer from intensity inhomogeneities, and the segmentations obtained with local region descriptors and boundary detectors do not correctly locate the cortical layer.

To overcome these limitations and obtain meaningful segmentations, it is necessary to introduce a distance constraint in the segmentation model that incorporates information about brain anatomy. For simplicity, we modify the model proposed in Zeng et al. [1999], but other models for cortex segmentation could be adopted, see e.g., MacDonald et al. [2000], Goldenberg et al. [2002]. We use a coupled surface model, where a functional is minimized when  $C^1$  captures the interface between cerebral spinal fluid and grey matter,  $C^2$  the grey-white matter boundary, and the distance between them is close to the expected cortical thickness  $d$  (about 3 mm). The problem is written in terms of the LSM, making use of the SDF  $\phi_1$  and  $\phi_2$  to define the bounding surfaces  $C^1$  and  $C^1$ . The functional to minimize is then given by

$$\min_{\phi_1, \phi_2} \mathcal{E}(w_b, w_{r1}, \phi_1) + \mathcal{E}(w_b, w_{r2}, \phi_2) + \frac{c}{2} \int_{\Omega} (\phi_1 - \phi_2 - d)^2 \quad \text{subject to} \quad \begin{cases} |\nabla \phi_1| = 1 \\ |\nabla \phi_2| = 1 \end{cases} \quad (4.31)$$

The term  $(\phi_1 - \phi_2 - d)^2$  penalizes segmentations where the distance between the bounding surfaces differs from the expected cortical thickness  $d$ . Indeed, when  $\phi_1$  and  $\phi_2$  are the SDFs

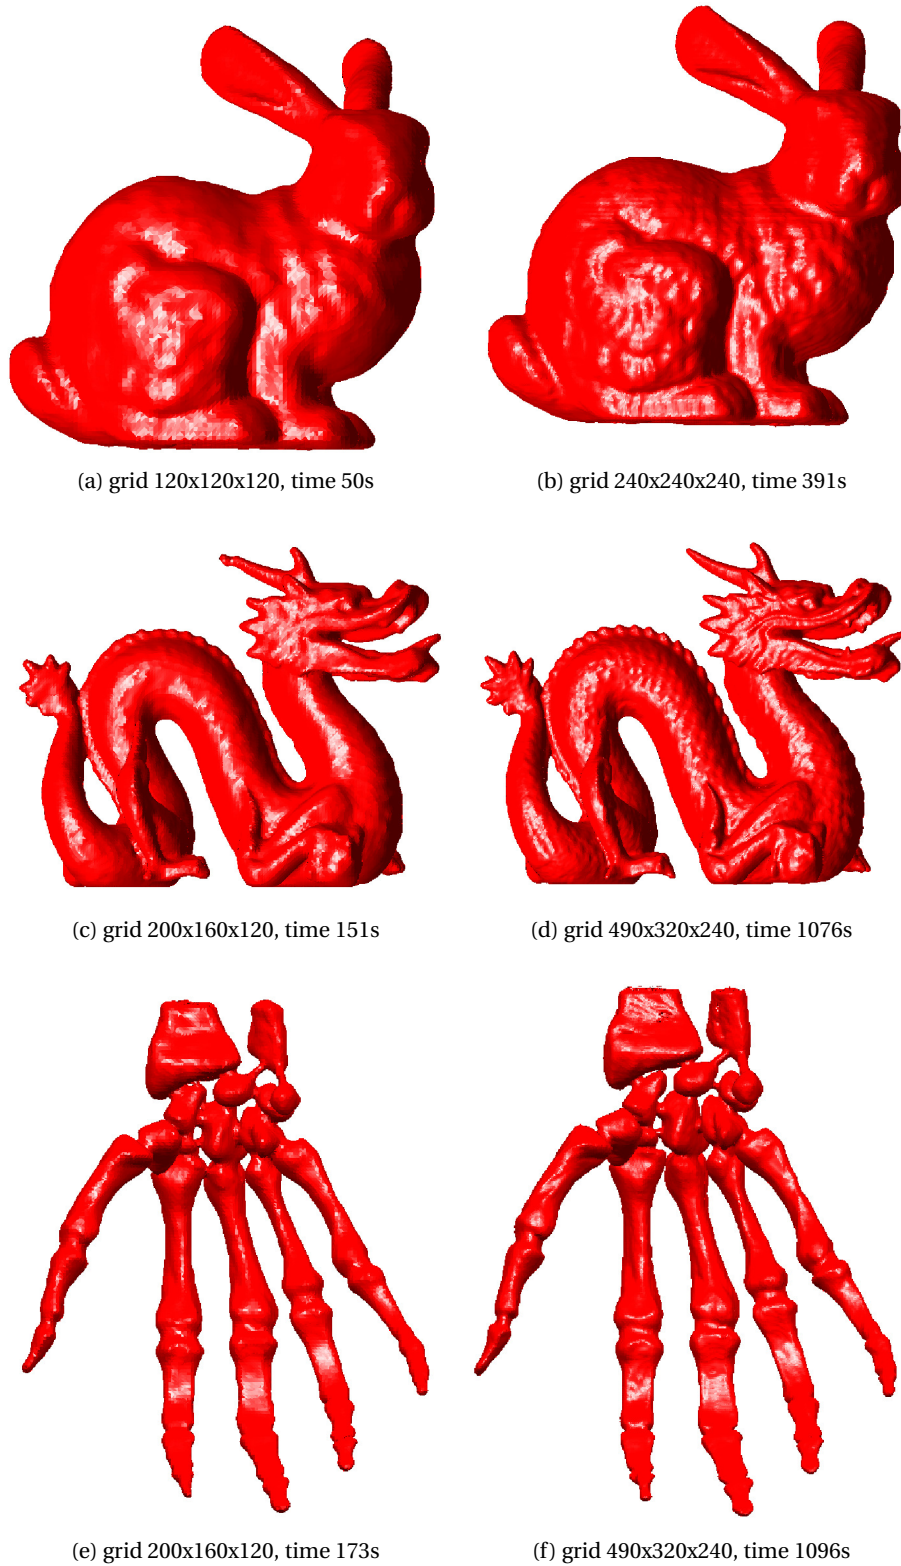


Figure 4.4: Reconstruction of the Stanford bunny, dragon, and hand surfaces from scattered data points at different resolutions.




Figure 4.5: Reconstruction of the Stanford Budha and horse surfaces from scattered data points at different resolutions.




Figure 4.6: Segmentation of grey-white matter interface on MRI images of human cortex. The segmentation obtained with the proposed method 4.6b is clearly closer to the ground truth 4.6d than the results obtained when no coupling terms is considered and the segmentation is performed independently for the inner and outer cortical surfaces 4.6c. The segmentation 4.6e obtained with 3 phases of the k-mean algorithm fails because region descriptors alone cannot segment grey-matter.

defined by  $C^1$  and  $C^2$ , the distance between the surfaces can be measured on the whole domain by  $\phi_1 - \phi_2$  and, consequently, the term  $(\phi_1 - \phi_2 - d)^2$  drives the segmentation to solutions where the distance between the surfaces is consistent with the cortical structure.

The minimization technique presented in Section 4.4 can be directly applied to this problem. The same variable-splitting and Lagrange multipliers are now defined and solved for each level set function independently, and only the alternate minimization with respect to  $\phi_1$  and  $\phi_2$  are modified. The minimization problems with respect to  $\phi_1$  (the analogous applies to  $\phi_2$ ) is now

$$\min_{\phi_1} \int_{\Omega} \frac{r_1}{2} \left( \phi_1 - \left( \phi_1 + \frac{\lambda_{1,1}}{r_1} \right) \right)^2 + \frac{c}{2} (\phi_1 - \phi_2 - d)^2 + \frac{r_4}{2} \left| \nabla \phi_1 - \left( \mathbf{p}_1 + \frac{\lambda_{1,4}}{r_4} \right) \right|^2, \quad (4.32)$$

and its associated Euler-Lagrange equation

$$(-r_4 \Delta + r_1 + c) \phi_1 = -r_4 \operatorname{div} \mathbf{p} - \operatorname{div} \lambda_4 + r_1 \left( \phi + \frac{\lambda_1}{r_1} \right) + c (\phi_2 + d) \quad (4.33)$$

can be solved efficiently by the fast Fourier transform as we have already explained in Section 4.4.

We have applied this technique in an illustrative experiment to segment the cortical layer in different slices of MRI images. Results are shown in Figure 4.6. We observe that coupling the level set functions and constraining the expected cortical thickness in the segmentation produces a segmentation in Figure 4.6b close to the ground truth, while the segmentation obtained without the coupling in Figure 4.6c is not able to follow the fine structures of the brain.



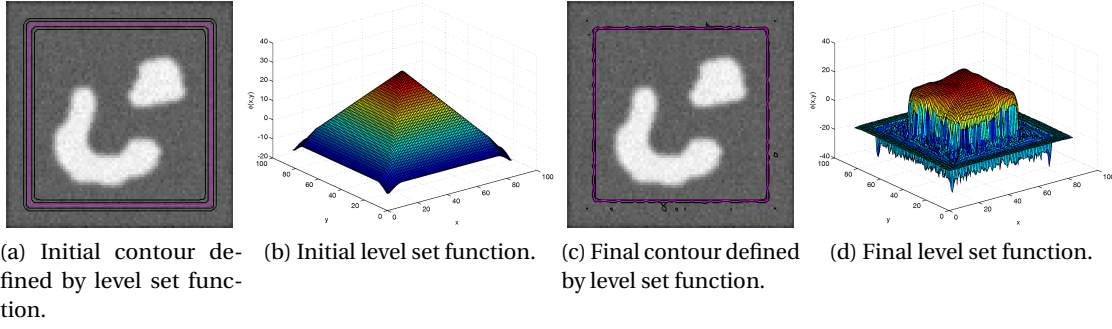


Figure 4.7: The level set function ( $\phi = 0$  in pink and  $\phi = \pm 1, \pm 2$  in dark) develops irregularities during the propagation and stops after a few iterations if the LSF is not constrained to be (or at least close to) an SDF.

### 4.5.3 Comparison with other level set methods preserving distance functions

In this section, we compare our algorithm to the other techniques designed to preserve the SDF in the LSM in the context of image segmentation. To that purpose, we first remind the necessity of preserving the LSF as an SDF with a simple example. We initialize the LSF as an SDF and evolve it with the PDE flow of the original LSM for the image segmentation problem (4.4). The level set function becomes too steep around its zero level set after a few iterations and stops its evolution because the geometric features of curvature and normals are not correctly estimated, see Figure 4.7. Two common approaches have been introduced to overcome this problem: periodically re-distancing the LSF, or maintaining the SDF during the evolution of the contours.

We will compare three different techniques: the standard re-distancing approach and two methods that preserve the SDF in the evolution of the contours, our method and the methods of Li et al. [2005, 2010]. The re-distancing process is carried out with the fast marching method of Adalsteinsson and Sethian [1999], while the methods of Li et al. are defined by introducing a penalty energy  $P$  in the objective functional to encourage the LSF to be close to an SDF.

The methods of Li et al., to be precise, are defined by the following minimization problem

$$\min_{\phi} \int_{\Omega} w_b |\nabla H(\phi)| + w_r H(\phi) + \frac{\mu}{2} P(|\nabla \phi|), \quad (4.34)$$

where the penalty  $P(x)$  is a potential function with a minimum at  $x = 1$ . Li et al. [2005] defines the potential  $P(x) = (x-1)^2$  penalizing deviations from SDF and encouraging the SDF property  $|\nabla \phi| = 1$  on the whole domain. This penalty is modified in Li et al. [2010] to encourage the SDF property only close to the zero level set and a flat level set function elsewhere. This is accomplished re-defining the penalty  $P$  as a double-well potential with two unique equal minima at  $x = 1$  and  $x = 0$ . In both cases, Li et al. solve the minimization problem with a descent PDE flow, with a time step limited by the value of  $\mu$  according to the CFL condition.

Figures 4.8h-4.8l show the re-distancing method of Adalsteinsson and Sethian [1999]. Although the final segmentation and LSF provide the desired results, the periodic re-initialization procedures produce non-smooth minimizations that result in jumps in the objective functional that appear each time we re-initialize the level set function. These annoying jumps appear because re-distancing is not guaranteed to exactly preserve the location of the zero level set of the level set function, this produces small abrupt changes in the location of the contour and, consequently, in the value of the objective functional. Besides, we do not know in general when to re-distance the LSF and must choose heuristically; in our experiments we re-distance every 5 iterations because re-distancing more often slows the algorithm, and re-distancing less often causes irregularities in the level set function. Other heuristics are possible and equally arguable.

Next, we consider the method of Li et al., which is closer to the method developed in this chapter. Li et al. introduces a penalty in the minimization problem that encourages the LSF to be close to a distance function. A penalty term, however, does not constrain exactly the LSF to be an SDF, and the value of the penalty constant  $\mu$  in (4.34) is limited by a trade-off between speed and accuracy. Indeed, a large value of  $\mu$  in Figures 4.8r-4.8v preserves the LSF close to a distance function, but it slows down significantly the minimization process. On the other hand, a small  $\mu$  speeds up convergence, but the LSF differs from an SDF and presents peaks and valleys close to the zero level set; this effect is shown in Figures 4.8m-4.8q, where the small island-like contours of the  $\pm 1, 2$  level sets (in black) might lead to instabilities that affect the location of the zero level contours (in pink). To avoid this issue, the penalty function is re-defined as a double-well potential in Li et al. [2010], encouraging the LSF to be an SDF close to the zero level set and flat elsewhere. The resulting method avoids the instabilities associated with the method of Li et al. [2005] when a small  $\mu$  is used, but it does not really provide a distance function. In terms of speed, the method of Li et al. [2010] is as fast as the original method of Li et al. [2005], and it does not develop potential instabilities. Both methods, however, are still limited by the CFL condition and are slower than the method proposed here.

Our method overcomes the limitations of the LSM of Li et al. [2005, 2010] because our formulation constrains the LSF to be an SDF, and the proposed algorithm is fast because the splitting strategy allows us to exploit data sparsity and closed-form solutions to develop efficient minimizations. Indeed, Figures 4.8q, 4.8v, and 4.8g show that our method keeps more faithfully the LSF as an SDF than Li et al.'s method, and results in a penalty energy  $\int_{\Omega} (|\nabla \phi| - 1)^2$  smaller by at least one order of magnitude. Observe that the minimum of this penalty is around 1 with our method, around 5 with Li et al.'s method with a large  $\mu$ , and around 100 with Li et al.'s method with a small value of  $\mu$ . In other words, our method is almost as accurate as re-distancing when it comes to preserving the SDF and overall faster than the other two methods. These advantages are provided by the augmented Lagrangian approach, which can preserve accurately the constraint while keeping a good minimization speed.

To quantify the improvement obtained with our method with respect to the other distance

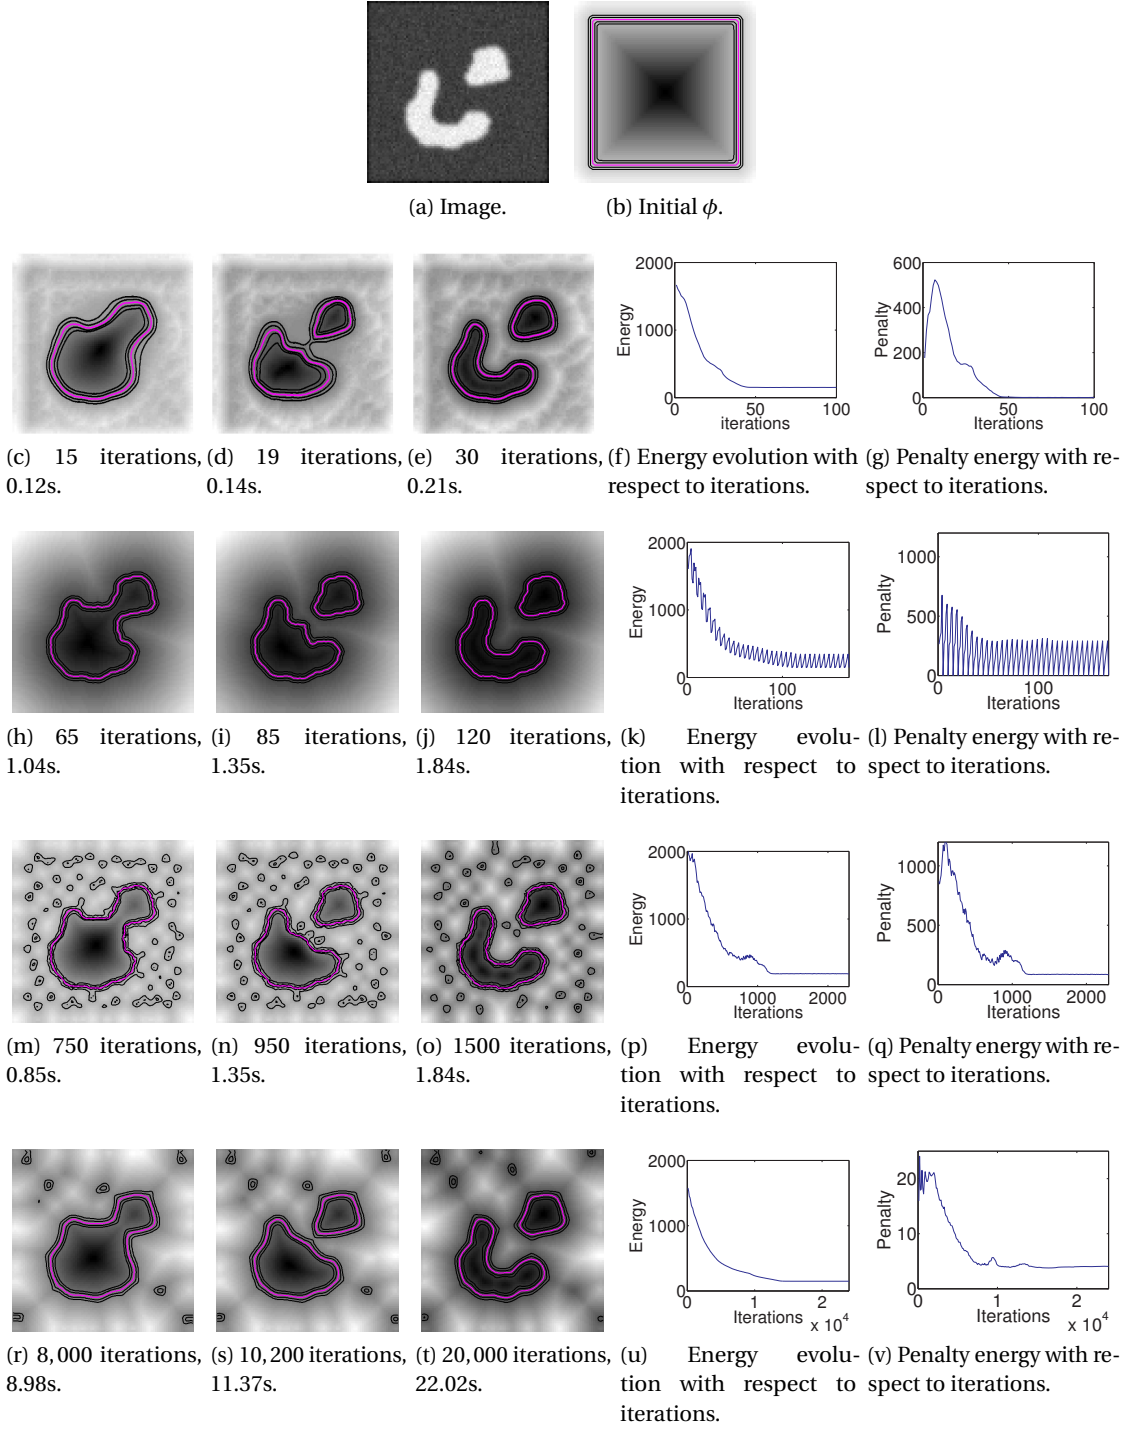


Figure 4.8: Comparison of the evolution of the level set function  $\phi$ : 2nd row, our algorithm; 3rd row, re-distancing of Adalsteinsson and Sethian [1999]; 4th row algorithm of Li et al. [2005] with a big penalty, and 5th row with small penalty. Each row shows the evolution of the LSF  $\phi$  at intermediate and final iterations ( $\phi = 0$  in magenta,  $\phi = \pm 1, \pm 2$  in black), the evolution of energy  $\int_{\Omega} w_b |\nabla H(\phi)| + w_r H(\phi) + \frac{1}{2} (|\nabla \phi| - 1)^2$  and penalty  $\int (|\nabla \phi| - 1)^2$ . Our method provides the best trade-off between speed and preservation of the distance function.



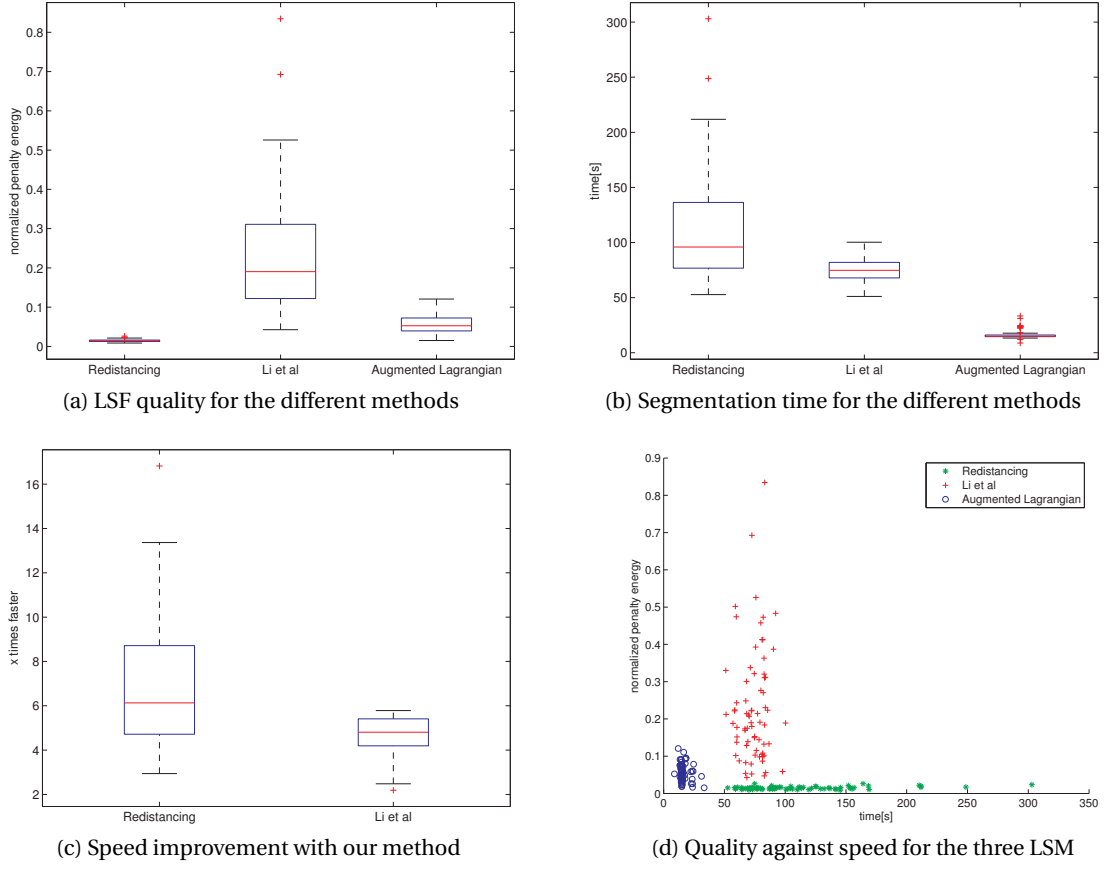



Figure 4.9: Comparison of quality of segmentation and speed for the different methods on a dataset of 72 images. Quality of the LSF is measured in terms of the penalty term at convergence  $\frac{1}{|\Omega|} \int_{\Omega} (|\nabla \phi| - 1)^2$  when the obtained contours are equivalent. Our method preserves the SDF almost as well as re-distancing, clearly better than Li et al.’s method, and is faster than any of them.

preserving LSM techniques, we have used the previous algorithms to segment 72 images from the Berkeley, Weizmann and GrabCut databases. For Li et al.’s method, we have chosen a relatively small penalty  $\mu$  that allows large time steps. We compare in Figure 4.9 the different algorithms in terms of the quality of preserving the SDF and speed. Figure 4.9a presents the values of the penalty  $\int_{\Omega} (|\nabla \phi| - 1)^2$  at convergence for all the images, showing that our method preserves the SDF almost as well as re-distancing and clearly better than Li et al.. Indeed, the penalty values for our method and re-distancing are similar, both being an order of magnitude smaller than with Li et al.’s method. We also compare the time for each method to reach convergence, which is assumed when  $\frac{\|\phi - \phi^*\|_2}{\|\phi^*\|_2} < 10^{-4}$ . We observe in Figures 4.9b and 4.9c that our algorithm is on average 5 to 6 times faster than the other two methods. Finally, we present in Figure 4.9d a scatter plot of the penalties obtained at convergence against the time required to converge, which shows that our method presents a good trade-off between accuracy of SDF preservation and speed.



(a) Redistancing: 5.2% misclassification in 17s. (b) Li et al.'s [2005]: 5.8% misclassification in 4s. (c) Li et al.'s [2010]: 6.2% misclassification in 4s. (d) Our method: 5.8% misclassification in 3.8s.

Figure 4.10: Segmentation of grey-white matter interface on MRI images of human cortex. The segmentation model (4.31) is solved with different approaches: re-distancing in Figure 4.10a, with the methods of Li et al. [2005] in Figure 4.10b, with the method of Li et al. [2010] in Figure 4.10c, and with the proposed algorithm in Figure 4.10d.

Finally, in order to compare the different methods in terms of accuracy of segmentation, we have tested the surface coupling of Section 4.5.2 with the previous level set methods for a synthetic image. In Figure 4.10 we show the resulting segmentations with the different methods, and measure their performance in terms of grey-white matter tissue classification. While re-distancing obtains the best performance with 5.24% of errors, it is the slower option with 17 s. The two methods of Li et al.'s take 4 s to converge, but Li et al. [2005] performs better than Li et al. [2010] because it encourages an SDF everywhere, not only close to the desired contours, and considers the topology of the brain over the whole image domain. Finally, our method provides the best trade-off between speed and accuracy as it takes 3.8 s to converge and produces less errors than any of Li et al. [2005, 2010]'s methods.

## 4.6 Conclusions

This chapter has introduced an efficient algorithm for the level set method that overcomes the main numerical limitations of the original level set algorithm, that is, the speed and the preservation of the signed distance function. The gain in speed is due to the splitting and augmented Lagrangian techniques used in the minimization, which decompose the original problem into sub-optimization tasks that are easier to solve and have closed-form solutions. Preserving the distance function, which is essential in many applications, is achieved through a constraint in the minimization problem that is handled by the augmented Lagrangian and merges seamlessly with the rest of the algorithm.

Extended numerical experiments show the three main properties of the algorithm: it is fast, it preserves the level set function as a signed distance function, and it provides the best trade-off between speed and accuracy of distance function preservation among the state-of-the-art.

Due to the wide use of level set functions, this technique can have an impact in a large community. In this dissertation, however, it only presents the numerical techniques used to solve the minimization problems efficiently. The underlying principle behind all these algorithms, dual variables, and equations, is beautifully simple: divide and conquer.



## 5 Conclusions and Future Work

This dissertation has explored the use of geometry in variational methods, proposing new solutions for different inverse problems in imaging that improve state-of-the-art techniques.

### 5.1 Conclusions

Geometry, in ancient Greek, literally means measuring the Earth; geometry originated from the necessity to model the physical world in an organized manner, and its concepts and constructions reflect how humans understand the universe. It is only natural, then, that when we try to capture, reconstruct, and analyse the physical world from two-dimensional images, geometry plays an important role in the models and techniques we develop.

In the context of variational methods and inverse problems in imaging, geometry goes a step further than regularization; it does not only assume that an image or a curve is smooth, but that it obeys certain properties characteristic of the physical world. These additional assumptions lead to more accurate variational models, but also to more complex minimization problems. Two logic questions arise at this point: when and how should we include geometry in a variational model?

Occam's razor<sup>1</sup> can guide the answer to our first question: "Everything should be made as simple as possible, but not simpler". In our case, we can use the razor to bias the choice of our inversion methods: when two competing variational models provide the same solutions, or qualitatively equivalent ones, the simpler model of the two is preferable. For instance, the use of the popular total variation regularization usually results in variational methods with convex functionals; in terms of minimization, then, our problems are over<sup>2</sup>. Indeed, convex problems have unique minimizers, or multiple equally good ones, necessary optimality conditions are

---

<sup>1</sup>This interpretation of Occam's razor is attributed to Einstein, in defence of his theory of relativity in *The Meaning of Relativity*.

<sup>2</sup>This idea is nicely expressed by Rockafellar: "...in fact, the great watershed in optimization isn't between linearity and non-linearity, but convexity and non-convexity". A detailed explanation of the implications of convexity in minimization can be found in Boyd and Vandenberghe [2004].

sufficient, and there is a large collection of numerical methods to efficiently solve them. The restriction to convex functionals, of course, limits the form of the variational energy and the accuracy of the underlying model, but convex models are usually enough for simple images and even desirable to avoid over-fitting. When additional accuracy is required or the structure of the data is more complex, then geometry comes into play.

The second question, the *how*, offers only an unhappy answer: *it depends*. It depends on the inverse problem under consideration, and it depends on the nature of the object being imaged. Geometry is an additional prior for the reconstruction of images in compressive sensing in Chapter 1, it is exploited in the coupling of both measurements (the image) and the solution (contour) in the segmentation method of Chapter 2, and it is intrinsic to the measurements in the shape-from-shading problem of Chapter 3. In this dissertation, therefore, a natural answer to the *how* requires a stroll through its chapters.

### 5.1.1 On Geometric Models

Chapter 1 proposed a recovery method for compressed sensing that iteratively aligns the normal vectors of the level curves of the image with the reconstructed signal. In compressed sensing, these normal vectors are buried in a reduced set of measurements and cannot be directly exploited in the recovery model. For this reason, the proposed technique alternates between the estimation of the normals to the level set curves of the image, and reconstruction of an image that matches the estimated normals, the compressed sensing measurements, and the sparsity constraint of total variation. This two-step procedure reduces the problem to a series of convex minimizations that can be efficiently solved, but it lacks a proof of convergence and requires an heuristic stopping criterion. The geometric alignment term of the normals is also extended to non-local operators in graphs for the recovery of textured images, where it exploits the geometric properties of images in terms of their non-local graphs. The inclusion of this alignment term as an image prior improves signal recovery in two ways: it is more robust to noise and a reduced number of measurements, and it is able to recover sharp edges as well as smoothly varying regions, overcoming the staircase effect of total variation.

The hamonic active contours developed in Chapter 2 interpret each possible segmentation of an image as a 2-dimensional manifold embedded in a higher dimensional space. With the metric induced by the proposed embedding, a minimal surface manifold corresponds to a segmentation that minimizes the usual region, edge and regularity terms of an active contour model, and it exploits the directional information of the gradients of the image and the level set function to align the contours of the segmentation with the edges of the image. The proposed formulation, moreover, naturally takes into account the coupling of the different image channels, and it is generalized to higher dimensional spaces and non-flat images where usual segmentation methods cannot be applied. As a result, the proposed harmonic active contour is able to segment grey-scale and color images with strong inhomogeneities and fragmented edges.

Chapter 3 investigated the problem of shape-from-shading for a very particular application, the recovery of the surface of a silicon wafer from an image of a printed circuit taken with a scanning electron microscope. In this case, geometry is part of the problem, not of the solution. Indeed, the irradiance equation characteristic to shape-from-shading is purely geometric because it relates the orientation of the unknown surface, in terms of its normal vectors, with the shading observed in the image. This geometric problem is ill-posed, and prior knowledge is necessary to obtain a correct solution; in this case, it takes the form of a shape-prior and a deformation field that relate the pattern in the circuit design with the reconstructed surface. The result is a robust shape-from-shading method that exploits prior knowledge from circuit design to reconstruct surfaces of silicon wafers and measures the deformation field associated with the manufacturing process. Above all, this chapter is a proof-of-concept that opens the doors to new investigations on shape-from-shading in the context of optical lithography.

### 5.1.2 On Variational Methods and Minimization

A geometric variational model is, *per se*, a minimization problem; as such, it requires a solution. Variational problems only have closed-form solutions in a few cases, and usually we have to develop numerical algorithms that find a solution for each instance of the given data. In the same way that simplifying the algebraic expression of a closed-form solution speeds its computation and reduces numerical errors, developing an efficient algorithm reduces the computational cost and errors associated with a numerical minimization problem. If it is an accepted convention that closed-form expressions should be simplified, then efficient numerical algorithms should be equally developed for each problem. This is the goal of Chapter 4: to present the techniques used in this dissertation to design efficient numerical algorithms for the previous variational models.

In particular, Chapter 4 develops an efficient algorithm for the level set method that overcomes the main numerical limitations of the original level set algorithm, that is, the speed and the preservation of the signed distance function. The distance function is preserved by the introduction of an explicit constraint in the minimization problem, while the gain in speed is due to the variable-splitting and augmented Lagrangian techniques used in the minimization. These techniques introduce additional variables, constraints, and Lagrange multipliers in the original minimization problem and decompose it into sub-optimization tasks that are easier to solve. The underlying principle behind the resulting algorithm, with the dual variables, update rules, and sub-minimizations, is beautifully simple and convincing: divide and conquer.

## 5.2 Future Work

The logic continuation of Chapter 1 would try to answer two open questions. First, how can we choose a meaningful stopping criterion for the proposed iterative two-step method? The discrepancy principle of the inverse scale space proposed in Burger et al. [2006] offers a

possible answer to this question and a starting point for future investigations. Second, how can we regularize the non-local normals? The main difficulty lies on the fact that non-local normals do not correspond to standard vector fields: they have a different number of components for each pixel, and the direction associated with each component depends on the relative position of the non-local neighbours to the pixel. A direct answer involves the second order non-local derivatives of Gilboa and Osher [2008], but threatens to be computationally intractable. An ad-hoc option is to consider the non-local gradient at each pixel as a function in the circle, with the angular coordinates given by relative position of the non-local neighbours to the pixel. With this interpretation, spherical harmonics can be used to obtain a vector representation of the non-local gradients that is easily regularized with standard techniques.

The strength of the harmonic active contours of Chapter 2 lies on the natural treatment of multi-channel images and the generalization to non-flat spaces. Therefore, two possible extensions of this work are the application of the proposed technique to texture segmentation and spherical images. In this case, the segmentation techniques proposed by Sagiv et al. [2006] and Bogdanova et al. [2007] could serve as kick-off points.

The shape-from-shading technique of Chapter 3 could be improved in several ways. First, the approximation of the reflectance map could be refined to account for proximity effects by introducing simulations of electron-sample interactions. The resulting reconstruction method would be more accurate by providing a more faithful model of the data acquisition system, but it would also be computationally more expensive. Second, the estimation of the deformation field could be substituted by a statistical shape model of the pattern being imaged, but to obtain such a model multiple measurements of the surface associated with each pattern are required. In this case, the reconstruction method would gain on speed, but the construction of a shape model by atomic force microscopy would be extremely expensive.

### Ending note

It would be difficult, and certainly out of my power, to trace a more general path for future work on the use of geometry in image processing; this dissertation reflects only a personal expedition on this vast field, and my conclusions are limited to the experiences of each chapter. Some of them, however, can be summarized in the following quote from another journey through space and geometry:

*“Distress not yourself if you cannot at first understand the deeper mysteries of Spaceland. By degrees they will dawn upon you.”*

Edwin Abbott Abbott, *Flatland: A Romance of Many Dimensions*, 1884.



# A Minimization Algorithm of Chapter 1

## Local Recovery Method

In this appendix, we provide the details of the minimization algorithm of Chapter 1 for the local recovery method proposed for compressed sensing.

### A.0.1 Notation

We discretize the image domain  $\Omega \subset \mathbb{R}^2$  with a regular grid of size  $n = n_x \times n_y$ . In  $\Omega$  we consider the images as scalar functions with  $u(i) \in \mathbb{R}$ , for each pixel  $i \in \Omega$ , and their gradients as vector-valued functions with  $\nabla u(i) \in \mathbb{R}^2$ . We use forward differences to compute the discrete gradients and backward differences for the divergence in order to preserve the adjoint relationship  $\text{div} = -\nabla^*$  in the discrete setting.

The discrete TV regularizer is then given by

$$J(u) = \sum_i |\nabla u(i)| = \sum_i \sqrt{\nabla_x u(i)^2 + \nabla_y u(i)^2}, \quad (\text{A.1})$$

where we denote the pixel-wise norm of vectors as  $|\mathbf{d}|(i) = \sqrt{d_x^2(i) + d_y^2(i)}$ . Our discretized TV is then the  $\ell_1$ -norm of the function computing the pixel-wise norm of the gradient, i.e.,  $J(u) = \|\nabla u\|_1$ .

For vector fields  $\mathbf{d} = (d_x, d_y)$ , we discretize the TV regularizer as follows

$$J(d_x, d_y) = \sum_i \sqrt{|\nabla d_x(i)|^2 + |\nabla d_y(i)|^2}. \quad (\text{A.2})$$

In that case we observe that it corresponds to the  $\ell_1$ -norm of the function computing the pixel-wise norm of the vector of combined gradients, i.e.,  $J(d_x, d_y) = \|\nabla \mathbf{d}\|_1$ . With that observation it is easy then to extend it to a weighted TV regularizer as  $J_w(d_x, d_y) = \|W \nabla \mathbf{d}\|_1$ , where  $W$  is a diagonal matrix with the weights associated to each pixel.

In the vector notation used in CS, we can efficiently compute the spatial derivatives multiplying the discrete functions arranged as a column vector with the sparse finite difference

matrices  $\nabla_x u = D_x u$ ,  $\nabla_y u = D_y u$ . Similarly, the discretization of the  $L_2$  inner product in  $\Omega$  corresponds to the usual dot product of vectors, i.e.  $\langle v, u \rangle = v^T u$ .

## A.1 Estimate $u$ matching measurements, sparsity prior, and normals

To reconstruct the image we need to solve the following convex minimization problem:

$$\min_u \|\nabla u\|_1 + \gamma v^T u + \frac{\alpha}{2} \|Au - f\|_2^2 \quad \text{with } v = \text{div } n. \quad (\text{A.3})$$

We propose an iterative algorithm to solve (A.3) based on splitting and constraint minimization techniques. The main idea is to split the original problem into sub-optimization problems which are well-known and easy to solve, and combine them together using an augmented Lagrangian. The proposed algorithm is guaranteed to converge thanks to the convexity of (A.3).

Let us consider the following constrained minimization problem, which is equivalent to (1.11):

$$\min_{u, \mathbf{d}} \|\mathbf{d}\|_1 + v^T u + \frac{\alpha}{2} \|Au - f\|_2^2 \quad \text{subject to } \mathbf{d} = \nabla u. \quad (\text{A.4})$$

The proposed splitting approach makes the original problem (1.11) easier to solve because (A.4) decouples the  $\ell_1$ -norm and the gradient operator  $\nabla$ .

Next, we reformulate this constrained minimization problem as an unconstrained optimization task. This can be done with an augmented Lagrangian approach, which translates the constraints into pairs of a Lagrangian multiplier and a penalty term. Let us define the augmented Lagrangian energy associated to (A.4):

$$\begin{aligned} \mathcal{L}_1(u, \mathbf{d}, \boldsymbol{\lambda}) = & \|\mathbf{d}\|_1 + v^T u + \frac{\alpha}{2} \|Au - f\|_2^2 + \lambda_x^T (d_x - D_x u) + \lambda_y^T (d_y - D_y u) \\ & + \frac{r}{2} \|d_x - D_x u\|_2^2 + \frac{r}{2} \|d_y - D_y u\|_2^2 \end{aligned} \quad (\text{A.5})$$

where  $\boldsymbol{\lambda} = (\lambda_x, \lambda_y)$  are the Lagrange multipliers and  $r$  is a positive constant. The constraint minimization problem (A.4) reduces to finding the saddle-point of the augmented Lagrangian energy  $\mathcal{L}_1$ . The solution to the saddle point problem (A.5) can be approximated iteratively by the following algorithm: initialize the variables and Lagrange multipliers to zero; at each iteration find an approximate minimizer of  $\mathcal{L}_1(u, \mathbf{d}, \boldsymbol{\lambda}^{k-1})$  with respect to the variables  $u, \mathbf{d}$ , and update the Lagrange multipliers with the residuals associated to the constraints; stop the process when  $u$  remains fix. As the Lagrangian  $\mathcal{L}_1$  is convex with respect to  $u, \mathbf{d}$ , we can find a minimizer by alternating the minimization with respect to each variable. The resulting method is equivalent to the alternating direction method of multipliers. The iterative method is summarized in Algorithm 8

The next step is to determine the solutions of the two sub-minimization problems (A.6), (A.7),

### A.1. Estimate $u$ matching measurements, sparsity prior, and normals

**Algorithm 3** Augmented Lagrangian method to solve (A.4), estimating  $u$  from CS measurements and normal matching

- 1: Initialize  $u, \mathbf{d}, \boldsymbol{\lambda}$
- 2: For each iteration  $l = 1, 2, \dots$ , find an approximate minimizer of  $\mathcal{L}_1$  with respect to variables  $(u, \mathbf{d})$  with fixed Lagrange multipliers  $\boldsymbol{\lambda}^l$ :

$$u^l = \underset{u}{\operatorname{argmin}} \mathcal{L}_1(u, \mathbf{d}^{l-1}, \boldsymbol{\lambda}^l) \text{ solved in Fourier domain} \quad (\text{A.6})$$

$$\mathbf{d}^l = \underset{\mathbf{d}}{\operatorname{argmin}} \mathcal{L}_1(u^l, \mathbf{d}, \boldsymbol{\lambda}^l) \text{ solved by shrinkage} \quad (\text{A.7})$$

- 3: Update Lagrange multipliers

$$\lambda_x^{l+1} = \lambda_x^l + r(d_x^l - D_x u^l)$$

$$\lambda_y^{l+1} = \lambda_y^l + r(d_y^l - D_y u^l)$$

- 4: Stop the iterative process when  $\frac{\|u^l - u^{l-1}\|}{\|u^l\|} < \epsilon$ .

which can be computed efficiently.

#### A.1.1 Minimization problem with respect to $u$

The sub-minimization problem (A.6) can be written as follows:

$$\min_u v^T u + \frac{\alpha}{2} \|Au - f\|_2^2 + \frac{r}{2} \|d_x + \frac{1}{r} \lambda_x - D_x u\|_2^2 + \frac{r}{2} \|d_y + \frac{1}{r} \lambda_y - D_y u\|_2^2. \quad (\text{A.8})$$

We see that it reduces to a quadratic minimization, with positive semi-definite Hessian matrix  $H = \alpha F^T R^T R F + r(D_x^T D_x + D_y^T D_y)$ . The optimality conditions read

$$Hu = b \quad \text{with} \quad b = \alpha F^T R f + D_x^T (r d_x + \lambda_x) + D_y^T (r d_y + \lambda_y).$$

Actually as  $R$  is a row selector,  $R^T R$  is a sparse diagonal matrix with non-zero entries on the selected Fourier coefficients, and we cannot assure the invertibility of  $H$ . We find an approximate solution defining the positive definite matrix  $H_\epsilon = H + \epsilon I_n$ , with a small  $\epsilon > 0$ , and solving the approximate system

$$H_\epsilon u = b + \epsilon \hat{u}, \quad (\text{A.9})$$

where we use the value of  $u$  from the previous augmented Lagrangian iteration to estimate  $\hat{u} = u^{l-1}$ . In the resulting system,  $H_\epsilon$  is block circulant and we can use the Fourier transform to decompose it as  $H_\epsilon = \mathcal{F}^T C \mathcal{F}$ , with  $C = \alpha R^T R + r \mathcal{F} \left( D_x^T D_x + D_y^T D_y \right) \mathcal{F}^T + \epsilon I_n$  a diagonal matrix. Consequently, the system (A.9) can easily be solved in the Fourier domain inverting the diagonal matrix  $C$ . In practice we use the FFT transform instead of doing the matrix multiplications with  $\mathcal{F}$  and  $\mathcal{F}^T$ , which gives us a solution of complexity  $\mathcal{O}(n \log n)$ .

### A.1.2 Minimization problem with respect to $\mathbf{d}$

The minimization problem with respect to  $\mathbf{d}$  corresponds to an  $\ell_1$  -  $\ell_2$  norm problem that can be solved by shrinkage. If we define  $\mathbf{z} = \frac{1}{r} \boldsymbol{\lambda} - \nabla u$ , Equation (A.7) is equivalent to

$$\min_{\mathbf{d}_x, \mathbf{d}_y} \sum_i |\mathbf{d}(i)| + \frac{r}{2} \sum_i |\mathbf{d}(i) - \mathbf{z}(i)|^2. \quad (\text{A.10})$$

The minimization of (A.10) can be done pixel-wise, and the solution is given by the shrinkage operator  $\mathcal{S}(\mathbf{z}, 1/r)$ , that is,

$$\mathbf{d}(i) = \max \left\{ |\mathbf{z}(i)| - \frac{1}{r}, 0 \right\} \frac{\mathbf{z}(i)}{|\mathbf{z}(i)|} \quad i = 1, \dots, n. \quad (\text{A.11})$$

## A.2 Regularization of normal vectors

To regularize the normals at each iteration we have to solve

$$\min_{|\mathbf{n}| \leq 1} \| |W(\nabla n_x, \nabla n_y)| \|_1 + \frac{\mu}{2} \|n_x - \hat{n}_x\|_2^2 + \frac{\mu}{2} \|n_y - \hat{n}_y\|_2^2,$$

where  $W$  is a diagonal matrix with weights associated to weighted TV regularizer. We use the same combination of splitting and augmented Lagrangian techniques than in Section A.1. To avoid repetition, on the following we will simply write the form of the constraint minimization problem, the augmented Lagrangian, and each of the sub-minimizations.

Equivalent constraint minimization is

$$\min_{\substack{\mathbf{n}=\mathbf{m}, |\mathbf{m}| \leq 1 \\ \mathbf{d}=\nabla n_x, \mathbf{e}=\nabla n_y}} \| |W(\mathbf{d}, \mathbf{e})| \|_1 + \frac{\mu}{2} \|n_x - \hat{n}_x\|_2^2 + \frac{\mu}{2} \|n_y - \hat{n}_y\|_2^2,$$

with associated augmented Lagrangian

$$\begin{aligned} \mathcal{L}_2(\mathbf{n}, \mathbf{m}, \mathbf{d}, \mathbf{e}, \boldsymbol{\lambda}, \mathbf{v}, \boldsymbol{\xi}) = & \| |W(\mathbf{d}, \mathbf{e})| \|_1 + \frac{\mu}{2} \|\mathbf{n} - \hat{\mathbf{n}}\|_2^2 + \lambda_x^T (d_x - D_x n_x) + \lambda_y^T (d_y - D_y n_y) \\ & + \frac{r_d}{2} \|d_x - D_x n_x\|_2^2 + \frac{r_d}{2} \|d_y - D_y n_y\|_2^2 + \mathbf{v}_x^T (e_x - D_x n_y) + \mathbf{v}_y^T (e_y - D_y n_y) + \frac{r_e}{2} \|e_x - D_x n_y\|_2^2 \\ & + \frac{r_e}{2} \|e_y - D_y n_y\|_2^2 + \xi_x^T (n_x - m_x) + \frac{r_m}{2} \|n_x - m_x\|_2^2 + \xi_y^T (n_y - m_y) + \frac{r_m}{2} \|n_y - m_y\|_2^2. \end{aligned}$$

The resulting minimization method is presented in Algorithm 4.

### A.2.1 Minimization problem with respect to $\mathbf{n}$

The sub-minimization problem with respect to  $n_x$  can be written as follows:

$$\min_{n_x} \frac{\mu}{2} \|n_x - \hat{n}_x\|_2^2 + \xi_x^T (n_x - m_x) + \frac{r_m}{2} \|n_x - m_x\|_2^2 + \frac{r}{2} \|d_x + \frac{1}{r_d} \lambda_x - D_x n_x\|_2^2 + \frac{r_d}{2} \|d_y + \frac{1}{r_d} \lambda_y - D_y n_x\|_2^2.$$

---

**Algorithm 4** Augmented Lagrangian method to regularize normal vectors

---

- 1: Initialize  $\mathbf{n}, \mathbf{m}, \mathbf{d}, \mathbf{e}, \boldsymbol{\lambda}, \mathbf{v}, \boldsymbol{\xi}$
- 2: For each iteration  $l = 1, 2, \dots$ , find an approximate minimizer of  $\mathcal{L}_2$  with respect to variables  $(n_x, n_y, \mathbf{m}, \mathbf{d}, \mathbf{e})$  with fixed Lagrange multipliers  $\boldsymbol{\lambda}^l, \mathbf{v}^l, \boldsymbol{\xi}^l$ :

$$\mathbf{n}^l = \underset{\mathbf{n}}{\operatorname{argmin}} \mathcal{L}_2(\mathbf{n}, \mathbf{m}^{l-1}, \mathbf{d}^{l-1}, \mathbf{e}^{l-1}, \boldsymbol{\lambda}^{l-1}, \mathbf{v}^{l-1}, \boldsymbol{\xi}^{l-1})$$

$$\mathbf{m}^l = \underset{\mathbf{m}}{\operatorname{argmin}} \mathcal{L}(\mathbf{n}^l, \mathbf{m}, \mathbf{d}^{l-1}, \mathbf{e}^{l-1}, \boldsymbol{\lambda}^{l-1}, \mathbf{v}^{l-1}, \boldsymbol{\xi}^{l-1})$$

$$\mathbf{d}^l = \underset{\mathbf{d}}{\operatorname{argmin}} \mathcal{L}(\mathbf{n}^l, \mathbf{m}^l, \mathbf{d}, \mathbf{e}^{l-1}, \boldsymbol{\lambda}^{l-1}, \mathbf{v}^{l-1}, \boldsymbol{\xi}^{l-1})$$

$$\mathbf{e}^l = \underset{\mathbf{e}}{\operatorname{argmin}} \mathcal{L}(\mathbf{n}^l, \mathbf{m}^l, \mathbf{d}^l, \mathbf{e}, \boldsymbol{\lambda}^{l-1}, \mathbf{v}^{l-1}, \boldsymbol{\xi}^{l-1})$$

- 3: Update Lagrange multipliers

$$\lambda_x^l = \lambda_x^{l-1} + r_d(d_x^l - D_x n_x^l)$$

$$\lambda_y^l = \lambda_y^{l-1} + r_r(d_y^l - D_y n_x^l)$$

$$v_x^l = v_x^{l-1} + r_e(e_x^l - D_x n_y^l)$$

$$v_y^l = v_y^{l-1} + r_e(e_y^l - D_y n_y^l)$$

$$\xi^l = \xi^{l-1} + r_m(\mathbf{n} - \mathbf{m})$$

- 4: Stop the iterative process when  $\frac{\|\mathbf{n}^l - \mathbf{n}^{l-1}\|}{\|\mathbf{n}^l\|} < \epsilon$ .
- 

We see that it reduces to a quadratic minimization, with positive definite Hessian matrix  $H = (\mu + r_m)I_n + r_d D_x^T D_x + r_d D_y^T D_y$ . The optimality conditions read

$$H \mathbf{n}_x = \mu \hat{\mathbf{n}}_x + r_m \mathbf{m}_x + D_x^T (r_d \mathbf{d}_x + \boldsymbol{\lambda}_x) + D_y^T (r_d \mathbf{d}_y + \boldsymbol{\lambda}_y) - \boldsymbol{\xi}_x.$$

As before, the resulting  $H$  is block circulant and we can use the Fourier transform to decompose it as  $H = \mathcal{F}^T C \mathcal{F}$ , with  $C = (\mu + r_m)I_n + r_d \mathcal{F} \left( D_x^T D_x + D_y^T D_y \right) \mathcal{F}^T$  a diagonal matrix. We solve the linear system in the Fourier domain efficiently with the FFT transform. Observe that the minimization problem with respect to  $n_y$  has the same form and can be solved with the same technique.

### A.2.2 Minimization problem with respect to $\mathbf{d}$ and $\mathbf{e}$

The minimization problem with respect to  $\mathbf{d}$  corresponds to the  $\ell_1 - \ell_2$  problem

$$\min_{d_x, d_y} \sum_i |w(i) \mathbf{d}(i)| + \frac{r_d}{2} \sum_i |\mathbf{d}(i) - \mathbf{z}(i)|^2, \quad (\text{A.12})$$

where  $\mathbf{z} = \frac{1}{r_d} \boldsymbol{\lambda} - \nabla n_x$ . As  $w(i) > 0$ , this minimization is equivalent to

$$\min_{d_x, d_y} \sum_i |\mathbf{d}(i)| + \frac{r_d}{2w(i)} \sum_i |\mathbf{d}(i) - \mathbf{z}(i)|^2. \quad (\text{A.13})$$

## Appendix A. Minimization Algorithms of Chapter 1: Local Recovery Method

---

A similar problem has already been solved in Section A.1 with the shrinkage operator, which is now adapted to include the weights  $w$ . The solution to this sub-minimization is given by

$$\mathbf{d}(i) = \max \left\{ |\mathbf{z}(i)| - \frac{w(i)}{r}, 0 \right\} \frac{\mathbf{z}(i)}{|\mathbf{z}(i)|} \quad i = 1, \dots, n. \quad (\text{A.14})$$

Due to the symmetry of the problems, the same minimization technique is used for  $\mathbf{e}$ .

### A.2.3 Minimization problem with respect to $\mathbf{m}$

Finally, the minimization problem with respect to  $\mathbf{m}$  reads

$$\min_{|\mathbf{m}(i)|=1} \frac{r_m}{2} \sum_i |\mathbf{m}(i) - \mathbf{z}(i)|^2, \quad \text{with} \quad \mathbf{z} = \mathbf{n} + \frac{1}{r_m} \boldsymbol{\xi} \quad (\text{A.15})$$

and can be solved pixel-wise. For each pixel we have the following 2D problem: given a point in space with coordinates  $\mathbf{z}(i) \in \mathbb{R}^2$  we want to find the point in the unit ball minimizing its distance to  $\mathbf{z}(i)$ . It is clear that the solution corresponds to the projection of the unconstrained minimizer  $\mathbf{z}(i)$  into the unit ball, i.e.,

$$\mathbf{m}(i) = \begin{cases} \mathbf{z}(i) & |\mathbf{z}(i)| \leq 1 \\ \frac{\mathbf{z}(i)}{|\mathbf{z}(i)|} & \text{otherwise} \end{cases}. \quad (\text{A.16})$$

This finalizes the sub-optimization tasks, and therefore this appendix.

# B Minimization Algorithm of Chapter 1

## Non-Local Recovery Method

In this appendix, we provide the details of the minimization algorithm of Chapter 1 for the non-local recovery method proposed for compressed sensing.

### B.0.4 Notation

In the discrete setting, the non-local gradient is a linear operator, and by arranging the image as a column vector it can thus be computed as a sparse matrix multiplication  $\nabla_G u = Du$ . The matrix  $D \in \mathbb{R}^{N \times n}$  ( $N = |E|$  indicates the number of edges in the graph) is derived from the weights associated to the edges and is usually sparse. Consequently  $d = Du \in \mathbb{R}^N$  is also a vector column, with as many components associated to node  $i$  as neighbours this node has. With the vector notation, the inner product of two vectors fields  $\mathbf{d}, \mathbf{e}$  defined in  $G$  is then computed as  $\langle \mathbf{d}, \mathbf{e} \rangle_G = \mathbf{d}^T \mathbf{e}$ . As in the continuous setting, the NL divergence  $\text{div}_G$  is derived from its adjoint relation with the non-local gradient  $\nabla_G^* = -\text{div}_G$  and, consequently, in matrix notation it corresponds to  $\text{div}_G \mathbf{d} = -D^T \mathbf{d}$ .

Since the minimization associated to (1.33) has already been explained for the vectorial case, in this appendix we focus on the minimizations associated to the non-local operators of Equations (1.30) and (1.34).

### B.1 Estimate $u$ matching measurements, non-local sparsity prior, and non-local normals

With the previous notation, the minimization problem (1.30) reads

$$u = \arg \min_u \|Du\|_1 + \gamma v^T u + \frac{\alpha}{2} \|Au - f\|_2^2 \quad (\text{B.1})$$

with  $v = \text{div}_G \mathbf{n}_G$ . This minimization is also reformulated as a constraint minimization and solved efficiently with an augmented Lagrangian. Compared to the local minimizations, in the splitting step we require an additional variable,  $s$ , to have efficient solutions for the posterior

---

## Appendix B. Minimization Algorithms of Chapter 1: Non-local Recovery Method

---

sub-minimization problems. The resulting constraint minimization of (B.1) is given by

$$\min_{u,s,\mathbf{d}} \| |\mathbf{d}|_G \|_1 + v^T u + \frac{\alpha}{2} \|As - f\|_2^2 \quad \text{subject to} \quad \begin{cases} \mathbf{d} = Du \\ s = u \end{cases}. \quad (\text{B.2})$$

The Lagrangian in that case reads

$$\begin{aligned} \mathcal{L}_3(u, s, \mathbf{d}, \boldsymbol{\lambda}_d, \lambda_u) = & \| |\mathbf{d}|_G \|_1 + v^T u + \frac{\alpha}{2} \|As - f\|_2^2 + \boldsymbol{\lambda}_d^T (\mathbf{d} - Du) + \frac{r_d}{2} \|\mathbf{d} - Du\|^2 \\ & + \lambda_u^T (u - s) + \frac{r_u}{2} \|u - s\|^2. \end{aligned}$$

The resulting minimization method is presented in Algorithm 5, where we have also hinted the solution to each of the sub-minimization problems.

---

**Algorithm 5** Augmented Lagrangian method for CS reconstruction matching normals by (B.2)

---

- 1: Initialize  $u, s, \mathbf{d}, \boldsymbol{\lambda}_d, \lambda_u$
- 2: For each iteration  $l = 1, 2, \dots$ , find an approximate minimizer of  $\mathcal{L}_3$  with respect to variables  $(u, s, \mathbf{d})$  with fixed Lagrange multipliers  $\boldsymbol{\lambda}_d^l, \lambda_u^l$ :

$$\begin{aligned} u &= \arg \min_u \mathcal{L}_3(u, s^{l-1}, \mathbf{d}^{l-1}, \boldsymbol{\lambda}_d^{l-1}, \lambda_u^{l-1}) \quad \text{conjugate gradients} \\ s &= \arg \min_s \mathcal{L}_3(u^l, s, \mathbf{d}^{l-1}, \boldsymbol{\lambda}_d^{l-1}, \lambda_u^{l-1}) \quad \text{solved in Fourier domain} \\ \mathbf{d} &= \arg \min_{\mathbf{d}} \mathcal{L}_3(u^l, s^l, \mathbf{d}, \boldsymbol{\lambda}_d^{l-1}, \lambda_u^{l-1}) \quad \text{solved by non-local shrinkage} \end{aligned}$$

- 3: Update Lagrange multipliers

$$\begin{aligned} \boldsymbol{\lambda}_d^l &= \boldsymbol{\lambda}_d^{l-1} + r_d (\mathbf{d}^l - Du^l) \\ \lambda_u^l &= \lambda_u^{l-1} + r_u (u^l - s^l) \end{aligned}$$

- 4: Stop the iterative process when  $\frac{\|u^l - u^{l-1}\|}{\|u^l\|} < \epsilon$ .
- 

### Minimization problem with respect to $u$

The minimization with respect to  $u$  corresponds to the following quadratic positive definite problem

$$\min_u v^T u + \boldsymbol{\lambda}_d^T (\mathbf{d} - Du) + \frac{r_d}{2} \|\mathbf{d} - Du\|^2 + \lambda_u^T (u - s) + \frac{r_u}{2} \|u - s\|^2. \quad (\text{B.3})$$

We find its minimizer solving its optimality conditions, which provide the following system of linear equations

$$K_u u = -\gamma v - \lambda_u + r_u s + D^T (\boldsymbol{\lambda}_d + r_d \mathbf{d}). \quad (\text{B.4})$$



### B.1. Estimate $u$ matching measurements, non-local sparsity prior, and non-local normals

Matrix  $K_u = r_u I + r_d D^T D$  is sparse, symmetric, and positive definite, and we have efficient algorithms to invert it. We choose an iterative method that we initialize with the solution to the previous augmented Lagrangian iterate. In particular we use conjugate gradients to exploit the symmetry and positive definition of  $K$ , with preconditioning matrix given by its incomplete Cholesky factorization. The resulting method is very fast, converging to enough precision with 2-3 iterations of conjugate gradients.

#### B.1.1 Minimization problem with respect to $s$

The minimization with respect to  $s$  is also a quadratic problem which can be efficiently solved, in that case in the Fourier domain. The problem reads

$$\min_s \frac{\alpha}{2} \|As - f\|^2 + \lambda_u^T (u - s) + \frac{r_u}{2} \|u - s\|^2. \quad (\text{B.5})$$

The optimality conditions in that case are

$$(\alpha A^T A + r_u I_n) s = \alpha A^T f + \lambda_u + r_u u. \quad (\text{B.6})$$

As before, the matrix  $K_s = \alpha A^T A + r_u I_n = F^T C F$  is block-circulant and the resulting system is diagonal in the Fourier domain with  $C = R^T R + r_u I_n$ . Therefore (B.6) can be efficiently solved with the FFT.

The introduction of the additional splitting variable  $s = u$  allows us to split the inversion of the full matrix  $\alpha A^T A + r_d D^T D$ , which would result from the use of a single variable for  $s$  and  $u$ , into the inversion of two matrices  $K_u$  and  $K_s$ . The sparse matrix  $K_u$  can be efficiently solved with a sparse incomplete Cholesky factorization, while the full matrix  $K_s$  is easily inverted in the Fourier domain. The original matrix  $\alpha A^T A + r_d D^T D$ , on the other hand, does not present an evident sparsity pattern or a direct decomposition with fast transforms.

#### B.1.2 Minimization problem with respect to $d$

The minimization with respect to  $d$  is equivalent to

$$\min_d \| |d|_G \|_1 + \frac{r_d}{2} \|d - z\|^2 \quad \text{with} \quad z = Du - \frac{\lambda_d}{r_d} \quad (\text{B.7})$$

As in the local case, this minimization is decoupled for each pixel  $i$  as follows

$$\min_{d(i,j)} \sqrt{\sum_{j \sim i} d^2(i,j)} + \sum_{j \sim i} \frac{r_d}{2} (d(i,j) - z(i,j))^2 \quad (\text{B.8})$$

and can be solved by a direct extension of the shrinkage operator to the graph. That is, for each node in the graph that is a neighbour of node  $i$ , the solution is given by

$$\mathbf{d}^*(i, j) = \max \left\{ |z|_G(i) - \frac{1}{r_d}, 0 \right\} \frac{z(i, j)}{|z|_G(i)}. \quad (\text{B.9})$$

## B.2 Regularization of non-local divergence of normal vectors

With the previous notation, the minimization problem (1.34) reads

$$\min_v \| |Dv|_G \|_1 + \frac{\mu}{2} \|v - \hat{v}\|_2^2. \quad (\text{B.10})$$

As in the local case, we decouple the  $\ell_1$  and the  $\ell_2$  problems by defining an additional variable  $\mathbf{d} = Du$  and rewriting (B.10) as the following constraint minimization problem

$$\min_{v, \mathbf{d}} \| |\mathbf{d}|_G \|_1 + \frac{\mu}{2} \|v - \hat{v}\|_2^2 \quad \text{subject to} \quad \mathbf{d} = Du \quad (\text{B.11})$$

with associated augmented Lagrangian

$$\mathcal{L}_4(u, \mathbf{d}, \boldsymbol{\lambda}_d) = \| |\mathbf{d}|_G \|_1 + \frac{\mu}{2} \|v - \hat{v}\|_2^2 + \boldsymbol{\lambda}_d^T (\mathbf{d} - Du) + \frac{r_d}{2} \|\mathbf{d} - Du\|_2^2. \quad (\text{B.12})$$

To minimize the Lagrangian  $\mathcal{L}_4$  with respect to  $u$  and  $\mathbf{d}$ , we alternate the direction of minimization with respect to each variable and proceed as indicated by Algorithm 6, where we have also hinted the solution to each one of the sub-minimization problems.

---

**Algorithm 6** Augmented Lagrangian method to regularize non-local divergence of normals from (B.11)

---

- 1: Initialize  $u, \mathbf{d}, \boldsymbol{\lambda}_d$
- 2: For each iteration  $l = 1, 2, \dots$ , find an approximate minimizer of  $\mathcal{L}_4$  with respect to variables  $(u, \mathbf{d})$  with fixed Lagrange multipliers  $\boldsymbol{\lambda}_d^l$ :

$$v = \arg \min_v \mathcal{L}_4(v, \mathbf{d}^{l-1}, \boldsymbol{\lambda}_d^{l-1}) \text{ solved with conjugate gradient}$$

$$\mathbf{d} = \arg \min_{\mathbf{d}} \mathcal{L}_4(v^l, \mathbf{d}, \boldsymbol{\lambda}_d^{l-1}) \text{ solved by non-local shrinkage}$$

- 3: Update Lagrange multipliers

$$\boldsymbol{\lambda}_d^l = \boldsymbol{\lambda}_d^{l-1} + r_d(\mathbf{d}^l - Du^l)$$

- 4: Stop the iterative process when  $\frac{\|v^l - v^{l-1}\|}{\|v^l\|} < \epsilon$ .
-

### **B.2.1 Minimization problem with respect to $v$**

The minimization with respect to  $v$  corresponds to the following quadratic positive definite problem

$$\min_v \frac{\mu}{2} \|v - \hat{v}\|_2^2 + \lambda_d^T (\mathbf{d} - Dv) + \frac{r_d}{2} \|\mathbf{d} - Dv\|_2^2. \quad (\text{B.13})$$

We find its minimizer by solving its optimality conditions, which provide the following system of linear equations

$$(\mu I + r_d D^T D) v = \mu \hat{v} + D^T (\lambda_d + r_d \mathbf{d}). \quad (\text{B.14})$$

We find the same form of matrix  $K = \mu I + r_d D^T D$  than in (B.4) and, therefore, we use the same conjugate gradient method to invert the system.

### **B.2.2 Minimization problem with respect to $d$**

The minimization with respect to  $d$  is equivalent to (B.7) changing  $u$  for  $v$ , in particular we have

$$\min_d \| |\mathbf{d}|_G \|_1 + \frac{r_d}{2} \|\mathbf{d} - \mathbf{z}\|^2 \text{ with } \mathbf{z} = Du - \frac{\lambda_d}{r_d}, \quad (\text{B.15})$$

which is solved with the same adaptation of the shrinkage operator to the graph. For each node neighbour to  $i$ , the solution is given by

$$d^*(i, j) = \max \left\{ |z|_G(i) - \frac{1}{r_d}, 0 \right\} \frac{z(i, j)}{|z|_G(i)}. \quad (\text{B.16})$$

This finalizes the sub-optimization tasks, and therefore this appendix.



## C Minimization Algorithm of Chapter 2

The optimization problem associated with the proposed HAC model is given by

$$\min_{\phi} \int_{\Omega} \sqrt{g}, \quad (\text{C.1})$$

with the surface element given by

$$\begin{aligned} g = & 1 + \alpha_i \|\nabla I^i\|^2 + \gamma_i \|\nabla f^i\|^2 + \beta \|\nabla \phi\|^2 + \frac{1}{2} \alpha_i \alpha_j [\nabla I^i, \nabla I^j]^2 \\ & + \alpha_i \gamma_j [\nabla I^i, \nabla f^j]^2 + \frac{1}{2} \gamma_i \gamma_j [\nabla f^i, \nabla f^j]^2 + \alpha_i \beta [\nabla I^i, \nabla \phi]^2 + \gamma_i \beta [\nabla f^i, \nabla \phi]^2. \end{aligned} \quad (\text{C.2})$$

The key idea of the proposed decomposition algorithm is to split the original problem into sub-optimization problems that are easier to solve. To that purpose, we rewrite the original problem (C.1) as a constraint minimization, and we use augmented Lagrangians to solve it. The resulting Lagrangian is minimized with respect to each variable independently, and the multipliers are then updated in a cyclic way. Since all the minimizations can be analytically solved or are decoupled for each pixel, the resulting algorithm is fast and easy to implement.

Let us consider the following constrained minimization problem

$$\min_{\phi, \mathbf{p}, \varphi, f^1, \dots, f^k, \mathbf{q}^1, \dots, \mathbf{q}^k} \int_{\Omega} \sqrt{g_c} \quad \text{subject to} \quad \begin{cases} \mathbf{p} = \nabla \phi \\ \mathbf{q}^i = \nabla f^i \quad 1 \leq i \leq k \\ \varphi = \phi \\ f^i = s^i H(\varphi) \quad 1 \leq i \leq k \end{cases}, \quad (\text{C.3})$$

which is equivalent to the original problem (C.1) with

$$\begin{aligned} g_c = & 1 + \alpha_i \|\nabla I^i\|^2 + \gamma_i \|\mathbf{q}^i\|^2 + \beta \|\mathbf{p}\|^2 + \frac{1}{2} \alpha_i \alpha_j [\nabla I^i, \nabla I^j]^2 + \alpha_i \gamma_j [\nabla I^i, \mathbf{q}^j]^2 \\ & + \frac{1}{2} \gamma_i \gamma_j [\mathbf{q}^i, \mathbf{q}^j]^2 + \alpha_i \beta [\nabla I^i, \mathbf{p}]^2 + \gamma_i \beta [\mathbf{q}^i, \mathbf{p}]^2. \end{aligned}$$

## Appendix C. Minimization Algorithm of Chapter 2

Next, we reformulate the constrained minimization of problem (C.3) as an unconstrained optimization task. This can be done with an augmented Lagrangian approach, which translates the constraints into pairs of a Lagrangian multiplier and a penalty term, see e.g., Glowinski and Le Tallec [1989]. Let us define the augmented Lagrangian energy  $\mathcal{L}(\phi, \mathbf{p}, f^1, \dots, f^k, \mathbf{q}^1, \dots, \mathbf{q}^k, \varphi, \Lambda)$  associated with the previous problem (C.3):

$$\begin{aligned} \mathcal{L} = \int_{\Omega} \sqrt{g_c} + \lambda_1 \cdot (\mathbf{p} - \nabla \phi) + \frac{r_1}{2} \|\mathbf{p} - \nabla \phi\|^2 + \lambda_{2,i} \cdot (\mathbf{q}^i - \nabla f^i) + \frac{r_{2,i}}{2} \|\mathbf{q}^i - \nabla f^i\|^2 \\ + \lambda_3(\varphi - \phi) + \frac{r_3}{2}(\varphi - \phi)^2 + \lambda_{4,i} \cdot (f^i - s^i H(\varphi)) + \frac{r_{4,i}}{2} (f^i - s^i H(\varphi))^2, \end{aligned} \quad (\text{C.4})$$

where the Lagrange multipliers  $\Lambda = (\lambda_1, \lambda_{2,1}, \dots, \lambda_{2,k}, \lambda_3, \lambda_{4,1}, \dots, \lambda_{4,k})$  are functions in  $\Omega$ , and  $r_1, \dots, r_{4,k}$  are positive constants.

The constraint minimization problem (C.3) reduces then to finding the saddle-point of the augmented Lagrangian  $\mathcal{L}$ . The solution to this saddle-point problem can be approximated by the following iterative algorithm: initialize the Lagrange multipliers to zero and the split variables according to the constraints; at each iteration, find an approximate minimizer of  $\mathcal{L}$  with respect to the variables  $\phi, \mathbf{p}, f^1, \dots, f^k, \mathbf{q}^1, \dots, \mathbf{q}^k$  by alternating the minimization with respect to each variable, and update the Lagrange multipliers with the residuals associated with each constraint; stop the process when  $\phi$  remains fixed. This process is summarized in Algorithm 8.

The next step is to determine the solutions of the sub-minimization problems (C.5)-(C.9), which can actually be computed efficiently. In the following, we simplify notation by omitting the super-index in the different sub-minimizations.

### C.0.3 Notation

We discretize the image domain  $\Omega \subset \mathbb{R}^2$  with a regular grid of size  $n = n_x \times n_y$ . We use forward differences to compute the discrete gradients and backward differences for the divergence in order to preserve the adjoint relationship  $\text{div} = -\nabla^*$  in the discrete setting. In matrix-vector notation, we can efficiently compute the spatial derivatives multiplying the discrete functions arranged as a column vector with the sparse finite difference matrices  $\nabla_x u = D_x u$ ,  $\nabla_y u = D_y u$ . Similarly, the discretization of the  $L_2$  inner product in  $\Omega$  corresponds to the usual dot product of vectors.

### C.0.4 Minimization associated with $\phi$ and $f^i$

The sub-minimization problem (C.5) can be written as follows:

$$\min_{\phi} \int_{\Omega} \frac{r_1}{2} \|\mathbf{v} - \nabla \phi\|^2 + \frac{r_3}{2} (z - \phi)^2,$$

---

**Algorithm 7** Lagrangian method for minimization of HAC method.

---

- 1: Initialize  $\phi, \mathbf{p}, f^1, \dots, f^k, \mathbf{q}^1, \dots, \mathbf{q}^k, \varphi, \Lambda$
- 2: At iteration  $l$ , find a minimizer of  $\mathcal{L}$  with respect to variables  $(\phi, \mathbf{p}, f^1, \dots, f^k, \mathbf{q}^1, \dots, \mathbf{q}^k, \varphi)$  with fixed Lagrange multipliers  $\Lambda^{l-1}$ :

$$\phi^l = \argmin_{\phi} \mathcal{L}(\phi, \mathbf{p}^{l-1}, f^{1^{l-1}}, \dots, f^{k^{l-1}}, \mathbf{q}^{1^{l-1}}, \dots, \mathbf{q}^{k^{l-1}}, \varphi^{l-1}, \Lambda^{l-1}) \quad (\text{C.5})$$

$$\mathbf{p}^l = \argmin_{\mathbf{p}} \mathcal{L}(\phi^l, \mathbf{p}, f^{1^{l-1}}, \dots, f^{k^{l-1}}, \mathbf{q}^{1^{l-1}}, \dots, \mathbf{q}^{k^{l-1}}, \varphi^{l-1}, \Lambda^{l-1}) \quad (\text{C.6})$$

$$f^{i^l} = \argmin_{f^i} \mathcal{L}(\phi^l, \mathbf{p}^l, f^1, \dots, f^k, \mathbf{q}^{1^{l-1}}, \dots, \mathbf{q}^{k^{l-1}}, \varphi^{l-1}, \Lambda^{l-1}) \quad (\text{C.7})$$

$$\mathbf{q}^{i^l} = \argmin_{\mathbf{q}^i} \mathcal{L}(\phi^l, \mathbf{p}^l, f^{1^l}, \dots, f^{k^l}, \mathbf{q}^1, \dots, \mathbf{q}^k, \varphi^{l-1}, \Lambda^{l-1}) \quad (\text{C.8})$$

$$\varphi^l = \argmin_{\varphi} \mathcal{L}(\phi^l, \mathbf{p}^l, f^{1^l}, \dots, f^{k^l}, \mathbf{q}^{1^l}, \dots, \mathbf{q}^{k^l}, \varphi, \Lambda^{l-1}) \quad (\text{C.9})$$

- 3: Update the Lagrange multipliers and region terms

$$\begin{aligned} \lambda_1^l &= \lambda_1^{l-1} + r_1(\mathbf{p}^l - \nabla \phi^l) \\ \lambda_{2,i}^l &= \lambda_{2,i}^{l-1} + r_{2,i}(\mathbf{q}^{i^l} - \nabla f^{i^k}) \quad i = 1, \dots, k \\ \lambda_3^l &= \lambda_3^{l-1} + r_3(\varphi^l - \phi^l) \\ s^i &= \left(I^i - \mu_+^i\right)^2 \text{ with } \mu_+^i = \frac{\int_{\Omega} I^i H(\varphi)}{\int_{\Omega} H(\varphi)} \quad i = 1, \dots, k \\ \lambda_{4,i}^l &= \lambda_{4,i}^{l-1} + r_{4,i}(f^{i^l} - s^i H(\varphi^l)) \quad i = 1, \dots, k \end{aligned}$$

- 4: Stop the iterative process when  $\|\phi^l - \phi^{l-1}\| < \epsilon$ .
-

## Appendix C. Minimization Algorithm of Chapter 2

---

where  $z = \varphi + \frac{\lambda_3}{r_3}$  and  $\mathbf{v} = \mathbf{p} - \frac{\lambda_1}{r_1}$ . The corresponding Euler-Lagrange equation is:

$$r_3 \phi - r_1 \Delta \phi = r_3 z - r_1 \operatorname{div} \mathbf{z}. \quad (\text{C.10})$$

With the notation explained above, Equation (C.10) is discretized as a linear system of equations  $A_\phi \phi = b_\phi$ , with matrix  $A_\phi$  and vector  $b_\phi$  given by

$$\begin{aligned} A_\phi &= r_3 I_n + r_1 D_x^T D_x + r_1 D_y^T D_y \\ b_\phi &= r_3 z - r_1 D_x^T \mathbf{v}_x - r_1 D_y^T \mathbf{v}_y. \end{aligned}$$

Matrix  $A_\phi$  is symmetric, definite positive, and block-circulant, and we can use the Fourier transform  $\mathcal{F}$  to decompose it as  $A_\phi = \mathcal{F}^T D_\phi \mathcal{F}$ , with  $D_\phi$  a diagonal matrix. Consequently, the system  $A_\phi \phi = b_\phi$  can easily be solved in the Fourier domain. In practice we use the FFT transform instead of doing the matrix multiplications with  $\mathcal{F}$  and  $\mathcal{F}^T$ , which gives us the following solution  $\phi = \mathcal{F}^T (D_\phi^{-1} \mathcal{F} b_\phi)$  of complexity  $\mathcal{O}(n \log n)$

The minimization problem associated with the features  $f^1, \dots, f^k$  have all the same form and can also be solved in the frequency domain. Without loss of generality, we present here the minimization associated with  $f^1$ , which reads

$$\min_{f^1} \int_{\Omega} \frac{r_{2,1}}{2} \|\nabla f^1 - \mathbf{v}\|^2 + \frac{r_{4,1}}{2} (f^1 - z)^2,$$

where we have defined the auxiliary variables  $\mathbf{v} = \mathbf{q}^1 + \frac{\lambda_{2,1}}{r_{2,1}}$  and  $z = s^1 H - \frac{\lambda_{4,1}}{r_{4,1}}$ . The corresponding Euler-Lagrange equation is

$$r_{4,1} f^1 - 2r_{2,1} \Delta f^1 = r_{4,1} z - r_{2,1} \operatorname{div} \mathbf{v},$$

which is discretized also as a linear system of equations  $A_f f^1 = b_f$ . Matrix  $A_f$  and vector  $b_f$  are given by

$$\begin{aligned} A_f &= r_{4,1} I_n + r_{2,1} D_x^T D_x + r_{2,1} D_y^T D_y \\ b_f &= r_{4,1} z - r_{2,1} D_x^T \mathbf{v}_x - r_{2,1} D_y^T \mathbf{v}_y, \end{aligned}$$

where  $A_f$  is also symmetric, positive definite, and block-circulant, and the system is solved again in the Fourier domain.

### C.0.5 Minimization associated with vector fields $\mathbf{p}$ and $\mathbf{q}^1, \dots, \mathbf{q}^k$

We first note that the minimization problem associated with the vector fields  $\mathbf{p}$  and each  $\mathbf{q}^1, \dots, \mathbf{q}^k$  are decoupled for each pixel and can be solved by point-wise minimization of the functions being integrated. As the problems associated with the vector fields  $\mathbf{q}^1, \dots, \mathbf{q}^k$  have



all the same form, we develop only the expressions for  $\mathbf{q}^1$  and  $\mathbf{p}$ , which read

$$\min_{\mathbf{p}} \int_{\Omega} \sqrt{g_c} + \lambda_1 (\mathbf{p} - \nabla \phi) + \frac{r_1}{2} \|\mathbf{p} - \nabla \phi\|^2 \quad (\text{C.11})$$

$$\min_{\mathbf{q}^1} \int_{\Omega} \sqrt{g_c} + \lambda_{2,1} (\mathbf{q}^1 - \nabla f^1) + \frac{r_{2,1}}{2} \|\mathbf{q}^1 - \nabla f^1\|^2. \quad (\text{C.12})$$

The main difficulty involved in (C.11) and (C.12) is the square root affecting the surface element  $\sqrt{g_c}$ , which does not allow for a closed-form solution. To overcome this issue, we use the *iterative re-weighted least squares* (IRLS) technique, which has been proposed in the context of Beltrami minimizations in Rosman [2010].

IRLS iteratively minimizes the square root term  $\sqrt{g_c}$  by the following process: at iteration  $l$  the square root is approximated by the weighted surface element  $\sqrt{g_c}^l \approx \frac{g_c}{\sqrt{g_c^{l-1}}} = \frac{g_c}{w^l}$  with fixed weight  $w^l$ , the resulting weighted least squares problem is solved, and the weight is updated for the next iteration. The problem reduces then to a series of a quadratic minimizations on  $\mathbf{p}$  or on  $\mathbf{q}^1$ , for which closed-form solutions are available.

For instance, applied to (C.11), at each iteration the minimization problem is equivalent to

$$\min_{\mathbf{p}} \int_{\Omega} \frac{\beta \|\mathbf{p}\|^2 + \alpha_i \beta [\nabla I^i, \mathbf{p}]^2}{w} + \frac{r_1}{2} \|\mathbf{p} - \mathbf{v}\|^2,$$

where  $\mathbf{v} = \nabla \phi - \frac{\lambda_1}{r_1}$ . By simple differentiation with respect to each component of  $\mathbf{p} = (\mathbf{p}_x, \mathbf{p}_y)$  we obtain the following  $2 \times 2$  linear system of equations for each pixel

$$A_p \mathbf{p} = r_1 w \mathbf{v} \text{ with } A_p = 2\beta \begin{pmatrix} 1 + \alpha_i [I_y^i]^2 + \gamma_i [\mathbf{q}_y^i]^2 + \frac{r_1 w}{2\beta} & \alpha_i I_x^i I_y^i + \gamma_i \mathbf{q}_x^i \mathbf{q}_y^i \\ \alpha_i I_x^i I_y^i + \gamma_i \mathbf{q}_x^i \mathbf{q}_y^i & 1 + \alpha_i [I_x^i]^2 + \gamma_i [\mathbf{q}_x^i]^2 + \frac{r_1 w}{2\beta} \end{pmatrix}.$$

Matrix  $A_p$  is symmetric, positive definite, and its determinant is given by

$$|A| = \left( \frac{2\beta}{w} \right)^2 g^f + r_1 \frac{2\beta}{w} (2\alpha_i \|\nabla I^i\|^2 + \gamma_i \|\mathbf{q}^i\|^2),$$

where  $g^f$  is the induced metric associated with the embedding  $(x, y, I^1, \dots, I^k, f^1, \dots, f^k)$ . With a  $2 \times 2$  linear system to solve at each pixel, we have an equivalent to a closed-form solution for each IRLS update. These formulas can be further simplified by ignoring the coupling of the  $x, y$  components of  $\mathbf{p}$  and updating each component iteratively.

The same procedure is used to solve (C.12). In this case we have the following linear system for each pixel  $B^1 \mathbf{q}^1 = \frac{r_{2,1}}{2\gamma_1} w \nabla f^i - \lambda_{2,1}$ .

Matrix  $B^1$  is symmetric, positive definite, and has entries<sup>1</sup>

$$\begin{aligned} B_{11} &= 1 + \sum_{i=1}^k \alpha_i \left[ I_y^i \right]^2 + \sum_{i=2}^k \gamma_i \left[ \mathbf{q}_y^i \right]^2 + \beta \left[ \mathbf{p}_y \right]^2 + \frac{w r_{2,1}}{2\gamma_1} \\ B_{12} &= B_{21} = \sum_{i=1}^k \alpha_i I_x^i I_y^i + \sum_{i=2}^k \gamma_i \mathbf{q}_x^i \mathbf{q}_y^i + \beta \mathbf{p}_x \mathbf{p}_y \\ B_{22} &= 1 + \sum_{i=1}^k \alpha_i \left[ I_x^i \right]^2 + \sum_{i=2}^k \gamma_i \left[ \mathbf{q}_x^i \right]^2 + \beta \left[ \mathbf{p}_x \right]^2 + \frac{w r_{2,1}}{2\gamma_1} \end{aligned}$$

The determinant of  $B^1$  reads

$$|B^1| = \frac{1}{w^2} g^{f_1} + r_{2,1} \frac{1}{2\gamma_1 w} \left( 2 \sum_{i=1}^k \alpha_i \|\nabla I^i\|^2 + \sum_{i=2}^k \gamma_i \|\mathbf{q}^i\|^2 \right),$$

where  $g^{f_1}$  is the induced metric associated with the embedding  $(x, y, I^1, \dots, I^k, f^2, \dots, f^k)$ . We have then also a closed-form for each IRLS update.

In practice, we have found that 3-5 iterations of IRLS are enough for both minimization problems.

### C.0.6 Minimization associated with $\varphi$

If we define the auxiliary variables  $z = \phi - \frac{\lambda_3}{r_3}$ ,  $r = r_{4,i}(s^i)^2$  and  $w = \frac{1}{r} \lambda_{4,i} s^i + \frac{r_{4,i}}{r} s^i f^i$ , the minimization problem associated with  $\varphi$  is equivalent to the following minimization

$$\min_{\varphi} \int_{\Omega} F(\varphi) = \min_{\varphi} \int_{\Omega} \frac{r_3}{2} (\varphi - z)^2 + \frac{r}{2} (H(\varphi) - w)^2 \quad (\text{C.13})$$

and can again be solved by pixel-wise minimization of  $F(\varphi)$ . Observe that for practical implementations, this minimization involves a smooth approximation  $H_{\epsilon}$  of the Heaviside function. We propose two steps to find quickly a minimizer of (C.13).

1) Find a solution  $\varphi^0$  of (C.13) for  $\epsilon = 0$  (i.e. for the distributional/non-smooth Heaviside function). A closed-form solution exists for this problem and can be computed as follows. The first term of  $F(\varphi)$  is minimized by  $\varphi^0 = z$ . As the Heaviside function can take only values 0 or 1, the second term is minimized by  $\varphi^0 < 0$  when  $w < \frac{1}{2}$  and by  $\varphi^0 \geq 0$  when  $w \geq \frac{1}{2}$ . That means that both terms, and therefore the function, are minimized by  $\varphi^0 = z$  if either  $w < \frac{1}{2}$  and  $z < 0$  or  $w \geq \frac{1}{2}$  and  $z \geq 0$ . Otherwise we must choose to minimize the greater of these terms and set  $\varphi^0 = 0$  if  $F(0) < F(z)$  and  $\varphi^0 = z$  otherwise.

2) Find a solution  $\varphi$  of (C.13) for  $\epsilon > 0$  using the standard Newton's method with  $\varphi^0$  as initialization. The iterative Newton's method for finding the minimizer of (C.13) reads as

---

<sup>1</sup>Einstein's convention is not used because the summations always exclude one index.

---

follows:

$$\varphi^{m+1} = \varphi^m - \frac{r_3 (\varphi^m - z) + r (H(\varphi^m) - w) \delta (\varphi^m)}{r_3 + r (H(\varphi^m) - w) \delta'(\varphi^m) + r \delta (\varphi^m)^2}.$$

Each iteration, Newton's method finds a second-order polynomial approximation to the function around the current iterate and minimizes it. Initialized close to a minimum, as our first step assures, Newton's method converges fast because the second order approximation is accurate, and the resulting minimization method requires only 3-5 iterations to converge.

This finalizes the sub-optimization tasks, and therefore this appendix.



## D Minimization Algorithm of Chapter 3

The variational problem associated with the reconstruction of a surface in Chapter 3 results in the following non-convex minimization problem

$$\min_{z, \mathbf{u}} \alpha A(\nabla z) + \beta B(z, \mathbf{u}) + \gamma G(\nabla z) + W(\nabla \mathbf{u}). \quad (\text{D.1})$$

The objective functional of this problem has terms involving the surface height map  $z$ , its spatial gradients  $\mathbf{d} = \nabla z$ , the deformation field  $\mathbf{u}$ , and its spatial derivatives  $V = \nabla \mathbf{u}$ <sup>1</sup>. To efficiently solve the problem, we design an iterative algorithm that considers these variables as independent, and we use an augmented Lagrangian to ensure that the relationships  $\mathbf{d} = \nabla z$  and  $V = \nabla \mathbf{u}$  hold.

The resulting Lagrangian is minimized with respect to each variable independently and the multipliers are then updated in a cyclic way. The independent minimizations for each variable are easier to solve because they reduce to convex problems involving only first-order derivatives or are decoupled for each pixel. The initial complex minimization of Equation (D.1) is thus divided in small and easier subtasks that are efficiently solved, and the resulting algorithm is faster than the descent PDE flow of problem (D.1). In image processing, similar techniques have been recently applied to the variational minimization problems associated with image denoising, image segmentation, and image reconstruction that were previously solved with PDE methods, see e.g., Wang et al. [2008], Goldstein and Osher [2009], Goldstein et al. [2009].

### D.0.7 Notation

We discretize the image domain  $\Omega \subset \mathbb{R}^2$  with a regular grid of size  $n = n_x \times n_y$ . In  $\Omega$  we consider images and height maps as scalar functions with  $z(i) \in \mathbb{R}$ , and their gradients are

---

<sup>1</sup>  $\mathbf{u}$  is a 2-dimensional vector field in  $\mathbb{R}^2$  and we denote as  $V = \nabla \mathbf{u}$  the  $2 \times 2$  matrix field obtained of applying the gradient operator to each component of  $\mathbf{u}$ .

therefore vector-valued functions with  $\nabla z(i) \in \mathbb{R}^2$ . We use forward differences to compute the discrete gradients and backward differences for the divergence in order to preserve the adjoint relationship  $\text{div} = -\nabla^*$  in the discrete setting. In vector notation, we can efficiently compute the spatial derivatives multiplying the discrete functions arranged as a column vector with the sparse finite difference matrices  $\nabla_x z = D_x z$ ,  $\nabla_y z = D_y z$ .

For numerical purposes we also relax the Heaviside function to a smooth approximation  $H_\epsilon(x) = \frac{1}{2} \left(1 + \frac{2}{\pi} \arctan\left(\frac{x}{\epsilon}\right)\right)$ , with its corresponding derivative  $\delta_\epsilon$  approximating the Dirac distribution.

### D.1 Numerical minimization algorithm

To reconstruct the height map  $z$  associated with the surface and estimate the deformation field  $u$ , we need to solve the minimization problem of Equation (D.1). To this purpose, let us consider the following constrained minimization problem, which is equivalent to (D.1):

$$\min_{\substack{\mathbf{d}=\nabla z \\ \mathbf{V}=\nabla u}} \alpha A(\mathbf{d}) + \beta B(z, \mathbf{u}) + \gamma G(\mathbf{d}) + W(\mathbf{V}). \quad (\text{D.2})$$

Next, we reformulate this constrained minimization problem as an unconstrained optimization task. This can be done with an augmented Lagrangian approach, which translates the constraints into pairs of a Lagrangian multiplier and a penalty term. The augmented Lagrangian associated with (D.2) reads

$$\mathcal{L}(z, \mathbf{d}, \mathbf{u}, \mathbf{V}, \boldsymbol{\lambda}_z, \boldsymbol{\lambda}_u) = \mathcal{E}_c(z, \mathbf{d}, \mathbf{u}, \mathbf{V}) + \int_{\Omega} \boldsymbol{\lambda}_z \cdot (\mathbf{d} - \nabla z) + \frac{r_z}{2} |\mathbf{d} - \nabla z|_2^2 + \boldsymbol{\lambda}_u \cdot (\mathbf{V} - \nabla \mathbf{u}) + \frac{r_u}{2} |\mathbf{V} - \nabla \mathbf{u}|_2^2$$

where the vectorial functions  $\boldsymbol{\lambda}_z, \boldsymbol{\lambda}_u$  are Lagrange multipliers, and  $r_z, r_u$  are positive constants.

The constraint minimization problem (D.2) reduces to finding the saddle-point of the augmented Lagrangian energy  $\mathcal{L}$ . To find a saddle-point, the following algorithm can be adopted: initialize the variables and Lagrange multipliers to zero, at each iteration  $k$  find an approximate minimizer of  $\mathcal{L}(z, \mathbf{d}, \mathbf{u}, \mathbf{V}, \boldsymbol{\lambda}_z^{k-1}, \boldsymbol{\lambda}_u^{k-1})$  with respect to the variables  $z, \mathbf{d}, \mathbf{u}, \mathbf{V}$ , and update the Lagrange multipliers with the residuals associated with each constraint. See, e.g., Glowinski and Le Tallec [1989].

In general, it is difficult to find the exact minimizer of the Lagrangian  $\mathcal{L}$  with respect to the variables  $z, \mathbf{d}, \mathbf{u}, \mathbf{V}$  simultaneously. However, experiments show that a good approximation can be found by alternating the minimization of  $\mathcal{L}$  with respect to each variable while considering the others fixed. The resulting method is equivalent to the alternating direction method of multipliers. The resulting iterative method is summarized in Algorithm 8.

The next step is to determine efficient solutions for the sub-minimization problems (D.3)-(D.6). To simplify notation we omit the super-index in the sub-minimization problems.

---

**Algorithm 8** Augmented Lagrangian method to solve the constraint minimization (D.2)

---

- 1: Initialize  $z, \mathbf{d}, \mathbf{u}, \mathbf{V}, \boldsymbol{\lambda}_z, \boldsymbol{\lambda}_u$
- 2: For each iteration  $l = 1, 2, \dots$ , find an approximate minimizer of  $\mathcal{L}$  with respect to variables  $(z, \mathbf{d}, \mathbf{u}, \mathbf{V})$  with fixed Lagrange multipliers  $\boldsymbol{\lambda}_z^l, \boldsymbol{\lambda}_u^l$ :

$$z^l = \underset{z}{\operatorname{argmin}} \mathcal{L}(z, \mathbf{d}^{l-1}, \mathbf{u}^{l-1}, \mathbf{V}^{l-1}, \boldsymbol{\lambda}_z^{l-1}, \boldsymbol{\lambda}_u^{l-1}) \quad (\text{D.3})$$

$$\mathbf{d}^l = \underset{\mathbf{d}}{\operatorname{argmin}} \mathcal{L}(z^l, \mathbf{d}, \mathbf{u}^{l-1}, \mathbf{V}^{l-1}, \boldsymbol{\lambda}_z^{l-1}, \boldsymbol{\lambda}_u^{l-1}) \quad (\text{D.4})$$

$$\mathbf{u}^l = \underset{\mathbf{u}}{\operatorname{argmin}} \mathcal{L}(z^l, \mathbf{d}^l, \mathbf{u}, \mathbf{V}^{l-1}, \boldsymbol{\lambda}_z^{l-1}, \boldsymbol{\lambda}_u^{l-1}) \quad (\text{D.5})$$

$$\mathbf{V}^l = \underset{\mathbf{V}}{\operatorname{argmin}} \mathcal{L}(z^l, \mathbf{d}^l, \mathbf{u}^l, \mathbf{V}, \boldsymbol{\lambda}_z^{l-1}, \boldsymbol{\lambda}_u^{l-1}) \quad (\text{D.6})$$

- 3: Update Lagrange multipliers

$$\boldsymbol{\lambda}_z^l = \boldsymbol{\lambda}_z^{l-1} + r_z(\mathbf{d}^l - \nabla z^l) \quad (\text{D.7})$$

$$\boldsymbol{\lambda}_u^l = \boldsymbol{\lambda}_u^{l-1} + r_u(\mathbf{V}^l - \nabla \mathbf{u}^l) \quad (\text{D.8})$$

- 4: Stop the iterative process when  $\frac{\|\mathbf{u}^l - \mathbf{u}^{l-1}\|}{\|\mathbf{u}^l\|} < \epsilon, \frac{\|z^l - z^{l-1}\|}{\|z^l\|} < \epsilon$ .
- 

### D.1.1 Minimization with respect to $z$

If we have an OPC model  $z_0$  of the printed surface, the sub-minimization problem with respect to  $z$  reads

$$\min_z \int_{\Omega} \frac{\beta}{2} (z - \bar{z}_0)^2 + \boldsymbol{\lambda}_z \cdot (\mathbf{d} - \nabla z) + \frac{r_z}{2} |\mathbf{d} - \nabla z|_2^2, \quad (\text{D.9})$$

where  $\bar{z}_0 = z_0(\mathbf{x} + \mathbf{u})$  does not depend on  $z$ .

We observe that it is a convex minimization problem, and therefore we can find a minimizer by solving its optimality conditions. Euler-Lagrange gives us the following optimality condition

$$\beta z - r_z \Delta z = \boldsymbol{\lambda}_z + \beta \bar{z}_0 - \operatorname{div}(\boldsymbol{\lambda}_z + r_z \mathbf{d}). \quad (\text{D.10})$$

If we discretize Equation (D.10), we obtain a system of linear equations  $A_z z = b$  with a positive definite matrix. In particular we have

$$\begin{aligned} A_z &= \beta I_n + r_z D_x^T D_x + r_z D_y^T D_y \\ b &= \beta \bar{z}_0 + D_x^T (r_d d_x + \lambda_{z,x}) + D_y^T (r_d d_y + \lambda_{z,y}). \end{aligned} \quad (\text{D.11})$$

Note that matrix  $A_z$  is block circulant and we can use the Fourier transform  $\mathcal{F}$  to decompose it as  $A_z = \mathcal{F}^T D \mathcal{F}$ , with  $D$  a diagonal matrix. Consequently, the system (D.11) can easily be solved in the Fourier domain. In practice we use the FFT transform instead of doing the matrix multiplications with  $\mathcal{F}$  and  $\mathcal{F}^T$ , which gives us the solution  $z = \mathcal{F}^T (D^{-1} \mathcal{F} b)$  with a

complexity  $\mathcal{O}(n \log n)$ .

If we compare the surface to the expected mask in terms of the Heaviside function, we have the following minimization problem

$$\min_z \int_{\Omega} \frac{\beta}{2} (H_{\theta}(z) - h)^2 + \boldsymbol{\lambda}_z \cdot (\mathbf{d} - \nabla z) + \frac{r_z}{2} |\mathbf{d} - \nabla z|_2^2, \quad (\text{D.12})$$

where  $h = H(\phi(\mathbf{x} + \mathbf{u}))$  does not depend on  $z$ , and  $H_{\theta}(x) = H(x - \theta)$  denotes the composition of the Heaviside function with the translation by  $\theta$ .

In this case we do not have a closed-form solution for the minimization and require an iterative algorithm. In practical terms, we derive a fast semi-implicit minimization algorithm, that we initialize with the value of the previous iterate in the augmented Lagrangian Algorithm 8. The resulting minimization with respect to  $z$  converges within 2-5 iterations.

The Euler-Lagrange equations associated with Equation (D.12) are

$$-r_z \Delta z = \beta(H_{\theta}(z) - h)\delta_{\theta}(z) + \boldsymbol{\lambda}_z + \beta \bar{z}_0 - \text{div}(\mathbf{v}), \quad (\text{D.13})$$

where  $\mathbf{v} = \boldsymbol{\lambda}_z + r_z \mathbf{d}$ , and  $\delta(\cdot)$  is the derivative of the Heaviside function, that is, the Dirac distribution. Given a small  $\epsilon > 0$ , the system is equivalent to

$$\epsilon z - r_z \Delta z = \epsilon z + \beta(H_{\theta}(z) - h)\delta_{\theta}(z) + \boldsymbol{\lambda}_z + \beta \bar{z}_0 - \text{div} \mathbf{v} \quad (\text{D.14})$$

and can be solved with the following fixed-point iterative method (iterations on  $k$ )

$$\epsilon z^{k+1} - r_z \Delta z^{k+1} = \epsilon z^k + \beta(H_{\theta}(z^k) - h)\delta_{\theta}(z^k) + \boldsymbol{\lambda}_z + \beta \bar{z}_0 - \text{div} \mathbf{v}. \quad (\text{D.15})$$

At each iteration we have a linear system equivalent to (D.10) that we solve again in the Fourier domain.

### D.1.2 Minimization with respect to $\mathbf{d}$

If we define  $\mathbf{v} = \nabla z - \frac{\boldsymbol{\lambda}_z}{r_z}$ , the minimization problem with respect to  $\mathbf{d} = (d_x, d_y)$  reads

$$\min_{\mathbf{d}} \int_{\Omega} \alpha (R(\mathbf{d}) - I_0)^2 + \gamma w |\nabla \mathbf{d}|^2 + r_z |\mathbf{d} - \mathbf{v}|_2^2. \quad (\text{D.16})$$

The ideal reflectance function for SEM is given by Equation (3.4), but in fact, at the nanometre scale the reflectance map is considerably more complex and has to account for proximity effects. As a result, Equation (3.4) is only an approximation of the SEM image acquisition system, and it can be equally modelled by its first order Taylor expansion in the minimization process. Such an approach has also been shown to improve convergence in Zheng and Chellappa [1991]. At iteration  $k + 1$  we approximate the reflectance map around the previous



iterate by

$$R(\mathbf{d}) \approx R^k + (d_x - d_x^k) R_x^k + (d_y - d_y^k) R_y^k \quad \text{with} \quad \begin{cases} R^k = R(\mathbf{d}^k) \\ R_x^k = \partial_{d_x} R(\mathbf{d}^k) \\ R_y^k = \partial_{d_y} R(\mathbf{d}^k) \end{cases} \quad (\text{D.17})$$

Defining  $\bar{R}^k = R(\mathbf{d}^k) - R_x^k d_x^k - R_y^k d_y^k - I_0$ , at each iteration the minimization can be rewritten as

$$\min_{\mathbf{d}} \int_{\Omega} \alpha (R_x^k d_x + R_y^k d_y + \bar{R}^k)^2 + \gamma w |\nabla \mathbf{d}|^2 + r_z |\mathbf{d} - \mathbf{v}|^2. \quad (\text{D.18})$$

For simplicity we minimize independently for each component of  $\mathbf{d}$  and correct this approximation through the iterative method. A simultaneous minimization could also be performed efficiently solving a  $2 \times 2$  linear system of equations for each pixel in the frequency domain, see e.g., Hahn et al. [2011].

The Euler-Lagrange equations associated with the minimization with respect to  $d_x$  and  $d_y$  are

$$\begin{aligned} r_z d_x + \alpha R_x^k d_x + \gamma \nabla w \cdot \nabla d_x - \gamma w \Delta d_x &= r_z v_x - \alpha R_x^k (R_y^k d_y + \bar{R}^k) \\ r_z d_y + \alpha R_y^k d_y + \gamma \nabla w \cdot \nabla d_y - \gamma w \Delta d_y &= r_z v_y - \alpha R_y^k (R_x^k d_x + \bar{R}^k). \end{aligned}$$

If we discretize the optimality conditions, we have a linear system of equations  $A_x d_x = b_x$  and  $A_y d_y = b_y$  with

$$\begin{aligned} A_x &= r_z I_n + \alpha D_{R_x k} + \gamma D_x^T D_w D_x + \gamma D_y^T D_w D_y \\ b_x &= \alpha D_{R_x k} (D_{R_y k} d_y - D_{\bar{R}^k}) + r_z v_x \\ A_y &= r_z I_n + \alpha D_{R_y k} + \gamma D_x^T D_w D_x + \gamma D_y^T D_w D_y \\ b_y &= \alpha D_{R_y k} (D_{R_x k} d_x - D_{\bar{R}^k}) + r_z v_y, \end{aligned}$$

where  $D_{R_x k}$ ,  $D_{R_y k}$  and  $D_w$  are diagonal matrices, with entries associated with the values of functions  $R_x^k$ ,  $R_y^k$  and  $w$  in  $\Omega$ . Matrices  $A_x$ ,  $A_y$  are then symmetric, sparse, and diagonally dominant, and the systems can be inverted with an iterative Gauss-Seidel method. As a consequence of the sparsity of the matrices, each iteration of the Gauss-Seidel method is extremely fast and convergence is achieved within a few updates.

### D.1.3 Minimization with respect to $\mathbf{u}$

If we have an OPC model  $z_0$  of the printed surface, the sub-minimization problem with respect to  $\mathbf{u}$  reads

$$\min_{\mathbf{u}} \int_{\Omega} \frac{\beta}{2} (z(\mathbf{x}) - z_0(\mathbf{x} + \mathbf{u}))^2 + \boldsymbol{\lambda}_u \cdot (V(\mathbf{x}) - \nabla \mathbf{u}(\mathbf{x})) + \frac{r_u}{2} |V(\mathbf{x}) - \nabla \mathbf{u}(\mathbf{x})|^2. \quad (\text{D.19})$$

The associated Euler-Lagrange equation is

$$\begin{aligned} -r_u \Delta u(\mathbf{x}) &= b(u) \\ b(u) &= -\beta z(\mathbf{x}) - \operatorname{div}(\boldsymbol{\lambda}_u(\mathbf{x}) + r_u V(\mathbf{x})) + \beta z_0(\mathbf{x} + \mathbf{u}) \nabla z_0(\mathbf{x} + \mathbf{u}). \end{aligned}$$

If we compare the surface to the expected mask in terms of the Heaviside function, we have the following minimization problem

$$\min_{\mathbf{u}} \int_{\Omega} \frac{\beta}{2} (h_{\theta}(z(\mathbf{x})) - H(\phi(\mathbf{x} + \mathbf{u})))^2 + \boldsymbol{\lambda}_u \cdot (V(\mathbf{x}) - \nabla \mathbf{u}(\mathbf{x})) + \frac{r_u}{2} |V(\mathbf{x}) - \nabla \mathbf{u}(\mathbf{x})|^2. \quad (\text{D.20})$$

The associated Euler-Lagrange equation is now

$$\begin{aligned} -r_u \Delta u(\mathbf{x}) &= b(u) \\ b(u) &= -\beta h_z(\mathbf{x}) - \operatorname{div}(\boldsymbol{\lambda}_u(\mathbf{x}) + r_u V(\mathbf{x})) + \beta H(\phi(\mathbf{x} + \mathbf{u})) \delta(\phi(\mathbf{x} + \mathbf{u})) \nabla \phi(\mathbf{x} + \mathbf{u}). \end{aligned}$$

In both cases, the optimality conditions reduce to

$$-r_u \Delta \mathbf{u} = \mathbf{b}(\mathbf{u}) \quad (\text{D.21})$$

and can be solved iteratively with the same scheme. To that purpose, we choose a small  $\epsilon > 0$  and initialize the following fixed-point method with the solution to the previous Lagrangian iterate

$$\epsilon \mathbf{u}^{t+1} - r_u \Delta \mathbf{u}^{t+1} = \mathbf{b}(\mathbf{u}^t) + \epsilon \mathbf{u}^t. \quad (\text{D.22})$$

This fixed-point scheme is discretized as a system of linear equations that is equivalent to (D.11) and can be efficiently solved by means of the FFT transform or a few iterations of Gauss-Seidel. We have experimentally observed that 2-5 iterations of this fixed-point algorithm are enough to achieve convergence.

#### D.1.4 Minimization with respect to $V$

Defining  $W = \nabla \mathbf{u} - \frac{\lambda \mathbf{u}}{r_u}$ , the minimization with respect to  $V$  reads

$$\min_V \int_{\Omega} |\mathbf{V} - \mathbf{W}|_F^2 + W \left( \frac{\mathbf{V} + \mathbf{V}^T + \mathbf{V}^T \mathbf{V}}{2} \right), \quad (\text{D.23})$$

where  $|\cdot|_F$  designates the Frobenius norm. In this case the minimization functional contains no derivatives and is decoupled for each pixel, that is, we only need to solve a scalar minimization problem independently for each pixel. To that purpose we use a semi-implicit gradient descent method with time step  $\tau$ . At each pixel we denote  $\mathbf{V} = \begin{pmatrix} V_{11} & V_{12} \\ V_{21} & V_{22} \end{pmatrix}$  and iteratively update  $\mathbf{V}$  pixel-wise with the following rule

$$\begin{aligned} V_{11}^{k+1} &= \frac{V_{11}^k + \tau (r_u W_{11} - \lambda c_{01}^k - \mu c_{02}^k - \mu c_{04}^k V_{12}^k)}{1 + \tau (r_u + \lambda c_{01}^k + \mu c_{02}^k)} \\ V_{12}^{k+1} &= \frac{V_{12}^k + \tau (r_u W_{12} - \mu c_{04}^k (1 + V_{11}^k))}{1 + \tau (r_u + \lambda c_{01}^k + \mu c_{03}^k)} \\ V_{21}^{k+1} &= \frac{V_{21}^k + \tau (r_u W_{21} - \mu c_{04}^k (1 + V_{22}^k))}{1 + \tau (r_u + \lambda c_{01}^k + \mu c_{02}^k)} \\ V_{22}^{k+1} &= \frac{V_{22}^k + \tau (r_u W_{22} - \lambda c_{01}^k - \mu c_{03}^k - \mu c_{04}^k V_{21}^k)}{1 + \tau (r_u + \lambda c_{01}^k + \mu c_{03}^k)}, \end{aligned}$$

where

$$\begin{aligned} c_{01} &= V_{11} + V_{12} + \frac{1}{2} (V_{11}^2 + V_{12}^2 + V_{21}^2 + V_{22}^2) \\ c_{02} &= 2V_{11} + V_{11}^2 + V_{21}^2 \\ c_{03} &= 2V_{22} + V_{12}^2 + V_{22}^2 \\ c_{04} &= V_{12} + V_{21} + V_{11} V_{12} + V_{21} V_{22}. \end{aligned}$$

A similar scheme has been proposed in Guyader and Vese [2009] in the context of image registration, but the relation  $\mathbf{V} = \nabla \mathbf{u}$  is here considered a constraint instead of a penalty in the minimization. Experimentally we have found that few iterations are enough to reach the accuracy required for the Lagrangian update rule of Algorithm 8.

This finalizes the sub-optimization tasks, and therefore this appendix.



## Bibliography

- David Adalsteinsson and James A. Sethian. The Fast Construction of Extension Velocities in Level Set Methods. *Journal of Computational Physics*, 148(1):2–22, January 1999. ISSN 0021-9991. doi: <http://dx.doi.org/10.1006/jcph.1998.6090>. 96, 97, 98
- Coloma Ballester, Marcelo Bertalmio, Vincent Caselles, Guillermo Sapiro, and Joan Verdera. Filling-in by joint interpolation of vector fields and gray levels. *IEEE Transactions on Image Processing*, 10(8):1200–11, January 2001. ISSN 1057-7149. doi: 10.1109/83.935036. 19
- Wolfgang Beil and I. C. Carlsen. Surface reconstruction from stereoscopy and “shape from shading” in SEM images. *Machine vision and applications*, 4(4):271–285, September 1991. ISSN 0932-8092. doi: 10.1007/BF01815304. 60, 62, 63
- José M. Bioucas-dias and Mário A. T. Figueiredo. A New TwIST : Two-Step Iterative Shrinkage/Thresholding Algorithms for Image Restoration. *IEEE transactions on Image Processing*, 16(12):2992–3004, 2007. 82
- Michael J. Black, Guillermo Sapiro, David H. Marimont, and David Heeger. Robust anisotropic diffusion. *IEEE Transactions on Image Processing*, 7(3):421–432, January 1998. ISSN 1057-7149. doi: 10.1109/83.661192. 17, 26
- Peter Blomgren and Tony F. Chan. Color TV: total variation methods for restoration of vector-valued images. *IEEE Transactions on Image Processing*, 7(3):304–309, January 1998. ISSN 1057-7149. doi: 10.1109/83.661180. 16
- Thomas Blumensath and Mike E. Davies. Iterative hard thresholding for compressed sensing. *Applied and Computational Harmonic Analysis*, 27(3):265–274, 2009. 19
- Iva Bogdanova, Xavier Bresson, Jean-Philippe Thiran, and Pierre Vandergheynst. Scale space analysis and active contours for omnidirectional images. *IEEE transactions on image processing*, 16(7):1888–901, July 2007. ISSN 1057-7149. 44, 45, 106
- Stephen P. Boyd and Lieven Vandenberghe. *Convex optimization theory*, volume 25. Cambridge University Press, New York, 7 edition, June 2004. ISBN 9780521833783. doi: 10.1080/10556781003625177. 4, 15, 103

## Bibliography

---

- Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11):1222–1239, 2001. 80
- Kristian Bredies, Karl Kunisch, and Thomas Pock. Total generalized variation. *SIAM Journal on Imaging Sciences*, 3(3):492–526, 2010. 5
- Xavier Bresson, Pierre Vandergheynst, and Jean-Philippe Thiran. Multiscale Active Contours. *International Journal of Computer Vision*, 70(3):197–211, December 2006a. ISSN 0920-5691. doi: 10.1007/s11263-006-7462-3. 45, 53
- Xavier Bresson, Pierre Vandergheynst, and Jean-Philippe Thiran. A Variational Model for Object Segmentation Using Boundary Information and Shape Prior Driven by the Mumford-Shah Functional. *International Journal of Computer Vision*, 68(2):145–162, March 2006b. ISSN 0920-5691. doi: 10.1007/s11263-006-6658-x. 65
- Xavier Bresson, Selim Esedoglu, Pierre Vandergheynst, Jean-Philippe Thiran, and Stanley Osher. Fast Global Minimization of the Active Contour/Snake Model. *Journal of Mathematical Imaging and Vision*, 28(2):151–167, July 2007. ISSN 0924-9907. doi: 10.1007/s10851-007-0002-0. 42
- Chaim Broit. *Optimal registration of deformed images*. PhD thesis, University of Pennsylvania, 1981. 68
- Antoni Buades, Bartomeu Coll, and Jean-Michel Morel. A review of image denoising algorithms, with a new one. *SIAM Journal on Multiscale Modelling and Simulation*, 4(2):490–530, 2005. 22, 24
- Martin Burger, Guy Gilboa, Stanley Osher, and Jinjun Xu. Nonlinear inverse scale space methods. *SIAM Journal on Imaging Sciences*, 4(1):179–212, 2006. ISSN 19364954. doi: 10.1137/070689954. 105
- Emmanuel J. Candès and Justin Romberg. Practical signal recovery from random projections. *Proceedings of SPIE*, 5674:76–86, 2005. 15
- Emmanuel J. Candès, Michael B. Wakin, and Stephen P. Boyd. Enhancing Sparsity by Reweighted  $l_1$  Minimization. *Journal of Fourier Analysis and Applications*, 14(5-6):877–905, October 2008. ISSN 1069-5869. doi: 10.1007/s00041-008-9045-x. 19, 20
- Emmanuel J. Candès, Yonina C. Eldar, Deanna Needell, and Paige Randall. Compressed sensing with coherent and redundant dictionaries. *Applied and Computational Harmonic Analysis*, 31(1):59–73, July 2011. ISSN 10635203. doi: 10.1016/j.acha.2010.10.002. 14
- Vincent Caselles, Ron Kimmel, and Guillermo Sapiro. Geodesic active contours. *International Journal of Computer Vision*, 22(1):61–79, January 1997. ISSN 1057-7149. doi: 10.1109/83.951533. 9, 11, 41, 42, 58, 65, 81, 82, 89

- Antonin Chambolle. An Algorithm for Total Variation Minimization and Applications. *Journal of Mathematical Imaging and Vision*, 20(1/2):89–97, January 2004. ISSN 0924-9907. doi: 10.1023/B:JMIV.0000011321.19549.88. 27
- Antonin Chambolle and Thomas Pock. A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging. *Journal of Mathematical Imaging and Vision*, 40(1): 120–145, December 2010. ISSN 0924-9907. doi: 10.1007/s10851-010-0251-1. 27
- Tony F. Chan and Luminita Vese. Active contours without edges. *IEEE Transactions on Image Processing*, 10(2):266–277, January 2001. ISSN 1057-7149. doi: 10.1109/83.902291. 3, 9, 11, 41, 42, 58, 81, 82, 89
- Tony F. Chan and Wei Zhu. Level Set Based Shape Prior Segmentation. *IEEE International Conference of Computer Vision and Pattern Recognition*, 2:1164–1170, 2005. doi: 10.1109/CVPR.2005.212. 65
- Tony F. Chan, Gene H. Golub, and Pep Mulet. A Nonlinear Primal-Dual Method for Total Variation-Based Image Restoration. *SIAM Journal on Scientific Computing*, 20(6):1964, 1999. ISSN 10648275. doi: 10.1137/S1064827596299767. 27
- Tony F. Chan, Selim Esedoglu, and Mila Nikolova. Algorithms for Finding Global Minimizers of Image Segmentation and Denoising Models. *SIAM Journal on Applied Mathematics*, 66(5):1632–1648, 2006. ISSN 00361399. doi: 10.1137/040615286. 3, 4, 42, 80
- Rick Chartrand. Exact Reconstruction of Sparse Signals via Nonconvex Minimization. *IEEE Signal Processing Letters*, 14(10):707–710, 2007. 19
- Yunmei Chen and Hemant D. Tagare. Using prior shapes in geometric active contours in a variational framework. *International Journal of Computer Vision*, 50(3):315–328, 2002. 65
- Gary E. Christensen, Richard D. Rabbitt, and Michael I. Miller. Deformable templates using large deformation kinematics. *IEEE Transactions on Image Processing*, 5(10):1435–1447, 1996. 68
- Richard Courant, Kurt Friedrichs, and Hans Lewy. On the Partial Difference Equations of Mathematical Physics. *IBM Journal of Research*, 3(March):215–234, 1967. 7, 80
- Daniel Cremers, Florian Tischhäuser, Joachim Weickert, and Christoph Schnörr. Diffusion Snakes: Introducing Statistical Shape Knowledge into the Mumford-Shah Functional. *International journal of computer vision*, 50(3):295–313, 2002. 65
- Daniel Cremers, Stanley Osher, and Stefano Soatto. Kernel density estimation and intrinsic alignment for knowledge-driven segmentation: Teaching level sets to walk. *Pattern Recognition*, pages 1–28, 2004. 65
- Reinhard Danzl and Stefan Scherer. Integrating Shape from Shading and Shape from Stereo for Variable Reflectance Surface Reconstruction from SEM Images. In *Workshop of the Austrian Association for Pattern Recognition*, pages 281–288, 2002. 60, 62, 63

## Bibliography

---

- Lorina Dascal, Guy Rosman, and Ron Kimmel. Efficient Beltrami filtering of color images via vector extrapolation. In *International Conference on on Scale Space and Variational Methods in Computer Vision*, pages 92–103. Springer-Verlag, 2007. 55
- Lorina Dascal, Guy Rosman, Xue-Cheng Tai, and Ron Kimmel. On semi-implicit splitting schemes for the Beltrami color flow. *Scale Space and Variational Methods in Computer Vision*, pages 259–270, 2009. 55
- Manfredo P. do Carmo. *Differential Geometry of Curves and Surfaces*. Prentice-Hall, Englewood Cliffs, NJ, 1976. 44
- Fangfang Dong, Zhen Liu, Dexing Kong, and Kefeng Liu. An Improved LOT Model for Image Restoration. *Journal of Mathematical Imaging and Vision*, 34(1):89–97, May 2009. ISSN 0924-9907. 19
- David L. Donoho. De-Noising by Soft-Thresholding. *IEEE Transactions on Information Theory*, 41(3):613–627, 1995. 89
- Jean-Denis Durou, Maurizio Falcone, and Manuela Sagona. Numerical methods for shape-from-shading: A new survey with benchmarks. *Computer Vision, Graphics, and Image Processing*, 109(1):22–43, January 2008. ISSN 10773142. doi: 10.1016/j.cviu.2007.09.003. 64
- Timothy P. Ellison and Christopher J. Taylor. Calculating the surface topography of integrated circuit wafers from SEM images. *Image and Vision Computing*, 9(1):3–9, February 1991. ISSN 02628856. doi: 10.1016/0262-8856(91)90042-N. 60, 62, 63
- Abderrahim Elmoataz, Olivier Lezoray, and Sébastien Bougleux. Nonlocal Discrete Regularization on Weighted Graphs: A Framework for Image and Manifold Processing. *IEEE Transactions on Image Processing*, 17(7):1047–1060, 2008. ISSN 10577149. doi: 10.1109/TIP.2008.924284. 22
- Virginia Estellers, Dominique Zosso, Xavier Bresson, and Jean-Philippe Thiran. Harmonic Active Contours for multichannel image processing. In *IEEE International Conference on Image Processing*, number 200021, pages 3141–3144, 2011. ISBN 9781457713033. 55
- Virginia Estellers, Dominique Zosso, Rongjie Lai, Stanley Osher, Jean-Philippe Thiran, and Xavier Bresson. Efficient algorithm for level set method preserving distance function. *IEEE Transactions on Image Processing*, 21(12):4722–4734, December 2012. ISSN 1941-0042. doi: 10.1109/TIP.2012.2202674. 11, 79
- Virginia Estellers, Jean-Philippe Thiran, and Xavier Bresson. Enhanced Compressed Sensing Recovery with Level Set Normals. *IEEE Transactions on Image Processing*, to appear:1–16, March 2013a. ISSN 1941-0042. doi: 10.1109/TIP.2013.2253484. 11, 13
- Virginia Estellers, Jean-philippe Thiran, and Maria Gabrani. Surface Reconstruction from Microscopic Images in Optical Lithography. *IEEE Transactions on Image Processing*, (Submitted to):1–12, 2013b. 11, 59



- Virginia Estellers, Dominique Zosso, Xavier Bresson, and Jean-philippe Thiran. Harmonic Active Contours. *IEEE Transactions on Image Processing*, under revision:1–12, 2013c. 11, 41
- Olivier Faugeras and Renaud Keriven. level set methods and the Stereo Problem. *INRIA Research Report*, 33:272–283, 1997. 81
- Bernd Fischer and Jan Modersitzki. Curvature based image registration. *Journal of Mathematical Imaging and Vision*, pages 81–85, 2003. 67
- Bernd Fischer and Jan Modersitzki. A unified approach to fast image registration and a new curvature based registration technique. *Linear Algebra and its Applications*, 380:107–124, March 2004. ISSN 00243795. doi: 10.1016/j.laa.2003.10.021. 67
- Alban Foulonneau, Pierre Charbonnier, and Fabrice Heitz. Affine-invariant geometric shape priors for region-based active contours. *IEEE transactions on pattern analysis and machine intelligence*, 28(8):1352–7, August 2006. ISSN 0162-8828. doi: 10.1109/TPAMI.2006.154. 65
- Alejandro Frangi, Wiro J. Niessen, Koen L. Vincken, and Max A. Viergever. Multiscale vessel enhancement filtering. In *Medical Image Computing and Computer-Assisted Intervention*, volume 1496, pages 130–137, 1998. 65
- Robert T. Frankot and Rama Chellappa. A method for enforcing integrability in shape from shading algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(4): 439–451, July 1988. ISSN 01628828. doi: 10.1109/34.3909. 64
- Guy Gilboa and Stanley Osher. Nonlocal operators with applications to image processing. *Multiscale Modeling and Simulation*, 7(3):1005–1028, 2008. ISSN 15403459. doi: 10.1137/070698592. 22, 23, 27, 106
- Roland Glowinski and Patrick Le Tallec. *Augmented Lagrangian and Operator-Splitting Methods in Nonlinear Mechanics*. SIAM, 1989. 82, 84, 85, 120, 128
- Roman Goldenberg, Ron Kimmel, Ehud Rivlin, and Michael Rudzsky. Cortex Segmentation : A Fast Variational Geometric Approach. *IEEE transactions on Medical Imaging*, 21(2): 1544–1551, 2002. 92
- Tom Goldstein and Stanley Osher. The Split Bregman Method for l-1 Regularized Problems. *SIAM Journal on Imaging Sciences*, 2(2):323–343, 2009. ISSN 19364954. doi: 10.1137/080725891. 7, 27, 43, 82, 84, 127
- Tom Goldstein, Xavier Bresson, and Stanley Osher. Geometric Applications of the Split Bregman Method: Segmentation and Surface Reconstruction. *Journal of Scientific Computing*, 45(1-3):272–293, November 2009. ISSN 0885-7474. doi: 10.1007/s10915-009-9331-z. 3, 42, 127
- Jose Gomes and Olivier Faugeras. Reconciling Distance Functions and Level Sets. *Scale Space Theories in Computer Vision*, pages 70–81, 1999. 80, 81, 90

## Bibliography

---

- Weihong Guo and Wotao Yin. EdgeCS: edge guided compressive sensing reconstruction. In *Proceedings of Visual Communications and Image Processing*, number 2, pages 17–23, 2010. 16, 19, 20, 27
- Carole Le Guyader and Luminita Vese. A combined segmentation and registration framework with a nonlinear elasticity smoother. *Scale Space and Variational Methods in Computer Vision*, 2:600–611, 2009. 66, 68, 133
- Jooyoung Hahn, Xue-Cheng Tai, Sofia Borok, and Alfred Marcel Bruckstein. Orientation-Matching Minimization for Image Denoising and Inpainting. *International Journal of Computer Vision*, 92(3):308–324, July 2010. ISSN 0920-5691. doi: 10.1007/s11263-010-0371-5. 19
- Jooyoung Hahn, Chunlin Wu, and Xue-Cheng Tai. Augmented Lagrangian Method for Generalized TV-Stokes Model. *Journal of Scientific Computing*, 50(2):235–264, April 2011. ISSN 0885-7474. doi: 10.1007/s10915-011-9482-6. 17, 19, 34, 131
- Lin He, Ti-Chiun Chang, Stanley Osher, Tong Fang, and Peter Speier. MR image reconstruction by using the iterative refinement method and nonlinear inverse scale space methods. *UCLA CAM Report*, 06-35, 2006. 20
- Hugues Hoppe, Tony DeRose, Tom Duchamp, John McDonald, and Werner Stuetzle. Surface Reconstruction from Unorganized Points. *ACM Computer Graphics*, 26(2):71–78, 1992. 83
- Berthold K. P. Horn. *Shape from shading: A method for obtaining the shape of a smooth opaque object from one view*. PhD thesis, Massachusetts Institute of Technology, 1970. 62
- Berthold K. P. Horn. Height and gradient from shading. *International Journal of Computer Vision*, 5(1):37–75, August 1990. ISSN 0920-5691. doi: 10.1007/BF00056771. 64
- Katsushi Ikeuchi and Berthold K. P. Horn. Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, 17(1-3):141–184, August 1981. ISSN 00043702. doi: 10.1016/0004-3702(81)90023-0. 64
- A Iserles. *A first course in the numerical analysis of differential equations*. Cambridge University Press, Cambridge; New York, 2009. ISBN 9780521734905 0521734908. v
- AG Jones and Christopher J. Taylor. Robust shape from shading. *Image and vision computing*, 12(7):411–421, September 1994. ISSN 02628856. doi: 10.1016/0262-8856(94)90025-6. 62
- AG Jones and Christopher J. Taylor. Scale space surface recovery using binocular shading and stereo information. *Proceedings of the BMVC*, pages 77–86, 1995. 60, 62, 63
- Ali Kayaalp, Ravishankar A. Rao, and Ramesh Jain. Scanning electron microscope-based stereo analysis. *International Conference on Computer Vision*, 3(4):231–246, September 1990. ISSN 0932-8092. doi: 10.1007/BF01211849. 60, 62, 63

- Satyanad Kichenassamy, Arun Kumar, Peter Olver, Allen Tannenbaum, and Anthony Yezzi. Gradient flows and geometric active contour models. In A Kumar, P Olver, A Tannenbaum, and A Yezzi, editors, *Computer Vision, IEEE International Conference on*, volume 0, pages 810–815. IEEE Comput. Soc. Press, June 1995. ISBN 0-8186-7042-8. doi: 10.1109/ICCV.1995.466855. 9, 42, 65, 81, 82
- Ron Kimmel. Fast Edge Integration. *Geometric Level Set Methods in Imaging Vision and Graphics*, pages 59–77, 2003. doi: 10.1007/b97541. 43, 54, 57
- Ron Kimmel and Alfred Marcel Bruckstein. On edge detection, edge integration and geometric active contours. In *Proceedings of International Symposium on Mathematical Morphology, ISMM*, pages 37–45, 2002. 9
- Ron Kimmel and Alfred Marcel Bruckstein. Regularized Laplacian zero crossings as optimal edge integrators. *International Journal of Computer Vision*, 53(3):225–243, 2003. 11, 41, 42, 43, 53, 54, 56, 57, 58
- Ron Kimmel and Nir Sochen. Orientation Diffusion or How to Comb a Porcupine. *Journal of Visual Communication and Image Representation*, 13(2000-02):238–248, March 2002. ISSN 10473203. doi: 10.1006/jvci.2001.0501. 44
- Ron Kimmel, Ravikanth Malladi, and Nir Sochen. Images as embedding maps and minimal surfaces: movies, color, and volumetric medical images. *IEEE International Conference of Computer Vision and Pattern Recognition*, 39(2):350–355, 1997. doi: 10.1109/CVPR.1997.609348. 45, 47, 52, 53
- Yvan G. Leclerc and Aaron F. Bobick. The Direct Computation of Height from Shading. In *IEEE International Conference of Computer Vision and Pattern Recognition*, pages 552–558, 1991. 64
- Victor Lempitsky and Yuri Boykov. Global Optimization for Shape Fitting. *IEEE International Conference of Computer Vision and Pattern Recognition*, pages 1–8, June 2007. doi: 10.1109/CVPR.2007.383293. 83, 90
- Michael E. Leventon, Eric W. Grimson, and Olivier Faugeras. Statistical shape influence in geodesic active contours. In *IEEE International Conference of Computer Vision and Pattern Recognition*, pages 1316–1323, 2000. 65, 81
- Chunming Li, Chenyang Xu, Changfeng Gui, and Martin D. Fox. Level Set Evolution Without Re-initialization : A New Variational Formulation. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 430–436, 2005. 81, 96, 97, 98, 100
- Chunming Li, Chenyang Xu, Changfeng Gui, and Martin D. Fox. Distance regularized level set evolution and its application to image segmentation. *IEEE Transactions on Image Processing*, 19(12):3243–54, December 2010. ISSN 1941-0042. doi: 10.1109/TIP.2010.2069690. 81, 96, 97, 99, 100

## Bibliography

---

- Tony Lindeberg. Edge detection and ridge detection with automatic scale selection. In *IEEE International Conference of Computer Vision and Pattern Recognition*, pages 465–470. IEEE Comput. Soc. Press, 1996. ISBN 0-8186-7258-7. doi: 10.1109/CVPR.1996.517113. 65
- Pierre-Louis Lions and B Mercier. Splitting Algorithms for the Sum of Two Nonlinear Operators. *SIAM Journal on Numerical Analysis*, 16(6):964–979, December 1979. ISSN 0036-1429. doi: 10.1137/0716071. 82
- Michael Lustig, J.H. Lee, David L. Donoho, and John M. Pauly. Faster Imaging with Randomly Perturbed , Undersampled Spirals and L-1 Reconstruction. In *Proceedings of the Annual Meeting of International Society for Magnetic Resonance in Medicine*, page 50, 2005. 15
- Marius Lysaker, Arvid Lundervold, and Xue-Cheng Tai. Noise removal using fourth-order partial differential equation with applications to medical magnetic resonance images in space and time. *IEEE Transactions on Image Processing*, 12(12):1579–90, January 2003. ISSN 1057-7149. doi: 10.1109/TIP.2003.819229. 19
- Marius Lysaker, Stanley Osher, and Xue-Cheng Tai. Noise Removal Using Smoothed Normals and Surface Fitting. *IEEE Transactions on Image Processing*, 13(10):1345–1357, 2004. 19
- Shiqian Ma, Yin Wotao, Yin Zhang, and Amit Chakraborty. An efficient algorithm for compressed MR imaging using total variation and wavelets. In *IEEE International Conference of Computer Vision and Pattern Recognition*, pages 1–8, June 2008. ISBN 978-1-4244-2242-5. doi: 10.1109/CVPR.2008.4587391. 15
- David MacDonald, Noor Kabani, David Avis, and Alan C. Evans. Automated 3-D extraction of inner and outer surfaces of cerebral cortex from MRI. *NeuroImage*, 12(3):340–56, September 2000. ISSN 1053-8119. doi: 10.1006/nimg.1999.0534. 92
- Chris Mack. extract\_CLbook.pdf. In 220, editor, *Fundamental principles of optical lithography: the science of microfabrication*, pages 121–124. Wiley, 2007. ISBN 978-0-470-01893-4. 60, 61, 66, 71
- Jan Modersitzki. *Numerical Methods for Image Registration (Numerical Mathematics and Scientific Computation)*. Oxford University Press, 2004. ISBN 0198528418. 66
- Hidetoshi Morokuma, Akiyuki Sugiyama, and Yasutaka Toyoda. A new matching engine between design layout and SEM image of semiconductor device. *Proceedings of SPIE*, 2005. 60
- David Mumford and Mubarak Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on pure and applied Mathematics*, XLII, 1989. 3, 4, 51
- Deanna Needell and Joel A. Tropp. CoSaMP: Iterative signal recovery from incomplete and inaccurate samples. *Applied and Computational Harmonic Analysis*, pages 1–30, 2009. 19

- Deanna Needell and Rachel Ward. Stable image reconstruction using total variation minimization. *arXiv preprint arXiv:1202.6429*, pages 1–23, 2012. 14
- Stanley Osher and Ronald P. Fedkiw. *Level set methods and dynamic implicit surfaces*, 2002. 80
- Stanley Osher and Nikos Paragios. *Geometric level set methods in imaging, vision, and graphics*. Springer, New York, 2006. ISBN 0387954880 9780387954882. 80
- Stanley Osher and James A. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics*, 79(1):12–49, November 1988. ISSN 00219991. doi: 10.1016/0021-9991(88)90002-2. 10, 42, 79, 80
- Stanley Osher, Martin Burger, Donald Goldfarb, Jinjun Xu, and Wotao Yin. An Iterative Regularization Method for Total Variation-Based Image Restoration. *Multiscale Modeling and Simulation*, 4(2):460, 2005. ISSN 15403459. doi: 10.1137/040605412. 20, 28
- Nikos Paragios. A level set approach for shape-driven segmentation and tracking of the left ventricle. *IEEE transactions on medical imaging*, 22(6):773–6, June 2003. ISSN 0278-0062. doi: 10.1109/TMI.2003.814785. 65
- Nikos Paragios and Rachid Deriche. Geodesic active contours and level sets for the detection and tracking of moving objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(3):266–280, March 2000. ISSN 01628828. doi: 10.1109/34.841758. 81
- Pietro Perona and Jitendra Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, July 1990. ISSN 01628828. doi: 10.1109/34.56205. 17
- Gabriel Peyré, Sébastien Bogleux, and Laurent Cohen. Non-local regularization of inverse problems. *European Conference on Computer Vision (ECCV)*, pages 57–68, 2008. 23
- Emmanuel Prados, Fabio Camilli, and Olivier Faugeras. A unifying and rigorous shape from shading method adapted to realistic data and applications. *Journal of Mathematical Imaging and Vision*, 3, 2006. 64
- Leili Baghaei Rad, Hanying Feng, Jun Ye, R. Fabian W. Pease, David G. Seiler, Alain C. Diebold, Robert McDonald, Michael C. Garner, Dan Herr, Rajinder P. Khosla, and Erik M. Secula. Computational Scanning Electron Microscopy. *International Conference on Frontiers of Characterization and Metrology*, 931(1):512–517, 2007. ISSN 0094243X. doi: 10.1063/1.2799427. 62
- Bhaskar D. Rao and Kenneth Kreutz-Delgado. An affine scaling methodology for best basis selection. *IEEE Transactions on Signal Processing*, 47(1):187–200, January 1999. 19

## Bibliography

---

- Eduard Reithmeier, Taras Vynnyk, and Thanin Schultheis. 3D-measurement using a scanning electron microscope. *Applied Mathematics and Computation*, 217(3):1193–1201, October 2010. ISSN 00963003. doi: 10.1016/j.amc.2010.01.107. 60, 62, 63
- R. Tyrell Rockafellar. Lagrange multipliers and optimality. *SIAM review*, 35(2):183–238, 1993. 103
- Guy Rosman. Polyakov Action for Efficient Color Image Processing. *UCLA CAM Report*, pages 1–21, 2010. 55, 123
- Mikael Rousson and Nikos Paragios. Shape priors for level set representations. In *European Conference on Computer Vision (ECCV)*, pages 78–92, 2002. 65
- Mark Rudelson and Roman Vershynin. Sparse reconstruction by convex relaxation: Fourier and Gaussian measurements. In *Conference on Information Sciences and Systems*, pages 207–212, March 2006. 15
- Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1-4):259–268, November 1992. ISSN 01672789. doi: 10.1016/0167-2789(92)90242-F. 5, 16, 19, 27
- Chen Sagiv, Nir Sochen, and Yehoshua Y. Zeevi. Integrated active contours for texture segmentation. *IEEE Transactions in Image Processing*, 15(6):1633–1646, June 2006. ISSN 1057-7149. 45, 52, 53, 106
- Alessandro Sarti, Ravikanth Malladi, and James A. Sethian. Subjective Surfaces : A Geometric Model for Boundary Completion. *International Journal of Computer Vision*, 46(3):201–221, 2002. 41, 43, 49, 52, 53
- Adam Seeger. *Surface reconstruction from AFM and SEM images*. PhD thesis, University of North Carolina, 2004. 66
- James A. Sethian. A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences*, 93(4):1591–1595, February 1996. ISSN 00278424. doi: 10.1073/pnas.93.4.1591. 57, 81
- Simon Setzer. Operator Splittings , Bregman Methods and Frame Shrinkage in Image Processing. *International Journal of Computer Vision*, 92(3):265–280, 2011. 82
- T Shibahara, T Minakawa, Michio Oikawa, Hiroyuki Shindo, H Sugahara, and Yutaka Hojo. A CD-gap-free contour extraction technique for OPC model calibration. *Proceedings of SPIE*, 7971(April):79710O–79710O, March 2011. 60
- Hiroyuki Shindo, Akiyuki Sugiyama, Hitoshi Komuro, Yutaka Hojo, Ryoichi Matsuoka, John L. Sturtevant, Thuy Do, Ir Kusnadi, Germain Fenger, Peter De Bisschop, and Jeroen Van de Kerckhove. High-precision contouring from SEM image in 32-nm lithography and beyond. *Proceedings of SPIE*, 7275:72751F–72751F–9, 2009. doi: 10.1117/12.814430. 60

- Chie Shishido, Maki Tanaka, and Mayuka Osaki. CD bias reduction in CD-SEM of very small line patterns: sidewall shape measurement using model-based library matching method. *Proceedings of SPIE*, 7638(March):76383I–76383I, March 2010. 60
- Nir Sochen, Ron Kimmel, and Ravikanth Malladi. A general framework for low level vision. *IEEE Transactions on Image Processing*, 7(3):310–318, January 1998. ISSN 1057-7149. doi: 10.1109/83.661181. 6, 41, 43, 45, 52, 53
- Petter Strandmark, Fredrik Kahl, and Niels Chr Overgaard. Optimizing parametric total variation models. In *IEEE International Conference of Computer Vision*, pages 2240–2247. Ieee, September 2009. ISBN 978-1-4244-4420-5. doi: 10.1109/ICCV.2009.5459464. 3, 4
- Mark Sussman, Peter Smereka, and Stanley Osher. A level set approach for computing solutions to incompressible two-phase flow, 1994. 81, 83
- Richard Szeliski. Fast shape from shading. *CVGIP: Image Understanding*, 53(2):129–153, March 1991. ISSN 1049-9660. doi: [http://dx.doi.org/10.1016/1049-9660\(91\)90023-I](http://dx.doi.org/10.1016/1049-9660(91)90023-I). 64
- Cyrus E. Tabery, Hidetoshi Morokuma, Akiyuki Sugiyama, and Lorena Page. Evaluation of OPC quality using automated edge placement error measurement with CD-SEM. *Proceedings of SPIE*, 2006. 60
- Cyrus E. Tabery, Hidetoshi Morokuma, Ryoichi Matsuoka, Lorena Page, George E. Bailey, Ir Kusnadi, and Thuy Do. SEM image contouring for OPC model calibration and verification. *Proceedings of SPIE*, 2007. 60
- Xue-Cheng Tai, Stanley Osher, and Randi Holm. Image Inpainting Using a TV-Stokes Equation. In *Image Processing Based on Partial Differential Equations*, Mathematics and Visualization, pages 3–22. Springer, 2007. ISBN 978-3-540-33267-1. 19
- Xue-Cheng Tai, Sofia Borok, and Jooyoung Hahn. Image Denoising Using TV-Stokes Equation with an Orientation-matching Minimization. In *Scale Space and Variational Methods in Computer Vision*, pages 490–501, 2009. 19
- Joel A. Tropp and Anna C. Gilbert. Signal Recovery From Random Measurements Via Orthogonal Matching Pursuit. *IEEE Transactions on Information Theory*, 53(12):4655–4666, 2007. 19
- Andy Tsai, Anthony Yezzi, William Wells, Clare Tempany, Dewey Tucker, Eric W. Grimson, and Alan S. Willsky. Model-based curve evolution technique for image segmentation. In *IEEE International Conference of Computer Vision and Pattern Recognition*, volume 00, pages 2–7, 2001. ISBN 0769512720. 65
- J N Tsitsiklis. Efficient algorithms for globally optimal trajectories, 1995. 81
- Alexander Vasilevskiy and Kaleem Siddiqi. Flux maximizing geometric flows. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(12):1565–1578, December 2002. ISSN 0162-8828. doi: 10.1109/TPAMI.2002.1114849. 42

## Bibliography

---

- John S. Villarrubia and Zejun J. Ding. Sensitivity of SEM width measurements to model assumptions. *Proceedings of SPIE*, 7272(Cd):72720R–72720R–15, 2009. doi: 10.1117/12.814300. 63
- John S. Villarrubia, András E. Vladár, Benjamin D. Bunday, and Michael Bishop. Dimensional Metrology of Resist Lines using a SEM Model-Based Library Approach. In *Proceedings of SPIE*, volume 199, pages 199–209, May 2004. doi: 10.1117/12.536871. 63
- John S. Villarrubia, Ronald G. Dixon, and András E. Vladár. Proximity-associated errors in contour metrology. *Proceedings of SPIE*, 7638:76380S–76380S, March 2010. 63
- Yilun Wang, Junfeng Yang, Wotao Yin, and Yin Zhang. A new alternating minimization algorithm for total variation image reconstruction. *SIAM Journal on Imaging*, pages 1–24, 2008. 7, 27, 84, 86, 127
- Chunlin Wu and Xue-Cheng Tai. Augmented Lagrangian method, dual methods, and split Bregman iteration for ROF, vectorial TV, and high order models. *SIAM Journal on Imaging Sciences*, 3(3):300–339, 2010. 7, 27
- Jian Ye, Xavier Bresson, Tom Goldstein, and Stanley Osher. A Fast Variational Method for Surface Reconstruction from Sets of Scattered Points. *UCLA CAM Report*, pages 1–22, 2010. 83
- Wotao Yin, Stanley Osher, Donald Goldfarb, and Jerome Darbon. Bregman Iterative Algorithms for  $\ell_1$ -Minimization with Applications to Compressed Sensing. *SIAM Journal on Imaging Sciences*, 1(1):143–168, 2008. 27
- Xiaolan Zeng, Lawrence H. Staib, Robert T. Schultz, and James S. Duncan. Segmentation and Measurement of the Cortex from 3D MR Images Using Coupled Surfaces Propagation. *IEEE transactions on Medical Imaging*, 18(10):100–111, 1999. 92
- Ruo Zhang, Ping-Sing Tsai, James Edwin Cryer, and Mubarak Shah. Shape-from-shading: a survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8):690–706, 1999. 64
- Xiaoqun Zhang, Martin Burger, Xavier Bresson, and Stanley Osher. Bregmanized nonlocal regularization for deconvolution and sparse reconstruction. *SIAM Journal on Imaging Sciences*, 3(3):253–276, 2010. 23, 27
- Hong-Kai Zhao. Implicit and Nonparametric Shape Reconstruction from Unorganized Data Using a Variational Level Set Method. *Computer Vision, Graphics, and Image Processing*, 80(3):295–314, December 2000. ISSN 10773142. doi: 10.1006/cviu.2000.0875. 81, 83
- Qinfen Zheng and Rama Chellappa. Estimation of illuminant direction, albedo, and shape from shading. *IEEE International Conference of Computer Vision and Pattern Recognition*, pages 540–545, 1991. 65, 130



- Dengyong Zhou and Bernhard Schölkopf. Regularization on Discrete Spaces. In WalterG. Kropatsch, Robert Sablatnig, and Allan Hanbury, editors, *Pattern Recognition. Lecture Notes in Computer Science*, volume 80 of *Lecture Notes in Computer Science*, pages 361–368. Springer, 2005. ISBN 978-3-540-28703-2. doi: 10.1007/11550518\\_45. 22
- Mingqiang Zhu, Stephen J. Wright, and Tony F. Chan. Duality-based algorithms for total-variation-regularized image restoration. *Computational Optimization and Applications*, 47 (3):377–400, December 2010. ISSN 0926-6003. doi: 10.1007/s10589-008-9225-2. 27



## Virginia Estellers

---

### CONTACT INFORMATION

*Professional  
address:* LTS5, STI-IEL Station 11  
CH-1015 EPFL  
Switzerland  
*Tel.:* +41 (0)21 693 46 82  
*e-mail:* virginia.estellers@epfl.ch  
*Web:* <http://lts5srv2.epfl.ch/~estellers>

### RESEARCH INTERESTS

**Mathematical modelling and computational techniques for inverse problems in imaging.**  
My research interests are in mathematical modelling and computational techniques for inverse problems in imaging, particularly using variational methods, convex optimization and partial differential equations. I am interested in the theoretical and physical aspects of the acquisition of images, their mathematical representations, and the development of efficient computational models. I am therefore able to apply my theoretical interests to a wide range of inverse problems in computer vision, from image segmentation to compressive sensing or shape-from-shading.

### EDUCATION

**Ph.D, Electrical Engineering, (expected) June 2013.**  
Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland.  
Research Advisor: Jean-Philippe Thiran.

**B.Sc. and M.Sc. in Mathematics, February 2008.**  
Polytechnic University of Catalonia, Barcelona, Spain .

**B.Sc. and M.Sc. in Electrical Engineering, July 2008.**  
Polytechnic University of Catalonia, Barcelona, Spain.

### RESEARCH EXPERIENCE

**Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland.**  
*PhD student* **September 2008 – ongoing.**

- Variational methods in image processing with applications to image segmentation, surface reconstruction and compressive sensing.
- Audio-Visual Speech Recognition and multi-modal signal processing.

**IBM Research, Zurich, Switzerland.**  
*Research internship* **May 2012 – November 2012.**

- Surface reconstruction from SEM images of integrated circuits.
- Fast image segmentation and classification method for failure detection in SEM images of integrated circuits.

**Polytechnic University of Catalonia, Barcelona, Spain.**  
*Undergraduate research grant* **October 2006 – January 2008.**

- Implementation and analysis of distributed algorithms in European project WINSOC (Wireless Sensor Networks with Self-Organization Capabilities for Critical and Emergency Applications).

### TEACHING EXPERIENCE

**Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland.**  
*Teaching assistant*

- Teaching assistant for courses Image Processing I and II.
- Supervision of students during their master project.

PROFESSIONAL  
ACTIVITIES

- Reviewer: IEEE Transactions on Image Processing, IEEE International Conference on Image Processing (ICIP), IEEE International Conference on Computer Vision (ICCV).
- Writing two research proposal for the Swiss National Foundation, which were accepted and granted.

PUBLICATIONS

**Pending articles**

V. Estellers, J.-P. Thiran, and M. Gabrani. *Surface reconstruction from microscopic images in optical lithography*.

Submitted to *IEEE Transactions on Image Processing*, 2012.

V. Estellers, D. Zosso, X. Bresson, and J.-P. Thiran. *Harmonic Active Contours*.

Submitted to *IEEE Transactions on Image Processing*, 2013.

**Journal articles**

V. Estellers, J.-P. Thiran, and X. Bresson. *Enhanced Compressed Sensing Recovery with Level Set Normals*.

Accepted in *IEEE Transactions on Image Processing*, 2012.

V. Estellers, D. Zosso, R. Lai, S. Osher, J.-P. Thiran, and X. Bresson. *An Efficient Algorithm for Level Set Method Preserving Distance Function*.

In *IEEE Transactions on Image Processing*, vol. 21, no. 12, pp.4722–4734, 2012.

V. Estellers, M. Gurban, and J.-P. Thiran. *On dynamic stream weighting for Audio-Visual Speech Recognition*

In *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 20, no.4, pp.1145–1157, 2012.

V. Estellers and J.-P. Thiran. *Multi-pose lipreading and Audio-Visual Speech Recognition*.

In *EURASIP Journal on Advances in Signal Processing*, vol. 2012.

**Conference articles**

V. Estellers, D. Zosso, X. Bresson, and J.-P. Thiran. *Harmonic Active Contours for multichannel image segmentation*.

In *Proc. IEEE International Conference on Image Processing (ICIP)*, pp. 3141–3144, 2011.

V. Estellers and J.-P. Thiran. *Multipose Audio-Visual Speech Recognition*.

In *Proc. IEEE European Signal Processing Conference (EUSIPCO)*, pp. 1065–1069, 2011.

V. Estellers and J.-P. Thiran. *Overcoming Asynchrony in Audio-Visual Speech Recognition*.

In *Proc. IEEE International Workshop on Multimedia Signal Processing (MMSp)*, pp. 466–471, 2010.

V. Estellers, P. M. Baggenstoss, and J.-P. Thiran. *Class-specific classifiers in audio-visual speech recognition*.

In *Proc. IEEE European Signal Processing Conference (EUSIPCO)*, pp. 1998–2002, 2010.

V. Estellers, M. Gurban, and J.-P. Thiran. *Selecting relevant visual features for speechreading*.

In *Proc. IEEE International Conference on Image Processing (ICIP)*, pp. 1433–1436, 2009.

PATENTS

**Contour fitting for automatic SEM filtering of images of integrated circuits.** IBM patent filed.

COMPUTER SKILLS

- **Programming Languages:** C, java, python.
- **Operating Systems:** Linux, Windows, Mac.
- **Technical Software:** Matlab.

MISCELLANEOUS

- **Languages** Spanish and Catalan (native), English and French (proficient), German (fluent).