

Learning to sample in Cartesian MRI

Présentée le 3 juin 2022

Faculté des sciences et techniques de l'ingénieur
Laboratoire de systèmes d'information et d'inférence
Programme doctoral en informatique et communications

pour l'obtention du grade de Docteur ès Sciences

par

Thomas SANCHEZ

Acceptée sur proposition du jury

Prof. A. M. Alahi, président du jury
Prof. V. Cevher, directeur de thèse
Dr R. van Heeswijk, rapporteur
Dr Ph. Ciuciu, rapporteur
Prof. D. van de Ville, rapporteur

On God rests my salvation and my glory;
my mighty rock, my refuge is God.
— The Bible, *Psalms* 62:8

Acknowledgements

This thesis was a journey that would have been impossible without the help and support of many people.

First, I would like to thank my advisor Volkan Cevher. He never gave up on me despite some tumultuous steps, and fought for me to reach the end of my thesis. Few would have done this. I am also thankful that he let me pursue freely my research interests, and I am grateful to him for having built an amazing group with colleagues who are both brilliant researchers and great friends.

Second, I would like to express my gratitude to Prof. Alexandre Alahi, Prof. Dimitri van de Ville, Dr. Ruud van Heeswijk and Dr. Philippe Ciuciu for being part of my thesis committee and for the interest that they took in reading and evaluating my thesis, as well as for their useful suggestions. I would like to thank in particular Dr. Ruud van Heeswijk and Dr. Philippe Ciuciu for all the insight that they gave me about MRI, and Dr. Philippe Ciuciu hosting me at Neurospin, despite the visit being shortened by the start of a pandemic. I have very good memories of my short time there.

My thanks then naturally go to my friends and colleagues, who supported me during these four years and made them unforgettable. I can never repay them enough for their kindness. Thank you Baran, Paul, Ya-Ping, Ahmet, Fatih, Ali K., Kamal, Thomas, Fabian, Leello, Pedro, Igor, Chaehwan, Maria, Arda, Armin, Yura, Nadav, Grigoris, Ali R., Andrej, Luca, Stratis, Kimon, Fanghui, and all the others for your friendship and the moments that we have shared. Of course, a couple of special mentions are in order. I want to thank Baran for introducing me to the world of MRI, and teaching me much about research and life. I want to thank Igor for the close collaboration that we had throughout my PhD, for putting up with me through manifold argumentative and extended meetings, pesky bugs and seemingly incomprehensible experimental results, and for being so patient and supportive. I am also thankful for his healthy German honesty and for the good friend he is. I want also to thank Leello for being one of the most fun people on this planet – I don't think that someone made me laugh more than him – for his optimistic look on life, for his practical wisdom and overall for being such an amazing friend. Thank you for bringing joy and friendship to everyone around you. Of course, I ought to thank Gosia, the secretary of our lab, for her amazing work and

Acknowledgements

efficiency in every practical matter. I suspect that she is the reason that the lab has not yet collapsed into chaos.

I am also grateful to all my friends for their support, in particular to Clément and Nathan Dupertuis, Peter Lecomte, Marion Conus and Christian Cuendet, as well as everyone from church, from the house group and from GBU. They are too numerous to be listed individually, and I feel very fortunate to be surrounded by so many kind and generous people.

Of course, without the support from my family, this thesis would not have been possible. I have been blessed by the unconditional love and encouragement from my parents, Esther and David, from my sister Myriam and her fiancé Sam, and from my brother Nicolas. I also cannot thank enough my parents-in-law, Karin and Jean-Paul.

Finally, my words turn to my best friend and amazing wife, Aline. Thank you for who you are. Thank you for your love, your support and unfailing encouragement. Without you, I could not have completed this thesis. You make every day brighter and push me to become a better man.

Lausanne, May 12, 2022

T. S.

Abstract

Magnetic Resonance Imaging (MRI) is a non-invasive, non-ionizing imaging modality with unmatched soft tissue contrast. However, compared to imaging methods like X-ray radiography, MRI suffers from long scanning times, due to its inherently sequential acquisition procedure. Shortening scanning times is crucial in clinical setting, as it increases patient comfort, decreases examination costs and improves throughput.

Recent developments thanks to compressed sensing (CS) and lately deep learning allow to reconstruct high quality images from undersampled images, and have the potential to greatly accelerate MRI. Many algorithms have been proposed in the context of reconstruction, but comparatively little work has been done to find acquisition trajectories that optimize the quality of the reconstructed image downstream.

Although in recent years, this problem has gained attention, it is still unclear what is the best approach to design acquisition trajectories in Cartesian MRI. In this thesis, we aim at contributing to this problem along two complementary directions.

First, we provide novel *methodological* contributions to this problem. We first propose two algorithms that improve drastically the greedy learning-based compressed sensing (LBCS) approach of Gözcü et al. (2018). These two algorithms, called lazy LBCS and stochastic LBCS scale to large, clinically relevant problems such as multi-coil 3D MR and dynamic MRI that were inaccessible to LBCS. We also show that generative adversarial networks (GANs), used to model the posterior distribution in inverse problems, provide a natural criterion for adaptive sampling by leveraging their variance in the measurement domain to guide the acquisition procedure.

Secondly, we aim at deepening the *understanding* of the kind of structures or assumptions that enable mask design algorithms to perform well in practice. In particular, our experiments show that state-of-the-art approaches based on deep reinforcement learning (RL), which have the ability to adapt trajectories on the fly to patient and perform long-horizon planning, bring at best a marginal improvement over stochastic LBCS, which is neither adaptive nor does long-term planning.

Overall, our results suggest that methods like stochastic LBCS offer promising alternatives

Abstract

to deep RL. They shine in particular by their scalability and computational efficiency and could be key in the deployment of optimized acquisition trajectories in Cartesian MRI.

Keywords: magnetic resonance imaging, experiment design, inverse problems, compressed sensing, reinforcement learning, deep learning, generative adversarial networks

Résumé

L'imagerie par résonance magnétique (IRM) est une modalité d'imagerie non invasive, non ionisante et offrant un contraste inégalé des tissus mous. Cependant, en raison de procédures d'acquisition inhéremment séquentielles, la vitesse d'acquisition en IRM est lente comparé à des méthodes telles que la radiographie par rayons X. La réduction des temps d'acquisition est cruciale en milieu clinique, car elle profite au confort du patient, diminue les coûts d'examen et améliore le rendement.

Les innovations récentes grâce à l'acquisition comprimée (CS) et dernièrement l'apprentissage profond permettent de reconstruire des images de haute qualité à partir d'images sous-échantillonnées et ont le potentiel de grandement accélérer l'IRM. De nombreux algorithmes ont été proposés dans le contexte de la reconstruction, mais comparativement peu de travaux ont été réalisés afin de trouver des trajectoires d'acquisition qui optimisent la qualité de l'image reconstruite en aval.

Bien que ces dernières années, ce problème ait gagné en attention, il n'est toujours pas clair quelle est la meilleure approche pour concevoir des trajectoires d'acquisition en IRM cartésienne. Dans cette thèse, nous visons à contribuer à ce problème selon deux directions complémentaires.

Premièrement, nous proposons de nouvelles contributions *méthodologiques* à ce problème. Nous proposons tout d'abord deux algorithmes qui améliorent considérablement l'approche gloutonne d'acquisition comprimée basée sur l'apprentissage (LBCS) de (Gözcü et al., 2018). Ces deux algorithmes, appelés *LBCS paresseux* et *LBCS stochastique*, s'étendent à des problèmes importants et cliniquement pertinents, tels que l'IRM parallèle 3D et dynamique, qui étaient inaccessibles à LBCS. Nous montrons également que les réseaux antagonistes génératifs (GAN), utilisés pour modéliser la distribution *a posteriori* dans les problèmes inverses, fournissent un critère naturel pour l'acquisition adaptative en utilisant leur variance dans le domaine de la mesure pour guider la procédure d'acquisition.

Deuxièmement, nous cherchons à approfondir la *compréhension* du type de structures ou d'hypothèses qui permettent aux algorithmes de conception de masques d'être performants en pratique. En particulier, nos expériences montrent que les approches de pointe basées sur l'apprentissage par renforcement (RL) profond, qui ont la capacité

Résumé

d'adapter leurs trajectoires au patient à la volée et d'effectuer une planification à long terme, apportent au mieux une amélioration marginale par rapport à LBCS stochastique, qui n'est ni adaptatif ni ne fait pas de planification à long terme.

Dans l'ensemble, nos résultats suggèrent que des méthodes comme LBCS stochastique offrent des alternatives prometteuses au RL profond. Elles brillent notamment par leur extensibilité et leur efficacité de calcul et pourraient être déterminantes dans le déploiement de trajectoires d'acquisition optimisées en IRM cartésienne.

Mots-clé : imagerie par résonance magnétique, planification d'expériences, problèmes inverses, acquisition comprimée, apprentissage par renforcement, apprentissage profond, réseaux antagonistes génératifs.

Contents

Acknowledgements	i
Abstract (English/Français)	iii
1 Introduction	1
2 MRI Background	9
2.1 MRI physics	9
2.1.1 Magnetization	9
2.1.2 Excitation	10
2.1.3 Relaxation	11
2.1.4 Image acquisition	12
2.1.5 Parallel imaging	16
2.1.6 Dynamic MRI	18
2.2 Mathematical description	19
2.2.1 Accelerated MRI	20
3 Reconstruction methods	23
3.1 Model-driven approaches	23
3.1.1 Compressed Sensing	23
3.2 Data-driven approaches	25
3.2.1 Dictionary learning	25
3.2.2 Unrolled neural networks	26
3.2.3 Deep image prior	27
3.2.4 A Bayesian interpretation of reconstruction	28
3.2.5 Generative approaches	29
3.3 Trade-offs	29
4 Optimizing sampling patterns	31
4.1 A taxonomy of mask design methods	31
4.2 Compressed sensing and MRI: Model-based, model-driven approaches	32
4.2.1 The contribution of Lustig et al. (2007)	32
4.2.2 Model-based, model-driven sampling in the sequel	33
4.3 Towards learning-based, data-driven approaches	34

Contents

4.4	A matter of performance: how do these methods compare to each other?	35
4.5	The advent of deep learning: Learning-based, data-driven approaches . .	36
4.6	Formalizing learning-based, data-driven mask design	39
4.6.1	Non-adaptive sampling masks	40
4.6.2	Adaptive sampling masks	41
4.6.3	On the optimality of the discrete mask optimization problem . .	42
5	Scalable learning-based sampling optimization	47
5.1	Learning-based Compressive Sensing (LBCS)	47
5.2	LBCS motivation: Submodularity	49
5.3	Scaling up LBCS	52
5.3.1	Limitations of LBCS	52
5.3.2	Stochastic LBCS	52
5.3.3	Lazy LBCS	54
5.4	Experiments - Stochastic LBCS	55
5.4.1	Experimental setting	56
5.4.2	Comparison of greedy algorithms	58
5.4.3	Single coil results	61
5.4.4	Multicoil experiment	70
5.4.5	Large scale static results	71
5.5	Experiments - Lazy LBCS	72
5.5.1	2D multicoil setting	73
5.5.2	3D multicoil setting	74
5.6	Discussion and Conclusion	77
6	On the benefits of deep RL in accelerated MRI sampling	81
6.1	Reinforcement Learning to optimize sampling trajectories	82
6.1.1	A short RL primer	83
6.1.2	The special case of MRI	85
6.1.3	Double Deep Q-Networks	86
6.1.4	Policy Gradient	88
6.1.5	LBCS as a simple RL policy	89
6.1.6	The questions at hand	89
6.2	The MRI data processing pipeline	90
6.3	Re-examining deep RL for MRI sampling	93
6.3.1	Experimental setting	93
6.3.2	Main results on Bakker et al. (2020)	95
6.3.3	Main results on Pineda et al. (2020)	97
6.4	Discussion	98
6.5	Relation to other works	101
6.6	Conclusion	101
7	Uncertainty driven adaptive sampling via GANs	103

7.1	Generative adversarial networks (GANs) for posterior modeling	104
7.1.1	Generative adversarial networks (GANs)	104
7.1.2	Conditional GANs	105
7.1.3	Conditional GANs for posterior modeling	107
7.1.4	Learning the posterior or learning moments	108
7.2	GAN based adaptive sampling	110
7.3	Comparison with Zhang et al. (2019b)	112
7.4	Experiments	113
7.4.1	MRI adaptive sampling	113
7.4.2	Beyond Fourier sampling: image-domain conditional sampling . .	119
7.5	Related works	121
7.6	Discussion	123
7.7	Conclusion	125
8	Conclusion and future works	127
A	Appendix for Chapter 6	133
A.1	Implementation details	133
A.2	Further experiments and results on Bakker’s setting	135
A.2.1	Data processing differences	135
A.2.2	Exact reproduction of Bakker et al.	136
A.2.3	Further results in the Bakker experiment	137
A.3	LBCS vs long-range adaptivity	140
A.4	LBCS complexity and parameters	142
A.5	Visual comparisons	143
A.6	Detail on the practical recommendations	147
B	Appendix for Chapter 7	149
B.1	Implementation details	149
B.1.1	Knee experiments	149
B.2	Extended results	151
B.2.1	Additional MRI knee results	151
B.2.2	Brain experiment	151
B.2.3	Additional MNIST results	153
B.2.4	CIFAR10 experiment	155
	Bibliography	159
	Curriculum Vitae	177

1 Introduction

Since its inception in the 1970s, Magnetic Resonance Imaging (MRI) has deeply impacted radiology and medicine. It allows for a non-invasive, non-ionizing imaging of the body, and enables among other things to visualize anatomical structures, physiological functions and metabolic information with unmatched precision (Wright, 1997; Feng et al., 2016b). Magnetic Resonance Imaging, as its name indicates, constructs an image of an object of interest by exploiting magnetic resonance. Namely, when submitted to a strong static magnetic field B_0 , the protons in the object, mostly in water molecules, will resonate when excited through a radio frequency pulse. In order to encode spatial information in the observed signal instead of a receiving global response, additional time and spatially varying magnetic fields, referred to as gradient fields, are superimposed to the original magnetic field B_0 .

However, this process does not give direct access to an image, but rather informs us about the different *frequencies* that construct it. By varying the gradient fields, it is possible to obtain frequency information relating to different locations and represent them in what is known as the Fourier space, or *k-space*. Figure 1.1 shows the difference between an image and its Fourier space representation. The object of interest is excited several times using different varying gradient fields, which defines a *k-space trajectory* along which data are acquired. When sufficient information has been acquired, an image can be readily obtained by computing the inverse Fourier transform of the signal.

There is considerable flexibility in the design of k-space trajectories, but a particularly common approach is to cover the k-space in a grid-like fashion, known as *Cartesian* sampling. A grid-like covering of k-space implies that the image can be directly obtained by applying a fast Fourier transform (FFT). Other approaches, known as non-Cartesian, do not cover k-space in a grid-like fashion, but can consist of radial spokes (Lauterbur, 1973), spirals (Meyer et al., 1992) or space-filling curves (Lazarus et al., 2019). They can provide a faster acquisition than Cartesian sampling and benefits such as resistance to motion artifacts, but require a more involved and slower reconstruction through

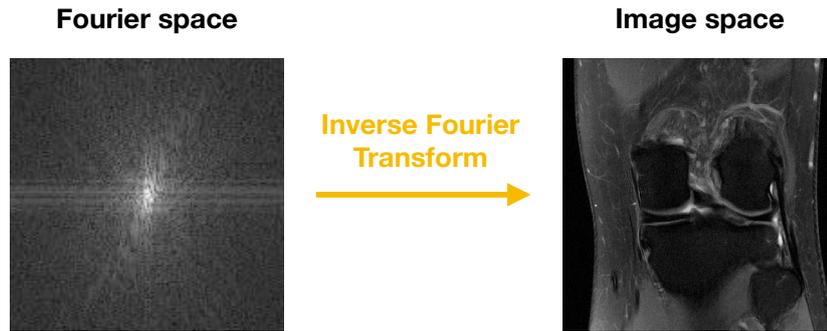


Figure 1.1 – Fourier domain and image domain representation of a knee MRI scan.

techniques such as gridding (O’sullivan, 1985; Jackson et al., 1991) or non-uniform fast Fourier transform (NUFFT) (Fessler and Sutton, 2003), as well as density compensation (Pipe and Menon, 1999; Pauly, 2007). In addition to this, non-Cartesian trajectories can be sensitive to errors in the magnetic field gradients, due to field imperfections of various sources (Vannesjo et al., 2013). As a result, Cartesian trajectories have been the most popular in practice (Lustig et al., 2008; Feng et al., 2016b).

However, MRI examinations suffer from being relatively slow compared to other imaging modalities, due to the sequential nature of the acquisition. For instance, MRI acquisition can easily exceed 30 minutes (Zbontar et al., 2019). Improving imaging speed has been a major concern for researchers of the community and has led to many innovations along the years, such as rapid imaging sequences (Frahm et al., 1986; Hennig et al., 1986) or parallel/multi-coil imaging (Sodickson and Manning, 1997; Pruessmann et al., 1999; Griswold et al., 2002a). But there is a need to further accelerate imaging. As scanning time is directly related to the number of acquisitions, collecting fewer of them directly enables faster imaging. In Cartesian imaging, this can be achieved by *undersampling* the Fourier space, which is implemented by *not* acquiring some lines in Fourier space. However, as we see on Figure 1.2, such undersampling degrades the quality of the image by introducing artifacts. This is called *aliasing*. As a result, an additional step of *reconstruction* is required in order to obtain a cleaner estimation of the original image.

The problem of reconstructing data from undersampled measurements is a typical instance of an *ill-posed inverse problem*. It is called an *inverse problem* because we aim at constructing an image of an object of interest from *indirect measurements*, in this case frequency information about the image. It is *ill-posed* because there is not sufficient information to directly retrieve the image of interest from the undersampled measurements. More precisely, there are infinitely many signals that could produce these partial measurements.

The typical way of solving this problem is by exploiting *prior knowledge* about the data, by leveraging some *structure* that is present in them. A particularly successful approach

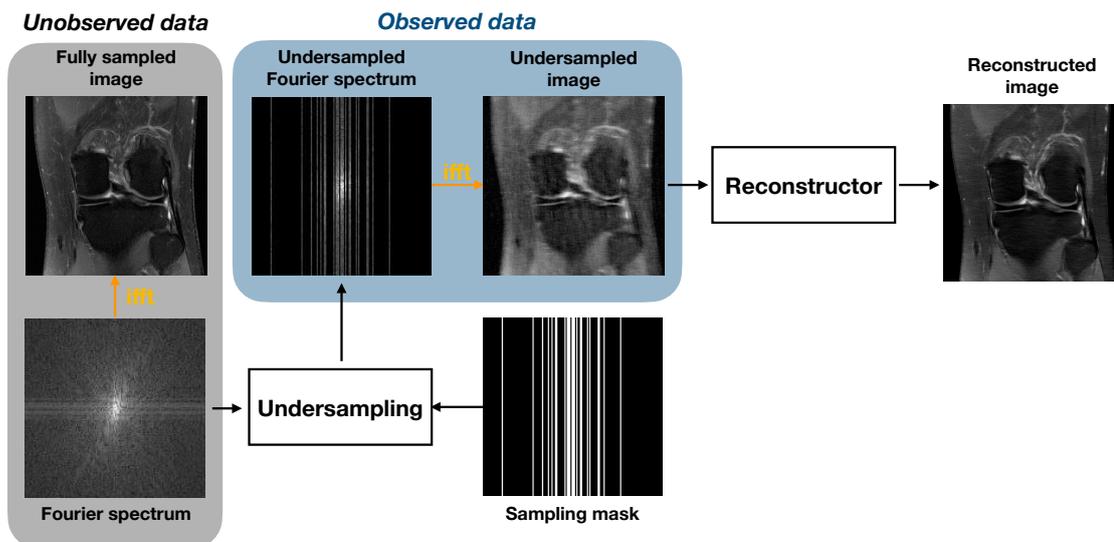


Figure 1.2 – Accelerated MRI by undersampling. Undersampling is governed by a sampling mask that tells what locations are acquired. The undersampled Fourier spectrum is then mapped back to image domain through an inverse Fourier transform, which gives an aliased image which needs to go through a reconstruction method in order to obtain a de-aliased estimation of the original, unobserved image. Here *ifft* stands for inverse fast Fourier transform.

relies on *sparsity*-based priors, where it is assumed that if the signal can be represented sparsely, or compactly in some transform domain, then, it can be reconstructed from undersampled measurements. This approach is known as Compressed Sensing (CS) (Donoho, 2006; Candès et al., 2006), and has been successfully applied to MRI (Lustig et al., 2008), for instance by representing the signal sparsely in the Wavelets domain. The signal is retrieved by solving an optimization problem with two components: a *data-fidelity* term that ensures that the reconstructed image is close to the measurements, and a *regularization* term that encourages the reconstructed to be sparse in some domain.

In the latter half of 2010s, deep learning-based approaches allowed pushing even further the boundaries of the state-of-the-art in MRI reconstruction (Sun et al., 2016; Schlemper et al., 2017; Mardani et al., 2018). By leveraging increasingly large amount of data, these methods are able to learn the structure directly from data, rather than imposing it through a mathematical model such as sparsity. These methods not only improved the quality of reconstruction, but also enabled fast reconstruction times, where CS methods often relied on slow iterative procedures.

However, in parallel to these developments, only few works tackled the problem of *how* the *k*-space should be undersampled. Although CS prescribed random sampling, practitioners quickly turned to heuristics favoring low frequency sampling, where a lot of energy and information about the structure of the signal is located (Lustig et al., 2008). Nevertheless, there were only a few works in the early 2010s aiming at directly optimizing the sampling

Introduction

mask¹, which can be attributed to the difficulty of the problem. Optimizing a sampling mask is a *combinatorial* problem, for which commonly used gradient-based optimization techniques are not readily available. As a result, two main trends emerged: approximate solutions to the combinatorial problem (Ravishankar and Bresler, 2011a; Liu et al., 2012), or the optimization of an underlying probability distribution from which a random mask is then sampled (Knoll et al., 2011b; Chauffert et al., 2013).

In more recent years, there has been a renewed interest in this problem, and a flurry of approaches have been proposed (Gözcü et al., 2018; Jin et al., 2019; Sherry et al., 2019; Bahadir et al., 2019; Huijben et al., 2020; Pineda et al., 2020; Bakker et al., 2020; Yin et al., 2021). This is mainly due to the wide availability of training data, the fast reconstruction times of deep learning-based methods and the wider use of auto-differentiation frameworks such as PyTorch and TensorFlow, which facilitate the routine computation of derivatives through complex models. While these methods are very diverse, they fit under two distinctive categories. Whereas previous approaches leveraged an underlying *model* from which candidate sampling masks were drawn, modern approaches are *learning-based* in the sense that they learn this model directly from data. Furthermore, earlier methods were *model-driven*, designing masks to minimize a mathematical criterion such as coherence (Lustig et al., 2008). In contrast, recent works are *data-driven*, as they learn policies that minimize the reconstruction error against a ground truth.

In this thesis, we aim at contributing to the problem of optimizing sampling masks along two complementary directions.

First, we will provide novel *methodological* contributions to the problem of sampling optimization. More specifically, we aim at developing mask optimization algorithms that do not depend on a specific reconstruction algorithm, as this will allow the method to remain usable in a field where the state-of-the-art evolves very quickly. In addition, we are concerned with designing *scalable* methods. Most clinically relevant settings involve high resolution, multi-coil images, and it is imperative for sampling optimization methods to scale up to such challenging settings in order to provide practical value.

Secondly, we aim at deepening the *understanding* of the kind of structures or assumptions that enable mask design algorithms to perform well in practice. It is not clear from the current literature what components are fundamental in a mask design algorithm. Should the mask design algorithm plan several steps of ahead? Should it adapt to each different patient? What kind of feedback should be available to the mask design algorithm in order to achieve the best performance? Amidst the race to develop ever more complex models, we aim at identifying the key components that drive the performance of modern mask design algorithms, in order to *efficiently* reach state-of-the-art performance.

¹We will refer to this problem interchangeably as the problem of mask design, sampling optimization, mask optimization, sampling mask optimization, sampling pattern optimization, or optimizing the acquisition trajectory.

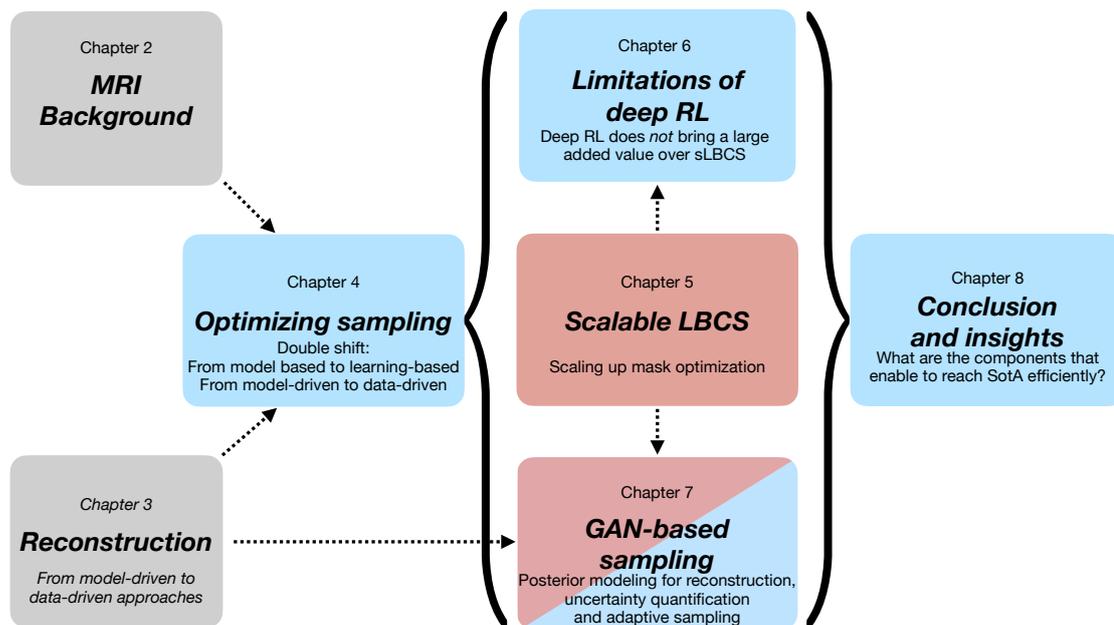


Figure 1.3 – Outline of the different chapters in this thesis. **Grey** boxes refer to chapters dealing with background material. **Red** boxes denote chapters with methodological contributions, and **blue** boxes denote chapters that contribute to a better understanding of the problem of optimizing sampling. Arrows show particular, direct dependencies between chapters.

Organization and contributions

Figure 1.3 illustrates the organization of this thesis, and each chapter is briefly described below, along with their respective contributions.

Chapter 2 – MRI Background

In this chapter, we provide a brief introduction to the physics underlying image acquisition in MRI, focusing on Cartesian sampling. We introduce multi-coil acquisition as well as dynamic imaging. We then provide the mathematical model that will be used to describe the acquisition in accelerated MRI.

Chapter 3 – Reconstruction methods

In this chapter, we review the algorithms used to reconstruct images from undersampled measurements. We describe the shift that has taken place from *model-driven* approaches based on Compressed Sensing (Lustig et al., 2007) towards *data-driven* approaches that construct a model from training data. We then focus on the different trends that have taken place within data-driven approaches, focusing on trained reconstruction models and generative approaches. Finally, we conclude by highlighting the trade-offs between the speed of reconstruction and the dependence on large training datasets.

Chapter 4 – Optimizing sampling patterns

In this last introductory chapter, we review the different approaches that have been taken to optimize sampling masks. We argue that, in a similar way that reconstruction methods moved from model-driven to data-driven approaches, a double shift has taken place in mask design. We show that mask design methods gradually moved from *model-based*, *model-driven* to *learning-based*, *data-driven* approaches. We discuss how this trend can be seen in the literature, and how it ties in with the empirical performance of these methods. Then, we proceed to formalizing the problem of learning-based, data-driven mask design and show how this problem changes when one aims at designing *fixed policies* – where a mask is trained and re-used as is at test time – or *adaptive policies* – that adapt to the current patient by integrating on-the-fly the information as it acquired. We conclude by proving the optimality of the discrete mask optimization problem, compared to the problem of finding the optimal sampling *distribution*.

Related publication.

Sanchez, T., Gözcü, B., van Heeswijk, R. B., Eftekhari, A., Ilıcak, E., Çukur, T., and Cevher, V. (2020a). Scalable learning-based sampling optimization for compressive dynamic MRI. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8584–8588.

Chapter 5 – Scalable learning-based sampling optimization

In this chapter, we extend the Learning-Based Compressive Sensing (LBCS) approach of Gözcü et al. (2018). We primarily aim at *drastically* improving their scalability of their method. We first discuss how their approach is rooted in submodular optimization (Krause and Golovin, 2014) and how algorithms proposed in this field could provide solutions to the limitation of LBCS. Their method scales as $O(P^2)$, where P is the resolution of the image, such a complexity is prohibitive in high dimensional settings such as 3D MRI or dynamic MRI, and is not convenient when reconstruction becomes slow, such as in multi-coil MRI. To address this, we first propose *lazy LBCS* (lLBCS), for the multi-coil and 3D settings. This algorithm leverages lazy evaluations (Minoux, 1978), where a list of upper bounds is *precomputed* and then traversed sequentially, drastically reducing the amount of computation over LBCS. Secondly, in the context of dynamic (multi-coil) MRI and large datasets, we propose *stochastic LBCS* (sLBCS) and show that it can achieve a very significant computational gain compared to LBCS, while retaining its performance.

Related publications.

Gözcü, B., Sanchez, T., and Cevher, V. (2019). Rethinking sampling in parallel MRI: A data-driven approach. In *27th European Signal Processing Conference (EUSIPCO)*.

Sanchez, T., Gözcü, B., van Heeswijk, R. B., Eftekhari, A., Ilıcak, E., Çukur, T., and Cevher, V. (2020a). Scalable learning-based sampling optimization for compressive dynamic MRI. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8584–8588.

Chapter 6 – On the benefits of deep RL in accelerated MRI sampling

In this chapter, we assess the performance of state-of-the-art (SotA) deep reinforcement learning (RL) methods for mask designs. These methods improve in principle over LBCS by designing policies that *adapt* to the patient at hand and leverage *long-horizon* planning in order to make better decisions. We mainly aim at making sense of seemingly contradictory conclusions. The work of Pineda et al. (2020) seems to indicate that long term planning could be the most important component in deep RL, as their results show that a non-adaptive, long term planning policy model trained on the dataset can perform as well as an adaptive, long-term planning policy. On the contrary, the contribution of Bakker et al. (2020) highlights the importance of adaptivity, as a greedy policy, that does not do long-term planning is found to closely match policies that do long term planning.

Our results, surprisingly show that sLBCS, a non-adaptive method that does not attempt long term planning, can perform as well as the state-of-the-art approaches of Pineda et al. (2020); Bakker et al. (2020). We further show that the current SotA RL methods only bring at best a limited added value, and that other factors such as the architecture of the reconstruction algorithm or the masks used to train it have a much larger impact on the overall performance. This work highlights the need for further discussions in the community about standardized metrics, strong baselines, and careful design of experimental pipelines to evaluate MRI sampling policies fairly.

Related preprint.

Sanchez, T.*, Krawczuk, I.* and Cevher, V. (2021). On the benefits of deep RL in accelerated MRI sampling. Available at <https://openreview.net/pdf?id=fRb9LBWUo56>.
* denotes equal contribution.

Chapter 7 – Uncertainty driven adaptive sampling via GANs

In this chapter, we take a step back from the question of designing the best sampling policy, and primarily aim at exploring how conditional Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) can be used to model inverse problems in a Bayesian fashion. GANs have been used mainly as reconstruction models in the context of inverse problems (Yang et al., 2018; Chen et al., 2022), but a few works took an additional step to show the value of learning the full posterior distribution in order to have the ability to provide uncertainty quantification as well (Adler and Öktem, 2018a).

Introduction

We argue that they do more and that conditional GANs for MRI reconstruction can naturally provide a criterion for adaptive sampling: by sequentially observing the location with the largest posterior variance in the measurement domain, the GAN provides an adaptive sampling policy without ever being trained for it.

Such a criterion is not restricted to acquisition in Fourier domain, like in MRI, but can also be readily used for image domain sampling. In this context, we show that our GAN-based policy can strongly outperform the non-adaptive greedy policy from sLBCS. However, in Fourier domain, the GAN-based policy does not match the performance of sLBCS. We provide an explanation to this phenomenon rooted in the concept of the information horizon used to make a decision at each step. We show that our model uses less information than LBCS to design its policy, and argue that this is the reason leading to its inferior performance. However, when compared with models that use the same amount of information to inform their policy, our model largely outperforms the competition.

Related preprint and workshop paper.

Sanchez, T., Krawczuk I., Sun, Z. and Cevher V. (2019). Closed loop deep Bayesian inversion: Uncertainty driven acquisition for fast MRI. *Preprint available at* <https://openreview.net/pdf?id=BJIPOIBKDB>.

Sanchez, T., Krawczuk, I., Sun, Z. and Cevher V. (2020). Uncertainty-Driven Adaptive Sampling via GANs. In *NeurIPS 2020 Workshop on Deep Learning and Inverse Problems*. Available at <https://openreview.net/pdf?id=IWLYCQmtvW>.

Chapter 8 – Conclusion and future works

In this last chapter, we review the contributions of the different chapters and summarize the insights that our contributions bring to the design of masks for Cartesian MRI sampling. We then provide an outlook of the opportunities for future works.

Remark (Bibliographic note). At the end of Chapters 4, 5, 6 and 7, we include a “bibliographic note” section, where we distinguish the contributions of the co-authors of the abovementioned publications.

Additional publication not included in this thesis.

Sun, Z., Latorre, F., Sanchez, T., and Cevher, V. (2021). A plug-and-play deep image prior. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 8103-8107.

2 MRI Background

Before diving in greater detail in the problems of reconstruction from undersampled measurements and optimization of the sampling masks, we first present the physical underpinnings of MRI. Although this survey will be very short, it hints at the remarkable depth and flexibility of MRI. This introduction is based upon the well-known textbook of Brown et al. (2014)¹.

2.1 MRI physics

2.1.1 Magnetization

Let us then consider a proton under a strong external magnetic field \mathbf{B}_0 . In practice, such a field is obtained by running a large electric current through a coil around the scanner. High current is required because typical magnetic fields in MRI are very large, around 3 T, and the strength of the magnetic field is directly proportional to the one of the electric current².

The spin of a proton, which is positively charged, can be represented as a gyroscope precessing about the field direction, as illustrated on Figure 2.1. The precession angular frequency can be quantified as the Larmor frequency ω_0 , with

$$\omega_0 = \gamma B_0 \tag{2.1}$$

where γ is known as the gyromagnetic ratio, and is a particle dependent constant. For instance, in water, the hydrogen proton has $\gamma \triangleq \gamma/(2\pi) \approx 42.6 \text{ MHz T}^{-1}$. Note that deriving this formula is involved and requires tools from relativistic quantum mechanics.

This description considers a single hydrogen nucleus in a static magnetic field, but in

¹Part of the content was presented in Sanchez (2018)

²For comparison, the strength of the magnetic field of the earth is typically between 40 and 70 μT .

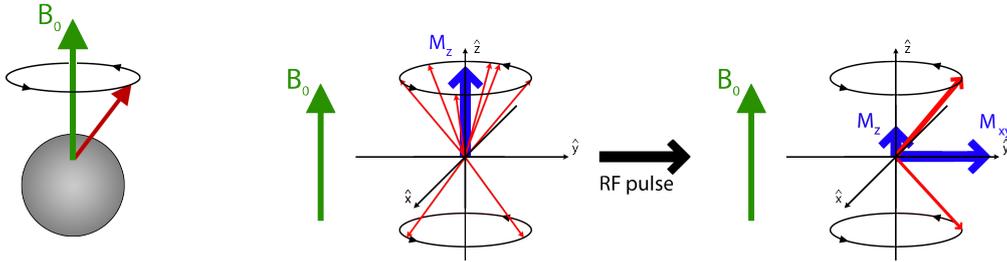


Figure 2.1 – Left: Illustration of precession of a particle when subject to an external field \mathbf{B}_0 (large green arrow). The small red arrow corresponds to the spin of the particle, and its temporal evolution is illustrated by the black arrow. Right: Illustration of the transverse magnetization appearing when a body is subject to an RF pulse, due to the protons precessing synchronously, and reaching the higher energy level, which results in arrows pointing downwards. This Figure was first used in Sanchez (2018).

practice, we are dealing with an entire population. Depending on their levels of energy, different protons of a population will either align in a parallel fashion against an external field \mathbf{B}_0 , while some will take an antiparallel position. Although there is only a very little amount of spins parallel to the magnetic field exceeding the number of spins antiparallel to it, this is sufficient to have a *spin excess* that yields a noticeable net magnetization \mathbf{M}_0 at a population level. \mathbf{M}_0 will be parallel to \mathbf{B}_0 because different protons will have incoherent phases, and will overall cancel each other out, leading to a net magnetization parallel to the static field. We call this resulting state the *equilibrium* of the system, which is illustrated in the middle of Figure 2.1, where in this case $\mathbf{M}_0 = \mathbf{M}_z$. However, this net magnetization cannot be detected in the equilibrium, as we need to excite the protons in order to induce a change in the magnetization, which in turn will induce an electrical current in a *receive coil*. This coil will allow to record a signal corresponding to the response of the protons to the excitation. The different components of the MRI that will be discussed throughout this chapter are illustrated on Figure 2.2.

2.1.2 Excitation

Let us denote by $\hat{\mathbf{z}}$ the direction of the external field \mathbf{B}_0 , and let us consider a body that we wish to image, for instance the knee of a patient. It is clear from the previous section that the hydrogen nuclei of the body will produce a net magnetization $\mathbf{M}_0 = M_0\hat{\mathbf{z}}$, where M_0 denotes the intensity of magnetization. The atoms will precess around $\hat{\mathbf{z}}$ at the Larmor frequency ω_0 .

In the excitation step, the system is destabilized by sending a radiofrequency (RF) pulse. It is easier to describe it and consider its effect by placing ourselves in a frame rotating at an angular frequency ω_0 around the axis $\hat{\mathbf{z}}$. This gives the following rotating unit

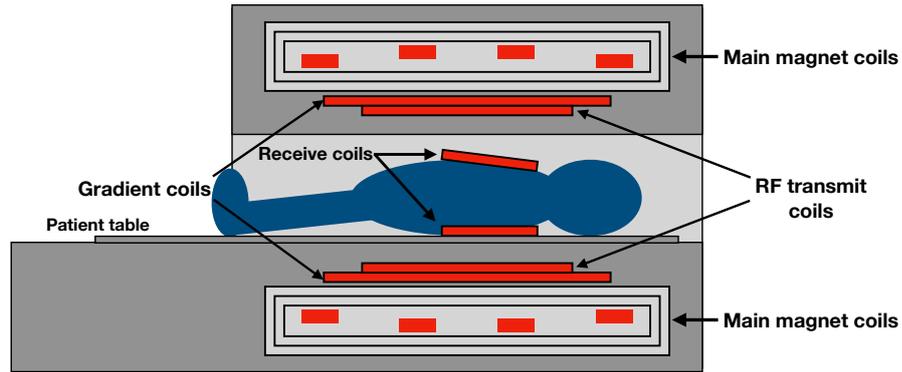


Figure 2.2 – Illustration of an MRI scanner and the different coils and components required to produce an image.

vectors $\hat{\mathbf{x}}'$, $\hat{\mathbf{y}}'$

$$\begin{aligned}\hat{\mathbf{x}}' &= \cos(\omega_0 t)\hat{\mathbf{x}} - \sin(\omega_0 t)\hat{\mathbf{y}} \\ \hat{\mathbf{y}}' &= \sin(\omega_0 t)\hat{\mathbf{x}} + \cos(\omega_0 t)\hat{\mathbf{y}}\end{aligned}$$

and $\hat{\mathbf{z}}$ remains unchanged. We define the RF pulse as a circularly rotating pulse of intensity B_1 , namely $\mathbf{B}_1 = B_1\hat{\mathbf{x}}'$. Applying this RF pulse will as a result cause the nuclei to resonate and flip the equilibrium magnetization \mathbf{M}_0 into a new net magnetization \mathbf{M} in the $\hat{\mathbf{x}}' - \hat{\mathbf{y}}'$ plane. \mathbf{M} will now have both a *transverse* component in the $\hat{\mathbf{x}}' - \hat{\mathbf{y}}'$ plane, and a *longitudinal* component along the $\hat{\mathbf{z}}$ axis. We refer to the angle between \mathbf{M} and the $\hat{\mathbf{z}}$ -axis as the *flip angle* (FA), and it quantifies how strongly the magnetization \mathbf{M} has been moved away from its equilibrium state \mathbf{M}_0 . This process is illustrated in the right of Figure 2.1. In particular, we see that the transverse component arises due to the synchronization of the precessing phases of the hydrogen nuclei, which cancelled each other out at equilibrium.

2.1.3 Relaxation

After applying the RF pulse, the system will gradually return to equilibrium, where \mathbf{M}_0 is aligned with \mathbf{B}_0 . This will require the nuclei to re-emit the energy that they absorbed during the excitation phase. The relaxation is governed by two main components, namely a *longitudinal* relaxation, characterized by a time constant \mathbf{T}_1 , and a *transverse* relaxation, characterized by \mathbf{T}_2 .

T_1 Relaxation. The T_1 relaxation describes the regrowth of the longitudinal component of the magnetization, i.e. how the $\hat{\mathbf{z}}$ component of \mathbf{M} grows back over the relaxation. This is also known as *spin-lattice* relaxation, as this describes a local process due to the

interaction of the spins with their surroundings. The process is formally described as

$$M_z(t) = M_0 \left(1 - \exp(-t/T_1)\right) + M_z(0) \exp(-t/T_1) \quad (2.2)$$

where $M_z(t)$ is the longitudinal magnetization, $M_z(0)$ describes the amount of longitudinal magnetization left after the RF pulse, and T_1 is a constant that depends on the type of tissue considered. We see that the longitudinal magnetization will be recovered at an exponential speed, and the speed will be determined by T_1 .

T_2 Relaxation. The T_2 relaxation process describes how clusters of spins gradually lose phase with each other after excitation. This dephasing will cause the transverse magnetization $M_{\hat{x}',\hat{y}'}$ or M_{\perp} to progressively disappear, as the equilibrium state is reached (cf. middle of Figure 2.1). We have

$$M_{\perp}(t) = M_0 \exp(-t/T_2) \quad (2.3)$$

where T_2 is a tissue-dependent time constant that describes how fast the transverse magnetization vanishes. In practice, $T_2 \ll T_1$, making the T_2 relaxation much faster. Also, in practice, the signal is likely to suffer additional suppression due to dephasing from inhomogeneities in the external field \mathbf{B}_0 .

As T_1 and T_2 are tissue dependent constants, constructing images by performing measurements at different steps of the relaxation allows to highlight different characteristics of tissues. They are one of the key components of the flexibility of MR imaging.

2.1.4 Image acquisition

Until now, we have discussed the global response of a sample to an RF pulse, but we have not addressed how to encode spatial information into the signal, in order to be able to construct an image. As previously discussed, the signal will be recorded by a receive coil, where an electric current is induced depending on the relaxation of the spins.

But the question of localizing the signal remains whole so far. To solve this issue, we will try to give to each location in the body a unique combination of phase and frequency, in order to be able to retrieve its contribution from the signal in the receive coil. This is achieved by introducing three additional coils to the scanner, known as gradient coils, whose role will be to produce some spatially varying magnetic fields in the \hat{x} , \hat{y} and \hat{z} directions respectively. As a result, the tissues of specific areas of the object will be excited, and this will allow to observe their local responses. Then, by repeating a measurement with different gradients, we will be able to sequentially construct a map of how the tissue responds to the RF pulses depending on its location. This is why an acquisition is often referred to as a *pulse sequence*, as we gather a sequence of

measurements in order to obtain a full image.

There exists several ways to construct an image with MRI, and we will illustrate how it works by considering a *2D gradient echo* sequence, and we will explain its different components. Sequences are frequently represented as diagrams, such as the one of Figure 2.3, where each line represents how different component vary in intensity through time. If we simply excite a body with an RF pulse, neglecting the relaxation effects, we can write the signal observed in the receive coil as

$$S(t) = \int \rho(\mathbf{x}) e^{i(\omega_0 t + \phi(\mathbf{x}, t))} d\mathbf{x} \quad (2.4)$$

where $\rho(\mathbf{x})$ is the *effective spin density*, and is the quantity that we wish to measure, and $\phi(\mathbf{x}, t)$ is the *accumulated phase*

$$\phi(\mathbf{x}, t) = - \int_0^t \omega(\mathbf{x}, t') dt'. \quad (2.5)$$

We see that in Equation 2.4, without additional considerations, the signal that we observe is integrated over the whole volume being imaged. In Equation 2.5, the main reason that ω is now a function of space and time is that we will add spatially varying gradient fields that in turn make ω dependent on space and time. One can also see that Equation 2.4 shows that measurements will not be provided directly in image space, but rather in Fourier space, often referred to as *k-space*: MRI allows us to obtain *frequency* information about our object of interest, and then an image can be obtained by mapping back these measurements into image space.

Slice Selecting (SS) Gradient. As we see on Figure 2.3, during the RF pulse, a constant \mathbf{z} -gradient is applied during a time τ_{RF} with intensity G_{SS} , and reversed afterwards. We will not detail here why this gradient is reversed, but this has the effect of making all components in the slice in phase, with a common accumulated phase $\phi = 0$. This enables to select the slice of interest and make sure that the signal will be strong for the acquisition. As a result, for a slice at location \mathbf{z}_0 with thickness Δz , the signal will be integrated out on the slice only, i.e. we will have the spin density

$$S(t_1) = \int \int \left[\int_{z_0 - \frac{\Delta z}{2}}^{z_0 + \frac{\Delta z}{2}} \rho(x, y, z) dz \right] dy dx \approx \int \int \rho(x, y, z_0) dy dx.$$

Phase Encoding (PE) Gradient. Let us now add the phase encoding \mathbf{y} -gradient to the picture. Assume that it is applied during a time τ_{PE} at intensity $G_{\text{min}} \leq G_y \leq G_{\text{max}}$. This will result in gradients acquiring a \mathbf{y} -dependent phase, and as a result

$$S(t_2) = \int \left[\int \rho(x, y, z_0) e^{-i2\pi \gamma G_y \tau_{\text{PE}} y} dy \right] dx$$

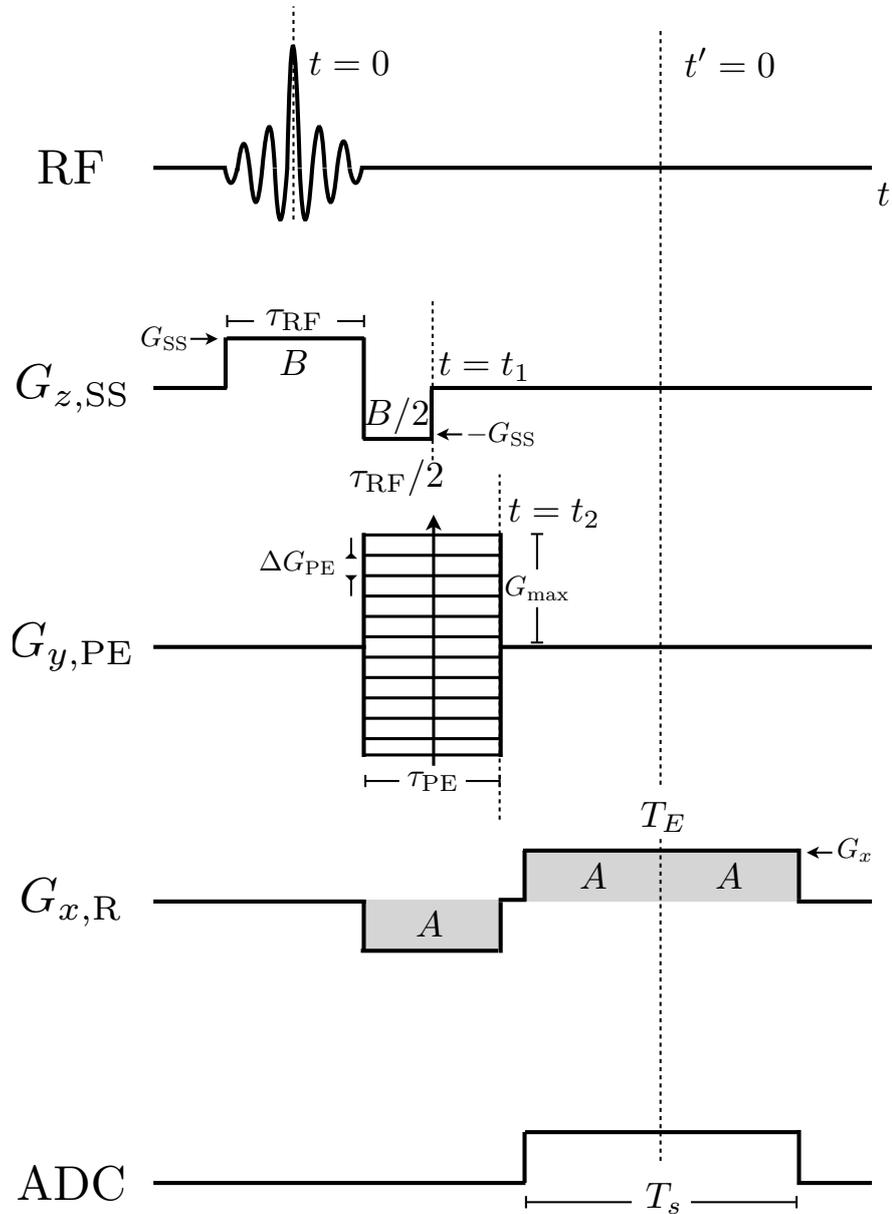


Figure 2.3 – A 2D gradient echo sequence diagram. Each line describes how different components interact during the time of a single excitation. On the first line, we see the RF pulse enabled at $t = 0$. On the second line, we have the slice selecting gradient $G_{z,SS}$, on the third line we have the phase encoding gradient $G_{y,PE}$, on the fourth we have the readout gradient $G_{x,R}$ and finally the analog-to-digital converter (ADC).

Readout (R) Gradient. Finally, the \mathbf{x} -gradient is first turned on negatively, dephasing the phases, followed by a rephasing stage where the observation through the ADC is carried out. It is enabled at the strength G_x and for a duration T_S . Given the shifted variable $t' = t - T_E$, we have as a result

$$S(t') = \int \left[\int \rho(x, y, z_0) e^{-i2\pi \varphi G_y \tau_{PE} y} dy \right] e^{-i2\pi \varphi G_x t' x} dx, \quad -T_S/2 \leq t' \leq T_S/2. \quad (2.6)$$

Equation 2.6 is very interesting as one can rewrite the (\mathbf{x}, \mathbf{y}) -directions as spatial frequencies depending on t and τ_y , namely

$$(k_x, k_y) = (\varphi G_x t', \varphi G_y \tau_{PE}).$$

This in turns enables us to rewrite Equation 2.6 as

$$S(k_x, k_y) = \int \int \rho(x, y, z_0) e^{-i2\pi(k_x x + k_y y)} dx dy \quad (2.7)$$

which turns out to be the Fourier transform of the spin density. More precisely, the signal obtained in the presence of a gradient echo structure in the readout direction after its encoding by a fixed G_y describes a *line* in the k -space of the Fourier transform of $\rho(\mathbf{x})$. This implies that by acquiring several lines at different k_y by varying G_y , we can effectively cover the k -space, as illustrated in Figure 2.4. As a result, when the k -space has been sufficiently densely covered, the spin-density can be recovered by taking the inverse Fourier transform of Equation 2.7, i.e.

$$\rho(x, y, z_0) = \int \int S(k_x, k_y) e^{+i2\pi(k_x x + k_y y)} dk_x dk_y \quad (2.8)$$

This approach is referred to as *Cartesian sampling* because it covers k -space in a grid-like fashion, as illustrated on Figure 2.4. Note that there exist different sequences that can perform Cartesian sampling, and even more sequences that are *not* Cartesian. We refer to such approaches as non-Cartesian sampling approaches, and they can include radial sampling (Lauterbur, 1973), spiral sampling (Meyer et al., 1992) or even some more flexible coverages of k -space (Lazarus et al., 2019).

Cartesian MRI has the advantage over these methods of enabling the use of the fast Fourier transform to compute the inversion of Equation 2.8, which is *not* the case for non-Cartesian methods. These approaches require some more sophisticated inversion, involving gridding (O'sullivan, 1985) and non-uniform Fourier transforms (Fessler and Sutton, 2003). The differences are however deeper than merely the inversion technique, as radial trajectories are known to be more robust to motion artifacts, while spiral acquisitions enable fast imaging but are prone to other artifacts, and Lustig et al. (2008) discuss these issues in greater depth.

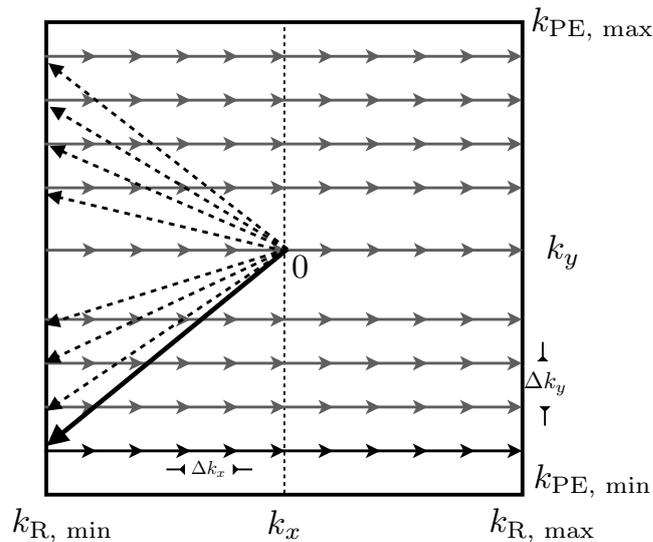


Figure 2.4 – Traversal of k -space for a gradient echo pulse sequence. The combined application of the phase encoding gradient $G_{y,PE}$ and readout gradient $G_{x,R}$ makes the signal move in k -space to the desired location in a diagonal trajectory. The phase encoding gradient is then stopped, and the readout gradient is reversed, bringing the spins back through a complete set of k_x value while keeping k_y constant. The experiment is then repeated by returning to a different k_y value until the whole of k -space has been covered in a grid-like fashion.

2.1.5 Parallel imaging

The need to cover k -space through sequential readouts makes MR imaging a slow imaging technique, and the quality of the images is greatly dependent on the absence of motion during the acquisition. This also implies that for body MRI, the patient should not breathe during the whole acquisition process, as this motion would yield a reconstruction image with many undesired artifacts (Glockner et al., 2005). This has motivated further research into ways to accelerate the acquisition.

A first step in this direction is the use of *parallel imaging*. This technique relies on a *multicoil* MRI acquisition, which was introduced at the beginning of the 1990s by Roemer et al. (1990). The idea was to use several receive coils positioned at different locations instead of a single large one. The coils record different signal intensities depending on their spatial location, which enables to achieve a larger signal to noise ratio (SNR) compared to a single coil acquisition. It was subsequently observed that instead of increasing the SNR compared to a single coil acquisition, exploiting the geometry of the coil array could help accelerate MRI by resolving the aliasing due to spacing out successive k -space phase encodes by a factor R known as the *acceleration factor* (Sodickson and Manning, 1997). This is what is known as *parallel imaging*. However, this acceleration introduces a tradeoff between acceleration and SNR as increasing R will reduce the SNR by a factor \sqrt{R} . Reconstructing the image becomes also non-trivial, as the images from each coil

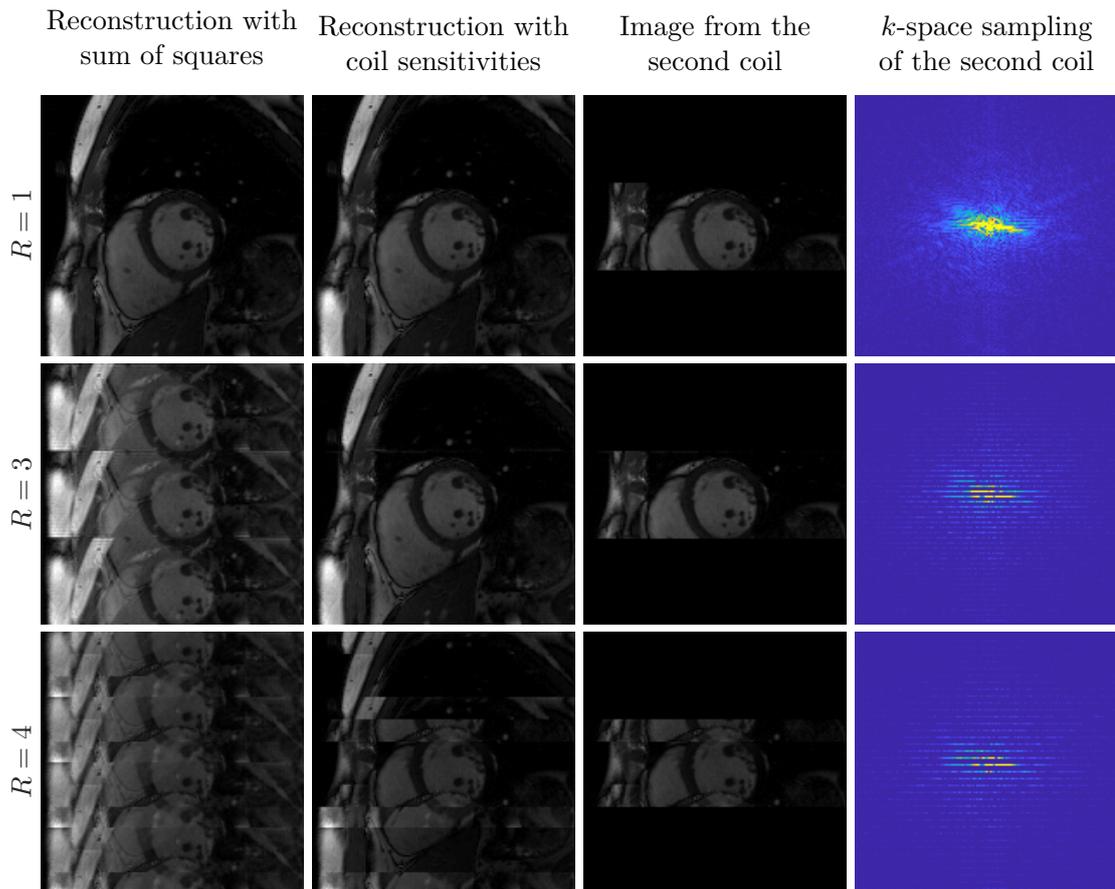


Figure 2.5 – Example of parallel imaging on a cardiac image, simulated with three coils, whose sensitivities are three superimposed rectangles, each covering one third of the image. First column: combination via the sum of squares (SoS) technique, where the images from each coil are simply summed weighted by their modulus. Second column: Reconstruction taking into account the sensitivity maps, allowing for precise reconstruction for an acceleration rate R up to 3. Third column: Image reconstructed for the second coil when taking into account its sensitivity. Fourth column: k -space of the second channel at different undersampling rates.

have a decreased field of view (FoV), resulting in aliasing when extended to the full FoV, as illustrated on Figure 2.5.

Various approaches have tackled this challenge, but they generally rely on some patient-dependent calibration. For instance, SENSE³ (Pruessmann et al., 1999) considered **coil sensitivity maps**, which describe how each coil responds to a signal at each location of the FoV. They are typically estimated from data from a prescan (Blaimer et al., 2004). GRAPPA⁴ (Griswold et al., 2002a) uses a small part of the signal near the center of k -space, called autocalibration signal, in order to compute how the data from different *channels* (coils) should be combined. Several reviews (Blaimer et al., 2004; Glockner et al.,

³Sensitivity encoding

⁴Generalized Autocalibrating Partially Parallel Acquisitions

2005; Deshmane et al., 2012) examine the differences and similarities of these techniques in greater detail and detail various clinical applications. Calibrationless methods have also been explored in more recent years (Trzasko and Manduca, 2011; Majumdar and Ward, 2012; El Gueddari et al., 2019).

An example of parallel imaging with coil sensitivities. We illustrate the role of sensitivity maps to provide additional information to resolve the aliasing due to the acceleration in Figure 2.5. This example assumes three coils with box-like sensitivities covering each a third of the image. We see that in the case $R = 1$, there is no aliasing, and reconstruction is simply achieved by an inverse Fourier transform. In the case $R = 3$. We see that adding the information of each coil allows to still resolve the aliasing, but this is not sufficient for $R = 4$, where the three coils do not provide enough information in order to resolve aliasing. We can see also on the last column how increasing R results in acquisition lines getting more spread out in Fourier space.

Note that this is an idealized example, as in this case the coil-sensitivities do not overlap, which means that summing up individual images multiplied by their respective sensitivities suffices to reconstruct the image. The SENSE algorithm is more involved and can resolve images with overlapping coil sensitivities.

2.1.6 Dynamic MRI

Before terminating this chapter, we will also cover the case of dynamic MRI, which will be relevant to Chapter 5 of this thesis. Dynamic MRI, as its name suggests aims at imaging dynamic, moving objects. It can be used for instance for cardiac imaging (Tsao et al., 2003) or vocal tract imaging (Echternach et al., 2010). The data are acquired in k-space at different times, resulting in k-t raw data, and the direct reconstruction is done with a frame by frame inverse Fourier transform.

In this setting, imaging speed is most critical for several reasons: if one wants to image a fast-moving object precisely, then a high temporal resolution will be required, which presents a challenge for MRI (Tsao et al., 2003). Accelerating the acquisition in this case will not only improve the comfort of the patient, but might also provide higher temporal resolution.

However, the motion between two frames should not be too large, and in practice, which would require breath-held examinations for Cartesian cardiac imaging (Gamper et al., 2008; Jung et al., 2009), which is often impractical for patients. In this case, techniques such as radial sampling can be greatly beneficial, as radial spokes continually sample the center of the k-space, and then motion can be retrospectively corrected in free-breathing samples (Winkelmann et al., 2007; Uecker et al., 2010; Feng et al., 2014, 2016b).

If one wishes to perform Cartesian sampling in this case, higher acceleration rates would

be required, and parallel imaging might not be sufficient to resolve aliasing. As a result, one might use a priori structure in the images to constrain the reconstruction of the image and obtain a good quality outcome. This is known as Compressed Sensing (Candès et al., 2006; Donoho, 2006), and will be discussed in Chapter 3.

2.2 Mathematical description

We close this chapter by introducing the mathematical framework that will be used in the sequel. So far, we have described our problem in a continuous fashion, where the signal is a function of a continuous time t or k-space location (k_x, k_y) . However, in a Cartesian acquisition, the data can be represented on a finite grid as a discrete set of values, depending on the total number of readouts. Then, within each readout, samples will be recorded at a time interval Δt , and so the resolution of k-space can be described as

$$\begin{aligned}\Delta k_x &= \gamma G_x \Delta t \\ \Delta k_y &= \gamma \Delta G_{\text{PE}\tau\text{PE}}\end{aligned}$$

Note that in order to guarantee an image without aliasing, Δt and ΔG_{PE} must be sufficiently small, and their precise value is dictated by the famous Nyquist-Shannon sampling theorem. Assuming a total of N_y phase encoding lines with each continuing N_x samples, we can represent the signal $S(k_x, k_y)$ as a matrix $\mathbf{S} \in \mathbb{C}^{N_x \times N_y}$, which can be written in a vectorized form $\mathbf{s} \in \mathbb{C}^P$, where $P = N_x \cdot N_y$.

In the Cartesian case, the spin density $\rho(x, y)$ will be obtained by computing an inverse Fourier transform on \mathbf{s} . It is common for the resulting image to be described as $\boldsymbol{\rho} \in \mathbb{C}^P$, with the same dimensions as \mathbf{s} . The inversion operation is performed by the inverse Fourier transform \mathcal{F}^{-1} , applied to \mathbf{s} , namely

$$\boldsymbol{\rho} = \mathcal{F}^{-1}(\mathbf{s}). \tag{2.9}$$

This operation is often represented as matrix vector multiplication, where the operator \mathcal{F}^{-1} is explicit as the FFT matrix \mathbf{F}^{-1} , and in addition, one can model the imperfections in the measurement as additive Gaussian noise $\boldsymbol{\epsilon} \in \mathbb{C}^P$, which yields

$$\boldsymbol{\rho} = \mathbf{F}^{-1} \mathbf{s} + \boldsymbol{\epsilon}. \tag{2.10}$$

Note that in practice, Equation 2.10 is not evaluated explicitly, but rather the inverse Fourier transform is directly computed using the Fast Fourier Transform (FFT) algorithm.

2.2.1 Accelerated MRI

In the case of accelerated MRI, the full k-space is not observed, and in practice, entire phase encoding lines are skipped at once, resulting in an acceleration of the scanning time proportional to the number of lines removed. We represent this operation as masking our unobserved coefficients, which is achieved by defining a diagonal masking matrix \mathbf{P}_ω that masks out coefficients indexed by the set ω , namely $\mathbf{P}_{\omega,ii} = 1$ if $i \in \omega$ and 0 otherwise. ω is defined as the sampling mask and has cardinality $|\omega| = N$. ω is typically structured so that entire lines are skipped at once, and a formal presentation of this structure will be discussed in Chapter 5. We refer to the ratio N/P as the *sampling rate*. This masking results in observing the partial image

$$\boldsymbol{\rho} = \mathbf{F}^{-1} \mathbf{P}_\omega \mathbf{s} + \boldsymbol{\epsilon} \quad (2.11)$$

Remark (Sampling mask and trajectories). There is a difference between a sampling mask and a sampling trajectory. The former only provides a fixed view of all locations acquired, while the latter comprises a *dynamic* component. The sampling trajectory tells not only what locations were acquired, but also the temporal ordering of the acquisition, detailing the order in which the samples are acquired throughout the multiple acquisition steps.

In the sequel, we will consider problems where we aim at learning how to reconstruct data from partial Fourier measurements, and how to optimize sampling masks to obtain the best reconstruction quality. In order to be aligned with the commonly used terminology, we will need to redefine some of the notation used until here. We assume that there exists some underlying ground truth image $\mathbf{x} \in \mathbb{C}^P$ from which we acquire partial Fourier measurements (or observations) \mathbf{y}_ω that can be noisy, i.e.

$$\mathbf{y}_\omega = \mathbf{A}_\omega \mathbf{x} + \boldsymbol{\epsilon} = \mathbf{P}_\omega \mathbf{F} \mathbf{x} + \boldsymbol{\epsilon}. \quad (2.12)$$

This will be our basic acquisition model for the whole of this thesis. \mathbf{A}_ω is defined as a shorthand for $\mathbf{P}_\omega \mathbf{F}$. Note that Equations 2.11 and 2.12 can be related by setting $\mathbf{s} = \mathbf{F} \mathbf{x}$ and $\boldsymbol{\rho} = \mathbf{F}^{-1} \mathbf{y}_\omega$. This reformulation allows to describe the image \mathbf{x} as the fundamental unknown, and the partial measurements to be described in the Fourier domain.

Remark. In the case of parallel or multicoil acquisition, Equation 2.12 is slightly adapted to account for acquisition on $j = 1, \dots, C$ coils with different spatial sensitivities

$$\mathbf{y}_{\omega,j} = \mathbf{A}_\omega \mathbf{S}_j \mathbf{x} + \boldsymbol{\epsilon}_j = \mathbf{P}_\omega \mathbf{F} \mathbf{S}_j \mathbf{x} + \boldsymbol{\epsilon}_j, \quad (2.13)$$

where the coil sensitivities can be represented as a $P \times P$ diagonal matrix where $S_{j,ii}$ describes the spatial sensitivity for the i -th pixel. We will also sometimes use the

compact notation

$$\mathbf{y}_\omega = \mathbf{A}_\omega \mathbf{S} \mathbf{x} + \boldsymbol{\epsilon} \quad (2.14)$$

for multicoil MRI, and in this case, $\mathbf{y}_\omega \in \mathbb{C}^{CP}$ will represent the stacked observation, \mathbf{A}_ω will be the coil-wise Fourier transform, $\mathbf{S} \in \mathbb{R}^{CP \times P}$ will be the stacked sensitivities.

Recovering the ground truth image \mathbf{x} from partial measurement \mathbf{y}_ω is challenging, as the problem of recovering an image from partial measurements is fundamentally ill-posed: there is mathematically infinitely many solutions that are consistent with \mathbf{y}_ω , but they are not physically meaningful. Such problems are often referred to as ill-posed *inverse* problems, where we aim at recovering a reference image from partial information that entail a loss of information. In this context, Equation 2.12 is referred to as a *forward* model. In the next chapter, we will discuss the approaches that have been studied in the literature to construct an estimate $\hat{\mathbf{x}}$ of \mathbf{x} from partial measurements \mathbf{y} .

3 Reconstruction methods

In this chapter, we consider the solutions that have been explored to the problem of *accelerated MRI* in the literature. We have defined this problem formally in Section 2.2.1: we consider partial measurements \mathbf{y}_ω obtained through the forward model of Equation 2.12, and aim at constructing an estimate $\hat{\mathbf{x}}$ of the underlying, unknown ground truth \mathbf{x} .

The difficulty of this problem resides in its ill-posedness: there exist infinitely many solutions consistent with the observations that are not physically meaningful. To counter it, it is necessary to impart additional information to the problem in order to reach a sensible solution. This was the approach taken in Compressed Sensing (CS) applied to MRI (Lustig et al., 2008). The ill-posedness of the problem is alleviated by imposing additional *structure* on it. All subsequent approaches rely on imposing some form of structure to estimate \mathbf{x} from \mathbf{y} , but have relied on increasingly complex, data-driven approaches.

We will trace back the evolution from CS approaches, that used simple mathematical models to obtain structure, to recent deep learning-based approaches that construct their models from training data: we moved from model-driven approaches to data-driven approaches. A more exhaustive presentation of these approaches can be found in Ravishankar et al. (2019); Doneva (2020).

3.1 Model-driven approaches

3.1.1 Compressed Sensing

It has been long known that a band-limited signal can be exactly reconstructed from a set of uniformly sampled frequencies, provided that their density is at least twice the highest frequency of the original signal. This result is famously known as the Nyquist-Shannon sampling theorem. However, sampling at the Nyquist rate can remain expensive, and there has been extensive research to achieve sub-Nyquist rate sampling and reconstruction.

Chapter 3. Reconstruction methods

In their seminal works, Donoho (2006); Candès et al. (2006) introduced the formalism of Compressed Sensing (CS), where they proved that by leveraging the structure present in natural images, such as sparsity or low-rankness, one can perfectly reconstruct a signal sampled much below the Nyquist rate. Achieving this required using a non-linear reconstruction method, framed as the following prototypical problem

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}_\omega \mathbf{x} - \mathbf{y}_\omega\|_2^2 + \lambda \mathcal{R}(\mathbf{x}) \quad (3.1)$$

where the first term, known as *data-fidelity*, enforces consistency with the measurement, and the second term, called *regularization*, defines a statistical model, capturing the desired structure of the signal. In the literature, $\mathcal{R}(\mathbf{x})$ has taken various forms, but typical models include sparsity in the wavelet domain, $\mathcal{R}(\mathbf{x}) = \|\mathbf{W}\mathbf{x}\|_1$ (Candès et al., 2006; Donoho, 2006; Lustig et al., 2007), or sparsity with Total Variation (TV) kind of constraints (Block et al., 2007; Knoll et al., 2011a). An alternative approach would rather impose a low-rank structure on the resulting image, typically using $\mathcal{R}(\mathbf{x}) = \|\mathbf{x}\|_*$ (Lingala et al., 2011), or more sophisticated structures such as the Hankel matrix (Jin et al., 2016). The common trend in all of these methods is the idea that *the structure of the signal implies a redundancy that can be exploited to represent the signal in a compact fashion* when transforming it appropriately. As expressing the structure in a mathematical fashion is not a straightforward task, these methods are the expression of various attempts to capture this idea.

Note that all these approaches typically introduce a trade-off between acquisition speed (by lowering the number of k-space lines acquired) and reconstruction speed. Solving Equation 3.1 requires using an iterative method, and while the computation is carried out offline, the reconstruction can last up to hours for specific settings with large accelerations (Feng et al., 2016a).

We note also for completeness that CS relies not only on leveraging additional structure in the image to be reconstructed, but also on incoherent sampling (Lustig et al., 2008). In practice, this often involves random-like sampling strategies, which enable the aliasing to be incoherent, i.e. to feature a noise-like structure, as illustrated on Figure 3.1. We will defer a more complete discussion of sampling to the next chapter.

Overall, it is clear that such approaches rely on *statistical* models, constructed to implement the designer’s assumptions, rather than being learned from data. While such methods enabled great progress in the last decade, the statistical models restrict the expressivity of the approach: they express abstract mathematical structures encoding our assumptions rather than focusing on the real data distribution expressed through the training set. This is why we will refer such approaches as *model-driven*.

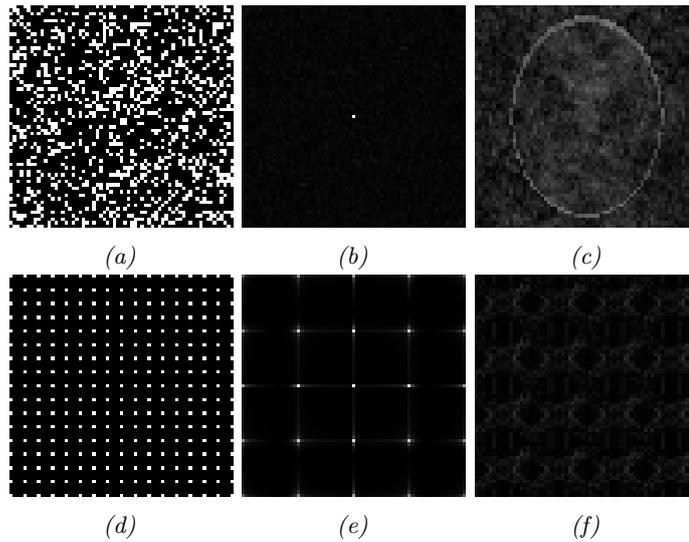


Figure 3.1 – Random (a) and regular (d) sampling masks, acquiring 25% of the total locations, and their corresponding point spread functions (PSF) in (b) and (e). The convolution of the PSF with the Shepp–Logan phantom is shown respectively in figures (c) and (f). We see that the random undersampling produces incoherent or noise-like aliasing, which still allows us to distinguish the original image. On the other hand, the aliasing produced by regular, grid-based sampling (known as fold-over artifacts) yields a low-intensity result where the original image cannot be distinguished.

3.2 Data-driven approaches

Many approaches tried to learn their model directly from data, and in this Section, we survey a couple of the most significant representatives of data-driven methods.

3.2.1 Dictionary learning

Dictionary learning approaches rely on the assumption that an image can be represented as a sparse combination of atoms from a dictionary, i.e. $\mathbf{x} = \mathbf{D}\mathbf{z}$ where \mathbf{D} is a dictionary and \mathbf{z} the sparse coefficient vector. This is sometimes referred to as a synthesis formulation, as opposed to 3.1, which shows an analysis formulation. Synthesis models are often applied patch-wise to images, and they aim at simultaneously learning a reconstruction \mathbf{x} , a dictionary \mathbf{D} and an encoding \mathbf{Z} (Ravishankar and Bresler, 2011b; Caballero et al., 2014)

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{D}, \mathbf{Z}} \quad & \frac{1}{2} \|\mathbf{A}_\omega \mathbf{x} - \mathbf{y}_\omega\|_2^2 + \beta \sum_{j=1}^N \|\mathbf{P}_j \mathbf{x} - \mathbf{D} \mathbf{z}_j\|_2^2 \\ \text{s.t.} \quad & \|\mathbf{z}_j\|_0 \leq s, \|\mathbf{d}_i\|_2 = 1, \forall i, j \end{aligned} \quad (3.2)$$

3.2.2 Unrolled neural networks

Unrolled neural networks are a data-driven, deep learning-based approach inspired by the iterative algorithms used to solve Equation 3.1. This problem is known as a composite optimization problem, and can be solved by various approaches. A famous approach is known as the Iterative Soft-Thresholding Algorithm (ISTA) (Daubechies et al., 2004), or Proximal Gradient method (Combettes and Pesquet, 2011), solves the problem by iteratively computing

$$\mathbf{x}^{t+1} = \text{prox}_{\lambda\mathcal{R}} \left(\mathbf{x}^t - \alpha \mathbf{A}^\dagger (\mathbf{A} \mathbf{x}^t - \mathbf{y}_\omega) \right). \quad (3.3)$$

The proximal operator is defined as $\text{prox}_g(\mathbf{x}) \triangleq \arg \min_{\mathbf{v}} g(\mathbf{v}) + \frac{1}{2} \|\mathbf{v} - \mathbf{x}\|_2^2$, and for some functions, can have a closed-form solution. This is typically the case when $\mathcal{R}(\mathbf{x}) = \|\mathbf{W}\mathbf{x}\|_1$, where the proximal takes the form of a soft-thresholding operator, i.e. $\text{prox}_{\lambda\|\mathbf{W}\cdot\|_1}(\mathbf{x}) = \text{sign}(\mathbf{W}\mathbf{x}) \odot \max(\mathbf{W}\mathbf{x} - \lambda, 0)$.

Motivated by the formulation of Equation 3.3, unrolled neural networks parameterize the proximal operator with a neural network $\mathbf{f}(\mathbf{y}_\omega, \omega; \theta)$, and then train a model of the form

$$\mathbf{x}^{t+1} = \mathbf{f} \left(\mathbf{x}^t - \alpha \mathbf{A}^\dagger (\mathbf{A}_\omega \mathbf{x}^t - \mathbf{y}_\omega); \theta_{(t)} \right) \quad (3.4)$$

for a fixed, small number of iterations $t = 1, \dots, T$. The output of the model is thus $\mathbf{f}(\mathbf{y}_\omega, \omega; \theta) = \mathbf{x}^T$ where $\theta = \{\theta_1, \dots, \theta_T\}$. As the proximal term capture the statistical properties of the signal that wish to recover, parametrizing it with a neural network enables to learn a model directly from data. The model being more flexible, it allows for improved reconstruction performance over methods using statistical models.

A flurry of approaches has been motivated by unrolling iterative algorithms, and we highlight a few examples hereafter. A more complete discussion of unrolled neural networks can be found in Liang et al. (2020). In Mardani et al. (2018), the parameters of the proximal mapping can be shared across iterations, or different at each iteration (Schlemper et al., 2017). The unrolling can leverage the structure of a different optimization algorithm, like the Alternating Direction Method of Multipliers (ADMM) in Sun et al. (2016), or the Primal-Dual Hybrid Gradient in Adler and Öktem (2018b). In the noiseless case, it is also common to replace the inner step $\mathbf{x}^t - \alpha \mathbf{A}^\dagger (\mathbf{A}_\omega \mathbf{x}^t - \mathbf{y}_\omega)$ with a step called *Data Consistency* (DC) (Schlemper et al., 2017; Zhang et al., 2019b), that replaces the reconstruction with the exact values at observed locations. Essentially, the inner step becomes

$$\text{DC}(\mathbf{x}^t, \mathbf{y}_\omega) = \mathbf{F}^\dagger (\mathbf{P}_{\omega^c} \mathbf{F} \mathbf{x}^t + \mathbf{y}_\omega) = \mathbf{F}^\dagger (\mathbf{P}_{\omega^c} \mathbf{F} \mathbf{x}^t + \mathbf{P}_\omega \mathbf{F} \mathbf{x}). \quad (3.5)$$

We see that at the observed locations ω , we simply keep the observation, while at the non-observed locations ω^c . Note that here, $\mathbf{P}_{\omega^c} = \mathbf{I} - \mathbf{P}_\omega$.

These unrolled approaches require training the models f_θ on a dataset of reference images and observations $\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^m$ by minimizing a loss $\ell(\cdot, \cdot)$:

$$\theta^* = \arg \min_{\theta} \frac{1}{m} \sum_{i=1}^m \ell(\mathbf{x}_i, \mathbf{f}(\mathbf{y}_i; \theta)). \quad (3.6)$$

Various losses have been used in the literature. Initial works used ℓ_2 loss (Schlemper et al., 2017; Hammernik et al., 2018), but lately practitioners have turned to using ℓ_1 loss, SSIM or compounds of these Muckley et al. (2021). ℓ_1 was indeed found to outperform ℓ_2 on vision tasks Zhao et al. (2016). The structural similarity (SSIM) (Wang et al., 2004) is a metric developed to match more closely the human perception, and recent work have proposed that a weighted sum between ℓ_1 loss and the multiscale SSIM (MS-SSIM) as a compromise (Zhao et al., 2016).

Deep learning-based reconstruction methods further shift the burden of computation, compared to compressed sensing approaches. While these enabled faster acquisition at the cost of a slower, iterative reconstruction procedures, deep learning-based methods retain the benefit of enabling sub-Nyquist sampling, but also enable fast reconstruction, in the order of 10^{-1} second for a slice (Jin et al., 2017; Hammernik et al., 2018). However, the computational burden is only shifted an additional step: these methods need to be trained with large, fully-sampled¹ datasets (Zbontar et al., 2019).

3.2.3 Deep image prior

In a spirit closer to dictionary learning and classical CS, Ulyanov et al. (2018) have shown that untrained CNN can provide a sufficiently strong prior to enable image reconstruction, even without training data. Given an untrained CNN \mathbf{f}_θ a single measurement instance \mathbf{y}_ω , the idea in deep image prior (DIP) is simply to minimize

$$\min_{\theta} \|\mathbf{y}_\omega - \mathbf{A}_\omega \mathbf{f}_\theta(\mathbf{z})\|_2^2 \quad (3.7)$$

where \mathbf{z} is a fixed, randomly sampled vector, and where the parameters θ are randomly initialized. Then, using an iterative optimization algorithm, fitting the weights to the observations enables the network to output a high-quality reconstructed image without supervision or training data. This method was successfully applied to MRI in (Darestani and Heckel, 2021; Yoo et al., 2021). A common pitfall of DIP is that this flexibility comes at the price of an increase inference time: the network weights must be retrained for each individual image. In addition, DIP can suffer from overfitting and so heuristics such as early stopping are used to counteract this (Ulyanov et al., 2018; Sun et al., 2021).

¹In their recent work, Yaman et al. (2020) proposed approaches to train only using undersampled data, but having access to a dataset of undersampled images is still required.

3.2.4 A Bayesian interpretation of reconstruction

So far, the discussion has been framed as moving from mathematically-defined to data-driven regularization in order to improve the recovery of Problem 3.1. However, this approach can also be framed from a Bayesian perspective, where the data consistency term $\frac{1}{2}\|\mathbf{A}_\omega \mathbf{x} - \mathbf{y}_\omega\|_2^2$ is viewed as a likelihood and the regularization $\lambda\mathcal{R}(\mathbf{x})$ as a prior. This was initially discussed by Ji et al. (2008) for compressed sensing. The idea revolves around Bayes rule, where the posterior $p(\mathbf{x}|\mathbf{y})$ can be expressed as

$$p(\mathbf{x}|\mathbf{y}) = \frac{\overbrace{p(\mathbf{y}|\mathbf{x})}^{\text{likelihood}} \overbrace{p(\mathbf{x})}^{\text{prior}}}{\underbrace{p(\mathbf{y})}_{\text{marginal}}} \propto p(\mathbf{y}|\mathbf{x})p(\mathbf{x}). \quad (3.8)$$

Under Gaussian noise, we can derive an explicit likelihood model for Equation 2.12, where, assuming $\epsilon \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$, i.e assuming additive Gaussian noise, we get $p(\mathbf{y}|\mathbf{x}) = \frac{1}{(2\pi\sigma^2)^{P/2}} \exp\left(-\frac{1}{2\sigma^2}\|\mathbf{y} - \mathbf{A}_\omega \mathbf{x}\|_2^2\right) \sim \mathcal{N}(\mathbf{y}; \mathbf{P}_\omega \mathbf{F} \mathbf{x}, \sigma^2 \mathbf{I})$. The regularization term can then be viewed as a prior from which the ground truth \mathbf{x} is sampled, in the case of classical CS, the prior can be specified as the Laplace density function $p(\mathbf{x}) = \left(\frac{\lambda}{2}\right)^P \exp(-\lambda\|\mathbf{W}\mathbf{x}\|_1)$ (Ji et al., 2008). Such an approach enables to view Problem 3.1 as a maximum a posteriori (MAP) estimation, where one aims to find the reconstruction $\hat{\mathbf{x}}$ that maximizes posterior probability

$$\begin{aligned} \hat{\mathbf{x}} &= \arg \max_{\mathbf{x}} p(\mathbf{x}|\mathbf{y}) = \arg \max_{\mathbf{x}} \frac{p(\mathbf{y}|\mathbf{x})p(\mathbf{x})}{p(\mathbf{y})} && \text{by definition} \\ &= \arg \max_{\mathbf{x}} p(\mathbf{y}|\mathbf{x})p(\mathbf{x}) && \text{as the optimization is indep. from } \mathbf{y} \\ &= \arg \min_{\mathbf{x}} -\log(p(\mathbf{y}|\mathbf{x})p(\mathbf{x})) && \text{by taking the log and negative of the problem} \\ &= \arg \min_{\mathbf{x}} \frac{1}{2\sigma^2}\|\mathbf{y} - \mathbf{A}_\omega \mathbf{x}\|_2^2 + \lambda\|\mathbf{W}\mathbf{x}\|_1 && \text{explicit the functions and discard the constants} \end{aligned}$$

The Bayesian perspective gives a different interpretation of reconstruction and of the role of regularization as a prior, and lends itself naturally to an extension to data-driven models, which act as learned prior models. The Bayesian framing of Equation 3.1 has also the advantage of naturally defining a criterion to acquire observations that have large uncertainty according to the model (Ji et al., 2008). We will discuss these implications in greater depth in the next Chapter as well as in Chapter 7. For now, we turn to our last category of deep learning-based approaches which make use of this Bayesian insight and attempt to explicitly learn meaningful, data-driven priors.

3.2.5 Generative approaches

The role of the reconstruction methods can be seen as learning a data-driven prior which has proven to enable higher quality of reconstruction than mathematically motivated models. While reconstruction methods embed the full iterative resolution of Problem 3.1 into a deep architecture, other approaches, relying on deep generative models, propose to explicitly parametrize the prior with a neural network (Bora et al., 2017; Jalal et al., 2021; Patel and Oberai, 2021). This has the advantage of keeping the optimization procedure more transparent, as it will be about explicitly finding a sample from a deep generative prior that is consistent with the observation. The approach, proposed initially by Bora et al. (2017) relies on training a generative adversarial network (GAN) to capture the prior distribution. We first briefly introduce GANs.

Generative adversarial networks. GANs (Goodfellow et al., 2014) are a specific type of deep learning models where a generator $\mathbf{g}_\theta(\mathbf{z})$ and a discriminator $d_\phi(\mathbf{x})$ are trained in a competitive fashion: starting from random noise $\mathbf{z} \sim p(\mathbf{z})$, the generator tries to construct samples that imitate a distribution of references images $p(\mathbf{x})$ from which we are given a large set of samples $\{\mathbf{x}_i\}_{i=1}^m$. The discriminator, on the other hand, tries to classify a sample based on how likely it is to originate from the reference distribution. The problem can formally be described as a game between the generator and the discriminator, where we solve

$$\min_{\theta} \max_{\phi} \mathbb{E}_{p(\mathbf{x})} [d_\phi(\mathbf{x})] - \mathbb{E}_{p(\mathbf{z})} [d_\phi(\mathbf{g}_\theta(\mathbf{z}))]. \quad (3.9)$$

Generative priors. Given a trained generator $\mathbf{g}_{\theta^*}(\mathbf{z})$ approximating the image distribution of interest $p(\mathbf{x})$, the approach of Bora et al. (2017), applied also to MRI in Patel and Oberai (2021), aims at finding a sample from the generator consistent with the observations by solving

$$\min_{\mathbf{z}} \|\mathbf{A}_\omega \mathbf{g}_{\theta^*}(\mathbf{z}) - \mathbf{y}_\omega\|_2 \quad (3.10)$$

While this approach does not enable to directly construct a maximum a posteriori (MAP) estimate, due to the GAN only yielding samples and no probabilities, the low-dimensionality of the latent space \mathbf{z} enables to efficiently carry out Markov Chain Monte Carlo (MCMC) computations (Patel and Oberai, 2021) to construct posterior samples.

3.3 Trade-offs

While all the methods presented in this chapter enable reconstruction from sub-Nyquist sampling rates, they feature different advantages and drawbacks. The state-of-the-art in terms of reconstruction performance is occupied by deep learning-based reconstruction methods, such as unrolled algorithms. They allow for the best quality and fastest inference

among all methods presented here, but require large training datasets and are not robust to distribution shifts (Jalal et al., 2021; Darestani et al., 2021). Generative priors also require large amounts of data to train the generative model, and generally perform less well than trained reconstruction networks, while also having a slower inference time (Darestani et al., 2021). They have the advantage however of being robust to some distribution shifts like change in the distribution of masks ω and their inference is more interpretable than deep learning methods, as the conditioning on the observations is explicitly computed from the prior, rather than through a black box like in an end-to-end reconstruction algorithm.

Compressed sensing (CS), dictionary learning (DIL) and deep image prior (DIP) are generally slower at inference, as they all rely on iterative reconstruction, and require the optimization of increasingly large amounts of variable: from the image in CS to the image and dictionary in DIL to the parameters of a deep network in DIP. However, DIP is found to generally outperform CS-based approaches. The trade-offs between DIL, DIP and Generative priors is however less clear, as these methods have not been extensively compared in the literature.

Overall, this survey highlights a couple of the general trends observed in the literature and is not exhaustive. It shows however a clear trend where different methods try to exploit the redundancy due to the structure in the data. Earlier works attempted to explicitly embed this structure by compactly representing the signal in a variety of different bases, such as Wavelet domain (Lustig et al., 2007), or through representing the signal as low-rank (Lingala et al., 2011). Later, approaches aimed rather at *learning* how to leverage this structure from data, and use the powerful tools of deep learning to represent the signal in complex ways, following a data-driven paradigm. All these methods rely however on the idea that the structure of the signal allows to alleviate the ill-posedness of the inverse problem, and retrieve a signal from undersampled observations, because of its underlying structure.

In the next chapter, we will detail the evolution in mask designs ω used to generate the observation and how they have been optimized in order to provide the most informative observations to enable the best reconstruction downstream.

4 Optimizing sampling patterns

In the previous chapter, we discussed the evolution of reconstruction methods for accelerated MRI, where a shift from model-driven approaches to data-driven approaches was observed. A similar shift can be observed in the design of sampling patterns.

4.1 A taxonomy of mask design methods

While the evolution of reconstruction methods can generally be described along the axis of moving from model-based approaches towards learning-based ones, the picture is less clear for the mask design methods. It is nonetheless to consider two axes of discussion that run in parallel to the evolution witnessed in reconstruction methods.

These two axes pertain to the two main components needed to optimize sampling patterns (or masks or trajectories). One needs **i)** a prior distribution from which candidates patterns are drawn, and **ii)** a criterion, a metric to quantify its performance. We will see that at the onset of compressed sensing, sampling patterns were drawn from heuristically designed parametric distributions, a setting widely known as *variable density sampling*, and gradually moved towards distributions tailored for a specific dataset. Similarly, the criteria used to establish the quality of a mask initially were mathematical quantities such as *coherence*, but the criteria became increasingly related to the reconstruction performance enabled by subsampling data using a given mask. One can clearly see a shift from **i)** *model-based* approaches towards *learning-based* ones, and from **ii)** *model-driven* criteria towards *data-driven* ones.

Remark. We chose to denote the distribution from which candidate sampling patterns are drawn as the *base* for the sampling optimization procedure, while the criterion used as what *drives* it.

In the rest of this chapter, we will survey the literature from the first application of compressed sensing to MRI onwards to show how mask design evolved from model-

based, model-driven approaches towards learning-based, data-driven approaches. We will conclude by formalizing the problems linked with optimizing mask design in this context.

4.2 Compressed sensing and MRI: Model-based, model-driven approaches

4.2.1 The contribution of Lustig et al. (2007)

With the onset of Compressed Sensing (CS), random sampling patterns were used, as they allowed to obtain noise-like, incoherent aliasing, as illustrated on Figure 3.1. This was indeed theoretically motivated, as random sampling easily enables to achieve incoherent sampling, which in turn meant that a small number of measurements were sufficient to exactly solve the problem 3.1 for $\mathcal{R}(\mathbf{x}) = \|\mathbf{W}\mathbf{x}\|_1$ with high probability.

Definition. Let (\mathbf{A}, \mathbf{W}) be two orthonormal bases of \mathbb{C}^p , where \mathbf{A} is the observation domain, and \mathbf{W} is the sparsifying (representation) basis. Then, the **coherence** between these bases is defined as (Candès and Wakin, 2008)

$$\mu(\mathbf{A}, \mathbf{W}) = \sqrt{p} \cdot \max_{1 \leq i, j \leq p} |\langle \mathbf{a}_i, \mathbf{w}_j \rangle|. \quad (4.1)$$

The coherence measures the maximal correlation between any element of \mathbf{A} and \mathbf{W} . $\mu(\mathbf{A}, \mathbf{W}) \in [1, \sqrt{p}]$, where coherence 1 implies maximal incoherence between the bases, and \sqrt{p} maximal coherence.

Coherence then allows to bound the minimal number of measurements N required for exactly solving 3.1 with high probability. Indeed, for a given signal $\mathbf{x} \in \mathbb{C}^p$ observed in domain \mathbf{A} , assuming that \mathbf{x} is S -sparse in the transformed domain \mathbf{W} , then a typical compressive sensing result (Candes and Romberg, 2007) dictates that, given N *random* measurements, exact reconstruction will occur with high probability when

$$N \geq C \cdot S \cdot \mu^2(\mathbf{A}, \mathbf{W}) \cdot \log P \quad (4.2)$$

for some positive constant C .

While the theory of CS obtains results for *random* measurements, practitioners have quickly sought to optimize the measurements, and a natural initial criterion has been to minimize the coherence between measurements, i.e. to select a sampling pattern ω with $|\omega| = N$ that enables to minimize $\mu(\mathbf{A}_\omega, \mathbf{W})$ (Lustig et al., 2007). However, when applying compressed sensing to MRI, it was quickly observed that purely random undersampling of the Fourier space is not practical, for several reasons (Lustig et al., 2007, 2008): **i**) it does not satisfy hardware and physical constraints that require smooth trajectories in Fourier space and the use of trajectories robust to system imperfections,

4.2. Compressed sensing and MRI: Model-based, model-driven approaches

ii) it does not consider the energy distribution of the signal in k-space, which is very concentrated around the center

For this reason, the seminal work of Lustig et al. (2007) proposed to use a variable-density sampling approach, where the k-space is less undersampled near the origin (low frequency, high energy region) and more towards the periphery (high frequency but low energy)¹. They used a distribution centered at the origin of k-space, where the probability of acquiring a location decayed at a polynomial rate as a distance of the origin, namely

$$P(k_x, k_y) = \left(1 - \frac{2}{\sqrt{p}} \sqrt{k_x + k_y}\right)^d \quad (4.3)$$

with decay $d > 1$, and $-p/2 \leq k_x, k_y \leq p/2$. This heuristic was proposed to imitate the decay in the spectrum observed on real data. The mask to be used is then estimated by a Monte-Carlo procedure: a set of masks are randomly sampled from Equation 4.3, and the mask with minimal coherence is retained.

The idea proposed by Lustig et al. (2007) of using variable-density sampling for MRI was not totally novel, as variable-density sampling had previously been used in the MRI community to design undersampling trajectories with incoherent aliasing when using linear reconstructions (Marseille et al., 1996; Tsai and Nishimura, 2000). However, Lustig et al. (2007) were the first to integrate this heuristic within the framework of CS.

It should be clear from the presentation that the approach of Lustig et al. (2007) is a *model-based approach*, as the masks are sampled from a heuristic probability distribution (eq. 4.3) and also a *model-driven approach*, as the criterion used to select the masks is *coherence*, a mathematical structure prescribed by the theory of compressed sensing.

4.2.2 Model-based, model-driven sampling in the sequel

A similarly model-based, model-driven was followed by Gamper et al. (2008), which considers dynamic MRI, although from a different initial mask: they attempted at randomizing a deterministic mask design proposed by Tsao et al. (2003) to achieve incoherent sampling. The *prior* in this case is given by a structured mask.

Some refinements to variable-density sampling (VDS) were developed for instance in (Wang et al., 2009), but in the sequel, two broad approaches were followed to improve the design of sampling patterns. The first approach leveraged learning-based approaches, and will be discussed in the next section, while the second attempted at developing theoretical models that more closely captured the structure of real-world signals.

¹The name *variable-density sampling* originates from the very fact that k-space is not uniformly subsampled, but rather with a variable density (Tsai and Nishimura, 2000).

The work of Wang and Arce (2009) is a first step in this direction, where improved VDS functions are proposed by exploiting a priori statistical information of real-world images in the wavelets domain. One can notice that this approach chose to tie the sampling pattern specifically to a type of regularization, namely sparsity in the wavelet domain, and obtains better sampling density for this problem but lose in generality over competing approaches.

The work of Puy et al. (2011) proposes a coherence-driven optimization of the variable density sampling procedure. Instead of relying on a heuristic distribution such as equation 4.3, the procedure proposes an objective that yields a VDS distribution that minimizes the coherence between the measurement and representation bases.

A large body of subsequent works (Krahmer and Ward, 2013; Roman et al., 2014; Adcock et al., 2015, 2017) rely on some additional theoretical tools such as *local coherence*, *asymptotic sparsity*, and *multilevel sampling* to give a solid ground to heuristics such as VDS that show good performance but lack theoretical justification.

4.3 Towards learning-based, data-driven approaches

In parallel to these works, there was a motivation to move towards learning-based approaches, due to some weaknesses in the VDS approach of Lustig et al. (2007). One of them is that the VDS distribution of Equation 4.3 is a heuristic that requires tuning to work well on different types of data, and that heuristically tuned variable density distributions generally performed better than the ones solely prescribed by CS theory (Chauffert et al., 2013).

Learning-based, model-driven sampling. A first work of Knoll et al. (2011b) aimed at solving this issue by constructing the VDS distribution from data rather than from a tuned heuristic. They simply average the spectra of several similar images, normalized it and sampled their mask from the resulting distribution, minimizing coherence to select the final mask, as in Lustig et al. (2007). This approach was then improved by Zhang et al. (2014); Vellagoundar and Machireddy (2015).

Early learning-based, data-driven approaches. Ravishankar and Bresler (2011a); Liu et al. (2012) proposed similar methods where they aimed at designing a sampling pattern that minimized the *reconstruction error* in Fourier space, rather than coherence. Both algorithms partition k-space in blocks, and iteratively design a mask by removing samples from low-error regions in order to re-assign them in high-error regions. Liu et al. (2012) made the important point that the design of a mask should reconstruction-based and not observation-based. Until then, approaches did not include the reconstruction algorithm into the design of the mask, but relied on model-based criteria relying on the observations, such as coherence. This insight proved to be impactful in the performance

4.4. A matter of performance: how do these methods compare to each other?

of the methods, and while no quantitative comparison was available in the paper, further works confirmed the impact of including the reconstruction algorithm in the mask design (Zijlstra et al., 2016).

While these methods capture the paradigm that will be overwhelmingly dominant in the deep learning era, there was little follow-up work until much later, around 2018.

The case of Seeger et al. (2010). The work of Seeger et al. (2010) lies in its own category. This approach builds on a sequential Bayesian experiment design perspective for sparse linear models (Seeger et al., 2007; Seeger and Nickisch, 2008). First a model of the posterior distribution $p(\mathbf{x}|\mathbf{y}_\omega)$ is built, and then, a sampling pattern is iteratively built by successively acquiring the readouts that have the highest entropy. This approach is able to sequentially tailor the sampling pattern to the need of an unseen image, given the previously acquired measurements. However, the approach features a quite high computational cost for a single image. Liu et al. (2012) raise that is not practical, as the matrix inversion required to compute the next measurement that must be acquired is too slow to match the readout time of a scanner.

While methods to optimize sequentially the sampling mask are common in the deep learning era, the approach of Seeger et al. (2010) is unique in its time, as all other works focused on optimizing the sampling pattern for a fixed sampling rate, and not in a sequential fashion. While more principled than other approaches, and using maximum entropy as the criterion to optimize the acquisition trajectory, this method remains model-based (using a Laplace prior) and model-driven (maximum entropy).

4.4 A matter of performance: how do these methods compare to each other?

Very few of the works discussed until then feature an extensive comparison of performance of their proposed approaches. Most works rely on visual comparisons, or state the weaknesses of previous methods, but do not compare against them. In this regard, the contributions that bring a comparison of various approaches to mask design are particularly valuable, and in this section, we will discuss two of them, namely Chauffert et al. (2013) and Zijlstra et al. (2016).

The work of Chauffert et al. (2013) aims at designing a better VDS density, while also providing a comparison between theoretically informed sampling patterns, and the ones obtained using the polynomial of Equation 4.3 tuned at various degrees $d \in \{1, 2, \dots, 6\}$, and a hybrid distribution mixing some theoretically prescribed components and heuristics. It is interesting to observe that in their data, the tuned hybrid distribution performs best, followed by the tuned polynomial and the optimal theoretically predicted distribution. This theoretically optimal distribution is based on the arguments from Rauhut (2010) uses

rejection of already acquired locations (Puy et al., 2011) and minimizes coherence, while the polynomial was tuned on the data using directly the final evaluation metric, namely PSNR. The authors conclude that the theory would need to capture additional structure of real-world images for sampling patterns obtaining better practical performance.

The article of Zijlstra et al. (2016) compares the performance of coherence-based VDS using Equation 4.3, of Knoll et al. (2011b), of Liu et al. (2012) and of VDS optimized to minimize normalized root mean squared error (NRMSE). The results show that coherence-based VDS is generally outperformed by Knoll et al. (2011b), which is matched by the optimized VDS. Both are outperformed by the method of Liu et al. (2012). The conclusion is however careful:

When putting the presented results in perspective, it is important, to remember that data-driven optimization of an undersampling pattern for CS is just one aspect of the CS implementation. Data-driven approaches should be taken to optimize other aspects of the CS implementation, such as the sparse domain and the reconstruction algorithm.

Overall, while these works do not exhaustively compare the existing methods in the literature, they give us reasons to view learning-based, data-driven methods as the most direct avenue towards improvement in mask design optimization.

4.5 The advent of deep learning: Learning-based, data-driven approaches

Interest for learning-based, data-driven approaches to mask optimization was renewed around 2018, with the advent of the first deep learning-based reconstruction methods applied to MRI (Sun et al., 2016; Schlemper et al., 2017; Mardani et al., 2018). This enabled the possibility to scale up the complexity of mask design methods, as deep learning methods enable fast reconstruction times, which was not the case for CS methods.

The first work that we discuss in this section comes slightly before this time and is the one of Baldassarre et al. (2016), where the authors propose a first principled, learning-based approach to optimizing a deterministic sampling pattern given a set of training signals. Note that in our taxonomy, their approach is both learning-based *and* data-driven. While the optimization only considers the energy of the signal in the case of Fourier sampling, and while the paper on *linear* reconstruction, the authors show improvement about the sampling patterns prescribed in Lustig et al. (2007); Roman et al. (2014). However, as they optimize only the sampling pattern to capture the *energy* of the signal and do not include the reconstruction algorithm in their optimization procedure, they fail to obtain significant improvement by turning to non-linear reconstruction methods.

4.5. The advent of deep learning: Learning-based, data-driven approaches

However, the authors published a sequel to this work with Gözcü et al. (2018), that introduced a sequential optimization of the sampling pattern in a fully data-driven fashion. The principle consists, at each step, to evaluate the improvement for a given performance metric of including a candidate location in the final sampling pattern. Once all candidate locations have been evaluated, the algorithm picks the one that brings the largest improvement and moves to the next step. This method was shown to outperform existing approaches, such as the one of Lustig et al. (2007) and Knoll et al. (2011b) by a large margin, due to its approach being both learning-based and data-driven. Part of the improvement is due also to not solely adding sampling locations based on the current error, but computing how adding a location would *actually* decrease the error in the reconstruction. The distinction is important as the method of Gözcü et al. (2018) is able to take into account how the newly acquired location affects the behavior of the non-linear reconstruction algorithm, rather than acquiring locations based on the current reconstruction error. We will describe the setting of Gözcü et al. (2018) in greater depth in Chapter 5 as this is the basis upon which this thesis was built.

At roughly the same time as Gözcü et al. (2018), Haldar and Kim (2019) published another work aiming at optimizing sampling, based on an experiment design approach to minimize the Cramér-Rao bound (CRB). The CRB provides a lower-bound to the covariance matrix that is based on the measurement matrix \mathbf{A}_ω . This approach is independent on the reconstruction algorithm, and provides a different take on a learning-based, model-driven approach, aiming at minimizing the covariance of an estimator. In the context of MRI sampling, this amounts to finding the measurement matrix that will encode the most information about the original signal \mathbf{x} .

The last non-fully deep learning-based work that we will discuss is the one of Sherry et al. (2019), where a bilevel optimization procedure is proposed: the first level solves an empirical risk minimization problem for the mask, while the second level consists in an iterative algorithm to solve a Problem similar to 3.1. The work takes an original approach whereby relaxing the mask optimization problem into a continuous optimization problem, they are able to tackle the mask optimization and the reconstruction under a unified framework. While the reconstruction remains model-based, the mask stage is clearly learning-based and data-driven, as training data are used, and aim at minimizing some loss between the original image and the reconstruction.

Due to the overwhelming performance of deep learning methods, most subsequent works focus exclusively on deep learning-based reconstruction methods, and all follow a learning-based, data-driven paradigm. All methods initialize the prior from which masks are drawn randomly, and aim at refining it throughout training in order to maximize some performance metric.

There are however some trends that can be highlighted within these methods. Some of them aim at *jointly* train a reconstructor and a sampling pattern for some fixed

sampling budget. Others tackle a problem of designing a mask *sequentially*, by adding new sampling locations based on the currently acquired information. Finally, some methods aim at training a model that is able to propose a *patient-adaptive* mask design. These methods often leverage tools from reinforcement learning (RL) and this is why we refer to such models as *policy* models.

The works of Bahadir et al. (2019); Aggarwal and Jacob (2020); Weiss et al. (2020); Huijben et al. (2020) relax the problem of optimizing the mask to a continuous optimization problem, and train the sampling pattern along with the neural network using backpropagation. After training, the mask is then mapped back as a boolean vector. Other approaches use a pre-trained reconstruction model and train a policy model that performs patient-adaptive sequential mask optimization using reinforcement learning (Pineda et al., 2020; Bakker et al., 2020). The approach of Zhang et al. (2019b) is unique, as they aim at training an evaluator that tries to estimate the current reconstruction error for each location in k-space, while using this evaluator to inform the training of the reconstruction model. Finally, a fourth category of approaches tackles the challenging problem jointly learning a reconstruction model and a policy model for patient-adaptive sampling (Jin et al., 2019; Van Gorp et al., 2021; Yin et al., 2021).

Remark (A short discussion of non-Cartesian sampling methods). The problem of optimizing sampling trajectories becomes more complex when moving to the non-Cartesian realm. The problem is inherently continuous, and there is a lot of freedom in designing non-Cartesian trajectories as long as hardware constraints such as bounded first and second derivatives of the trajectory (Boyer et al., 2016).

As a result, there is great diversity in the methods proposed. Traditional approaches such as radial spokes (Lauterbur, 1973) and spirals (Meyer et al., 1992) are typically *model-based*, *model-driven*, but more recent space filling trajectories like SPARKLING (Lazarus et al., 2019) fall under this category. Weiss et al. (2019); Chaithya et al. (2022) attempted to train non-Cartesian sampling and reconstruction jointly, in an end-to-end fashion, yielding *learning-based*, *data-driven* approaches, while Wang et al. (2021) proposed to parametrize the trajectories with B-spline to facilitate learning, resulting in a *model-based data-driven* method. Finally, (Chaithya et al., 2021) also proposed to refine the SPARKLING algorithm by learning from data the underlying variable-density distribution, yielding a *learning-based*, *model-driven* approach.

These approaches were compared in Chaithya and Ciuciu (2022), where additional challenges due to non-Cartesian sampling are highlighted. In this setting, it appears that the algorithm used to optimize the sampling trajectory greatly matters, as it is possible for trajectories to remain stuck close to initialization. As a result, merely being a learning-based, data-driven approach does not guarantee an improved performance over model-based model-driven alternatives. Nonetheless, the best

performing method was found to be the hybrid learning approach of Chaithya et al. (2022).

4.6 Formalizing learning-based, data-driven mask design

Recall the observation setting described in Section 2.2.1. We consider the problem of accelerated MRI, where we obtain partial information $\mathbf{y}_\omega \in \mathbb{C}^P$ about a ground truth signal $\mathbf{x} \in \mathbb{C}^P$

$$\mathbf{y}_\omega = \mathbf{A}_\omega \mathbf{x} + \boldsymbol{\epsilon} = \mathbf{P}_\omega \mathbf{F} \mathbf{x} + \boldsymbol{\epsilon}$$

where $\boldsymbol{\epsilon}$ a complex Gaussian noise vector. \mathbf{F} is the Fourier transform operator, and \mathbf{P}_ω is our masking operator, implemented as a diagonal matrix with diagonal entries $\mathbf{P}_{\omega,ii} = 1$ if $i \in \omega$ and 0 otherwise. ω denotes the sampling mask and $|\omega| = N$ denotes its cardinality.

As described in Chapter 3, we construct an estimate $\hat{\mathbf{x}}$ of the ground truth \mathbf{x} using a reconstruction algorithm f_θ such that $\hat{\mathbf{x}} = f_\theta(\mathbf{y}, \omega)$, where θ denote the parameters of the reconstruction algorithm. This can correspond to the regularization parameter of a compressed sensing method, or denote all the parameters of a neural network.

A learning-based approach requires access to a data distribution $p(\mathbf{x})$ that can give us access to reference images, while a data-driven approach requires a way to quantify how close an estimation $\hat{\mathbf{x}}$ is from the ground truth \mathbf{x} . This is done by using a performance metric $\eta(\cdot, \cdot) : \mathbb{C}^P \times \mathbb{C}^P \rightarrow \mathbb{R}$, which we will aim to maximize². The process is illustrated on Figure 4.1.

An ideal sampling algorithm would tailor the mask to each instance of $\mathbf{x} \sim p(\mathbf{x})$, solving

$$\max_{\omega: |\omega| \leq N} \eta(\mathbf{x}, \hat{\mathbf{x}}_\theta(\mathbf{y}_\omega = \mathbf{P}_\omega \mathbf{A} \mathbf{x})), \quad (4.4)$$

However, it is evident that such an approach is infeasible in practice, because solving Equation 4.4 would require access to the ground truth image \mathbf{x} at the time of evaluation. This is not the case in reality, where we only have access to fully sampled training signals for training. We turn now to two main approaches to this problem that have been discussed before. They rely on either **i)** using *non-adaptive* masks, designed on a training set offline and deployed at testing time or **ii)** using *adaptive* masks, where a sampling *policy* or *heuristic* π_ϕ is trained offline and then takes patient-adaptive decision during the acquisition at evaluation time.

²A similar reasoning would hold if one would consider a loss $\ell(\cdot, \cdot) : \mathbb{C}^P \times \mathbb{C}^P \rightarrow \mathbb{R}$ instead of a performance metric. In this case, one would then need to find the sampling mask that minimizes it.

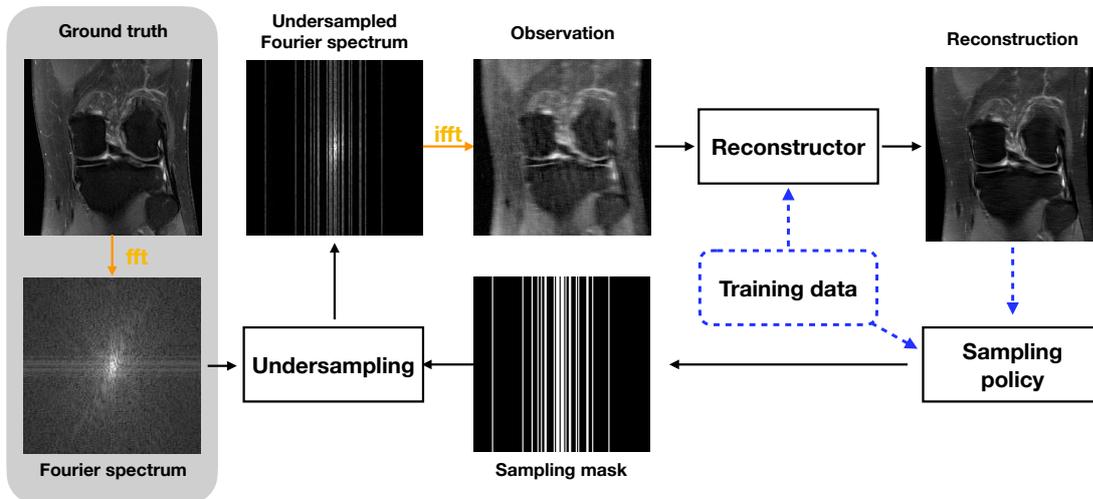


Figure 4.1 – Overview of sampling in accelerated MRI, using a knee image. Acquisition physically happens sequentially in Fourier space (full readout lines acquired at once), and a sampling policy can decide ahead of time what sampling mask to use (non-adaptive sampling), or integrate information from the current reconstruction to decide on the next location to acquire (adaptive sampling).

Blue dashed lines indicate optional relations: not all policies and reconstruction methods rely on training data, and not all policies are patient-adaptive. The light grayed area indicates that we do not have access to the full ground truth and spectrum at evaluation time, but that we rather directly acquire undersampled data.

4.6.1 Non-adaptive sampling masks

In this first setting, given access to a set of training data $\{\mathbf{x}_1, \dots, \mathbf{x}_m\} \sim p(\mathbf{x})$, we aim at finding the mask that exhibits the best average performance, namely

$$\max_{\omega: |\omega| \leq N} \frac{1}{m} \sum_{i=1}^m \eta(\mathbf{x}_i, \mathbf{f}_\theta(\mathbf{y}_{\omega,i}; \omega)) \text{ s.t. } \mathbf{y}_{\omega,i} = \mathbf{A}_\omega \mathbf{x}_i + \epsilon. \tag{4.5}$$

However, solving Equation 4.5 remains challenging, due to the combinatorial nature of the problem. It also brings questions of *generalization*, of whether a mask ω that shows good performance on the training set would also show good performance on a testing set of unseen samples from the same distribution $p(\mathbf{x})$. While standard learning-theoretic can give answer on the generalization problem (Gözcü et al., 2018), the combinatorial nature of the problem dooms to failure any naive approach, as the total number of different masks grows exponentially, with a complexity $O(2^N)$.

Two main types of approaches have been proposed for Problem 4.5, relying either on tackling the combinatorial problem in an approximate fashion (Gözcü et al., 2018; Zibetti et al., 2020) or relaxing the mask to a continuous variable that is then solved using gradient-based optimization (Bahadir et al., 2019; Aggarwal and Jacob, 2020; Weiss et al., 2020; Huijben et al., 2020; Sherry et al., 2019). Combinatorial approaches typically

4.6. Formalizing learning-based, data-driven mask design

tackle Problem 4.5 as a *sequential* problem, starting from some initial mask ω_0 , and gradually growing the mask until the sampling budget N is exhausted. On the contrary, *continuous* approaches choose to directly optimize the mask at the maximal sampling budget N , a setting which we refer as *fixed* sampling.

4.6.2 Adaptive sampling masks

In contrast to Equation 4.5, adaptive sampling generally relies on a two-step procedure, where one must first train a *policy* model π_ϕ , and subsequently evaluate it on testing data. The problem of training the policy reads

$$\max_{\phi} \frac{1}{m} \sum_{i=1}^m \eta(\mathbf{x}_i, \mathbf{f}_\theta(\mathbf{y}_{\omega_T, i})), \text{ where } \begin{cases} \omega_t = \omega_{t-1} \cup v_t \\ v_t \sim \pi_\phi(\mathbf{f}_\theta(\mathbf{y}_{\omega_{t-1}})) \\ |\omega_t| \leq N \end{cases} \quad t = 1, \dots, T. \quad (4.6)$$

Here, T describes the number of acquisition rounds that the policy does. While the problem may look similar to Problem 4.5, there are *two* fundamental differences to be noted. First, the optimization here is on the *parameters* ϕ of the policy model, rather than on the mask itself. Secondly, the problem is inherently sequential, as the mask is gradually built from atoms v_t , $t = 1, \dots, T$ that are all given from the policy π_ϕ at different stages in the acquisition.

Solving the problem remains challenging, as gradient-based optimization of ϕ remains challenging, due to the sampling operation $v_t \sim \pi_\phi(\mathbf{f}_\theta(\mathbf{y}_{\omega_{t-1}}))$ preventing a direct computation of the gradients. As a result, two main approaches can be found to tackle this problem, relying either again on relaxing the mask to a continuous variable (Van Gorp et al., 2021; Yin et al., 2021), or leveraging reinforcement learning (Bakker et al., 2020; Pineda et al., 2020; Jin et al., 2019), which is naturally suited to consider the problem of searching the policy that yields the best performance given a changing environment. In our case, the environment is described by the unknown underlying ground truth.

The second stage of adaptive sampling is the *inference* step, where the trained policy π_ϕ is then used to guide the acquisition of an unseen image \mathbf{x}_{test} :

$$v_{t+1} \in \arg \max_{v \in [P]} [\pi_\phi(\mathbf{f}_\theta(\mathbf{y}_{\omega_t, \text{test}}))]_v \quad (4.7)$$

where $[\cdot]_v$ denotes the v -th entry of the vector, $\omega_t = \omega_{t-1} \cup v_t$, for $t = 1, \dots, T$. At inference time, we see that the test image is gradually sampled at the locations v_t provided by the policy and based on the reconstruction at the previous step.

Remark (Using heuristics). The inference of Equation 4.7 is not restricted to policies trained following Equation 4.6, but can be paired with different heuristics.

For instance, Zhang et al. (2019b) trained an *evaluator* that estimates the error made by the reconstruction model at the different locations to be acquired. They then use this heuristic for sampling, by acquiring at each step the location that is estimated to have the highest reconstruction error.

We refer to this approach as a *heuristic* rather than a *policy* because it is not specifically trained to estimate what would be the best location to answer next; it only trains a model that estimates the current reconstruction error, and then uses it in a greedy fashion.

We will expand upon the use of heuristics for adaptive sampling in Chapter 7, where we propose to use the variance of a generative adversarial network (Goodfellow et al., 2014) to perform adaptive sampling.

Remark (Non-sequential adaptive sampling). Although a sequential acquisition is a natural choice for adaptive sampling, it is possible to perform patient-adaptive sampling with a fixed (non-sequential) mask.

This was done for instance by Bakker et al. (2021), where in the case of multi-coil MRI, the authors used the auto-calibration signal (ACS), a small set of measurements that is systematically acquired, to design a patient-adaptive fixed mask in a one-shot fashion. They inputted the measurements into a policy model that outputs a single mask for the rest of the acquisition as a result. However, their results surprisingly show that their best performing sampling policies explicitly learn to be non-adaptive, and their results provide then a state-of-the-art *non-adaptive* sampling mask.

4.6.3 On the optimality of the discrete mask optimization problem

In the previous sections, we defined Problems 4.4 and 4.5 as optimization problems over masks. However, compressed sensing approaches Puy et al. (2011); Chauffert et al. (2013) rather optimized the probability *distribution* from which masks are subsequently drawn. These arguments were also constructed in the context of model-based, model-driven methods, aiming at minimizing coherence. Can we relate the optimization of a probability distribution to the optimization of a discrete mask such as Problem 4.5? How does the problem of finding an optimal probability distribution change when moving from a model-based, model-driven context to a learning-based, data-driven context? These are the questions that we will answer in this section.

While in Chauffert et al. (2013), the optimization was done on the distribution π from which masks are subsequently sampled, in our case, the performance metric will be optimized directly for a fixed mask rather than a distribution. We will formally show below that in a data-driven context, we can construct an optimal sampling distribution π

4.6. Formalizing learning-based, data-driven mask design

from an optimal mask. The optimal mask describes the support of an optimal probability distribution.

We model the mask designing process as finding a probability mass function (PMF) $\pi \in S^{P-1}$, where $S^{P-1} := \{\pi \in [0, 1]^P : \sum_{i=1}^P \pi_i = 1\}$ is the standard simplex in \mathbb{R}^P . π assigns to each location i in the k -space a probability π_i to be acquired. The mask is then constructed by drawing without replacement from π until the cardinality constraint $|\omega| = N$ is met. We denote this dependency as $\omega(\pi, N)$. The problem of finding the optimal sampling distribution is subsequently formulated as

$$\max_{\pi \in S^{P-1}} \eta(\pi), \quad \eta(\pi) := \mathbb{E}_{\omega(\pi, N)} [\eta(\mathbf{x}, \mathbf{f}_\theta(\mathbf{y}, \omega))], \quad (4.8)$$

where the index set $\omega \subset [P]$ is generated from π and $[P] := \{1, \dots, P\}$. This problem corresponds to finding the probability distribution π that maximizes the expected performance metric with respect to the data $p(\mathbf{x})$ and the masks drawn from this distribution. To ease the notation, we will use $\eta(\mathbf{x}, \mathbf{f}_\theta(\mathbf{y}, \omega)) \equiv \eta(\mathbf{x}; \omega)$.

In practice, we do not have access to $\mathbb{E}_{p(\mathbf{x})} [\eta(\mathbf{x}; \omega)]$ and instead have at hand the training images $\{\mathbf{x}_i\}_{i=1}^m$ drawn independently from $p(\mathbf{x})$. We therefore maximize the *empirical* performance by solving

$$\max_{\pi \in S^{P-1}} \eta_m(\pi), \quad \eta_m(\pi) := \frac{1}{m} \sum_{i=1}^m \mathbb{E}_{\omega(\pi, N)} [\eta(\omega, \mathbf{x}_i)]. \quad (4.9)$$

Note that Problem 4.9, while being an empirical maximization problem, is distinct from Problem 4.5 above, as it seeks to optimize the distribution from which masks are drawn, and as a result, performs an expectation over all the masks drawn from it.

Given that Problem 4.9 looks for masks that are constructed by sampling N times without replacement from π , the following holds.

Proposition 1. *There exists a maximizer of Problem 4.9 that is supported on an index set of size at most N .*

Proof. Let the distribution $\hat{\pi}_n$ be a maximizer of Problem 4.9. We are interested in finding the support of $\hat{\pi}_n$. Because $\sum_{|\omega|=N} \Pr[\omega] = 1$, note that

$$\begin{aligned} \max_{\pi \in S^{P-1}} \eta_m(\pi) &:= \max_{\pi \in S^{P-1}} \sum_{|\omega|=N} \frac{1}{m} \sum_{i=1}^m \eta(\mathbf{x}_i; \omega) \cdot \Pr[\omega|\pi] \\ &\leq \max_{\pi \in S^{P-1}} \max_{|\omega|=N} \frac{1}{m} \sum_{i=1}^m \eta(\mathbf{x}_i; \omega) \\ &= \max_{|\omega|=N} \frac{1}{m} \sum_{i=1}^m \eta(\mathbf{x}_i; \omega). \end{aligned}$$

Let $\hat{\omega}_N$ be an index set of size N that maximizes the last line above. The above

Chapter 4. Optimizing sampling patterns

holds with equality when $\Pr[\hat{\omega}_N] = 1$ and $\Pr[\omega] = 0$ for $\omega \neq \hat{\omega}_N$ and $\pi = \hat{\pi}_N$. This in turn happens when $\hat{\pi}_N$ is supported on $\hat{\omega}$. That is, there exists a maximizer of Problem 4.9 that is supported on an index set of size N . \square

While this observation does not indicate how to find this maximizer, it nonetheless allows us to simplify Problem 4.9. More specifically, the observation that a distribution $\hat{\pi}_N$ has a compact support of size N implies the following:

Proposition 2.

$$\text{Problem 4.9} \equiv \max_{|\omega|=N} \frac{1}{m} \sum_{i=1}^m \eta(\mathbf{x}_i; \omega) \quad (4.10)$$

Proof. Proposition 1 tells us that a solution of Problem 4.9 is supported on a set of size at most n , which implies

$$\text{Problem 4.9} \equiv \max_{\pi \in S^{P-1}; |\text{supp}(\pi)|=N} \eta_m(\pi). \quad (4.11)$$

That is, we only need to search over compactly supported distributions π . Let S_Γ denote the standard simplex on a support $\Gamma \subset [P]$. It holds that

$$\begin{aligned} \text{Problem 4.11} &\equiv \max_{|\Gamma|=N} \max_{\pi \in S_\Gamma} \eta_m(\pi) \\ &= \max_{|\Gamma|=N} \max_{\pi \in S_\Gamma} \frac{1}{m} \sum_{i=1}^m \eta(\mathbf{x}_i; \Gamma) \cdot \Pr[\Gamma|\pi] \\ &= \max_{|\Gamma|=N} \max_{\pi \in S_\Gamma} \frac{1}{m} \sum_{i=1}^m \eta(\mathbf{x}_i; \Gamma) \\ &= \max_{|\Gamma|=N} \frac{1}{m} \sum_{i=1}^m \eta(\mathbf{x}_i; \Gamma). \end{aligned} \quad (4.12)$$

To obtain the second and third equalities, one observes that all masks have a common support Γ with N elements, i.e. $\pi \in S_\Gamma$ allows only for a single mask ω with N elements, namely $\omega = \Gamma$. \square

The framework of Problem 4.9 captures most variable-density based approaches of the literature that are defined in a learning-based fashion (Knoll et al., 2011b; Ravishankar and Bresler, 2011a; Vellagoundar and Machireddy, 2015; Haldar and Kim, 2019; Bahadir et al., 2019; Sherry et al., 2019), and Proposition 2 shows that Problem 4.12, that we tackled in Gözcü et al. (2018); Gözcü et al. (2019); Sanchez et al. (2020a) and develop in Chapter 5, also aims at solving the *same* problem as these probabilistic approaches.

Note that while the present argument considers sampling *points* in the Fourier space, it is readily applicable to the Cartesian case, where full lines are added to the mask at once.

4.6. Formalizing learning-based, data-driven mask design

Several additional observations and remarks can be made. The argument presented here is thus mainly a justification of the objective of Equation 4.5 that we use for a greedy mask optimization: we do not aim at solving a fundamentally different problem than the one of finding the optimal sampling density in a data-driven approach when trying to find the best performing mask. Indeed, Proposition 2 shows that we do not need to find a density, but that finding the support of cardinality N of an optimal distribution is sufficient. And proposition 1 shows the existence of an optimal distribution with such a support.

The argument does not formally prove the suboptimality of variable-density sampling, although one can argue for this from the intuition that these results bring. Indeed, any distribution not compactly supported on a set of N optimal location will open the possibility to a suboptimal mask being sampled by doing sampling without replacement. The heuristic of picking sampling locations at random without replacement could be a cause of the worse practical performance of VDS compared to LBCS, as we will see in Section 5.4 of next Chapter.

Finally, these propositions do not prescribe any particular algorithm that should be used to solve Problem 4.5. Rather, they are limited to establishing an equivalence between *problems*, and do not speak either about the expected *generalization* of these solutions. Generalization is however addressed by standard learning theoretic arguments (Gözcü et al., 2018).

Bibliographic note

Propositions 1 and 2 are due to Armin Eftekhari.

5 Scalable learning-based sampling optimization

In this chapter, we present our extensions to the framework of Gözcü et al. (2018). We primarily aim at drastically improving the scalability of their method. We begin by presenting in detail their learning-based CS (LBCS) framework in Section 5.1. In Section 5.2, we take a short detour in the field of submodular function maximization, from which Gözcü et al. (2018) took inspiration, and discuss how algorithms that have been proposed in this context could help improve the LBCS method. In Section 5.3, we present our two improvements over LBCS, namely *lazy* LBCS and *stochastic* LBCS. We then carry out extensive validation of these methods in Sections 5.4 and 5.5, and discuss the results in Section 5.6¹.

5.1 Learning-based Compressive Sensing (LBCS)

LBCS aims at designing a non-adaptive mask in a sequential fashion, and tackles Problem 4.5. Indeed, solving this problem in a naive, brute-force approaches would only be possible in very simple cases, as the total number of masks grows exponentially, with a rate $O(2^P)$. Therefore, one must use *approximate* solutions to this problem. Gözcü et al. (2018) proposed a greedy, parameter-free approach, where a mask is grown sequentially by adding elements to the mask in an iterative fashion. This approach, called Learning-Based Compressed Sensing (LBCS) allows to reduce the number of configurations evaluated to a complexity $O(P^2)$. The simplified procedure is illustrated on Figure 5.1.

A couple of elements must be introduced in order to formally define their algorithm. To implement the constraints of sampling in MRI, typically acquiring full lines in the

¹The work in this chapter is based on the following publications:

Gözcü, B., Sanchez, T., and Cevher, V. (2019). Rethinking sampling in parallel MRI: A data-driven approach. In *27th European Signal Processing Conference (EUSIPCO)*.

Sanchez, T., Gözcü, B., van Heeswijk, R. B., Eftekhari, A., Ilıcak, E., Çukur, T., and Cevher, V. (2020a). Scalable learning-based sampling optimization for compressive dynamic MRI. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8584–8588.

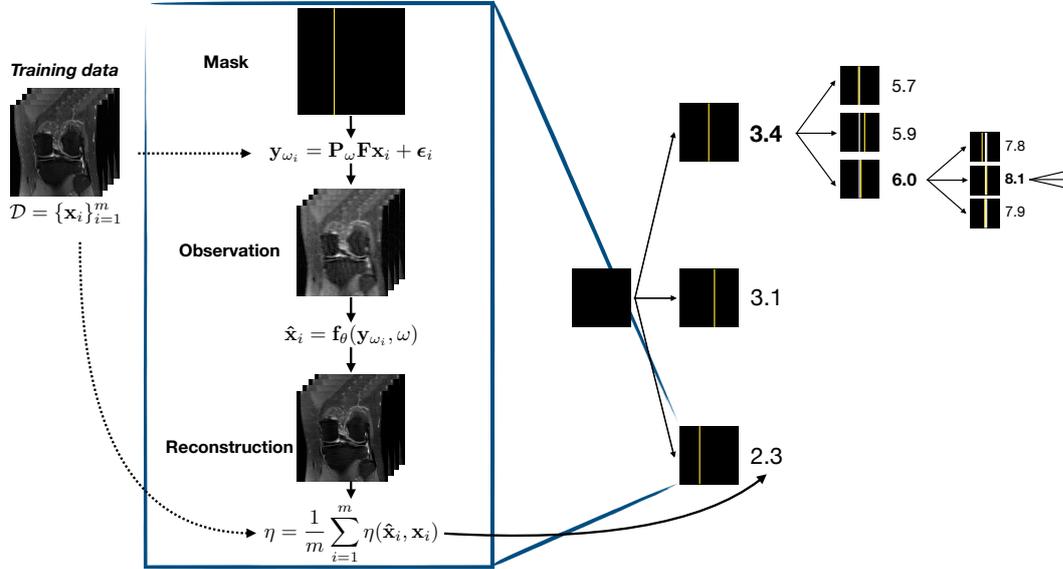


Figure 5.1 – Illustration of the LBCS greedy approach. On the right, we start from an empty mask, and select the line that yields to the best performance. Once it is acquired, we keep this mask, try adding to it the line that yields the best performance given what has been acquired, and continue until the final sampling budget is reached. On the left, we see how the performance is assessed: we use a set \mathcal{D} of training data, that are subsampled according to the current mask, then reconstructed and evaluated using the performance metric η . The computation is repeated for each different mask.

Fourier space at once, we define a *set of subsets* \mathcal{S} of $\{1, \dots, P\}$ and assume that the final mask will be constructed as

$$\omega = \bigcup_{j=1}^t S_j, \quad \text{s.t. } S_j \in \mathcal{S} \quad (5.1)$$

for some $t > 0$. Gözcü et al. (2018) also introduce a *cost function* $c(\omega) \geq 0$ and a *maximal cost* Γ that the final mask ω must not exceed. In the case of Problem 4.5, we have $c(\omega) = |\omega|$ and $\Gamma = N$. With these two additions, the LBCS method (Gözcü et al., 2018) is displayed in Algorithm 1. The procedure is very simple: while budget is available (1.2), the algorithm tests all feasible candidates $S \in \mathcal{S}$ (1.3-7) and adds permanently to the sampling mask the candidate that yields the largest performance gain (1.8).

The greedy approach of LBCS brings multiple advantages over previous heuristics such as variable-density sampling. First, the mask ω has a *nested* structure: by recording the order in which elements were added to the mask, it is possible to immediately adapt to different costs Γ . This feature is not possible for many approaches that directly optimize the sampling mask for a given sampling rate. In addition, this approach does not have any parameter that requires to be tuned, making it particular easy to implement in

Algorithm 1 LBCS

Input: Training data $\mathbf{x}_1, \dots, \mathbf{x}_m$, decoder \mathbf{f}_θ ,
sampling subsets \mathcal{S} , cost function c , maximum cost Γ
Output: Sampling pattern ω

- 1: $\omega \leftarrow \emptyset$
- 2: **while** $c(\omega) \leq \Gamma$ **do**
- 3: **for** $S \in \mathcal{S}$ such that $c(\omega \cup S) \leq \Gamma$ **do**
- 4: $\omega' = \omega \cup S$
- 5: For each j , set $\mathbf{y}_j \leftarrow \mathbf{P}_{\omega'} \mathbf{F} \mathbf{x}_j$, $\hat{\mathbf{x}}_j \leftarrow \mathbf{f}_\theta(\mathbf{y}_j, \omega')$
- 6: $\eta(\omega') \leftarrow \frac{1}{m} \sum_{j=1}^m \eta(\mathbf{x}_j, \hat{\mathbf{x}}_j)$
- 7: **end for**
- 8: $\omega \leftarrow \omega \cup S^*$, $S^* = \arg \max_{S: c(\omega \cup S) \leq \Gamma} \eta(\omega \cup S) - \eta(\omega)$
- 9: **end while**
- 10: **return** ω

practice. It also does not rely on any probability distribution from which candidate masks are drawn, which makes it fully learning-based. Finally, the algorithm is not tied to a specific reconstruction method or performance metric, which makes it easy to integrate in novel approaches.

5.2 LBCS motivation: Submodularity

The idea to apply a greedy approach to the problem of mask design was initially motivated by techniques used in the field of submodular function maximization (Baldassarre et al., 2016; Krause and Golovin, 2014). We discuss in greater depth these concepts here as the motivation for scaling up LBCS are also motivated by algorithms proposed in the context of submodular function maximization.

Submodular functions have been studied for a long time (Nemhauser et al., 1978; Minoux, 1978), and formalize among other things the idea of *diminishing returns*. This is the idea that adding an element to a smaller set increases the objective more than when it is added to a larger set. Formally

Definition (Submodularity). A set function $\eta(\omega)$ mapping subsets $\omega \in [P]$ to real numbers is said to be **submodular** if, for all $\omega_1, \omega_2 \subset [P]$ with $\omega_1 \subseteq \omega_2$ and all $i \in [P]$, it holds that

$$\eta(\omega_1 \cup \{i\}) - \eta(\omega_1) \geq \eta(\omega_2 \cup \{i\}) - \eta(\omega_2). \quad (5.2)$$

The function is said to be **modular** if the same holds true with equality in place of inequality

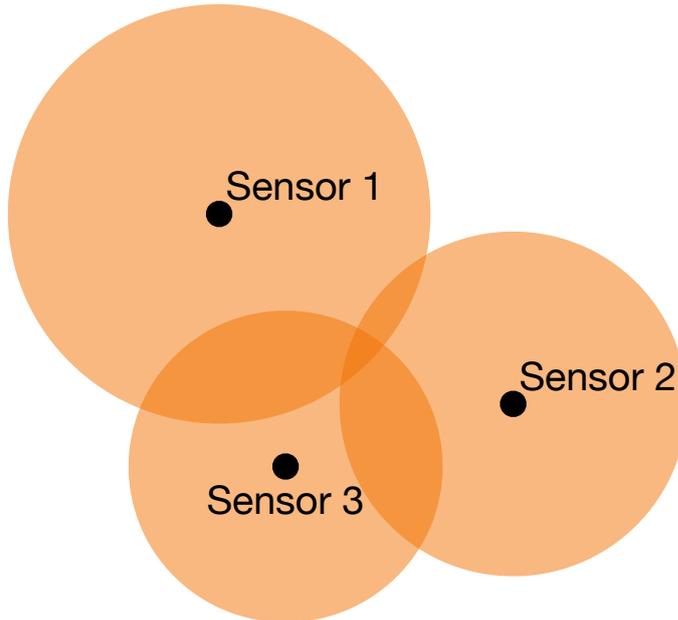
Chapter 5. Scalable learning-based sampling optimization

An example of this is sensor placement. As one places more sensors to cover an area more fully, often, the area covered by sensors will be overlapping, and so adding a new sensor will bring less information if a large set of sensors is already available. This is illustrated on Figure 5.2.

This typically translates into problems of submodular function maximization, where a common constraint lies on the cardinality of the solution:

$$\max_{\omega \subseteq [P]} \eta(\omega) \text{ subject to } |\omega| \leq N. \quad (5.3)$$

In the case of linear reconstruction in MRI, it can be shown that the problem of finding the sampling mask that gives minimal mean squared error is a *modular* optimization problem: there is no diminishing returns, but each component can be optimized individually (Baldassarre et al., 2016). In the case of a non-linear reconstruction in MRI, there is no formal proof that shows that submodularity exactly holds, but empirical results show a diminishing return throughout acquisition. This is what motivated Gözcü et al. (2018) to leverage tools from submodular optimization for mask design in MRI.



$$\eta(\{1\} \cup \{3\}) - \eta(\{1\}) \geq \eta(\{1, 2\} \cup \{3\}) - \eta(\{1, 2\})$$

Figure 5.2 – Illustration of the diminishing return property of submodular functions. If h quantifies the total coverage of sensors, we see that adding sensor 3 to sensor 1 increases the total coverage more than when adding it to the set $\{1, 2\}$, as there is more surface coverage surface overlapping in this case.

Indeed, greedy algorithms have a special place in submodular optimization because of the celebrated result by Nemhauser et al. (1978), who proved that the greedy algorithm

5.2. LBCS motivation: Submodularity

provides a good approximation to the optimal solution of NP-hard instances of submodular function maximization (Krause and Golovin, 2014).

Theorem 1 (Nemhauser et al. (1978)). *Let $\eta : 2^P \rightarrow \mathbb{R}_+$ be a nonnegative monotone submodular function. Let $\omega_t, t \geq 0$ be the sequence obtained by the greedy algorithm. Then, for all positive t, N ,*

$$\eta(\omega_t) \geq \left(1 - e^{-t/N}\right) \max_{\omega:|\omega|\leq N} \eta(\omega). \quad (5.4)$$

In particular, for $t = N$, $\eta(\omega_N) \geq (1 - 1/e) \max_{\omega:|\omega|\leq N} \eta(\omega)$.

For many classes of submodular functions h , this result is the best that can be achieved with any efficient algorithm (Krause and Golovin, 2014). In particular, Nemhauser and Wolsey (1978) proved that any algorithm that can only evaluate h on a polynomial number of set will not be able to obtain an approximation guarantee that is better than $(1 - 1/e)$. This property could be key to explaining the strong empirical performance of the LBCS algorithm of Gözcü et al. (2018) in the context of MRI.

Although the approximation guarantee cannot be generally improved, many works have proposed more *efficient* approaches than the simple greedy algorithm. Two works are of particular interest to us, as they serve as the motivation to the algorithms that we will describe in the next section. These algorithms are the *lazy greedy* (Minoux, 1978) and *stochastic greedy*, called *lazier* than lazy greedy (Mirzasoleiman et al., 2015).

The lazy greedy algorithm relies on the fact that if at the i -th iteration of the greedy algorithm, an element S is expected to bring a marginal benefit $\Delta(S|\omega_i) = \eta(\omega_i \cup S) - \eta(\omega_i)$, then it holds that $\Delta(S|\omega_j) \leq \Delta(S|\omega_i)$ for all $j \geq i$ for a submodular function η . Exploiting this fact, one can keep a list of upper bounds on the marginal benefits of each element S , called $\rho(S)$, initialized at $+\infty$. Then, at each iteration, the element S with largest value $\rho(S)$ is picked and updated as $\rho(S) = \Delta(S|\omega)$. If $\rho(S) \geq \rho(S') \forall S' \in \mathcal{S}$, the marginal benefit of this element S is larger than the upper bounds of each other marginal contribution, and is consequently added permanently to ω . This can speed up the algorithm by several orders of magnitude, since $\rho(S) \geq \rho(S')$ can be satisfied after a few trials, instead of trying each available S at every iteration (Krause and Golovin, 2014).

The stochastic greedy algorithm (Mirzasoleiman et al., 2015) samples at each step of the greedy algorithm a subset of candidates $\mathcal{S}_{iter} \subseteq \mathcal{S}$, with $|\mathcal{S}_{iter}| = k$ at line 3 of Algorithm 1. As a result, given a nonnegative monotone submodular function $\eta : 2^P \rightarrow \mathbb{R}_+$ and taking $k = \frac{P}{N} \log(1/\epsilon)$, the stochastic greedy algorithm achieves a $(1 - 1/e - \epsilon)$ approximation guarantee in expectation to the optimum, using only $O(P \log(1/\epsilon))$ evaluations.

Both of these algorithms inspired the methods that we proposed to scale up LBCS, and their implementation in the context of MRI will be respectively covered in Sections 5.3.3

and 5.3.2.

Remark (Submodular optimization under noise). Although the cases so far mostly discussed optimizing submodular functions without measurement noise, such a setting is not realistic in the case of prospective MRI acquisitions, where repeated lines will have a different noise, and resampling can improve the average signal. Nonetheless, submodular optimization under noise has been studied, and algorithms that adapt to such settings have been proposed (Singla et al., 2016; Hassidim and Singer, 2017).

5.3 Scaling up LBCS

5.3.1 Limitations of LBCS

As we have hinted above, in many practical scenarios, evaluating $O(P^2)$ configurations remains very expensive. This is also due to the fact that each configuration requires reconstructing m data points. In particular, the LBCS approach becomes prohibitively expensive when either of the following occurs:

1. When the number of candidate locations grows significantly. This is particularly the case, in 3D MRI, where one needs to undersample entire volumes, or dynamic MRI, where one undersamples multiple frames for a single 2D image.
2. When the size of the dataset grows, it becomes impractical to evaluate LBCS on the whole testing set at once.
3. When the running time of the reconstruction method becomes large. This is particularly the case for multi-coil MRI, which is the setting most commonly used in practice. This issue also occurs when dealing with larger data, from 3D or dynamic MRI.

As a result, we wish to develop variants of LBCS that allow to tackle these different settings, while retaining the performance of LBCS as much as possible. We will discuss two variants of LBCS that were respectively proposed in Gözcü et al. (2019) and Sanchez et al. (2020a), and that aim at tackling mostly the issues of multi-coil and 3D MRI as well as dynamic MRI. We start by explaining how they differ from LBCS.

5.3.2 Stochastic LBCS

With stochastic LBCS (sLBCS), we aim at addressing two fundamental drawbacks of the approach of Gözcü et al. (2018), namely that **1.** its scaling is quadratic with respect to the total number of readout lines, **2.** its scaling is linearly with respect to the size of the dataset. With this approach, we also aim at integrating a specific constraint of

dynamic MRI (dMRI), namely that the sampling mask should have a fixed number of readouts by temporal frame.

Scalability with respect to the size of the dataset is essential for the applicability of sLBCS to modern datasets, which can contain several thousands of images, making it impractical to go through the entire dataset at each inner loop of the algorithm.

Before describing the algorithm, let us briefly recall the dynamic MRI setting. The main difference with static MRI is that the signal will be a vectorized video, i.e. $\mathbf{x} \in \mathbb{C}^P$, where $P = H \times W \times T$ instead of $H \times W$. As a result, the operator \mathbf{A}_ω will be constructed differently, and in particular, the Fourier transform will be applied frame-wise. In addition, the mask ω will be constructed as a union of elements from a set \mathcal{S} that will represent lines at different time frames, and will be described as $\mathcal{S} = \mathcal{S}_1 \cup \dots \cup \mathcal{S}_T$. We also define the set $\mathcal{L} \subseteq \{1, \dots, m\}$ as a set of indices where the training data $\{\mathbf{x}\}_{i=1}^m$ will be subsampled.

Let us now detail *how* our proposed stochastic greedy method solves the limitations **1.** and **2.** of LBCS (Algo. 1). The issue **1.** is solved by picking uniformly at random at each iteration a batch possible readout lines \mathcal{S}_{iter} of constant size k from a given frame \mathcal{S}_t , instead of considering the full set of possible lines \mathcal{S} (line 3 in Alg. 2), and is inspired from the stochastic greedy of Mirzsoleiman et al. (2015); the number of candidate locations is further reduced by iterating through the lines to be added from each frame \mathcal{S}_t sequentially (lines 1, 3 and 11 in Algorithm 2); **2.** is addressed by considering a fixed batch of training data \mathcal{L} of size l instead of the whole training set of size m at each iteration (line 4 in Algorithm 2);

Remark. We will discuss the impact of introducing the batch size k for the readout lines and the subset \mathcal{L} to subsample training data in the results. However, we want to mention here that one has to be careful to pick the subset \mathcal{L} before the `for` loop of line 5 in Algorithm 2. Indeed, failing to do this would result in the different candidates S would be evaluated on different subsets of the data. As there is some variation within the data, for instance in terms of SNR, such an approach would yield performances that are *not* comparable among different candidates. In turn, this would skew the results towards picking the candidate location S with an associated subset \mathcal{L}_S with high signal-to-noise ratio, at the expense of selecting a readout location for the performance improvement that it actually brings. This would be especially the case if $|\mathcal{L}|$ were small. Picking \mathcal{L} before this `for` loop ensures that the readout line selected will indeed be the one that brings the largest performance gain among the candidates for the subset of data \mathcal{L} .

The improvements of sLBCS allow to reduce the computational complexity from $\Theta(mr(HT)^2)$ to $\Theta(lrkHT)$, effectively speeding up the computation by a factor $\Theta(\frac{m}{l} \frac{HT}{k})$, where r denotes the cost of a single reconstruction. Our results show that this is achieved

without sacrificing any reconstruction quality. We will see that the improvement is particularly significant for large datasets, as the batch size l will typically remain small, typically up to 100 samples, whereas MRI datasets can contain more than 15 000 data points (Zbontar et al., 2019).

Algorithm 2 sLBCS: Greedy mask optimization algorithm for (d)MRI

Input: Training data $\{\mathbf{x}\}_{i=1}^m$, recon. rule g , sampling set \mathcal{S} , max. cardinality n , samp. batch size k , train. batch size l

Output: Sampling pattern ω

```

1: (dMRI) Initialize  $t = 1$ 
2: while  $|\omega| \leq n$  do
3:   Pick  $\begin{cases} \mathcal{S}_{iter} \subseteq \mathcal{S} \text{ (MRI)} \\ \mathcal{S}_{iter} \subseteq \mathcal{S}_t \text{ (dMRI)} \end{cases}$  at random, with  $|\mathcal{S}_{iter}| = k$ 
4:   Pick  $\mathcal{L} \subseteq \{1, \dots, m\}$ , with  $|\mathcal{L}| = l$ 
5:   for  $S \in \mathcal{S}_{iter}$  such that  $|\omega \cup S| \leq \Gamma$  do
6:      $\omega' = \omega \cup S$ 
7:     For each  $\ell \in \mathcal{L}$  set  $\hat{\mathbf{x}}_\ell \leftarrow \mathbf{f}_\theta(\mathbf{P}_{\omega'} \mathbf{F} \mathbf{x}_\ell, \omega')$ 
8:      $\eta(\omega') \leftarrow \frac{1}{|\mathcal{L}|} \sum_{\ell \in \mathcal{L}} \eta(\mathbf{x}_\ell, \hat{\mathbf{x}}_\ell)$ 
9:   end for
10:   $\omega \leftarrow \omega \cup S^*$ , where  $S^* = \underset{\substack{S: |\omega \cup S| \leq n \\ S \in \mathcal{S}_{iter}}}{S}}{\arg \max} \eta(\omega \cup S)$ 
11:  (dMRI)  $t = (t \bmod T) + 1$ 
12: end while
13: return  $\omega$ 

```

5.3.3 Lazy LBCS

In cases where the cardinality of the candidate elements \mathcal{S} grows extremely large, even Algorithm 2 can become impractical. In 2D MRI, LBCS deals with images of size $H \times W$ while having a candidate set of size H , with H different lines. However, in the case of 3D MRI, the number of candidate locations explodes, as we move to images of size $P = H \times W \times S$, where S is the number of slices, with a candidate set of size $H \times W$, i.e., one is allowed to sample in both phase and frequency encoding directions. While in 2D dynamic MRI, the candidate set is of size $H \times T$ with T typically much smaller than H , in the case of 3D MRI, we often have $H \approx W$. In addition, we have less structure than in dMRI, as there is no option to cycle through frames. This implies that the batch of randomly sampled locations must be much larger to enable a good performance, which results in the cost remaining overall prohibitive. An additional reason for this increased cost is the reconstruction time of 3D methods. The number of slices can be much larger than the number of frames in dMRI, slowing down reconstruction considerably. Moving from single coil to multi-coil data further increases this cost, although this change does *not* increase the complexity of the sampling optimization problem.

Lazy LBCS proposes a solution to this limitation by leveraging another approach to speed up LBCS. It uses *lazy* evaluations (Minoux, 1978) which is equivalent to but faster than the greedy algorithm for submodular functions, and possess optimization guarantees for such functions (Nemhauser et al., 1978). While in our present setting, verifying submodularity is difficult because writing the performance of reconstructions in a closed form is generally not possible, we were motivated by the fact that Algorithm 3 works well in practice even if the objective function is not exactly submodular. Indeed, as we will observe in the experiments, lazy evaluations also perform well for mask optimization.

The implementation of lazy evaluations for LBCS is described in Algorithm 3. We see that a list of upper bounds is initially computed (l.2), and that we sequentially traverse the list of upper bounds until we find a candidate sample S that brings a larger improvement than any of the upper bounds in ρ (l.8-9).

Algorithm 3 Lazy Learning-based Compressive Sensing (ILBCS)

Input: Training data $\mathbf{x}_1, \dots, \mathbf{x}_m$, decoder \mathbf{f} , sampling subsets \mathcal{S} , cost function c , maximum cost Γ

Output: Sampling pattern ω

```

1:  $\omega \leftarrow \emptyset$ 
2:  $\rho(S) \leftarrow +\infty \forall S \in \mathcal{S} \text{ s.t. } c(\omega \cup S) \leq \Gamma$ 
3: while  $c(\omega) \leq \Gamma$  do
4:    $\omega' \leftarrow \omega \cup S$ , where  $S = \underset{S' \in \mathcal{S}: c(\omega \cup S') \leq \Gamma}{\arg \max} \rho(S')$ 
5:   For each  $j$ , set  $\mathbf{y}_j \leftarrow \mathbf{P}_{\omega'} \mathbf{F} \mathbf{x}_j$ ,  $\hat{\mathbf{x}}_j \leftarrow \mathbf{f}_{\theta}(\mathbf{y}_j, \omega')$ 
6:    $\eta(\omega') \leftarrow \frac{1}{m} \sum_{j=1}^m \eta(\mathbf{x}_j, \hat{\mathbf{x}}_j)$ 
7:    $\rho(S) \leftarrow \eta(\omega') - \eta(\omega)$ 
8:   if  $\rho(S) \geq \rho(S') \forall S' \in \mathcal{S} \text{ s.t. } c(\omega \cup S') \leq \Gamma$  then
9:      $\omega = \omega \cup S$ 
10:  end if
11: end while
12: return  $\omega$ 

```

5.4 Experiments - Stochastic LBCS

In this Section, we validate the performance of sLBCS, and will discuss ILBCS in Section 5.5. We carry out extensive experiments to validate our stochastic LBCS method. We first compare sLBCS against the vanilla LBCS methods of Gözcü et al. (2018), paying particular attention to how it compares to the original method in terms of performance and computational complexity. We study the influence on the batch sizes k and l on performance and on the design of the mask. Then, we carry out experiments on both single coil and multicoil data, as well as noise free and noisy data. Finally, we show the benefit of our method on large scale datasets, for which LBCS is prohibitively expensive.

5.4.1 Experimental setting

Dataset. Our experiments were carried out on three different datasets.

- *Cardiac dataset.* The data set was acquired in seven healthy adult volunteers with a balanced steady-state free precession (bSSFP) pulse sequence on a whole-body Siemens 3T scanner using a 34-element matrix coil array. Several short-axis cine images were acquired during a breath-hold scan. Fully sampled Cartesian data were acquired using a 256×256 grid, with relevant imaging parameters including $320 \text{ mm} \times 320 \text{ mm}$ field of view (FoV), 6 mm slice thickness, $1.37 \text{ mm} \times 1.37 \text{ mm}$ spatial resolution, 42.38 ms temporal resolution, 1.63/3.26 ms TE/TR, 36° flip angle, 1395 Hz/px readout bandwidth. There were 13 phase encodes acquired for a frame during one heartbeat, for a total of 25 frames after the scan.

The Cartesian cardiac scans were then combined to single coil data from the initial $256 \times 256 \times 25 \times 34$ size, using adaptive coil combination (Walsh et al., 2000; Griswold et al., 2002b). This single coil image was then cropped to a $152 \times 152 \times 17$ image. This was done because a large portion of the periphery of the images are static or void, and also to enable a greater computational efficiency. In the experiments, we used three volumes for training and four for testing.

- *Vocal dataset.* The vocal dataset that we used in the experiments comprised 4 vocaltract scans with a 2D HASTE sequence (T2 weighted single-shot turbo spin-echo) on a 3T Siemens Tim Trio using a 4-channel body matrix coil array. The study was approved by the local institutional review board, and informed consent was obtained from all subjects prior to imaging. Fully sampled Cartesian data were acquired using a 256×256 grid, with $256 \text{ mm} \times 256 \text{ mm}$ field of view (FoV), 5 mm slice thickness, $1 \text{ mm} \times 1 \text{ mm}$ spatial resolution, 98/1000 ms TE/TR, 150° flip angle, 391 Hz/px readout bandwidth, 5.44 ms echo spacing (256 turbo factor). There was a total of 10 frames acquired, which were recombined to single coil data using adaptive coil combination as well (Walsh et al., 2000; Griswold et al., 2002b).
- *fastMRI.* The fastMRI dataset was obtained from the NYU fastMRI initiative (Zbontar et al., 2019). The anonymized dataset comprises raw k-space data from more than 1,500 fully sampled knee MRIs obtained on 3 and 1.5 Tesla magnets. The dataset includes coronal proton density-weighted images with and without fat suppression.

Reconstruction algorithms. We consider two reconstruction algorithms, namely *k-t FOCUSS* (KTF) (Jung et al., 2009), and *ALOHA* (Jin et al., 2016). Their parameters were selected to maintain a good empirical performance across all sampling rates considered.

KTF aims at reconstructing a dynamic sequence by enforcing its sparsity in the x-f

domain. Formally, it aims at solving

$$\min_{\mathbf{z}} \|\mathbf{z}\|_1 \text{ subject to } \|\mathbf{y}_\omega - \mathbf{P}_\omega \mathbf{F} \mathbf{F}_t^{-1} \mathbf{z}\|_2 \leq \epsilon$$

where \mathbf{F}_t^{-1} denotes the inverse temporal Fourier transform. It tackles the problem by using a reweighted quadratic approach (Jung et al., 2007).

ALOHA (Jin et al., 2016) is an annihilating filter based low-rank Hankel matrix approach. ALOHA exploits the duality between sparse signals and their transform domain representations, cast as a low-rank Hankel matrix, and effectively transforms CS reconstruction into a low-rank matrix completion problem. We refer the reader to their detailed derivation for further details.

Mask selection baselines:

- *Coherence-VD* (Lustig et al., 2007): We consider a random *variable-density* sampling mask with Gaussian density and optimize its parameters to minimize coherence. This is an instance of a *model-based, model-driven* approach.
- *Golden* (Li et al., 2018): This approach designs a Cartesian mask by adding progressively lines according to the golden ratio. This is a *model-based, model-driven* approach.
- *LB-VD* (Gözcü et al., 2018; Gözcü et al., 2019): Instead of minimizing the coherence as in *Coherence-VD*, we perform a grid search on the parameters using the training set to optimize reconstruction according to the same performance metric as our method. This makes this approach *model-based, data-driven*.

We also consider several variants of LBCS and sLBCS that we will use throughout the experiments. We define abbreviations for conciseness:

- **G-v1** is the original LBCS algorithm.
- **G-v2** is a version of sLBCS that uses a batch of training data l and cycles through the frames.
- **SG-v1** is a version of the stochastic LBCS algorithm that considers a batch of sampling locations k but uses all training data, i.e. $l = |m|$.
- **SG-v2** is sLBCS, using a batch of sampling locations k , a batch of training data l and cycling through the frames.

Remark. In this work, we mainly focused on small scale datasets with a large set of candidate sampling locations. Our experiments therefore are carried out using SG-v1 and SG-v2. We did not observe any significant performance difference between these variants of sLBCS.

However, we note that G-v2 might be of interest in deep-learning based settings, where one deals with large datasets, such as fastMRI (Zbontar et al., 2019). This will be further discussed in Section 5.4.5.

5.4.2 Comparison of greedy algorithms

We first evaluate SG-v1 and SG-v2 against G-v1 to establish how sensitive our proposed method is to its parameters k and l . We carry out this experiment on the cardiac data, and report only the results on KTF, although similar trends can be observed with other reconstruction algorithms.

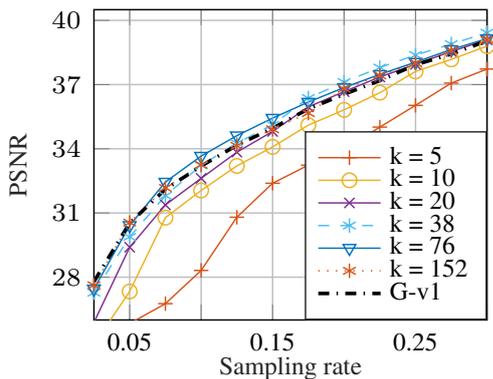


Figure 5.3 – PSNR as a function of the rate for KTF, comparing the effect of the batch size on the quality of the reconstruction for SG-v1. The result is averaged on 4 testing images of size $152 \times 152 \times 17$.

Influence of the batch size k on the mask design.

We first discuss the tuning of the batch size used in SG-v1, to specifically study the effect of different batch sizes. Figure 5.3 provides quantitative results of the performance, while Figure 5.4 shows the shapes of the masks obtained with various batch sizes. Unsurprisingly, small batch sizes yield poor results, as this algorithm has to choose from a very small set of randomly sampled candidates, being likely to miss out on important frequencies, such as the center ones. It is interesting to see however that the PSNR reaches the result from G-v1 with as few as 38 samples (out of the 152 candidates of a frame, and $152 \times 17 = 2584$ candidates overall).

In some cases, such as $k = 38$, SG-v1 can even yield an *improved* performance over G-v1, with 60 times less computation than G-v1. We hypothesize that it is due to the noise introduced by sLBCS allowing to avoid suboptimal configurations encountered by G-v1. Note that this might not be expected if the problem was to be exactly submodular, but this is *not* the case in practice.

Focusing on Figure 5.4, one can make several observations. First of all, as expected, taking a batch size of 1 yields a totally random mask, and taking a batch size of 5 yields a mask that is more centered towards low frequency than the one with $k = 1$ but it still

has a large variance. Then, as the batch size increases, resulting masks seem to converge to very similar designs, but those are slightly different from the ones obtained with G-v1. Overall, our initial experiments suggest, as a rule of thumb, that roughly sampling 25% of a frame enables sLBCS to match the performance of G-v1.

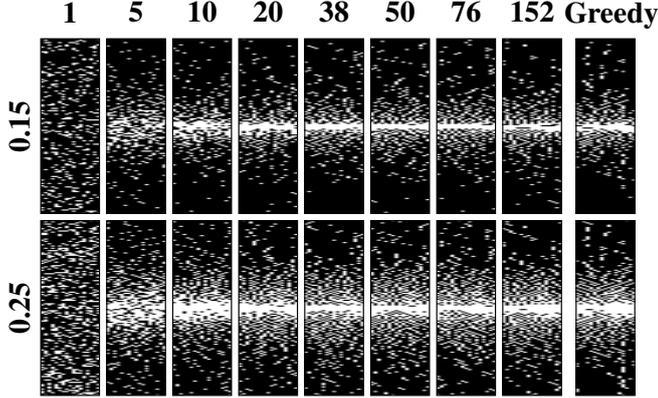


Figure 5.4 – Learning-based masks obtained with SG-v1 for different batch sizes k using KTF as a reconstruction algorithm, shown in the title of each column, for 15% and 25% sampling rate. The optimization used data of size $152 \times 152 \times 17$, with a total of 2584 possible phase encoding lines for the masks to pick from.

Influence of the batch size l . We then include the effect of using a smaller data batch size l into the equation, comparing the performances of G-v1 with SG-v1 and SG-v2. The results are presented on Figure 5.5. Both SG-v1 and SG-v2 behave similarly, and for $k = 38$, both methods end up outperforming G-v1, with respectively 60 times less and 180 times less computation. As seen before, using a *too* small batch size k (e.g. 10) yields a drop in performance. Using a batch of training images l (SG-v2) does not seem to significantly reduce the performance compared to SG-v1, while substantially reducing computations. As a result, in the sequel, we use $k = 38$ and $l = 1$ for SG-v2.

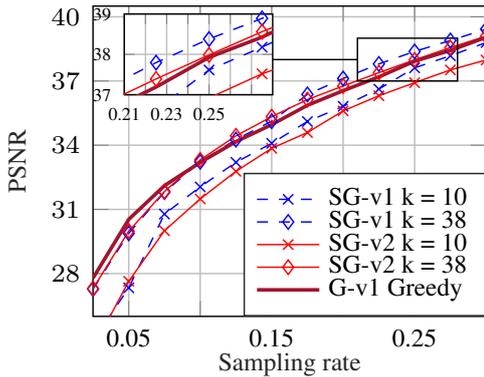


Figure 5.5 – PSNR as a function of the sampling rate for KTF, comparing the different reconstruction methods as well as the effect of the batch size on the quality of the reconstruction for SG.

Computational costs. We discuss now the computational costs for the different variations of the greedy methods used in the single coil experiments. Table 5.1 provides the running times and empirically measured speedup for the greedy variation, and Table 5.2 provides the computational times required to obtain the learning-based variable density (LB-VD) parameters through an extensive grid-search. The empirical speedup is computed as

$$\text{Speedup} = \frac{t_{\text{G-v1}} \cdot n_{\text{procs, G-v1}}}{t_{\text{SG-v2}} \cdot n_{\text{procs, SG-v2}}} \quad (5.5)$$

where $t_{\text{G-v1}}$, $t_{\text{SG-v2}}$ refer to the measured running times of the algorithms, and $n_{\text{procs, G-v1}}$, $n_{\text{procs, SG-v2}}$ to the number of CPU processes used to carry out the computations.

Chapter 5. Scalable learning-based sampling optimization

Algorithm	Setting	G-v1 (LBCS)		SG-v1 (sLBCS-v1)			SG-v2 (sLBCS-v2)		
		Time	n_{procs}	Time	n_{procs}	Speedup	Time	n_{procs}	Speedup
KTF	152×152×17×3	6d 23h	152	11h 40	38	58 (68)	3h 25	38	170 (204)
	256×256×10×2	~ 7d 8h*	256	12h 20	64	57 (68)	5h 20	64	173 (204)
IST	152×152×17×3	3d 11h	152	5h 30	38	60 (68)	1h 37	38	184 (204)
ALOHA	152×152×17×3	~ 25d 1h*	152	1d 14h 25	38	62 (68)	18h 13	38	133 (204)

Table 5.1 – Running time of the greedy algorithms for different decoders and training data sizes. The setting corresponds to n_x , n_y , n_{frames} , n_{train} . The number n_{procs} denotes how many parallel processes are used by each simulation. * means that the runtime was extrapolated from a few iterations. We used $k = n_{\text{procs}}$ for SG-v1 and SG-v2 and $l = 1$ for SG-v2. The speedup column contains the measured speedup and the theoretical speedup in parentheses.

Algo.	Setting	n_{pars}	n_{procs}	Time
KTF	152×152×17×3	1200	38	6h 30
	256×256×10×2*	2400	64	6h 45
IST	152×152×17×3	1200	38	3h 20
ALOHA	152×152×17×3	1200	38	1d 8h

Table 5.2 – Comparison of the learning-based random variable-density Gaussian sampling optimization for different settings. n_{pars} denotes the size of the grid used to optimize the parameters. For each set of parameters, the results were averaged on 20 masks drawn at random from the distribution considered. The n_{pars} include a grid made of 12 sampling rates (uniformly spread in $[0.025, 0.3]$), 10 different low frequency phase encodes (from 2 to 18 lines), and different widths of the Gaussian density (uniformly spread in $[0.05, 0.3]$) – 10 for the images of size 152×152 , 20 in the other case.

The main point of these tables is to show that the computational improvement is very significant in terms of resources, and that our approach improves greatly the efficiency of the method of Gözcü et al. (2018). This ratio might differ from the predicted speedup factor of $\Theta\left(\frac{m}{l} \frac{NT}{k}\right)$ due to computational considerations. Table 5.1 shows that we have roughly a factor 1.2 between the predicted and the measured speedup, mainly due to the communication between the multiple processes as well as I/O operations.

Moving to the grid search computations for LB-VD, displayed in Table 5.2, we ought to do several remarks. The number of parameters was chosen in order to ensure the same order of magnitude of computations to be carried out for both LB-VD and sLBCS. Note also that in opposition to SG-v2, the grid search cannot be sped up by using a batch of training data: all the evaluation should be done on the same data. In the case of SG-v2 however, one can draw a new batch of training data at each sampling round of the algorithm.

5.4.3 Single coil results

Main results. We see on Figures 5.6 and 5.7 that the SG-v2 brings a consistent improvement over all baselines considered, even though some variable-density techniques are able to provide good results for some sampling rates and algorithms. Compared to Coherence-VD, there is always at least 1 dB improvement at any sampling rate, and it can be as much as 6.7 dB at 5% sampling rate for ALOHA. For the rest, the improvement is over 0.5 dB. Figure 5.6 also clearly indicates that the benefits of our learning-based framework become more apparent towards higher sampling rates, where the performance improvement over LB-VD reaches up to 1 dB. Towards lower sampling rates, with much fewer degrees of freedom to achieve good mask designs, the greedy method and LB-VD yield similar performance, which is expected².

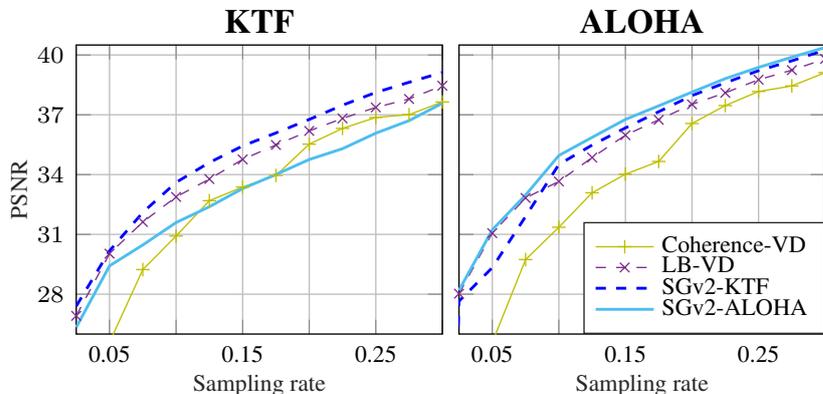


Figure 5.6 – PSNR as a function of sampling rate, evaluating the different masks (Coherence-VD, LB-VD, SG-v2-KTF, SG-v2-ALOHA) on both reconstruction methods (KTF and ALOHA). The results are averaged over 4 images.

Focusing on Figure 5.7, we observe that comparing reconstruction algorithms using the model-based, model-driven approach Coherence-VD does not allow for a faithful comparison of their performance. In this case, the performance difference between KTH and ALOHA is marginal. However, when moving to the model-based, data-driven approach LB-VD, the difference becomes quantitatively clearer. Moving to the learning-based, data-driven approach SG-v2 makes the difference even more noticeable: ALOHA with its corresponding mask clearly outperforms KTF, and this conclusion could not be made by looking solely at reconstructions with VD-based masks. This trend is also reflected on Figure 5.6.

This is an important observation, as most works have developed new reconstruction algorithms that were then tested using random masks designed with criteria such as minimal coherence. However, our results suggest that this approach does not reflect the *true* performance of a reconstruction method, and that one should look to optimize the

²At low sampling rates, learning-based approaches will prioritize low-frequency phase encoding lines, as those contain most of the energy at a given frame.

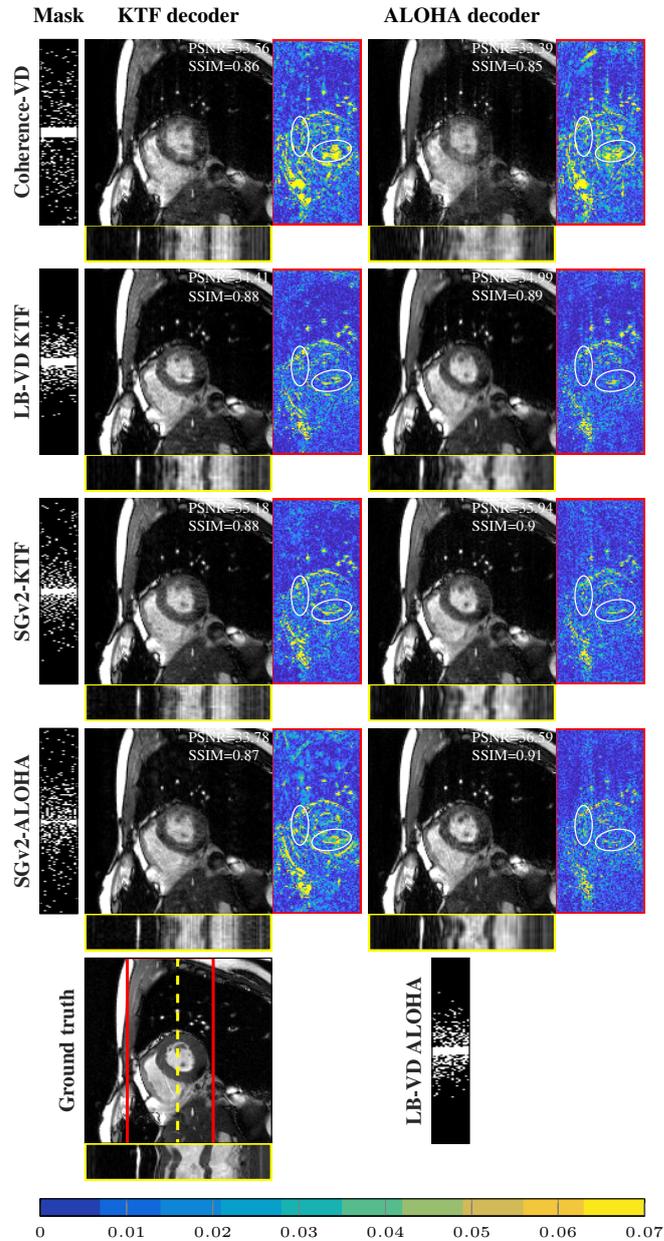


Figure 5.7 – Comparison of the different reconstruction masks and decoders, for a sampling rate of 15% on a single sample with its PSNR/SSIM performances.

mask and the reconstruction algorithm jointly.

The results on Figure 5.6 highlight also an interesting asymmetry between the performance of the SG-v2 mask optimized on KTF (SG-v2-KTF) or ALOHA (SG-v2-ALOHA): the one from KTF is more centered towards low frequencies than the one of ALOHA, and also performs well, and better than LB-VD on ALOHA. However, the reverse is not true: the mask from ALOHA is very scattered throughout k-space and leads to a very poor

performance when applied to KTF. This trend is also observed on multi-coil data, as highlighted on Table 5.3. This suggests that ALOHA, not relying explicitly on ℓ_1 sparsity, might succeed in exploiting varied information, from both low-frequency emphasizing masks from models promoting sparsity, but also higher frequency information.

Remark. This point does not hold directly for deep learning methods, that will be discussed in later chapters. Indeed, the reconstruction methods that we use in these experiments are all based on an underlying *model*, such as sparsity or low-rank. They are *not* data-driven approaches to reconstruction.

The question of what truly reflects the performance of a reconstruction method becomes even more complex in the case of deep learning-based reconstruction method, where the model is trained on a distribution of different sampling masks, making reconstruction and sampling even more intricate. We will come back to this question in Chapter 6.

Remark. The results in the sequel are obtained with the SG-v1 algorithm, rather than SG-v2. Recall that SG-v1 already accelerated G-v1 by a factor 60, and that we did not observe, in our experiments, a significant performance difference between SG-v1 and SG-v2³.

Cross-performances of performance measures. Up to here, we used PSNR as the performance measure, and we now compare it with the results of the greedy algorithm paired with SSIM, a metric that more closely reflects perceptual similarity (Wang et al., 2004). For brevity, we only consider ALOHA in this section. In the case where we optimized for SSIM, we noticed that unless a low-frequency initial mask is given, the reconstruction quality would mostly stagnate. This is why we chose to start the greedy algorithm with 4 low-frequency phase encodes at each frame in the SSIM case.

The reconstructions for PSNR and SSIM are shown on Figure 5.8, where we see that the learning-based masks outperform the baselines across all sampling rates except at 2.5% in the SSIM case. The quality of the results is very close for both masks, but each tends to perform slightly better with the performance metric for which it was trained. The fact that the ALOHA-SSIM result at 2.5% has a very low SSIM is due to the fact that we impose 4 phase encodes across all frames, and the resulting sampling mask at 2.5% is a low pass mask in this case.

A visual reconstruction is provided in Figure 5.9, we see that there is almost no difference in reconstruction quality, and that the masks remain very similar. Overall, we observe in this case that the performance metric selection does not have a dramatic effect on the

³The main reason of having these experiments on SG-v1 is that they were obtained from an earlier set of results. However, we believe that they remain insightful on the behavior of sLBCS, highlighting not only application on different settings, but we expect that such insights would also carry out to SG-v2.

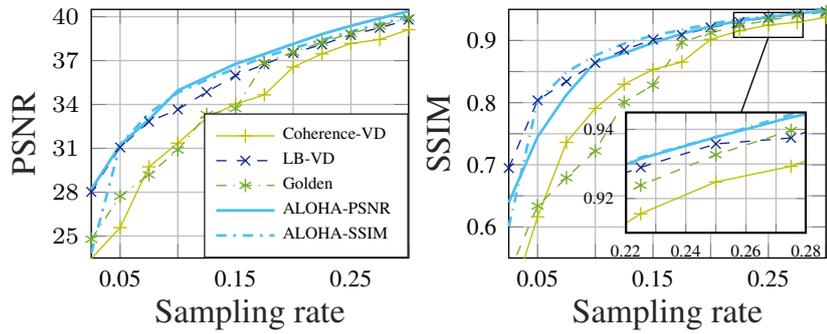


Figure 5.8 – PSNR and SSIM as a function of sampling rate for ALOHA, comparing the SG-v1 results optimized for PSNR and SSIM with the three baselines, averaged on 4 testing images of size $152 \times 152 \times 17$.

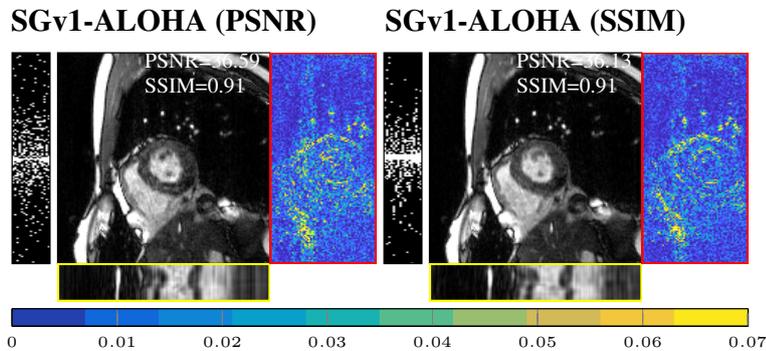


Figure 5.9 – Comparison of the sampling masks optimized for PSNR and SSIM with ALOHA, at 15% sampling. The images and masks can be compared to those of Figure 5.7, as the settings are the same.

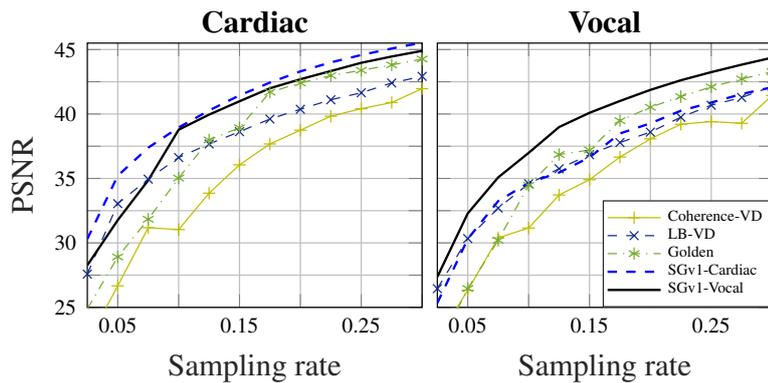


Figure 5.10 – PSNR as a function of sampling rate for KTF, comparing SG-v1 with both baselines, averaged on 2 testing images for both cardiac and vocal data sets of size $256 \times 256 \times 10$.

quality of reconstruction, and our greedy framework is still able to produce masks that outperform the baselines when optimizing SSIM instead of PSNR.

Experiments with different anatomies. In these experiments, we consider both the single coil cardiac dataset and the vocal imaging dataset, both of size $256 \times 256 \times 10$. The cardiac dataset was trained on 5 samples and tested on 2, using only the first ten frames of each scan, whereas the vocal one used 2 training samples and 2 testing samples. In this setup, the k-space of the cardiac dataset tends to vary more from one sample to another than the vocal one, making the generalization of the mask more complicated. This issue would require more training samples, but imposing SG-v1 algorithm to start with 4 central phase encoding lines on each frame was found to be sufficient to acquire the peaks in the k-space across the whole dataset. *SGv1-Cardiac* refers to the stochastic greedy algorithm using cardiac data, and *SGv1-Vocal* is its vocal counterpart. The algorithm used a batch of size $k = 64$ at each iteration⁴, and the results were obtained using only KTF.

The results are reported on the Figures 5.10 and 5.11, and we see that, for the both datasets, the greedy approach provides superior results against VD sampling methods across all sampling rates. It is striking that, in this setting, the SG-v1 approach outperforms even more convincingly all the baselines, and the LB-VD approach, in this case, is outperformed by more than 2dB by SG-v1, where it remained very competitive in the other settings. This difference is clear in the temporal fidelity of both reconstructions on Figure 5.11, where we see that the LB-VD approach loses sharpness and accuracy compared to SG-v1.

Comparison across anatomies. The main complication coming from applying the masks across anatomies is that the form of the k-space might vary heavily across datasets: the vocal spectrum is very sharply peaked, while the cardiac one is much broader. Comparing the cross-performances on Figures 5.11, we see that the and *SGv1-vocal* masks generalizes much better on the cardiac datasets than the other way around. This can be explained from the differences in the spectra: the cardiac one being more spread out, the cardiac mask less faithfully captures the very low frequencies of the k-space, which are absolutely crucial to a successful reconstruction on the vocal dataset, thus hindering the reconstruction quality. Also, we see that it is important for the trained mask to be paired with its anatomy to obtain the best performance.

Additional visual reconstructions for cardiac and vocal dataset. We show in Figures 5.12 and 5.13 reconstruction at different frames which provide clearer visual information to the quality of reconstruction compared to the temporal profiles.

For these images, the PSNR and SSIM are computed with respect to each individual frame, showing the quality of the reconstruction in a much more detailed fashion than before, where we considered each dynamic scan as a whole. Generally, we as previously observed, the mask trained for a specific anatomy will most faithfully capture the sharp contrast transitions in the dynamic regions of the images. For the vocal images, we see

⁴This follows the rule of thumb of using 25% of the number of lines in a given frame for k .

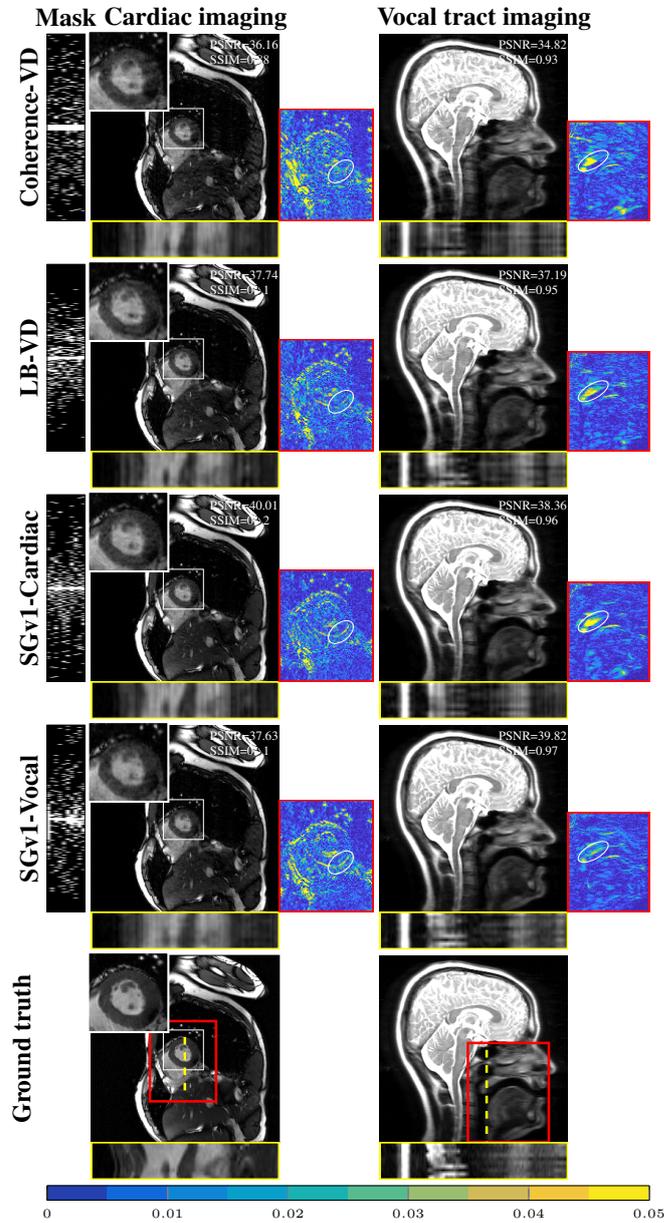


Figure 5.11 – Reconstruction for KTF at 15% sampling for the cardiac and vocal anatomies of size $256 \times 256 \times 10$. Figures showing different frames for the vocal and cardiac images are available in Figures 5.12 and 5.13.

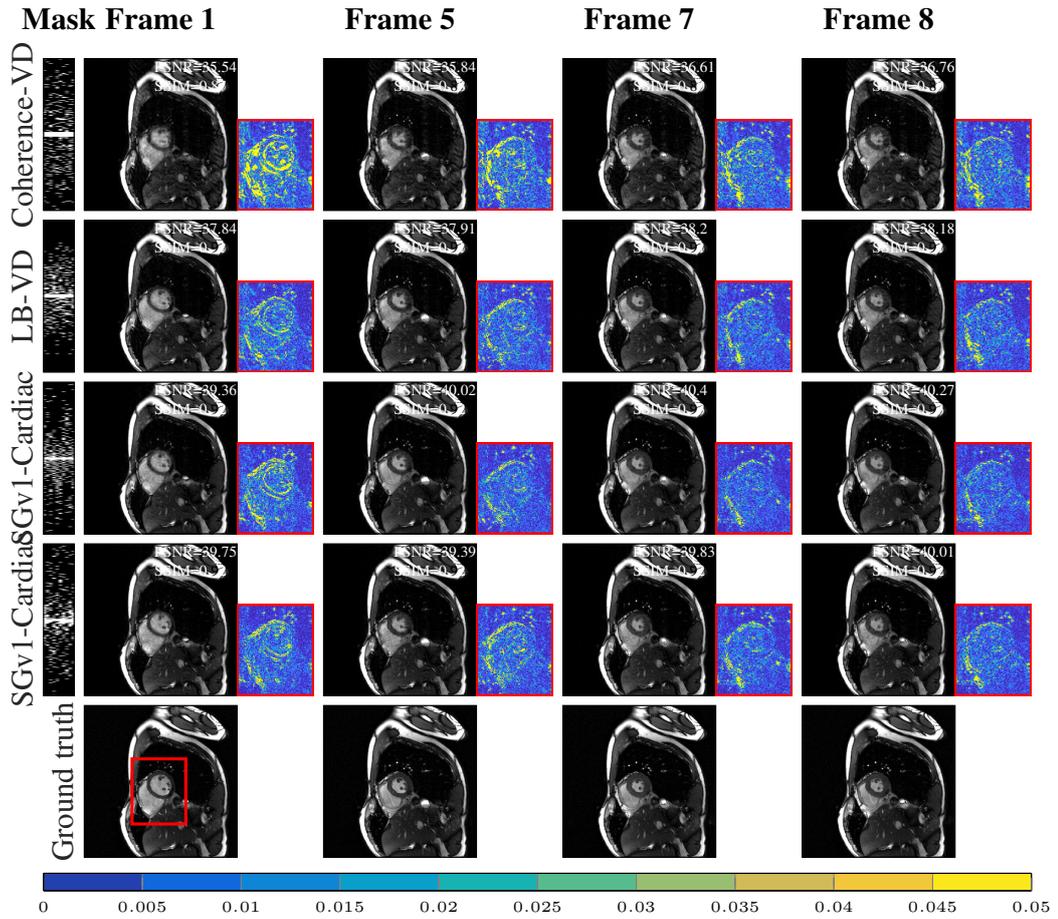


Figure 5.12 – Reconstruction with KTF (Jung et al., 2009) at 15% sampling rate for the cardiac anatomy of size $256 \times 256 \times 10$. It unfolds the temporal profile of Figure 5.11. The PSNR and SSIM displayed are computed for each image individually, and the overall PSNR for each image is the one of Figure 5.11. The ground truth is added at the end of each line for comparison.

that sampling the first frame more heavily is important in order to avoid having a very large PSNR discrepancy, as observed for the other masks. The PSNR remains quite stable across the frames otherwise.

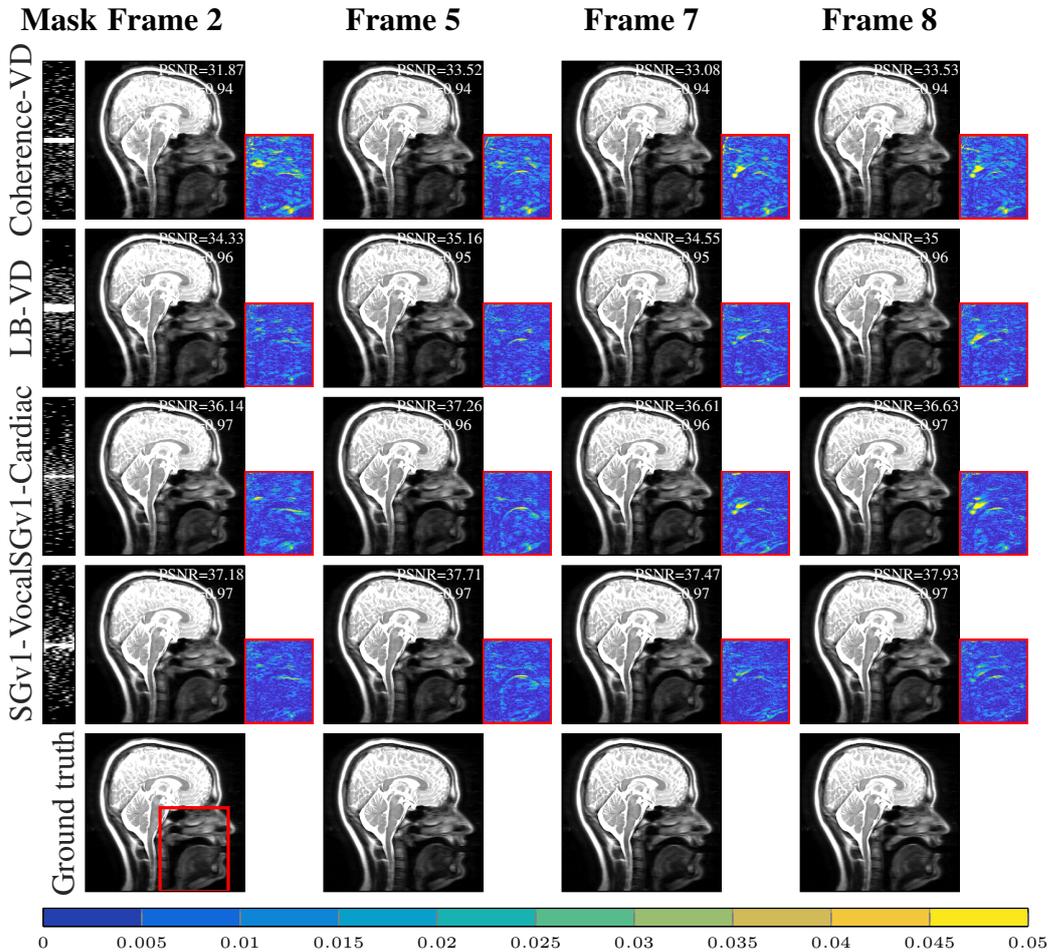


Figure 5.13 – Reconstruction with KTF at 15% (Jung et al., 2009) sampling rate for the vocal anatomy of size $256 \times 256 \times 10$. It unfolds the temporal profile of Figure 5.11. The PSNR and SSIM displayed are computed for each image individually, and the overall PSNR for each image is the one of Figure 5.11. The ground truth is added at the end of each line for comparison.

Noisy experiments. In order to test the robustness of our framework to noise, we artificially added bivariate circularly symmetric complex random Gaussian noise to the normalized complex images, with a standard deviation $\sigma = 0.05$ for both the real and imaginary components. We then tested to see whether the greedy framework is able to adapt to the level of noise by prescribing a different sampling pattern than in the previous experiments.

We chose to use V-BM4D (Maggioni et al., 2012) as denoiser with its default suggested mode using Wiener filtering and low-complexity profile, and provided the algorithm the standard deviation of the noise as the denoising parameter. The comparison between the fully sampled denoised images and the original ones yields an average PSNR of 24.95 dB across the whole dataset. Due to the fact that none of the reconstruction algorithms that we used have a denoising parameter incorporated, we simply apply the V-BM4D

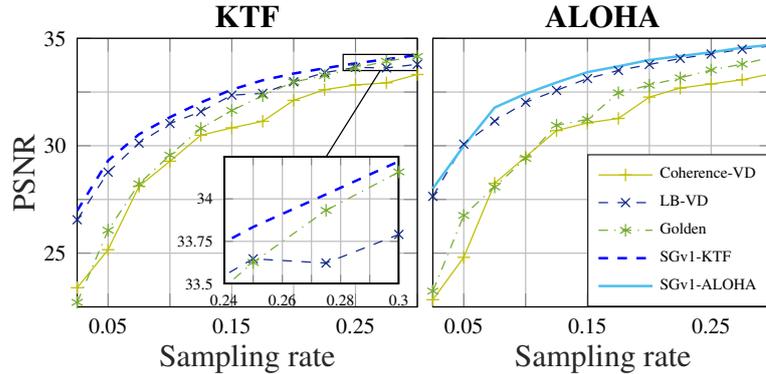


Figure 5.14 – PSNR as a function of sampling rate for both reconstruction algorithms considered, comparing SG-v1 with the three baselines, averaged on 4 noisy testing images of size $152 \times 152 \times 17$. The PSNR is computed between the denoised reconstructed image and the original (not noisy) ground truth.

respectively to the real and the imaginary parts of the result of the reconstruction. The results that we obtain are presented on the Figures 5.14 and 5.15.

It is interesting to notice on Figure 5.15 that the sLBCS framework outperforms the baselines that are not data-driven by a larger margin than in the noiseless case, and this is again especially true at low sampling rates. In this case however, the difference between SG-v1 and LB-VD methods is much smaller, and this might be explained by the fact that noise corrupts the high frequency samples, and thus the masks concentrate more around low-frequencies, leaving less room for designs that substantially differ.

We see a clear adaptation of the resulting learning-based mask, as shown by comparing Figures 5.7 and 5.15: the masks SGv1-KTF and SGv1-ALOHA, which are trained on the noisy data, are closer to low-pass masks, due to the high-frequency details being lost to noise, and hence, no very high frequency samples are added to the mask.

Also, notice that even if the discrepancy in PSNR is only around 0.8 – 1 dB between the golden ratio sampling and the optimized one, the temporal details are much more faithfully preserved by the learning-based approach, which is crucial in dynamic applications. The inadequacy of coherence-based sampling is highlighted in this case, as very little temporal information is captured in the reconstruction with both decoders. Also, for both decoders, there is a clear improvement on the preservation of the temporal profile when using learning-based masks compared to the baselines; the improvement of the SGv1-ALOHA mask of around 3dB also shows how well our framework is able to adapt to this noisy situation, whereas Coherence-VD yields results of unacceptable quality.

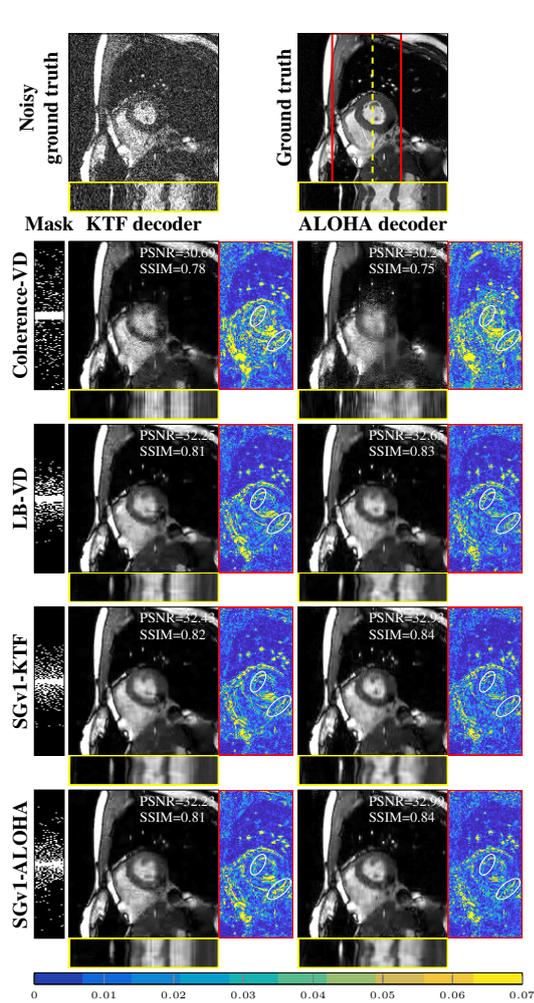


Figure 5.15 – Reconstructed denoised version from the noisy ground truth on the first line, at 15% sampling. The PSNR is computed with respect to the original ground truth on the top right.

5.4.4 Multicoil experiment

We now illustrate the applicability of our SG-v1 algorithm to multicoil data, a setting much more representative of routine clinical practice than single coil.

For the multicoil experiment, we used the previously described cardiac dataset, but we did not crop the images. We took the first 12 frames for all subjects, and selected 4

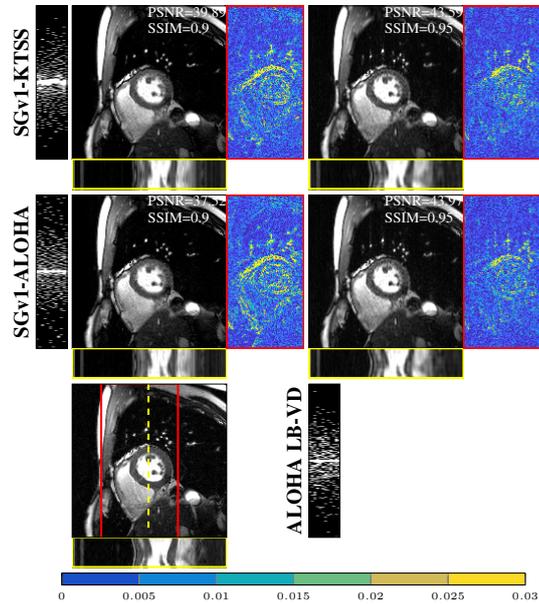


Figure 5.16 – Reconstruction with KTSS (Otazo et al., 2010) and ALOHA (Jin et al., 2016) at 15% sampling rate for a 4 coil parallel acquisition of cardiac cine size $256 \times 256 \times 12$. The setting is otherwise similar as the one presented in Figure 5 of (Sanchez et al., 2020a).

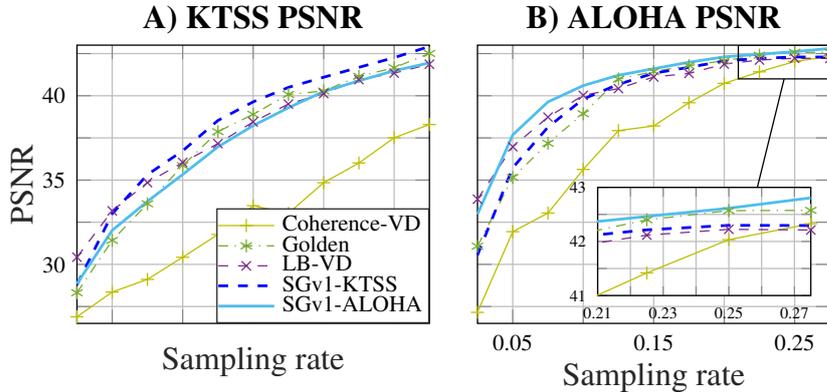


Figure 5.17 – PSNR as a function of sampling rate for KTSS (Otazo et al., 2010) and ALOHA (Jin et al., 2016) in the multicoil setting, comparing SG-v1 with the coherence-VD (Lustig et al., 2007), LB-VD and golden ratio Cartesian sampling (Li et al., 2018), averaged on 4 testing images of size $256 \times 256 \times 12$ with 4 coils.

coils that cover the region of interest. Each image was then normalized in order for the resulting sum-of-squares image to have at most unit intensity. When required, the coil sensitivities were self-calibrated according to the idea proposed by Feng et al. (2013), which averages the signal acquired over time in the k-space and subsequently performs adaptive coil combination (Walsh et al., 2000; Griswold et al., 2002b).

The advantage of using self-calibration is that the greedy optimization procedure can simultaneously take into account the need for accurate coil estimation as well as accurate reconstruction, thus potentially eliminating the need for a calibration scan prior to the acquisition. A more complete discussion of the accuracy of self-calibrated coil sensitivities is presented in Feng et al. (2013).

We used k-t SPARSE-SENSE (Otazo et al., 2010) and ALOHA (Jin et al., 2016) for reconstruction. While the first requires coil sensitivities, the second reconstructs the images directly in k-space before combining the reconstructed data.

5.4.5 Large scale static results

This last experiment shows the scalability of our method to very large datasets. We used the fastMRI dataset (Zbontar et al., 2019) consisting of knee volumes and trained the mask for reconstructing the 13 most central slices of size 320×320 , which yielded a training set containing 12 649 slices. For the sake of brevity, we only report computations performed using total variation (TV) minimization with NESTA (Becker et al., 2011). For mask design, we used the SG-v2 method with $k = 80$ and $l = 20$ (2 500 fewer computations compared to G-v1). The LB-VD method was trained using 80 representative slices and optimizing the parameters with a similar computational budget as SG-v2. The result on Figure 5.18 shows a uniform improvement of our method over the LB-VD approach.

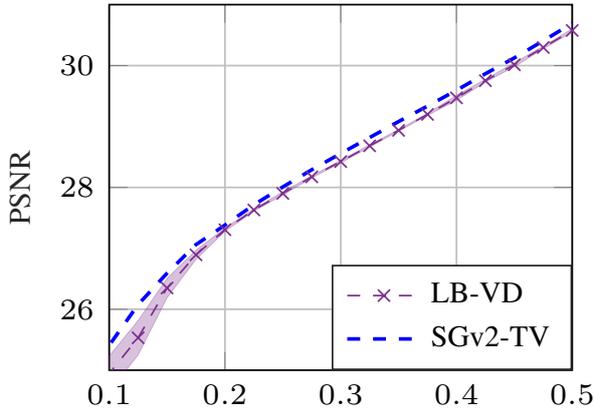


Figure 5.18 – PSNR as a function of the sampling rate for TV, averaged on the 13 most central slices of the fastMRI validation set (Zbontar et al., 2019) (2587 slices). SGv2 outperforms LB-VD over all sampling rates.

5.5 Experiments - Lazy LBCS

Dataset. Our experiments were carried out on two datasets.

- *Brain data.* The data were acquired on a 3T MRI system (Siemens), and the protocols were approved by the local ethics committee. The dataset consists of 2D T1-weighted brain scans of seven healthy subjects, which were scanned with a FLASH pulse sequence and a 12-channel receive-only head coil. In our experiments, we use 4 different slices of sizes 256×256 from two subjects as training data and 100 slices from 5 subjects as test data. To speed up the computations, we compressed our 12-channel data into 4-channel data.
- *Knee data.* The data were acquired by a 3T system (GE Healthcare) using FSE CUBE sequence with the proton density weighting. The number of channels is 8 and the matrix size is 320×320 . We used the data of 8 subjects as training data and remaining 12 subjects as test data. This dataset is publicly available at <http://mridata.org>.

We used PSNR as the image quality metric η , but also reported the results evaluated on SSIM (Wang et al., 2004).

Reconstruction algorithms. We consider the following decoders: Basis Pursuit (BP) (Donoho, 2006), Total Variation (TV), SENSE (Pruessmann et al., 1999), and ALOHA (Jin et al., 2016). For BP, we let the sparsifying transform to be the Shearlet transform (Kutyniok et al., 2016). For SENSE, we used ESPIRiT (Uecker et al., 2014) method to estimate coil sensitivities.

Basis pursuit is a method that directly minimizes the ℓ_1 norm of the image in a sparse domain. A similar type of convex optimization formulation is *total variation* (TV) minimization. SENSE is a method that requires the knowledge of coil sensitivities to retrieve an unaliased image from parallel acquisitions (Pruessmann et al., 1999). We use

its implementation in the BART toolbox (Uecker et al., 2015) to iteratively solve

$$\min_{\mathbf{z}} \|\mathbf{y}_\omega - \mathbf{P}_\omega \mathbf{F} \mathbf{S} \mathbf{z}\|_2^2 + \lambda \|\mathbf{W} \mathbf{z}\|_1 \quad (5.6)$$

where $\mathbf{y}_\omega \in \mathbb{C}^{CP}$ is the coil-wise subsampled Fourier measurements, $\mathbf{S} = [\mathbf{S}_1 \dots \mathbf{S}_C]^T \in \mathbb{R}^{CP \times P}$ are the stacked coil sensitivities, $\mathbf{P}_\omega \mathbf{F}$ is the subsampled Fourier transform applied coil-wise to the product of the image with the stacked coil sensitivities and \mathbf{W} is a sparsifying operator. Recall that the multicoil formulation was discussed in greater depth in Equation 2.14. Other compressed sensing multicoil algorithms include (Uecker et al., 2014; Chaâri et al., 2011). Recall that ALOHA exploits the duality between sparse signals and their transform domain representations cast as a low-rank Hankel matrix, and transforms the problem into a low-rank matrix completion.

5.5.1 2D multicoil setting

We first compare the performances of lLBCS with various baselines. As it can be seen in Figure 5.19 for a range of subsampling rates, the coherence-based VD masks can perform better than low-pass (LP) mask. However, they have weaker performance compared to the parameter sweep with PSNR as the objective function (LB-VD). This is expected as the LB-VD masks exploit the knowledge of the training set, whereas the coherence-based VD masks do not. The best performance is provided by our proposed approach, lLBCS (LB in the results). Indeed, the greedy LB algorithm is free from the constraint of a parametric distribution, which allows the data to decide which mask is best suited to a given decoder. Also, the VD-based masks result in artifacts especially visible in the ALOHA reconstructions that are suppressed in the LB-case.

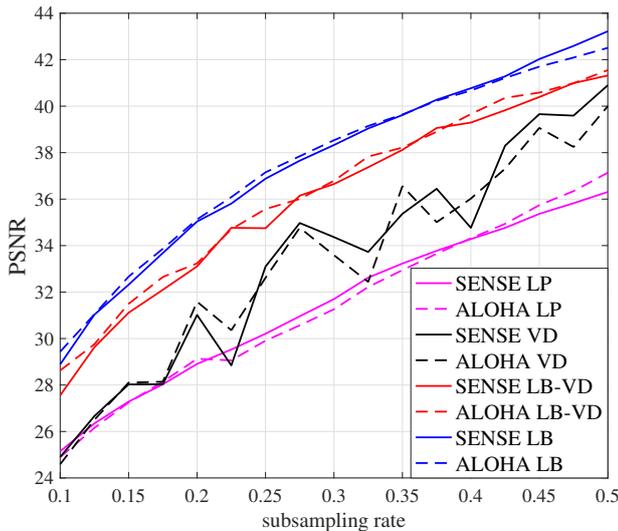


Figure 5.19 – PSNR performances of various masks averaged over 100 test images: LP (low-pass mask), VD (variable density), LB-VD and lLBCS (LB).

Cross-performances of single and multi-coil masks.

Mask	<i>Single-coil</i>		<i>Multicoil</i>	
	TV	BP	SENSE	ALOHA
Coherence	29.04	29.60	33.10	32.61
Low pass	30.59	30.78	30.21	29.91
TV LB	32.82	34.23	34.56	35.17
BP LB	32.90	34.27	35.02	35.07
SENSE LB	31.51	31.82	36.88	36.61
ALOHA LB	29.00	26.95	30.73	37.15

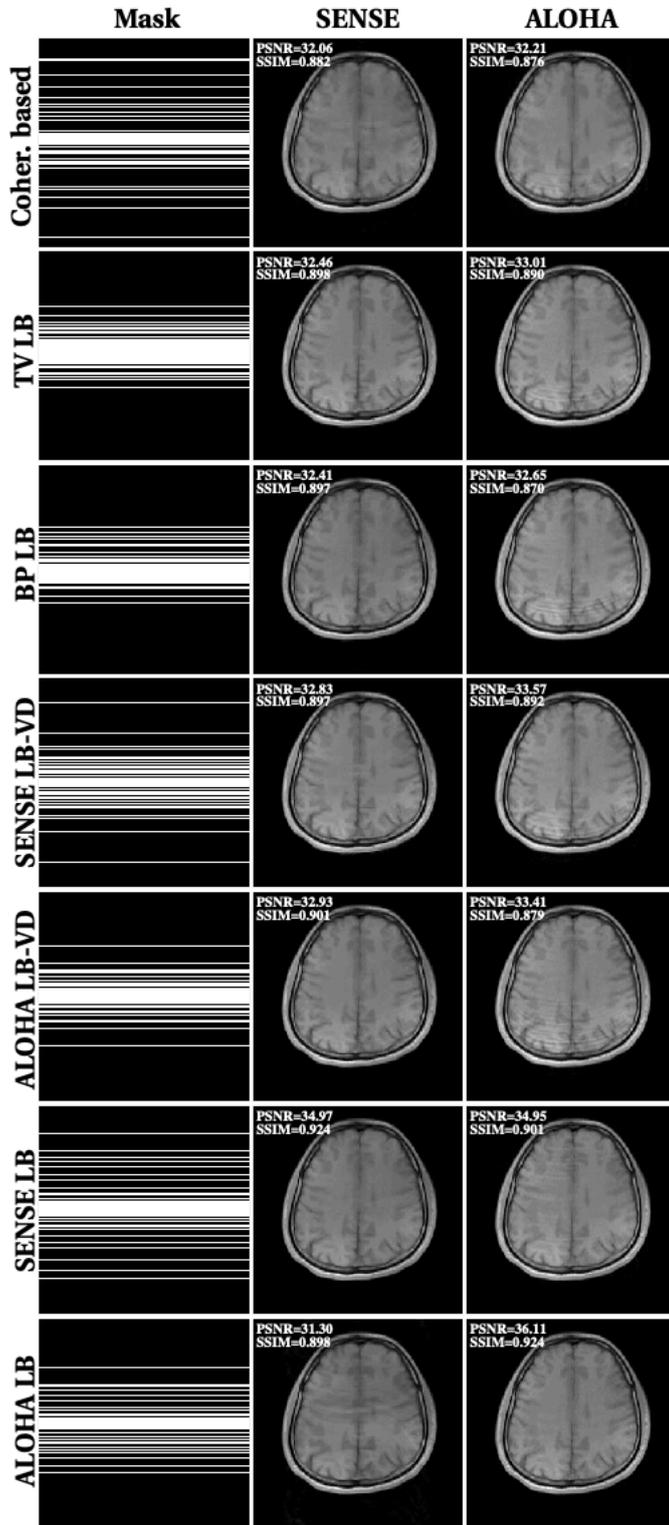
Table 5.3 – Cross performance of single and multi-coil masks on various reconstruction algorithms, at 20% subsampling rate, averaged over 100 test slices.

In Figure 5.20, we have also included single coil masks (TV and BP) obtained from using LBCS on single coil data. They result in inferior performance when used with the multi-coil SENSE and ALOHA reconstruction algorithms. Compared to these single coil masks, the optimal multi-coils masks are found to be more spread in k-space. As expected, one obtains the best performances in terms of visual quality and performance metrics when the LB-mask is paired with the reconstruction algorithm for which it was trained. In Table 5.3, we provide an average cross-performance of single/multi-coil masks and reconstruction algorithms. We see that it is important to optimize single and multi-coil masks separately as they are quite different from each other, which was also observed by Haldar and Kim (2019).

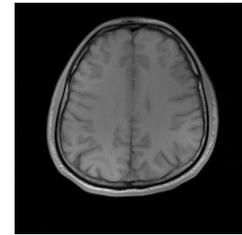
5.5.2 3D multicoil setting

In this section we demonstrate the effectiveness of ILBCS in providing masks for 3D MRI by comparing it to common 3D MRI subsampling masks: Controlled Aliasing in Parallel Imaging Results in Higher Acceleration (CAIPIRINHA) (Breuer et al., 2006), Poisson disc (PD) sampling (Bridson, 2007) and its variable density variant (VD-PD) (Vasanawala et al., 2011), and adaptive random sampling method (Knoll et al., 2011b). For brevity, we used only the SENSE algorithm as the multi-coil reconstruction algorithm and for all masks we took the central 24×24 region as calibration region to estimate the coil sensitivities. As seen in Figure 5.21, the mask given by ILBCS provides superior image quality compared to the benchmarks and comparable result to more recent VD-PD sampling whose polynomial decay parameter is optimized by a grid search (LB-VD).

Although this trend holds for low sampling rates, Figure 5.22 shows that the baselines catch up around 15% sampling and then outperform ILBCS. The trend is similar on both PSNR and SSIM evaluation, and several observations can be made. While our method does not uniformly outperform other baselines like in other cases, it remains competitive in the region of interest for such methods, namely low sampling rates. Indeed,



(a)



(b)

Figure 5.20 - Optimized masks and example reconstructions under SENSE and ALOHA decoding at 20% sampling rate. See Figure 5.19 for ground truth.

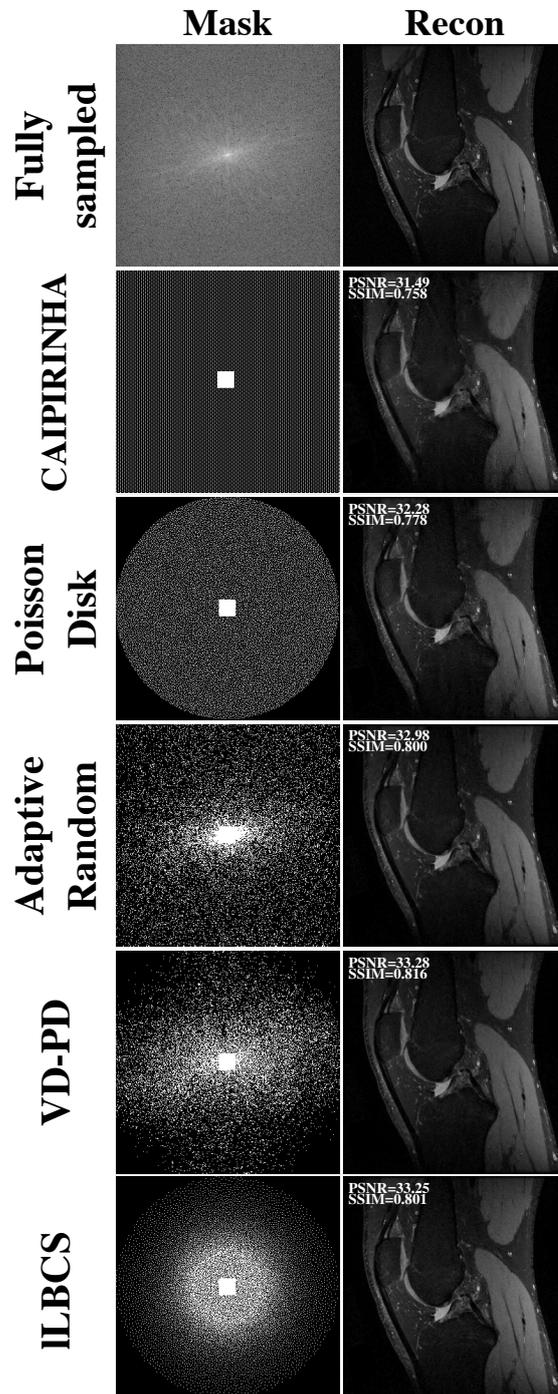


Figure 5.21 – Optimized masks and example reconstructions of knee images under SENSE decoding at 6-fold acceleration.

as acquiring these 3D volumes required 40 minutes of scanning (Sawyer et al., 2013), one aims at achieving gains by high accelerations. For instance, six-fold acceleration has been convincingly investigated in a prospective study for 3D knee MRI by Fritz et al. (2016).

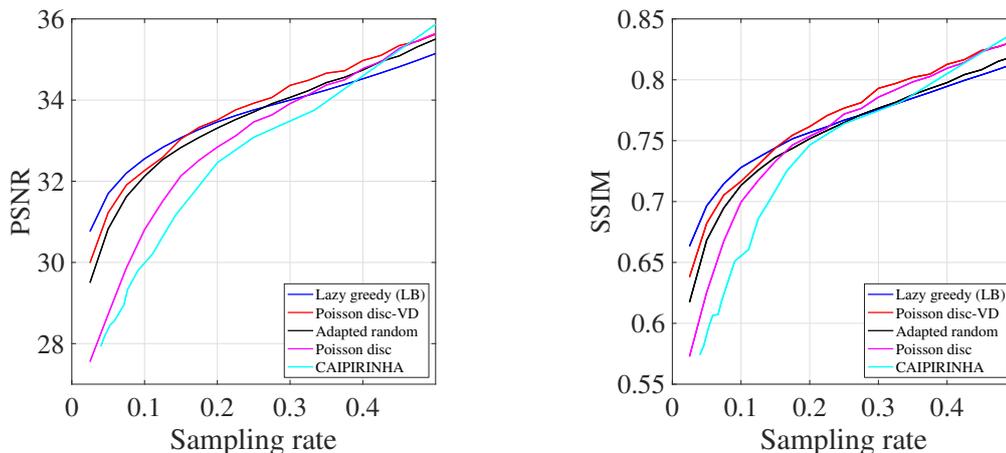


Figure 5.22 – PSNR and SSIM performance of various masks averaged over 120 testing 3D volumes. The ILBCS model was trained using PSNR.

Figure 5.23 shows the number of evaluations done per round of ILBCS, i.e. the number of times where the list of upper bounds $\rho(\cdot)$ is updated before adding an element to the mask ω in Algorithm 3. Two observations can be made: first, after the initial, expensive, step of evaluation, there are remarkably few updates in ILBCS, averaging at 2.644 per round, whereas the regular LBCS algorithm would have 76 368.5, which amounts to the enormous speedup of a factor 28 884. This testifies to the efficiency of the cost reduction achieved by ILBCS. However, and this is our second observation, it seems that the number of lazy evaluations are steadily decreasing throughout the iterations, and we hypothesize that the decrease of performance observed in Figure 5.22 can be attributed to the list of upper bounds getting increasingly unreliable as the number of iterations increase.

This suggests avenues to further improve the performance of ILBCS. The simplest approach would consist in periodically re-computing the list of upper bounds, to more accurately track the potentially useful sampling locations. Although this step would not be necessary for *submodular* functions, it could benefit our objective, which does not exactly satisfy submodularity. Alternatively, one could also re-evaluate a couple of elements from the list randomly at each iteration, spreading the computational load over a larger number of iterations.

5.6 Discussion and Conclusion

In this chapter, we presented two algorithmic contribution that broaden the scale of applications of LBCS (Gözcü et al., 2019) beyond its original limits, by successfully scaling to settings where the number of candidates grows drastically compared to 2D Cartesian MRI, to large dataset and to methods that present long-running times, as in multicoil and 3D MRI.

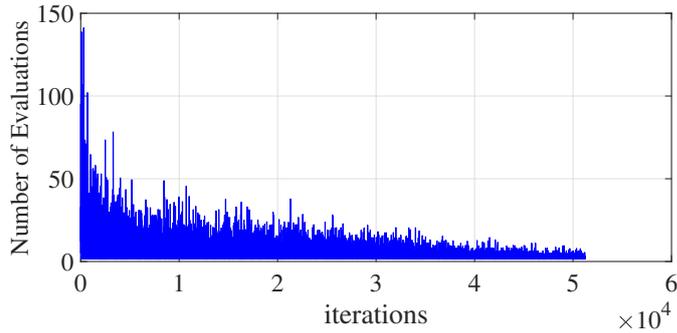


Figure 5.23 – Number of lazy evaluations per round of the lazy LBCS algorithm, until 50% sampling rate (50624 locations acquired, initialization with 24×24 centermost frequencies).

Adapting these algorithms to non-Cartesian settings such as radial sampling would be of great interest to clinicians, as radial sampling has seen a large success in body MRI due to the resistance to motion that it offers, enabling free breathing imaging (Coppo et al., 2015; Feng et al., 2016a). Although golden angle sampling has been widely used (Feng et al., 2014), we anticipate that adjusting the distribution of radial spokes more closely to the underlying data distribution would improve the overall quality of reconstructed images. We anticipate however that this would involve challenges in simulating a non-Cartesian acquisition and quantifying the improvement of adding new measurements to an image, but recent deep learning reconstruction methods could provide a solution to this issue (Ramzi et al., 2022).

Both sLBCS and lLBCS could certainly be improved by endowing them with additional heuristics. sLBCS would likely benefit from candidate lines being drawn from a probability favoring low frequency locations rather than fully at random. lLBCS could benefit from periodically re-computing the list of upper bounds ρ . Recent works such as Zibetti et al. (2020) also showed the benefits of heuristics for mask design. For instance, they bias the selection of candidates based on the current line-wise NMSE and avoid selected elements at complex-conjugated position in order to accelerate the optimization procedure.

This should however not be the main point that one takes away from this chapter. Our work illustrates the great benefit of moving from model-based, model-driven approaches for mask design to learning-based, data-driven ones: trying to maximize a performance metric on data (data-driven approach), rather than some abstract mathematical structure (model-driven approach) such as coherence enables to more efficiently use the limited available resources. Similarly, moving from model-based to learning-based approaches allows to remove the constraints on the shape of the masks that can be explored, and our results show that this greatly benefits the reconstruction quality. Stated differently, our results suggest that model-based and/or model-driven approaches *limit* the performance of CS methods applied to MRI, through their underlying model.

Bibliographic note

For the lazy LBCS work⁵, the work is mostly due to Baran Gözcü. The author of this dissertation contributed to the formalization of the lazy LBCS algorithm and to part of the 2D multi-coil experiments. Convinced of the potential of LBCS, Volkan Cevher provided the drive to further validate and extend the method.

For the stochastic LBCS work⁶, Ruud van Heeswijk acquired the cardiac data, Efe Ilıcak and Tolga Çukur provided the vocal tract data. Baran Gözcü was mostly involved in supervising the project. Armin Eftekhari provided Propositions 1 and 2, discussed in the previous chapter.

⁵Gözcü, B., Sanchez, T., and Cevher, V. (2019). Rethinking sampling in parallel MRI: A data-driven approach. In *27th European Signal Processing Conference (EUSIPCO)*.

⁶Sanchez, T., Gözcü, B., van Heeswijk, R. B., Eftekhari, A., Ilıcak, E., Çukur, T., and Cevher, V. (2020a). Scalable learning-based sampling optimization for compressive dynamic MRI. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8584–8588.

6 On the benefits of deep RL in accelerated MRI sampling

In parallel to the refinements to the Learning-based Compressive Sampling (LBCS) (Gözcü et al., 2018) approach that we studied in the previous chapter, the progress of deep learning (DL) applied to reconstruction motivated several researchers to turn their attention to the problem of optimizing sampling patterns in a learning-based, data-driven fashion.

In this Chapter, we wish to focus on the state-of-the-art (SotA) approaches that leverage deep reinforcement learning (RL) to perform adaptive mask design in a sequential fashion (Bakker et al., 2020; Pineda et al., 2020). These approaches are particularly relevant to us for two main reasons.

1. These two SotA approaches can be viewed as very natural extensions of (s)LBCS, which can be interpreted as a very basic RL policy. This makes it both easy and insightful to compare LBCS to these methods, as it will provide a baseline to quantify the benefits brought by moving from LBCS to deep RL.
2. More sophisticated RL policies promise to address two fundamental limitations of LBCS: they can perform adaptive sampling and incorporate long-term planning in their decision.

However, the SotA approaches of Pineda et al. (2020) and Bakker et al. (2020) draw seemingly conflicting conclusions from their results. The work of Pineda et al. (2020) seems to indicate that long term planning could be the most important component in deep RL, as their results show that a non-adaptive, long term planning policy model trained on the dataset can perform as well as an adaptive, long-term planning policy. On the contrary, the contribution of Bakker et al. (2020) highlights the importance of adaptivity, as a greedy policy, that does not do long-term planning is found to closely match policies that do long term planning.

Our results synthesize this apparent conflict: *We observe that sLBCS, a simple, easy-to-*

train method that does not rely on deep RL, does not attempt long term planning and is by definition not adaptive can perform as well as the state-of-the-art approaches of Pineda et al. (2020); Bakker et al. (2020).

This trend can be consistently observed on the fastMRI single-coil knee dataset (Zbontar et al., 2019), where we carry out evaluations across a variety of settings, whether full scale images, cropping the field of view, various mask designs used for training and various architectures. We observe that such small changes in the experimental pipeline can easily lead to different ordering of the RL methods and sLBCS in terms of performance or an invalidation of the conclusions drawn by the results. We refer to such occurrences as *reversals*. It is possible then to choose an experimental setting and performance metrics to support one’s desired conclusion: that RL outperforms LBCS, that LBCS outperforms RL or that their difference is not significant.

Together with the observation that current SotA RL methods only add marginal value over the simple baselines at best, the work in this chapter highlights the need for further discussions in the community about standardized metrics, strong baselines, and careful design of experimental pipelines to evaluate MRI sampling policies fairly¹.

6.1 Reinforcement Learning to optimize sampling trajectories

Before introducing how RL can be used to optimize sampling for MRI, let us briefly contextualize the works of Jin et al. (2019), Pineda et al. (2020) and Bakker et al. (2020). Jin et al. (2019) were the first to attempt at using an RL-based approach to mask design, trying to train jointly a reconstruction model and a policy model through self-supervised learning. The approach is based on the AlphaGo model (Silver et al., 2017), a model that successfully learned to play go through self-play and that eventually ended-up the beating the world champion of Go. In this approach, the policy model is trained to learn the predictions done by a Monte-Carlo Tree Search (MCTS) (Coulom, 2006).

Pineda et al. (2020) leverage Double Deep Q-Networks (DDQN) (Van Hasselt et al., 2016), a state-of-the-art deep RL approach that is particularly suited to the problem at hand and that has recently shown to be a reliable and stable method. They train two variants of the model, one to perform non-adaptive sampling (Dataset-specific, DS-DDQN), and the other that is adaptive (Subject-specific, SS-DDQN). Their results show that DS-DDQN outperforms SS-DDQN in nearly all cases.

Finally, Bakker et al. (2020) took a different approach, leveraging the policy gradient

¹The work in this chapter is based on the preprint: Sanchez, T.* , Krawczuk, I.* and Cevher, V. (2021). On the benefits of deep RL in accelerated MRI sampling. Available at <https://openreview.net/pdf?id=fRb9LBWUo56>. * denotes equal contribution.

6.1. Reinforcement Learning to optimize sampling trajectories

Sampling policy	Adaptive	Long horizon
LBCS (Gözcü et al., 2018)	✗	✗
AlphaGo (Jin et al., 2019)	✓	✓
DDQN (Pineda et al., 2020)	✗(✓)	✓
Policy Gradient (Bakker et al., 2020)	✓	✗(✓)

Table 6.1 – Methods that will be considered in the paper. Bakker et al. (2020) considered fixed vs. adaptive and greedy vs. non-greedy, Pineda et al. (2020) considered data specific vs subject specific policies and also compared against greedy methods.

theorem to enable the training of the policy network (Sutton and Barto, 2018, ch. 13). They both a greedy and a non-greedy (long horizon) policy, and find that, surprisingly, “a simple greedy approximation of the objective leads to solutions nearly on-par with the more general non-greedy approach.”

We summarize the features of the different models in Table 6.1, and for the reasons discussed above, we consider the method of Pineda et al. (2020) to be represented by their non-adaptive DS-DDQN, and the method of Bakker et al. (2020) to be represented by their greedy policy model.

6.1.1 A short RL primer

Assume that we have an agent in an unknown environment, and that the agent can interact with the environment and potentially obtain rewards based on its actions. In reinforcement learning, we would like the agent to learn to maximize its cumulative reward from its interaction with the environment. This means that the agent should learn a strategy that allows it to adapt to its environment in order to maximize its reward. In Go, this could mean train the agent to win as many games as possible, or in MRI, this could amount to selecting the next line to acquire to maximize the reconstruction quality.

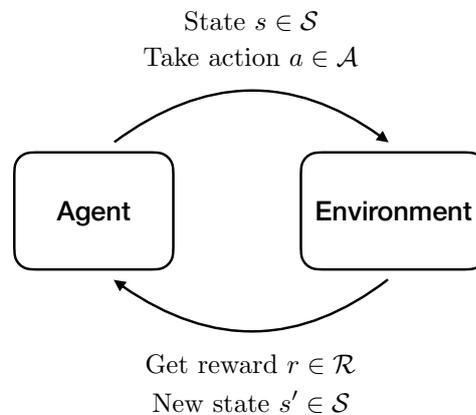


Figure 6.1 – Illustration of an agent interacting with its environment.

Chapter 6. On the benefits of deep RL in accelerated MRI sampling

Formally speaking, in RL, an agent acts in an *environment* with an underlying model that may be known or not. The agent will stay in one of the many states $s \in \mathcal{S}$ of the environment, and take an action $a \in \mathcal{A}$, which allow it to move from its current state to another state s' . As a result of taking the action a , the agent receives a reward $r \in \mathcal{R}$ as feedback.

This process is repeated several times, and produces a sequence of states, actions and rewards. This results in a tuple

$$(S_1, A_1, R_2, S_2, \dots, S_T)$$

where T is the total number of steps taken. The states, actions and rewards are capitalized here as they refer to random variables. This tuple is referred to as an *episode*. During an episode, the agent interacts with the environment in order to learn a better strategy. This strategy is described as a policy $\pi(s)$, which is a probability distribution over which actions taken in a state s are likely to get the maximal cumulative reward. In order to quantify this, each state is associated with a value function $V_\pi(s)$, which predicts the expected amount of future reward that we would receive by acting following the current policy π . In RL, one generally aims at learning the policy and/or the value function.

Modeling the environment. In order to do that, we need to further describe the environment by a model. A model usually relies on two quantities, the *transition probability* P and the reward R . Starting in state s and taking action a , the transition probability describe how likely it is that we end up moving to state s' and obtaining reward r as a result, namely

$$P(s', r|s, a) = \mathbb{P}[S_{t+1} = s', R_{t+1} = r | S_t = s, A_t = a]$$

where \mathbb{P} describes here a probability. The reward function is defined as the expected reward for a given state action pair

$$R(s, a) = \mathbb{E}[R_{t+1} | S_t = s, A_t = a] = \sum_{r \in \mathcal{R}} r \sum_{s' \in \mathcal{S}} P(s', r|s, a)$$

Value function and state-action function. The value function aims at predicted the expected amount of future reward, that is commonly referred to as the *return* and defined as

$$G_t \triangleq R_{t+1} + \gamma R_{t+2} + \dots \gamma^{T-t-1} R_T = \sum_{k=t}^T \gamma^{k-1} R_{t+k}$$

where $\gamma \in [0, 1]$ is called the *discounting factor*. The idea of the discounting factor is to penalize rewards in the future as a way to nudge the agent towards *immediate* benefit. Note that the return G_t always sums every future reward until the end of the episode, at time T .

6.1. Reinforcement Learning to optimize sampling trajectories

We can then use the return to define formally the value function

$$V_\pi(s) \triangleq \mathbb{E}_\pi[G_t | S_t = s] \quad (6.1)$$

which means that the value of the state s is the expected return given that we currently are in the state $S_t = s$. A related important quantity is the *state-action function*

$$Q_\pi(s, a) \triangleq \mathbb{E}_\pi[G_t | S_t = s, A_t = a] \quad (6.2)$$

which is similar to the value function except that it considers a state-action pair (s, a) . The two can be related as

$$V_\pi(s) = \sum_{a \in \mathcal{A}} Q_\pi(s, a) \pi(a|s) \quad (6.3)$$

Markov Decision Processes. Most RL problems can be expressed as Markov Decision Processes (MDPs). This implies in particular that the current state S_t and action A_t contain all the necessary information to future decisions, i.e.

$$\mathbb{P}(S_{t+1} | S_t, A_t) = \mathbb{P}(S_{t+1} | S_1, A_1 \dots, S_t, A_t).$$

In other words, this means that the future S_{t+1} is conditionally independent of the past $(S_1, A_1 \dots, S_{t-1}, A_{t-1})$ given the present S_t, A_t . An MDP consists of five elements that we already encountered, namely $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, P, R, \gamma \rangle$. We have a set of states, a set of actions, the transition probability, the reward function and the discount factor. This description comprises everything that we need to tackle a RL problem.

6.1.2 The special case of MRI

The problem of optimizing the sampling mask of Cartesian MRI fits the description of RL well, but is not exactly described by an MDP. If we look to learn a sampling pattern on an image \mathbf{x} , we can start from an empty mask and gradually make observations, trying to learn which are the best locations to acquire. It is natural to view the current mask ω_t as the current state s_t , and interpret the action a_t as the next location that we seek to acquire. The reward is easily quantified as the improvement in reconstruction obtained by adding to the sampling mask the location corresponding to a_t :

$$r(s_t, a_t) = \eta(\mathbf{x}, \hat{\mathbf{x}}_{\theta; \omega_{t+1}}) - \eta(\mathbf{x}, \hat{\mathbf{x}}_{\theta; \omega_t}) \quad (6.4)$$

where $\omega_{t+1} = \omega_t \cup a_t$ and $\hat{\mathbf{x}}_{\theta; \omega_{t+1}} = \hat{\mathbf{x}}_{\theta}(\mathbf{P}_{\omega_{t+1}} \mathbf{F} \mathbf{x})^2$. Recall also that $\eta(\cdot, \cdot) : \mathbb{C}^P \times \mathbb{C}^P \rightarrow \mathbb{R}$ is a performance metric. In the sequel, to avoid unnecessary clutter, we refer to the

²Note that we slightly abuse notation here and use the action a_t to also refer to its corresponding sampling location.

reconstruction at state s_t simply as $\hat{\mathbf{x}}_t$.

Note also that while the RL framework allows for transition *probabilities*, the case of MRI, they are deterministic: we know in what state we will end up should we take an action.

The main difference lies in the fact that although the reward is well defined, it depends on the unobserved underlying ground truth \mathbf{x} . As a result, our problem falls into the category of *Partially Observable* Markov Decision Processes (POMDP) (Monahan, 1982). To take this into account, we need to introduce the set of observations \mathcal{O} and the conditional observation function $O(s'|a, o)$. This leads us to define a POMDP as the 7-element tuple $\langle \mathcal{S}, \mathcal{O}, \mathcal{A}, P, O, R, \gamma \rangle$, where \mathcal{S} is the set of all ground truths and mask $\{\langle \mathbf{x}, \omega \rangle\}$, \mathcal{O} is the set of all observations $\{\langle \hat{\mathbf{x}}, \omega \rangle\}$, \mathcal{A} is the set of all actions, corresponding typically to all columns to be acquired. The transition probability function is deterministic as discussed above, and the conditional observation probability will also be deterministic, as taking action a_t from state s_t will result in observing the tuple $\langle \hat{\mathbf{x}}_{t+1}, \omega_{t+1} = \omega_t \cup a_t \rangle$. Finally, we have defined the reward function above, and the discount factor γ can be taken as a parameter.

6.1.3 Double Deep Q-Networks

Until now, we have defined the fundamental building blocks of RL and the environment corresponding to the problem of learning sampling policies. We turn now to practical methods to learn sampling policies³.

The state of possible state-action pairs in RL is very large, and in most modern RL applications, it is simply impossible to explore it in an exhaustive manner. Similarly, storing the value for each state-action pair is impractical, and in practice, one will choose to represent the state-value function with a parametric model $Q_\phi(s, a)$ (with parameters ϕ). Of course, deep neural networks have turned out to be prime candidates for such models, and led to the famous Deep Q-Networks (DQN) (Mnih et al., 2015). In Q-learning, the Q-network $Q_\phi(s, a)$ is learned by minimizing the objective

$$\mathcal{L}(\theta) = \mathbb{E}_{(s', r, a, s)} \left[\left(\overbrace{r + \gamma \max_{a'} Q_\phi(s', a')}^{\text{Target } Y^Q} - Q_\phi(s, a) \right)^2 \right].$$

One aims at solving this problem because $r + \gamma \max_{a'} Q_\phi(s', a')$ provides a more reliable estimate to the future return than $Q_\phi(s, a)$. Indeed, it can be proven that the optimal policy satisfies $Q_*(s, a) = r + \gamma \max_{a'} Q_*(s', a')$ given any tuple (s, a, r, s') , and thus the loss quantifies how far one is from optimality.

³This is a big jump forward, but the interested reader should read the excellent introduction to reinforcement learning (Weng, 2018).

6.1. Reinforcement Learning to optimize sampling trajectories

However, this objective makes the training unstable and slow, the sequence of iterates (s, a, r, s') from a single episode are heavily correlated, and because the target Y_Q depends on Q_ϕ that is updated each iteration. As a result, Mnih et al. (2015) proposed two mechanisms to alleviate these issues. Experience replay is the idea that instead of computing updates directly on the sequence of iterates from the current episodes, the tuples (s, a, r, s') are stored in a buffer from which a sample is drawn at random to perform the update. The second mechanism is periodic updating, where the network used in the target has its weights frozen and only periodically updated.

While these mechanisms allowed to fix some large issues of Q-learning, DQN was found to still suffer from overestimating the value of state-action pairs, and Van Hasselt et al. (2016) proposed to use Double-Q learning as a way to address this issue. As a result, the objective looks like

$$\mathcal{L}(\theta) = \mathbb{E}_{(s', r, a, s) \sim U(D)} \left[\overbrace{(r + \gamma \max_{a'} Q_{\phi^-}(s', a') - Q_\phi(s, a))^2}^{\text{Target } Y^{\text{DoubleQ}}} \right].$$

The changes here are the ones from DQN: $U(D)$ is a uniform distribution over the replay memory, and ϕ^- refer to the weights of the frozen network. The main difference however is to be found in the target Y^{DoubleQ} . Double Q-learning is based on the observation that in the Q-learning target Y^Q , the same values are used to select an action and to evaluate making it more likely to select overestimated values (Hasselt, 2010). This can be made explicit as

$$Y^Q = r + \gamma Q_\phi(s', \arg \max_{a'} Q_\phi(s', a'))$$

In double Q-learning, two values of weights ϕ and ϕ' are used to decouple the action selection and evaluation, resulting in the target

$$Y^{\text{DoubleQ}} = r + \gamma Q_\phi(s', \arg \max_{a'} Q_{\phi'}(s', a')).$$

After the Q -function has been trained, inference on new environments is done by following greedily the policy prescribed by the Q -function, namely

$$a_t = \arg \max_a Q_\phi(s_t, a) \text{ for } t = 1, \dots, T.$$

This is the model that Pineda et al. (2020) used in their article. They propose two variants of the model, a subject-specific DDQN (SS-DDQN) as well as a dataset-specific DDQN. The difference lies in the input given to the Q -function. In the SS-DDQN, the neural network will take as input $\hat{\mathbf{x}}_t$ and ω_t and output a $|\mathcal{A}|$ dimensional vector, with the probability to take each action $a \in \mathcal{A}$ as the next move. On the contrary, the DS-DDQN will only take as input ω_t , and as a result, the policy will only depend on the current

mask and will *not* adapt to the current subject.

6.1.4 Policy Gradient

Bakker et al. (2020) took a different approach to the problem, called policy gradient. Contrarily to Q-learning, this approach aims at directly learning a policy, without first learning a value function. The policy itself is parametrized as a function $\pi_\phi(s) : \mathcal{S} \rightarrow \mathcal{A}$ that associates to each state the probability of taking an action $a \in \mathcal{A}$. It aims finding the set of parameters ϕ that maximize

$$\max_{\phi} \left\{ \mathcal{J}(\phi) = V_{\pi_\phi}[s_1] = \sum_a \pi_\phi(a|s_1) Q_{\pi_\phi}(s_1, a) \right\}.$$

In this case, as the policy is parametrized by ϕ and not Q as we previously had, the function $Q_{\pi_\phi}(s, a)$ must be *computed* rather than simply evaluated on the input. This is achieved by using the relation connecting the value function to the state-action function and the definition of the state-action function (cf. Equations 6.3 and 6.2).

The policy gradient theorem aims at providing an analytical expression for $\nabla_{\phi} \mathcal{J}(\phi)$, in order to optimize ϕ using a gradient-based approach. We will directly use its result and refer the interested reader to Sutton and Barto (2018, p.324ff). The policy gradient theorem yields

$$\nabla_{\phi} \mathcal{J}(\phi) \propto \mathbb{E}_{\pi_\phi} [G_t \nabla \log \pi_\phi(A_t|S_t)]. \quad (6.5)$$

This result is significant because it allows to estimate the gradient of the policy with an *expectation* over the return and the gradient of the logarithm of policy. The expectation means that the gradient of the policy is amenable to a Monte-Carlo approximation, where one takes samples of the policy to approximate the expectation.

This is the basis used by Bakker et al. (2020), but the result in their case is slightly modified to have

$$\nabla_{\phi} \mathcal{J}(\phi) \propto \mathbb{E}_{\pi_\phi, \mathbf{x}_0} \left[\sum_{t=1}^T \nabla \log \pi_\phi(A_t|S_t) \sum_{t'=t}^T \left(\gamma^{t'-t} (r(S_{t'}, A_{t'}) - b(S_{t'})) \right) \right] \quad (6.6)$$

where the expectation over \mathbf{x}_0 is the expectation over the initial states, and $b(S_{t'})$ is a baseline that is aimed at reducing the high variance of the gradient estimator (Sutton and Barto, 2018, pp. 329-331).

The expression in Equation 6.6 is not directly computable as it still contains an expectation over the policy, and A_t, S_t depend on it. Bakker et al. (2020) propose both a *greedy* and

6.1. Reinforcement Learning to optimize sampling trajectories

a *non-greedy* approach to compute a Monte-Carlo (MC) estimate. They read

$$\nabla_{\phi} J(\phi) \approx \frac{1}{q-1} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \sum_{i=1}^q \sum_{t=1}^T \left[\nabla_{\phi} \log \pi_{\phi}(\mathbf{x}_t) \sum_{t'=t}^T \gamma^{t'-t} \left(r_{i,t'} - \frac{1}{q} \sum_{j=1}^q r_{j,t'} \right) \right] \quad (\text{Non-greedy}), \quad (6.7)$$

$$\nabla_{\phi} J(\phi) \approx \frac{1}{q-1} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}} \sum_{i=1}^q \sum_{t=1}^T \left[\nabla_{\phi} \log \pi_{\phi}(\mathbf{x}_t) \left(r_{i,t} - \frac{1}{q} \sum_{j=1}^q r_{j,t} \right) \right] \quad (\text{Greedy}) \quad (6.8)$$

where $r_{i,t}$ denotes the reward for the i -th (out of q) MC sample at time step t . We see here that the non-greedy approach requires an additional sum in the expectation: for each time step t , a full rollout must be computed until the last acquisition step in order to do a single gradient update. On the contrary, a greedy approach is much faster, as it there is no sum inside the expectation, except the averaging over the q different MC samples taken.

Note that, in the non-greedy case, Bakker et al. (2020) only sample multiple trajectories at the initial state, as sampling multiple ones at each state would lead to a combinatorial explosion. After training, inference on new environments is done by simply sampling from the obtained policy, i.e. $p(a|s) = \pi_{\phi}(a|s)$.

6.1.5 LBCS as a simple RL policy

It should be clear now to the reader that LBCS can easily be viewed as a simple greedy policy. Instead of training a sophisticated deep neural network, it simply relies on training a non-adaptive policy that does

$$\pi(a|s_t) = \mathbb{1}_{a=a_t}, \text{ where } a_t = \arg \max_a \left\{ r(s_t, a) = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [\eta(\mathbf{x}, \hat{\mathbf{x}}_{\theta; \omega_{t+1}}) - \eta(\mathbf{x}, \hat{\mathbf{x}}_{\theta; \omega_t})] \right\}.$$

where the policy is deterministic given the current state, and does not adapt to the current patient as the choice is based on the expectation over the training data. In this case, $\pi(a|s_t)$ really becomes $\pi(a|\omega_t)$. The inference is then simply done by following the policy, i.e. $\pi(a|\omega_t) = \mathbb{1}_{a=a_t}$ for all $\mathbf{x} \sim p(\mathbf{x})$.

6.1.6 The questions at hand

After this primer into RL that introduced the different existing methods for RL-based mask design, we come back again to our initial question: What component of RL does matter most? How does patient-adaptivity help? What is the contribution in long-term planning? Are there other components of the MR pipeline that affect the performance of the method more than the policy?

As the works of Pineda et al. (2020) and Bakker et al. (2020) do not compare to each other, and work in very different settings, it is not possible to answer these questions directly from their respective works.

We chose to use LBCS (Gözcü et al., 2018) as starting point as it is clear that it implements a learning-based, data-driven policy that is *not* adaptive *nor* long-term planning. This method will provide the baseline from which we can quantify the improvements brought by adaptivity and long-term planning.

Secondly, as the methods of Bakker et al. (2020) and Pineda et al. (2020) are preprocessed fairly differently before being used to train RL policies, we wish to understand how these policies are sensitive to changes in parameters such as the training of the reconstruction algorithm, or changes in the image field of view, or the way that k-space is undersampled. In particular, this will help us understand whether the seemingly contradictory conclusions originate from different processing of the data, or whether they reflect something more fundamental. In the sequel, we will first define the processing pipeline precisely, before carrying thorough evaluations of these methods.

6.2 The MRI data processing pipeline

In this section, we will step through the stages of the MRI data processing pipeline as shown in Figure 6.2. Such considerations are not reflected in the mathematical model representing the acquisition as presented in Equation 2.12. Recall that we consider an inverse problem, where we seek to recover a signal $\mathbf{x} \in \mathbb{C}^P$ from partial observations $\mathbf{y} \in \mathbb{C}^N$, $N \ll P$ obtained by subsampling a Fourier transform $\mathbf{F} \in \mathbb{C}^{P \times P}$:

$$\mathbf{y}_\omega = \mathbf{P}_\omega \mathbf{F} \mathbf{x} + \boldsymbol{\epsilon}, \tag{6.9}$$

where $\boldsymbol{\epsilon} \in \mathbb{C}^N$ is the noise, $\omega \subseteq [P] := \{1, \dots, P\}$ is the index of sampling locations, with cardinality N , and \mathbf{P}_ω is a diagonal matrix such that $(\mathbf{P}_\omega)_{ii} = 1$ if $i \in \omega$, 0 otherwise. This problem is inherently ill-posed, due to $N \ll P$, and we first construct an estimate $\hat{\mathbf{x}}_{\omega;\theta} = f_\theta(\mathbf{y}_\omega, \omega)$ of the original signal \mathbf{x} , where f_θ is a reconstruction method parametrized by θ .

We turn now to the description of the practical pipeline, which is independent of the sampling method and simply operates on a given mask. We will also highlight a few common caveats.

As we will show in Section 6.3.2, seemingly trivial changes at even a single of these stages can affect the results of the evaluation and lead to changes in the ordering of the performance of different policies, an observation that we will refer to as *reversal*. We already stress that the main reason for such reversals to occur is that the RL policies generally do not generally bring a large or consistent improvement over sLBCS. We mark

6.2. The MRI data processing pipeline

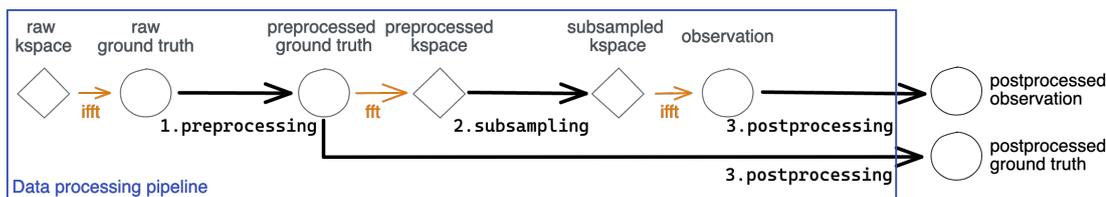


Figure 6.2 – Illustration of the data processing pipeline of MRI subsampling. Diamonds represent data in Fourier domain (k -space) and circles represent data in image domain. The postprocessed observation and ground truth are the data that are subsequently used for training the reconstruction and policy models. The pipeline features three main blocks, namely preprocessing, subsampling and postprocessing.

sections where we observed changes leading to reversals with * and sections here they only shifted the results but do not lead to reversals with *italics*.

Data sources and storage. In all cases, the ground truth signal \mathbf{x} is initially acquired as a complex signal in k -space, generally using multiple coils. The fastMRI dataset (Zbontar et al., 2019) provides this raw multicoil data, as well as a simulated single-coil k -space which we use throughout this work.

Preprocessing*. Due to computational constraints, it is common to resize the images by cropping and/or resizing or use magnitude images over the raw data. Cropping and resizing changes the ground truth distribution, which as seen in Table 6.3 can also lead to reversals in the final results.

Caveat #1. Converting the raw ground truth to or starting from a magnitude induces a conjugate symmetry in Fourier space that is not present in real data, a distortion that makes our modeling less faithful to the physical model.

Sampling. While in the real world the data is actually acquired by sampling k -space (prospective sampling), in practice, the acquisition is generally simulated by retrospectively undersampling fully sampled Cartesian data following Equation 2.12.

When training reconstructor and sampling method separately, training data are generally constructed with random masks that sample a certain fraction of center frequencies, and then the rest from a random distribution. There has been considerable variation in defining the parameters of these distributions, but to this day, no systematic study of their effect on the reconstruction quality have been carried out.

Post processing and reconstruction. After sampling, the observation will be processed according to the implementation details of the reconstruction algorithm (e.g., normalized or standardized) and the reconstruction is computed.

Typical normalization, as used for instance in Pineda et al. (2020), consist of dividing

Chapter 6. On the benefits of deep RL in accelerated MRI sampling

every input by a fixed value equal to the average energy of the dataset. Bakker et al. (2020) standardized their data and then clamped them to have standardized data in the range $[-6, 6]$.

We found that standardization used by Bakker et al. (2020) was necessary only when paired with non-residual models, which tracks with observations in the literature that residual networks can work without normalization, although usually this is studied for network internal normalization which we also leverage (see e.g. Zhang et al. (2019a) which removes normalization entirely by using a proper initialization).

The results of Appendix A.2.1 illustrate the impact of postprocessing.

Caveat #2. Normalization should occur consistently between ground truth and observation data, as failing to do this can lead to inconsistent values in the reconstruction, especially in parts where data consistency occurs. In Bakker et al. (2020), the authors normalize observations and ground truth using their respective statistics and denormalize the reconstructed image using the *ground truth* statistics, which is not compatible with the use of data consistency in the reconstructor. If one wishes to get a realistic estimate of the performance at deployment, it is also advisable to use only statistics available at test time (i.e. observation statistics, not ground truth statistics).

Evaluation metrics*. Finally, to judge the results, the most common metrics are peak signal-to-noise ratio (PSNR), mean squared error (MSE) and structural similarity index metric (SSIM).

All metrics tend to promote smooth images when used as a training loss (Muckley et al., 2021), but it is widely known that MSE and PSNR focus on low-frequencies/high-energy components, which was also noted by Bakker et al. (2020). However, while MSE and PSNR only operate on image differences and treats every pixel independently, SSIM is computed from local averages (Wang et al., 2004; Zbontar et al., 2019).

The metric is then reported as a curve plotted against sampling rate (Zhang et al., 2019b), its reciprocal, the acceleration rate (Pineda et al., 2020) or aggregated by computing the average performance, the performance at end of sampling, or the area under curve (AUC).

As discussed throughout the experiments, the choice of metric and its aggregation can greatly impact the conclusions drawn from results.

Remark (Area under curve (AUC) computation). We detail here the computation used to summarize performance curves with the AUC, used in Tables 6.3 and 6.4.

The area under curve is a numerical integration of the performance curve of the form $\{r_t, \ell(\mathbf{x}_i, \hat{\mathbf{x}}_\theta(\mathbf{y}_{\omega_{t,i}}))\}_{t=t_0}^T$ that relates the sampling rate or acceleration factor at the t -th step to its performance. We used the `sklearn` implementation, that implements it using a trapezoidal rule. We compute an individual AUC for each test sample and then aggregate the resulting set $\{AUC_i\}_{i=1}^{n_{\text{test}}}$ by computing its empirical mean and variance.

This measure is susceptible to changes in reparation of the area under the curve, and this is the reason that the result changes when moving from sampling rate to acceleration factor (1/sampling rate). Indeed, when representing acceleration factors, in the case of Pineda et al. (2020) (cf. Table 6.4) 20 out of 332 sampling locations cover 90% of the acceleration factor plot, biasing the AUC towards attributing most of its weight the high acceleration factors.

6.3 Re-examining deep RL for MRI sampling

Our initial impetus is the observation that the non-adaptive oracle used in Bakker et al. (2020) is highly reminiscent of the LBCS method Gözcü et al. (2018), which has been used as a strong baseline in the literature (Jin et al., 2019; Sanchez et al., 2020b). To our surprise, evaluating these methods against the greedy RL method in the data processing pipeline used in Zhang et al. (2019b) resulted in the fixed method not only coming close to the non-adaptive oracle, but actually closely matching the RL method (see Table 6.3). Wanting to perform a fair comparison as well as to assess the impact of the variability in the pipelines used across the literature, we closely replicated the pipeline of Bakker et al. (2020)⁴ as well as an extensive ablation study to understand this reversal and identify its source.

6.3.1 Experimental setting

In the sequel, the term *setting* then refers to a particular choice of preprocessing, subsampling and postprocessing of the data, independently of the data source (we always use the same dataset) and the model or sampling policy used. Preprocessing and postprocessing are relevant throughout the training of the reconstruction and the policy models, whereas changes in the sampling masks only directly affect the pretraining of the reconstructor, as the training the policy is done by successive rollouts using the policy model itself.

Dataset and preprocessing. Like the original paper of Bakker et al. (2020), we used

⁴We thank the authors for providing the original checkpoints and general help and responsiveness during this replication. Our results did not show the *exact* same numbers as Bakker et al. (2020), as we did not use the same train-val-test split and dataset undersampling. We discuss this further in Appendix A.2.1.

the fastMRI (Zbontar et al., 2019) single-coil knee dataset for the experiments. We slightly modify follow their preprocessing (using complex data instead of magnitude data, different data normalization) that do not affect the relative ordering of the different methods. We provide ablations and a detailed discussion of these changes in Appendix A.2.1.

Specifically, they *crop* the data to 128×128 , which we refer to as (c)). Zhang et al. (2019b) instead *resizes* the data to 128×128 which we replicate by first cropping to 256×256 and then resizing them to 128×128 . This preprocessing results in images with different fields of view. We refer to this preprocessing as (c+r) and integrate it to our ablations. We also evaluate horizontal (h) sampling masks, used in Gözcü et al. (2018); Jin et al. (2019) in addition to the vertical (v) sampling masks used by Bakker et al. (2020); Zhang et al. (2019b) and Pineda et al. (2020).

The deep reconstructors used in Pineda et al. (2020); Bakker et al. (2020) pre-trained by randomly sampling a mask from a set of distributions with different parameters. We ablated over two pretraining regimes which we abbreviate as b and z, respectively. We describe them in more detail in Appendix A.1

Reconstruction models. We ablate over the two reconstruction models used in the RL SotA works of Pineda et al. (2020) and Bakker et al. (2020), namely a cascade of residual networks (*cResNet*) architecture from Zhang et al. (2019b), used in Pineda et al. (2020) and the *UNet* baseline provided in the fastMRI dataset, used in Bakker et al. (2020). The training details and masks used for the training of the reconstruction models are described in Appendix A.1.

For the UNet, like Bakker et al. (2020) we take a single channel for magnitude only data as input, while for the cResNet we take the 2 complex channels as inputs. All models are trained using ℓ_1 loss and directly output the reconstruction without a final nonlinearity. The UNet has 837 635 parameters in total, while the cResNet is larger with 1 093 479 parameters.

Sampling methods In each ablation setting we compared the method of Bakker et al. (2020) to the following baselines and oracles (citations at the end of each item refer to prior works that also evaluate them):

- Random sampling (**Random**): Acquire a fixed proportion of low-frequency lines in Fourier and then randomly sample the remaining lines (Jin et al., 2019; Pineda et al., 2020; Bakker et al., 2020).
- Low-to-high frequencies (**LtH**): select k-space lines from low-to-high frequencies lines (Zhang et al., 2019b; Pineda et al., 2020; Jin et al., 2019).
- (Stochastic) Learning-based Compressive Sampling (**LBCS**) (Gözcü et al., 2018; Sanchez et al., 2020a): This method trains a non-adaptive, greedy sampling policy that selects as a measurement candidate in each acquisition step the column that

6.3. Re-examining deep RL for MRI sampling

leads to the greatest average improvement over a sample from the training dataset. We use the stochastic version that scales better to large dataset and images (Jin et al., 2019).

- **Non-adaptive Oracle (NA Oracle)** (Bakker et al., 2020): This oracle is computed by training and evaluating LBCS directly on the test set, and can serve to illustrate the benefit of adaptivity in greedy methods. This is the instance of a non-adaptive, greedy sampling method (Bakker et al., 2020).

For the training of the policy models of Bakker et al. (2020), we use the parameters of the greedy model in their paper, which we refer to simply as **RL** in the sequel. We excluded the non-greedy version, as Bakker et al. (2020) notice that the performance of the non-greedy model with discount factor $\gamma = 0.9$ is always close to one standard deviation of the greedy model, but significantly more computationally demanding.

We studied both the short and long horizon sampling regimes but only report on the long horizon for conciseness, with the short horizon results summarized in Appendix A.2.1. Except for the deterministic LtH and NA Oracle, we report the performance of each method averaged on three runs/separately trained RL policies, along with the standard deviation.

Summarizing, we start in the setting of Bakker et al. (2020) except for using complex data instead of magnitude data and then ablate over: **i**) cropped (**c**) vs cropped+resized (**c+r**), **ii**) vertical (**v**) vs horizontal (**h**) Cartesian sampling, **iii**) reconstruction with a pretrained UNet (Ronneberger et al., 2015; Bakker et al., 2020) or a cResNet (Zhang et al., 2019b) and **iv**) the training regime for the reconstructor proposed in Bakker et al. (2020) (**b**) or in Zhang et al. (2019b) (**z**). The base setting is referred to as **cvb**, for cropped (preprocessing) + vertical (mask) + Bakker et al. (sampling parameters) using a both a UNet and a cResNet reconstructor.

6.3.2 Main results on Bakker et al. (2020)

We first evaluated the method of Bakker et al. (2020) on two settings, namely **cvb** and **c+rhz**, and report it with two different aggregations. Table 6.2 shows the SSIM of the policies at the final sampling rate considered (25%), as used in as used in Bakker et al. (2020), whereas Table 6.3 has the SSIM aggregated over the whole acquisition trajectory by computing its AUC⁵.

Several observations can be made from these tables regarding *reversals*. Focusing first on Table 6.2 only, we can see that although the first column, corresponding to the processing done by Bakker et al. (2020), supports that RL outperforms LBCS, changing architecture

⁵A more complete ablation, using the setting **cvz** to ablate over the impact of different mask parameters, **c+rvz** to study the impact of different field of views and mask orientation is carried out in Appendix A.2.3.

Chapter 6. On the benefits of deep RL in accelerated MRI sampling

(second column) already invalidates these results, with LBCS matching RL. But the change is even more significant by changing *settings*. By moving from images where the field of view is cropped and undersampling occurs vertically (1st and 2nd column) to a different field of view with horizontal undersampling (3rd and 4th column), the relative ordering of LBCS and RL is *reversed*, and LBCS shows the best performance.

However, if we include Table 6.3 into the picture, then even the previous conclusions do not hold: on the `c+rhz` setting (3rd and 4th columns), RL matches or outperforms LBCS. *Conclusions drawn from the results can be reversed or invalidated by using a different way to aggregate the results or by changing the setting on which the evaluation is carried out.*

Although contradicting conclusions can be drawn by looking at individual results, a consistent trend from Tables 6.2 and 6.3, as well as the ablations in Appendix A.2.3, is that *the return on investment (ROI) of adopting RL over LBCS is generally marginal compared to what changes in the modeling pipeline can bring.*

We see for instance that improving the reconstruction architecture yields much more significant quantitative gains. Moving from a UNet to a cResNet brings an order of magnitude more improvement (around 0.01 SSIM difference) than what RL brings over LBCS (around 0.0015 at best).

We see also in Appendix A.2.3 that the distribution of masks used to pretrain the reconstruction algorithm has a similar impact over performance. The `b` masks of Bakker et al. (2020) are discretely distributed from sampling rates 12.5% to 25%, which matches exactly the short horizon experiment range. On the contrary the `z` masks of Zhang et al. (2019b) are distributed from 7% to 37.5%, whereas the long horizon experiment spans a range from 3% to 25% sampling rate. As a result, we observe that `b` masks consistently perform better than `z` in the short horizon experiment, whereas the opposite holds in the long horizon experiments.

As an example, in the long horizon case, using `z` masks over `b` masks brings a consistent improvement of ~ 0.004 in the SSIM AUC, whereas RL improves over LBCS by ~ 0.0006 in the SSIM AUC.

Main conclusion. These results suggest that steps such as matching the pretraining and evaluation mask distributions and using strong reconstructors has a significantly larger influence and consistent effect on the performance than moving from a greedy, fixed policy (LBCS) to a greedy, adaptive one (RL).

Additional observations. We also see that the performance of LBCS always remains close to the NA Oracle, testifying to the generalization ability of the fixed LBCS masks as indicated by theory. Other reversals can be observed in both cases when changing the data processing for both the SSIM at 25% and the AUC evaluations.

6.3. Re-examining deep RL for MRI sampling

To conclude, note that the large gap in performance between the `cvb` and the `c+rhz` settings is not due to a significant difference in performance of the reconstruction methods, but rather originates from the difference in field of view. This can be seen by comparing the SSIM value of the observations (before reconstruction) using the deterministic LtH policy in both cases: in the `cvb` setting, one gets an AUC of 0.4989 for the observation SSIM, whereas in the `c+rhz` setting, the AUC is 0.6534.

Policy	cropped, vert., Bakker		cropped+resized, horiz., Zhang	
	UNet	cResNet	UNet	cResNet
Random	0.5249 ± 0.0001	0.5432 ± 0.0004	0.6567 ± 0.0003	0.6725 ± 0.0006
LtH	0.5832	0.6197	0.7325	0.7714
LBCS	0.6294 ± 0.0009	0.6417 ± 0.0011	0.7768 ± 0.0000	0.7941 ± 0.0000
RL	0.6298 ± 0.0002	0.6415 ± 0.0002	0.7761 ± 0.0000	0.7935 ± 0.0001
NA Oracle	0.6301	0.6428	0.7771	0.7942

Table 6.2 – SSIM at 25% on the test set, on knee data comparing two models in the long horizon setting of Bakker et al. (2020), for the `cvb` (cropped, vertical, Bakker-type of masks) and `c+rhz` (cropped+resized, horizontal, Zhang-like masks) settings. A different aggregation of the results is shown on Table 6.3.

Policy	cropped, vert., Bakker		cropped+resized, horiz., Zhang	
	UNet	cResNet	UNet	cResNet
Random	0.4348 ± 0.0001	0.4432 ± 0.0002	0.5860 ± 0.0002	0.5885 ± 0.0002
LtH	0.4849	0.5107	0.6678	0.6902
LBCS	0.5134 ± 0.0004	0.5243 ± 0.0003	0.7035 ± 0.0001	0.7096 ± 0.0001
RL	0.5142 ± 0.0001	0.5242 ± 0.0002	0.7035 ± 0.0002	0.7111 ± 0.0009
NA Oracle	0.5140	0.5247	0.7038	0.7099

Table 6.3 – AUC on the test set using SSIM, on knee data comparing two models in the long horizon setting of Bakker et al. (2020), for the `cvb` (cropped, vertical, Bakker-type of masks) and `c+rhz` (cropped+resized, horizontal, Zhang-like masks) settings. A different aggregation of the results is shown on Table 6.2.

6.3.3 Main results on Pineda et al. (2020)

Given these results, we can observe that if there is indeed a benefit to patient-adaptive greedy policies, it is highly sensitive to the parameters used, and the improvement over the non-adaptive greedy baseline is marginal. One might wonder whether there will be a significant gain from adaptive RL policies if they are trained to perform long term planning on a longer horizon? To investigate this, we replicated the second SotA RL method described in Pineda et al. (2020).

Experimental setting. We used the pretrained checkpoints from Pineda et al. (2020)⁶, and compared to their Subject-Specific DDQN (SS-DDQN) as well as their Dataset-Specific DDQN (DS-DDQN).

This work uses a small discount factor of $\gamma = 0.5$ and operates on the full fastMRI image of size 640×332 , meaning it is both trained for a longer time horizon as well as having a lot more leeway for decision making.

We replicated the *extreme* setting which starts with 2 center frequencies (0.6% undersampling rate or $166\times$ acceleration) and is evaluated up to 100 frequencies (30.1% undersampling rate or $3.32\times$ acceleration)⁷. As can be seen in Table 6.4, the LBCS mask is very close to the performance of both the data and sample specific DDQN policies provided by Pineda et al. (2020).

Results. We see on Table 6.4 and Figure 6.3 other instances of reversals. On Table 6.4, for a given metric, different ways of aggregating the results give rise to *reversals* in interpretation. This is also the case when moving from SSIM to PSNR. Although there is a clear performance gain over Random and LtH masks, the gain from LBCS to DDQN is not consistent.

This discrepancy can be explained by the fact that, as displayed on Figure 6.3, the acceleration factor puts more emphasis on the sub 5% range (i.e. below 17 lines) where LBCS has not yet caught up on DDQN. This results in the AUC weighting disproportionately high acceleration factors: $\sim 90\%$ of the plot and of the corresponding AUC consists of 20 lines out of the 100 acquired. Similarly, considering the final sampling rate ignores the performance of the method throughout the acquisition procedure.

We also performed a more detailed comparison of the masks given by the policies, using a subset of the first 200 test set images, keeping the order fixed in across methods, which can be found in Appendix A.3. They can be summarized as follows: **i)** the adaptivity of SS-DDQN mainly affects the ordering of frequencies, a large section of acquired frequencies being shared across samples and in similar regions to the LBCS mask at the final sampling rate, and **ii)** it confirms that the adaptive masks have a small edge only until about a sampling rate of 5%, after which LBCS catches up and overtakes the RL policy.

6.4 Discussion

As we have seen in Section 6.3.2 and Section 6.3.3, the fixed mask of LBCS is competitive with or outperforms both the greedy policy gradient methods of Bakker et al. (2020) and

⁶The models can be found at <https://facebookresearch.github.io/active-mri-acquisition/misc.html>

⁷We thank the authors for providing us with the original scores and general responsiveness and helpfulness during this replication.

Policy	SSIM			PSNR		
	Samp. rate	Acc. factor	Final rate	Samp. rate	Acc. factor	Final rate
Random	0.5801	0.4497	0.6723	26.489	22.327	28.962
LtH	0.5636	0.4506	0.6686	27.169	23.133	29.360
LBCS	0.6079	0.4787	0.6886	28.491	23.799	30.211
DS-DDQN	0.6101	0.4797	0.6855	28.240	23.978	29.652
SS-DDQN	0.6139	0.4797	0.6882	28.424	23.918	29.929
Adaptive Oracle	0.6341	0.4910	0.7131	29.013	24.498	30.683

Table 6.4 – AUC on the test set when calculated against sampling rate and acceleration factor ($1/\text{sampling rate}$), as well as performance at the final sampling rate (100 lines acquired out of 332) on the knee dataset, using the processing of Pineda et al. (2020).

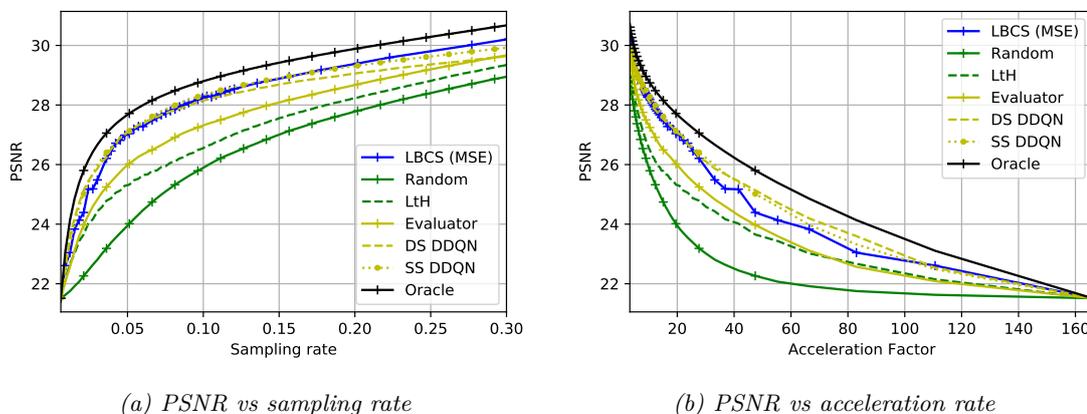


Figure 6.3 – PSNR performance plot on the test set, using the knee dataset with the processing from Pineda et al. (2020). The plots feature two ways to report the same result, which are also displayed in Table 6.4. The SSIM performance plot can be found in Figure A.3 in the Appendix.

the non-greedy DDQN methods of Pineda et al. (2020). Our ablations in Section 6.3.2 and Appendices A.2.3 and A.4, as well as the example of Figure 6.3 highlight that the benefit of current RL methods over the fixed baseline, if it exists at all, is so small that a change of field of view, architecture or mask distribution (cf. Tables 6.3 and A.3) can lead to a reversal in the relative performance ordering.

Note also that the computational resources required to train a deep RL model are much more significant than what is required to obtain a mask through sLBCS. sLBCS does not require to train a deep model and as a result does not need to perform any backpropagation. In the setting of Bakker et al. (2020), this results in being able to train the mask using 20 GB of RAM on a A-100 GPU in slightly less than 20 minutes. With the same computational budget, the greedy deep RL method of (Bakker et al., 2020), which is trained significantly more quickly than longer horizon policies, takes roughly 25 minutes for a single epoch out of 50. However, both methods could be accelerated by optimizing the data processing computations, which take a large part of the computation time in both cases. In the setting of Pineda et al. (2020), the comparison is even more

striking. In communication with the authors, we learned that it took them more than 20 days to train either of the DDQN models, while the training of sLBCS takes 30 minutes when distributed over 5 A-100 GPUs. This points to two main advantages of this testifies to the main advantage of sLBCS over RL-based approaches. Being a fixed policy, sLBCS does not require to do multiple iterations in order to converge, but is a one-shot approach. In addition, sLBCS can be *heavily parallelized*, while deep RL-based approaches are inherently sequential, and result in a slow training procedure.

Table 6.4 confirms the expectation that methods trained on SSIM can underperform with respect to PSNR and vice versa. In general agreement with the literature, we see that using several metrics is important to capture these trade-offs, as reporting only one particular metric might hide underperformance in the other.

There have been various ways of reporting the metrics, ranging from showing sampling curves as in Figure A.3a, reporting the area under curve (AUC) tables or simply reporting the metric at the final sampling rate. We see that reporting only the summarizing statistic can be deceitful and give relatively little information on the overall performance of the method. For instance, using only the last sampling rate, one could not distinguish between a method that performs well at all sampling rates leading to the final one, and one that has overall mediocre results and quickly improves on the performance at the end.

For this reason, we consider sampling curves the gold standard for assessing the quality of accelerated MRI sampling methods, as a good performance at all sampling rates coupled with some stopping criterion could mean that MRI scans could be stopped when sufficient information has been acquired, further increasing their efficiency. If one wishes to report a single number, we recommend using AUC unless only the final performance is of interest.

There is also the question of reporting acceleration factors or sampling rates, which put focus on very different regimes of sampling: acceleration factors, as shown in Pineda et al. (2020) focus on extreme undersampling rates. While is true that the high acceleration regimes are where it is most desirable to make improvements, results on sampling rates have the advantage of uniformly distributing the performance throughout a range of interest instead of a single acquisition determining a third of the metric as in Figure 6.3b.

Alternatives to using acceleration factors could consist in displaying sampling rates over a region of interest or computing the AUC by explicitly assigning more weight to low sampling rates instead of doing it implicitly and nonlinearly through acceleration factors.

Finally, a close study of the masks given by the SS-DDQN in Figure A.2, the edge of the SS-DDQN over the LBCS mask in Figure A.1a and the comparison of the LBCS and SSDQN edge over a fixed heuristic in Figures A.1b and A.1c reveals that the frequencies selected by the adaptive policy are very similar to those selected by the LBCS mask and

that the variability between samples is concentrated at very early sampling rates.

6.5 Relation to other works

While we focused strictly on RL trained sampling policies leveraging pre-trained reconstructors, there are other paradigms for accelerated MRI sampling present in the literature.

The first line is that of using end-to-end training of masks using stochastic relaxations of the sampling mask to be able to differentiate it and optimize the reconstructor and sampling method jointly. Bahadir et al. (2019); Huijben et al. (2020) directly optimize for a single sampling rate and sampling mask. More recent work extend this technique for jointly training an adaptive policy, doing away with the pretraining and making more efficient use of the capacity of the reconstructor (Yin et al., 2021; Van Gorp et al., 2021). Another line of joint training relies on self-supervised learning using Monte-Carlo Tree Search, as was done in Deepmind’s AlphaGo (Silver et al., 2017; Jin et al., 2019)

An open question regarding this line of research is whether the joint training enables to achieve a better sampling policy or simply serves as a curriculum and a way to specialize the sampling policy and reconstructor onto each other. To our knowledge, neither Van Gorp et al. (2021) nor Yin et al. (2021) investigated this. An interesting experiment could be to train a LBCS sampling mask on the co-trained reconstructor to see whether it can recover the sampling performance of the co-trained sampling network.

More similar to the spirit of our work, the recent study of Shimron et al. (2021) tackles subtle biases induced by the use of stored data and improper processing, Edupuganti et al. (2020) investigated specifically uncertainty methods for MRI reconstruction.

For works focusing on RL, Henderson et al. (2018) performed extensive ablation studies showing the sensitivity of RL methods to even minute variations in parameters. Engstrom et al. (2020); Gronauer et al. (2021) showed that the improvements of SotA RL methods can be traced to the exploitation of a subset of implementation details and algorithmic improvements and that even simple algorithms leveraging this subset can achieve SotA performance, similar to our work. This showcases that adapting to the data distribution seems to be sufficient to reach SotA in accelerated MRI sampling.

6.6 Conclusion

Taken together, our observations lead us to conclude the apparent conflict between the works of Bakker et al. (2020) and Pineda et al. (2020) is simply because at least *in their current state*, neither of the RL SotA methods do not offer significant benefits over fixed, greedily trained masks. Since greedy algorithms tend to perform near optimal

in settings with submodularity (Krause and Golovin, 2014), we conjecture a similar structure might be present in problems like the MRI sampling problem. Determining whether this conjecture holds or whether the lack of added value originates from specific RL algorithm and their training remains to be determined and should be the focus for any researcher set on applying RL to such a problem.

Recommendations. Our results also enable us to provide practical advice, summarized below. We provide a more extensive discussion of these statements in Appendix A.6.

- Focus on improvements in the reconstructor architecture, mask distribution and algorithms used for training the reconstructor.
- Compare against strong baselines, such as LBCS.
- Show sampling curves and use AUC to aggregate your results instead of performance at the final sampling rate.
- Be mindful about preprocessing settings when evaluating a policy model. We recommend using the cropped+vertical setting with the data normalization implemented by Zbontar et al. (2019).

We also want to emphasize that conducting the experiments in this paper would have been impossible if the RL methods had not been exemplary in terms of openness and reproducibility. Without access to the checkpoints and code, and without the authors' responsiveness we would not have been able to reproduce both works and add the missing baseline. Despite the theoretical guarantees of LBCS, we were surprised that it matched and sometimes simply outperformed more sophisticated methods. We therefore do not view our work as criticism of these works but rather as an extension and a synthesis, and urge any future work to follow their lead in publishing codes, checkpoints and data.

Bibliographic note

The idea of investigating the *limitations* of current deep RL methods originated from discussions with Igor Krawczuk. He also contributed the experiments on the setting of Pineda et al. (2020).

7 Uncertainty driven adaptive sampling via GANs

In the previous chapter, we considered methods to optimize sampling based on deep reinforcement learning, and showed that in their current state, they fail to deliver on their promises of improving performance by doing long-term planning and adapting to the patient.

In this chapter, we take a step back from the question of designing the best sampling policy, and primarily aim at exploring how conditional Generative Adversarial Networks (GANs) (Goodfellow et al., 2014; Mirza and Osindero, 2014) can be used to model inverse problems in a Bayesian fashion. We are particularly interested on how they can be used to learn the posterior $p(\mathbf{x}|\mathbf{y}_\omega)$, namely the distribution of ground truth images given observations.

In the context of inverse problems, GANs have mainly been used either as reconstruction models (Yang et al., 2018; Chen et al., 2022) or as generative priors (Bora et al., 2017; Jalal et al., 2021), as discussed in Section 3.2.5. However, the work of Adler and Öktem (2018a) showed that GANs trained to approximate the posterior can learn to act as reconstructors and simultaneously provide uncertainty quantification (UQ).

We argue that there is even more that they can do when trained to approximate $p(\mathbf{x}|\mathbf{y}_\omega)$. In this work, we show how conditional GANs naturally provide a criterion for adaptive sampling. Indeed, by sequentially observation the location with the largest posterior variance in the measurement domain, the GAN provides an adaptive sampling policy *without ever being trained for it*. In addition, our GAN-based sampling policy can easily be adapted to provide sampling in other measurement domains, contrarily to competing methods that are designed with Fourier sampling explicitly in mind (Zhang et al., 2019b).

In particular, we show that when considering image domain sampling, our GAN-based policy can strongly outperform the non-adaptive greedy policy from sLBCS. However, in Fourier domain, the GAN-based policy does not match the performance of sLBCS. Indeed, in image domain, it seems that there is a greater variability in the location

of the structure of interest, and as a result, adaptivity becomes ever more important. On the other side, in the previous chapter, we saw that adaptivity does not necessarily bring much in the Fourier domain. Nonetheless, we provide a more comprehensive explanation to the underperformance of our policy in Fourier, rooted in the concept of the *information horizon* used to make a decision at each step. We show that our model uses *less* information than LBCS to design its policy, and argue that this is the reason leading to its inferior performance. However, when compared to models that use the same amount of information to inform their policy, such as Zhang et al. (2019b), our model shines and largely outperforms its competition.

Overall, this work proposes an all-in-one approach for reconstruction, uncertainty quantification (UQ) and sampling, and shows the importance of incorporating relevant *information* into the policy for it to achieve a strong empirical performance.

The plan of this chapter is the following. In Section 7.1, we will first provide an introduction about GANs, and show how they can be used to model the posterior of inverse problems (Section 7.1.3). We will then show in Section 7.2 how this posterior can be used to perform adaptive sampling; this is our main contribution. We will then briefly discuss the paper of Zhang et al. (2019b) in Section 7.3, as they are the most conceptually similar work to ours, before moving to our experiments in Section 7.4. There, we validate that using the maximum posterior variance as a policy, GANs can be used to perform adaptive sampling in Fourier as well as image domains. We then briefly review relevant literature (Section 7.5), discuss our results (Section 7.6) before concluding (Section 7.7)¹.

7.1 Generative adversarial networks (GANs) for posterior modeling

7.1.1 Generative adversarial networks (GANs)

We will begin by describing what GANs are and how they work. GANs are *generative models*: given training data $\{\mathbf{x}_1, \dots, \mathbf{x}_m\} \sim p(\mathbf{x})$, they aim at learning to generate new samples that look like they have been generated by the distribution $p(\mathbf{x})$. Several deep learning approaches have been explored to tackle the problem of generation, which we will discuss in Section 7.5, and GANs have shined in particular by their ability to handle multi-modal outputs (Goodfellow, 2016).

GANs tackle the problem of generation as a two-player game, where two networks have

¹The work in this chapter is based on the following preprint and workshop paper: Sanchez, T., Krawczuk I., Sun, Z. and Cevher V. (2019). Closed loop deep Bayesian inversion: Uncertainty driven acquisition for fast MRI. *Preprint available at* <https://openreview.net/pdf?id=BJIPOIBKDB>. Sanchez, T., Krawczuk, I., Sun, Z. and Cevher V. (2020). Uncertainty-Driven Adaptive Sampling via GANs. In *NeurIPS 2020 Workshop on Deep Learning and Inverse Problems*. Available at <https://openreview.net/pdf?id=IWLYCQmtvW>.

7.1. Generative adversarial networks (GANs) for posterior modeling

opposite goals. The *generator* tries to generate candidate samples that resemble the true distribution $p(\mathbf{x})$, whereas the *discriminator* aims at distinguishing real samples from generated ones. The generator network learns to map data from a latent space to the data distribution, and is described as $\mathbf{g}_\theta : \mathbb{R}^L \rightarrow \mathbb{R}^P$. The discriminator performs a classification task, receiving a sample and outputting the probability that it has to be real, and so is described as $d_\phi : \mathbb{R}^P \rightarrow \mathbb{R}$.

The two-player game can then be formally framed as (Goodfellow et al., 2014):

$$\min_{\theta} \max_{\phi} \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [\log (d_{\phi}(\mathbf{x}))] + \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} [\log (1 - d_{\phi}(\mathbf{g}_{\theta}(\mathbf{z})))] \quad (7.1)$$

where $\mathbf{z} \sim p(\mathbf{z})$ is a random variable in the latent space follow a simple distribution such as a multivariate Gaussian $p(\mathbf{z}) = \mathcal{N}(\mathbf{z}; \mathbf{0}, \mathbf{I})$. The objective is simply a cross-entropy term for binary classification. The model is called *adversarial* because \mathbf{g}_θ and d_ϕ aim at achieving opposite goals, they are adversaries in the game.

In practice, the models are trained by minimizing Equation 7.1 in an alternating fashion, where d_ϕ is optimized for k steps and then \mathbf{g}_θ is optimized for a single step. This is due to the fact that the maximization over ϕ happens in the inner problem, and in theory, this would imply that each time that \mathbf{g}_θ change, d_ϕ should be trained until convergence. As this would incur prohibitive costs, the optimization is done over k steps only.

GANs have had tremendous practical success, and manage to generate photorealistic images (Karras et al., 2019) or transfer artistic styles (Isola et al., 2017) among other things. But their application are extremely widespread (Goodfellow, 2016; Pan et al., 2019; Kazemina et al., 2020).

7.1.2 Conditional GANs

However, in the present case, we are not interested in *unconditional* generation, where the model directly learns to generate samples from the data distribution $p(\mathbf{x})$. We rather want to be able to approximate *conditional* distributions of the form $p(\mathbf{x}|\mathbf{y})$, where we want to condition the data distribution on some specific observation \mathbf{y} .

The first conditional GANs aimed at doing class-conditional generation (Mirza and Osindero, 2014), where the model should learn to generate images from the data distribution that are consistent with the class given as input. For instance, if the model aimed at learning to represent cats and dogs, conditional generation would be the ability to integrate class information in the generative model in order to obtain samples from different classes on demand, i.e. have a generator of the form $\mathbf{g}_\theta(\mathbf{z}, c)$, where c is the label of the class of interest.

However, in the case of MRI and inverse problems, we are more interested in image-

conditional generation, where the generator is conditioned on a partial observation, and aims at generating a full image that is both consistent with the partial observation and likely to originate from the data distribution $p(\mathbf{x})$. This was explored later in tasks such as multimodal image-to-image translation, where a model aims at translating a same image between two domain (e.g. between a sketch and a realistic picture) (Zhu et al., 2017). Particularly relevant to us is the work of Adler and Öktem (2018a), which explicitly shows that learning a conditional Wasserstein GANs (WGANs) (Arjovsky et al., 2017) in an inverse problem amounts to approximating the posterior $p(\mathbf{x}|\mathbf{y})$. We discuss this work in greater depth as they were, until recently (Kovachki et al., 2020), the only work that formally connected conditional WGANs to learning the posterior in the context of inverse problems. Similarly, (Belghazi et al., 2019) show that a conditional GAN can serve to learn an exponential amount of conditionals, although without explicitly connecting it to learning a posterior distribution.

Conditional Wasserstein GANs

The Wasserstein distance, also known as the Earth-Mover distance, is defined as a distance between two probability distributions p and q

$$\mathcal{W}(p, q) \triangleq \inf_{\gamma \in \Pi(p, q)} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \gamma} [\|\mathbf{x} - \mathbf{y}\|] \quad (7.2)$$

where $\Pi(p, q)$ denotes the set of all joint distributions $\gamma(\mathbf{x}, \mathbf{y})$ that respectively marginalize to p and q . In Wasserstein GANs (WGANs), one aims to find the distribution \mathbf{g}_θ that minimizes the Wasserstein distance with the reference distribution p . Although Equation 7.2 does not look like a two-player game at first glance, it can be expressed in a minmax form by leveraging the Kantorovich-Rubinstein duality (Villani, 2008, Remark 6.5 on p. 95), which states that

$$\mathcal{W}(p, q) = \sup_{d: \|d\|_L \leq 1} \mathbb{E}_{\mathbf{x} \sim p} [d(\mathbf{x})] - \mathbb{E}_{\mathbf{z} \sim q} [d(\mathbf{z})] \quad (7.3)$$

where $\|d\|_L \leq 1$ refers to functions with a Lipschitz constant smaller or equal to 1. This results in the objective

$$\min_{\theta} \mathcal{W}(p(\mathbf{x}), \mathbf{g}_\theta) = \min_{\theta} \max_{\phi: \|d_\phi\|_L \leq 1} \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [d_\phi(\mathbf{x})] - \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} [d_\phi(\mathbf{g}_\theta(\mathbf{z}))] \quad (7.4)$$

where similarly to Equation 7.1, we aim at training a generator to produce samples that resemble $p(\mathbf{x})$, with the discriminator output being large for *real* samples and small for *generated* ones. Note that there are important differences in the objective, as we see that Equation 7.4 has no logarithms in the objective, and imposes a constraint on the Lipschitz constant of the discriminator. This has been enforced various ways, but the most common approach relies on regularizing the objective with a term called gradient penalty (Gulrajani et al., 2017) that penalizes large Lipschitz constants. One of the

7.1. Generative adversarial networks (GANs) for posterior modeling

original motivations of WGANs is to reduce the mode collapse behavior observed in GANs (Arjovsky et al., 2017). Mode collapse is an issue where a GAN does not learn all modes of the distribution, but remains stuck on a single mode.

7.1.3 Conditional GANs for posterior modeling

Most significant to this present work is however the extension of Equation 7.4 to an image conditional setting and the connection with a Bayesian modeling approach. In their work, Adler and Öktem (2018a) show that we can learn the posterior distribution $p(\mathbf{x}|\mathbf{y})$ using WGANs, and using Kantorovich-Rubinstein duality, they show that the problem of learning the posterior by minimizing

$$\min_{\theta} \mathbb{E}_{\mathbf{y} \sim p(\mathbf{y})} [\mathcal{W}(p(\mathbf{x}|\mathbf{y}), \mathbf{g}_{\theta}(\cdot, \mathbf{y}))]$$

can be framed as

$$\min_{\theta} \max_{\phi: \|d_{\phi}\|_L \leq 1} \mathbb{E}_{\substack{(\mathbf{x}, \mathbf{y}) \sim p(\mathbf{x}, \mathbf{y}) \\ \mathbf{z} \sim p(\mathbf{z})}} [d_{\phi}(\mathbf{x}, \mathbf{y}) - d_{\phi}(\mathbf{g}_{\theta}(\mathbf{z}, \mathbf{y}), \mathbf{y})] \quad (7.5)$$

where $p(\mathbf{x}, \mathbf{y})$ is the joint distribution of reference images and observations. We see that the main difference with Equation 7.4 lies in the fact that the generator and discriminator both take two inputs instead of one, and that we use a pair of samples from the joint distribution in order to learn the posterior. Note that it is easy to obtain samples from the joint distribution, as given a ground truth image \mathbf{x} , the corresponding observation can be obtained by applying the forward model of Equation 2.12 to it, with any mask ω .

Using such a model applied to Computed Tomography (CT), Adler and Öktem (2018a) show that the learned model can perform both reconstruction and UQ, by taking the empirical mean and variance using samples $\{\mathbf{g}_{\theta}(\mathbf{z}_i, \mathbf{y})\}_{i=1}^{n_s}$.

Deep Bayesian Inversion

In their work, Adler and Öktem (2018a) take an additional step from Equation 7.4, in order to further prevent the issue of *mode collapse*. To tackle the issue Adler and Öktem (2018a), propose to use a conditional mini-batch discriminator approach, inspired from Karras et al. (2017); Salimans et al. (2016). Their idea is to have the discriminator distinguish between pairs of conditional samples that can contain the original image or

generated ones. Their objective becomes:

$$\min_{\theta} \max_{\phi: \|d_{\phi}\|_L \leq 1} \mathbb{E}_{\substack{(\mathbf{x}, \mathbf{y}) \sim p(\mathbf{x}, \mathbf{y}) \\ \mathbf{z}_1, \mathbf{z}_2 \sim p(\mathbf{z})}} \left[\frac{1}{2} \left(d_{\phi} \left(\left(\mathbf{x}, \mathbf{g}_{\theta}(\mathbf{z}_2, \mathbf{y}) \right), \mathbf{y} \right) + d_{\phi} \left(\left(\mathbf{g}_{\theta}(\mathbf{z}_1, \mathbf{y}), \mathbf{x} \right), \mathbf{y} \right) \right) - d_{\phi} \left(\left(\mathbf{g}_{\theta}(\mathbf{z}_1, \mathbf{y}), \mathbf{g}_{\theta}(\mathbf{z}_2, \mathbf{y}) \right), \mathbf{y} \right) \right] \quad (7.6)$$

Neural Conditioner (Belghazi et al., 2019)

Although not framed as a Bayesian approach, the work of Belghazi et al. (2019) shows how to use a regular conditional GAN to learn an exponential amount of conditional distributions, which in the bottom line is the same task as learning the posterior distribution for various conditionals \mathbf{y} . As they are learning conditionals, their generator is called a *neural conditioner*. Contrarily to Adler and Öktem (2018a), Belghazi et al. (2019) focus on the problem of inpainting, where an image $\mathbf{x} \in \mathbb{R}^P$ is divided between *available* features indexed by the mask ω_a and *requested* features that are not observed, and are indexed with the mask ω_r . Their task then is to train a model to approximate the distribution $p(\mathbf{x}_{\omega_r} | \mathbf{x}_{\omega_a})$. This corresponds to an inverse problem where the observation is described as $\mathbf{y} = \mathbf{x}_{\omega_a} = \mathbf{P}_{\omega_a} \mathbf{x}$. In the sequel, for compactness, we write $\mathbf{P}_{\omega_a} \mathbf{x}$ as $a \cdot \mathbf{x}$. Overall, their objective reads

$$\min_{\theta} \max_{\phi} \mathbb{E}_{\mathbf{x}, a, r} [\log (d_{\phi}(r \cdot \mathbf{x}, a \cdot \mathbf{x}, a, r))] + \mathbb{E}_{\mathbf{x}, a, r, \mathbf{z}} [\log (1 - d_{\phi}(\mathbf{g}_{\theta}(a \cdot \mathbf{x}, a, r, \mathbf{z}) \cdot r, a \cdot \mathbf{x}, a, r))]. \quad (7.7)$$

Here, the discriminator and generator embed directly the mask information a and r , and as in the conditional model of Equation 7.6, the generator takes the observation $a \cdot \mathbf{x}$ and the noise \mathbf{z} , and the discriminator takes as input both the candidate sample $r \cdot \mathbf{x}$ or $r \cdot \mathbf{g}_{\theta}(a \cdot \mathbf{x}, a, r, \mathbf{z})$ along with the observation $a \cdot \mathbf{x}$.

7.1.4 Learning the posterior or learning moments

If one is primarily interested in using the posterior to compute its mean and variance for reconstruction and uncertainty quantification, then alternative approaches exist that rely on supervised learning instead of adversarial training. Indeed, adversarial training can often be unstable, whereas training in a supervised fashion is often easier.

The idea, used in particular by Zhang et al. (2019b) for MRI and proposed initially in Kendall and Gal (2017), assumes that the data is distributed normally around a mean $\boldsymbol{\mu} \in \mathbb{R}^P$ and a diagonal covariance $\text{diag}(\boldsymbol{\sigma})$, $\boldsymbol{\sigma} \in \mathbb{R}^P$, $\sigma_i \geq 0$. As discussed in Section 3.2.4, minimizing the reconstruction error $\|\mathbf{y}_{\omega} - \mathbf{A}_{\omega} \hat{\mathbf{x}}\|_2^2$ can be viewed as maximizing the likelihood of a Gaussian distribution with unit diagonal covariance. However, it is possible to model the uncertainty at a given pixel as a Gaussian centered at the reconstruction

7.1. Generative adversarial networks (GANs) for posterior modeling

mean $\hat{\mathbf{x}}_\theta$ with variance $\hat{\sigma}_\theta^2$. Both the reconstruction mean and variance can be output of two heads of a model with shared parameters, i.e. $(\hat{\mathbf{x}}_\theta, \hat{\sigma}_\theta^2) = \mathbf{f}_\theta(\mathbf{y}_\omega)$. And they can be learned through minimizing the conditional negative log-likelihood (NLL)

$$\min_{\theta} \ell_{\text{NLL}}(\theta) = \min_{\theta} \frac{1}{mP} \sum_{i=1}^m \sum_{j=1}^P \frac{(\mathbf{x}_{i,j} - \hat{\mathbf{x}}_{\theta,i,j})^2}{2\hat{\sigma}_{\theta,i,j}^2} + \frac{1}{2} \log(2\pi\hat{\sigma}_{\theta,i,j}^2) \quad (7.8)$$

where $\mathbf{x}_{i,j}$ denotes the j -th pixel of the i -th data point for a training set of size m . Such an approach of aiming at solving $\min_{\theta} -\log(p_{\theta}(\mathbf{x}|\mathbf{y}))$ is referred to as a *negative log-likelihood* in the machine learning (ML) literature, but in our context of inverse problems, Equation 7.8 amounts to assuming that the *posterior* takes the shape of a Gaussian distribution, and that we learn its moments by minimizing the parameters of the neural network $\mathbf{f}_\theta(\mathbf{y}_\omega)$.

Remark. There is here a considerable overlap with the terminology used in Bayesian Neural Networks (Kendall and Gal, 2017), where the aim is to learn a posterior distribution over *weights*,

$$p(\theta|\mathbf{x}, \mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{x}, \theta)p(\theta)}{p(\mathbf{y}|\mathbf{x})}.$$

We see that in our case, the notation is inverted compared to classical ML literature, as they denote the label or reference as \mathbf{y} and the data as \mathbf{x} , whereas in our case, the reference image is denoted \mathbf{x} , and the data (i.e. the observation) is denoted as \mathbf{y} . This is because in standard ML, we do not consider an inverse problem where we aim at retrieving the data from partial observations obtained through a forward model, but rather assume that \mathbf{y} contains some high-level structure (e.g. a class) that is hidden in some raw data \mathbf{x} .

We see then that $p(\mathbf{y}|\mathbf{x}, \theta)$ in the ML notation corresponds to $p_{\theta}(\mathbf{x}|\mathbf{y})$ in our notation, and serves in the ML case the role of likelihood, hence Equation 7.8 describing $p_{\theta}(\mathbf{x}|\mathbf{y})$ as a *negative log-likelihood*. However, in an inverse problem, $p_{\theta}(\mathbf{x}|\mathbf{y})$ plays the role of *posterior*, as seen in Section 3.2.4. We have

$$p_{\theta}(\mathbf{x}|\mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{x})p(\mathbf{x})}{p(\mathbf{y})}.$$

In the case of an inverse problem, $p(\mathbf{y}|\mathbf{x})$ is the likelihood defined by the forward model of Equation 2.12, $p(\mathbf{x})$ the prior defined by the ground truth data and $p(\mathbf{y})$ is the marginal. As computing probabilities from a discrete set of training data is difficult and marginalization is computationally prohibitive on such problems, Zhang et al. (2019b) circumvent the problem by directly assuming that the posterior is a Gaussian distribution, and simply train a model to learn its moments from data. However, as we will see in the sequel, such modeling simplifications incur a cost,

as this reductive assumption does not allow the posterior to model multi-modal distributions.

7.2 GAN based adaptive sampling

The initial motivation of this work was a belief that the benefits of generative methods have not been fully exploited in the Compressed Sensing (CS) setting. Although GANs have been used as priors (Bora et al., 2017; Narnhofer et al., 2019), as reconstructors (Yang et al., 2018) and approximate posteriors for UQ (Adler and Öktem, 2018a), their ability to naturally provide an adaptive policy has not been explored.

Although this connection between modeling the posterior and adaptive sampling was explored in the context of classical CS (Ji et al., 2008; Seeger et al., 2010), this has not been exploited in the context of GANs, even though they provide a very flexible, learning-based model of the posterior distribution.

This is precisely our approach, which is in some sense very natural. Given an observation $\mathbf{y}_\omega = \mathbf{P}_\omega \mathbf{A} \mathbf{x} + \boldsymbol{\epsilon}$ and a trained conditional GAN $g_{\theta^*}(\cdot, \mathbf{y}_\omega)$ approximating the posterior $p(\mathbf{x}|\mathbf{y}_\omega)$, we can compute the estimated conditional pixel-wise variance in the transformed domain $\text{Var}[\mathbf{A} \mathbf{x} | \mathbf{y}_\omega]$ by taking several samples from the generator, transforming them to the observation domain \mathbf{A} , and acquiring the location with the largest posterior variance, i.e.

$$v_t = \arg \max_{v: v \in [P]} \mathbf{P}_v \text{Var}[\mathbf{A} \mathbf{x} | \mathbf{y}_{\omega_t}]. \quad (7.9)$$

Departing from the idea of explicitly training a policy to guide adaptive sampling (Jin et al., 2019; Zhang et al., 2019b; Pineda et al., 2020; Bakker et al., 2020), our approach simply minimizes the variance of an approximate posterior. This turns out to yield a strong sampling policy *despite the networks involved being trained solely for conditional generation*.

Remark. In order to perform adaptive sampling using Equation 7.9, the model must fulfill some requirements. First of all, one must be able to *sample* for the model, and secondly, sampling must be efficient.

Although these requirements are readily filled by our approach, this is not the case of all models. In particular, this disqualifies methods that only learn the mean and pixel-wise variance of the posterior, by minimizing for instance the NLL (Equation 7.8). Such statistics do not capture the complex, multi-modal distributions that arise in the inverse problem, and sampling from the Gaussian distribution resulting from minimizing Equation 7.8 will not work in practice, as the samples will not originate from the true posterior, which can be multi-modal, but rather from the Gaussian that simplifies the underlying distribution. We confirm in our experiments that sampling from the

7.2. GAN based adaptive sampling

NLL, transforming the samples into Fourier and taking the maximum variance to guide the acquisition does not work.

In addition, if one wishes to deploy such adaptive sampling in a clinical setting, the inference time must be shorter than the repetition time (TR), which is typically in the order of 10^2 - 10^3 ms. While our approach and the one of Zhang et al. (2019b) satisfy this constraint, this is not the case for all generative approaches. In particular, autoregressive models (Van Oord et al., 2016) or score-based sampling approaches suffer from long inference times (Song et al., 2021).

An important distinction between our current approach and the ones that were discussed previously is the *information horizon* used by the different policies. The information horizon refers to how far ahead rollouts occur in order to train a given policy, and is described on Table 7.1 for different policy models. While it is clear that policies that integrate long horizon planning leverage N -step information, other methods can make use of 0-step or 1-step information.

Sampling policy	Adaptive	Long horizon	Information horizon
LBCS (Gözcü et al., 2018)	✗	✗	1-step
AlphaGo (Jin et al., 2019)	✓	✓	N -step
DDQN (Pineda et al., 2020)	✗(✓)	✓	N -step
Policy Gradient (Bakker et al., 2020)	✓	✗(✓)	1-step
Evaluator (Zhang et al., 2019b)	✓	✗	0-step
GAN (Ours)	✓	✗	0-step

Table 7.1 – Methods that will be considered in the paper. This table extends the perspective given in Table 6.1. Different methods make use of 0/1/ N -step information in their sampling policy.

The distinction lies in the fact that our GAN-based approach and the one of Zhang et al. (2019b) make their choice based on the location with the maximum error at the current step, while LBCS or the Policy Gradient approach (Bakker et al., 2020) integrate the actual error that remains when adding a location in the mask. Formally, given a test image \mathbf{x} , greedily trained 1-step policies aim at acquiring the location which yield to the best performance at the next step

$$v_{t+1} = \arg \max_{v \in [P]} \eta(\mathbf{x}, \hat{\mathbf{x}}_{\omega_t \cup v}). \quad (7.10)$$

Evaluating this for a given point requires *computing a reconstruction for each candidate* v , in order to compute the *actual* performance gain that would be achieved by adding v to the mask ω_t . 0-step information methods rely instead on acquiring the location feature the largest *current* error

$$v_{t+1} = \arg \max_{v \in [P]} \|\mathbf{P}_v \mathbf{F}(\mathbf{x} - \hat{\mathbf{x}}_{\omega_t})\|. \quad (7.11)$$

Evaluating Equation 7.10 requires only a *single reconstruction*, and then selection is performed by masking in the Fourier space in order to find the location with the largest current error.

As we will see in the results, 0-step methods do not outperform 1-step methods, and we hypothesize that this is because they do not incorporate any feedback in the design of their policy, contrarily to 1-step methods. 0-step methods simply they pick the location with the largest current error, but do not incorporate the *a posteriori* information on how good their action *actually* was. This is mainly because 0-step approaches use a heuristic that was not explicitly trained to be a policy, but to perform a surrogate task, such as estimating the variance of the posterior distribution in Fourier (in our case), or estimating the line-wise Fourier MSE (Zhang et al., 2019b).

7.3 Comparison with Zhang et al. (2019b)

We have mentioned several times until now the conceptual similarities with the approach of Zhang et al. (2019b) with our approach. We have shown in Section 7.1.4 that their approach is *not* generative, but aims at learning the mean and variance of the posterior distribution by minimizing a NLL (Equation 7.8). These statistics do not allow to directly perform adaptive sampling, as we show in Appendix B.2.2.

This is why, in order to perform adaptive sampling, Zhang et al. (2019b) also train with their reconstructor an additional model, called an *evaluator* $e_\phi(\hat{\mathbf{x}}, \omega)$, that tries to predict the line-wise error in Fourier domain. The role of the evaluator is not only to provide guidance for adaptive sampling, similarly to the variance in Fourier space in our case, but it also helps to refine the reconstruction, by adding an adversary that aims at discriminate between real reconstructed lines and reconstructed ones.

The idea of the evaluator is to provide an individual prediction of whether a line in k-space is likely to be real or reconstructed. Consider the k -th $v_k \in [P]^2$, and a reconstruction $\hat{\mathbf{x}}$. Zhang et al. (2019b) define the k -th *spectral map* as $\mathbf{m}^k(\hat{\mathbf{x}}) = \mathbf{F}^{-1}(\mathbf{P}_{v_k} \mathbf{F} \hat{\mathbf{x}})$, \mathbf{m}_i represents the inverse Fourier transform of a single k-space line, and has dimension $\mathbf{m}_i \in \mathbb{C}^P$. The evaluator then aims at fitting target scores using the following kernel

$$t(\hat{\mathbf{x}}, \mathbf{x})_k = \exp\left(-\gamma \|\mathbf{m}^k(\hat{\mathbf{x}}) - \mathbf{m}^k(\mathbf{x})\|_2^2\right) \quad (7.12)$$

where γ is a scalar hyperparameter. Note that $t(\hat{\mathbf{x}}, \mathbf{x}) \in [0, 1]$, where it is 1 when the spectral maps of $\hat{\mathbf{x}}$ and \mathbf{x} perfectly match. The evaluator $e_\phi(\hat{\mathbf{x}}, \omega)$ is then trained to minimize

$$\min_{\phi} \ell_E^E(\phi) = \min_{\phi} \frac{1}{m} \sum_{i=1}^m \sum_{k=1}^P |e_\phi(\hat{\mathbf{x}}_i, \omega)_k - t(\hat{\mathbf{x}}_i, \mathbf{x}_i)_k|^2. \quad (7.13)$$

²For simplicity, we simply rewrite here the set of all pixels $\{1, \dots, P\}$, as the argument is the same whether pixels or lines are considered.

The intuition is simple, the evaluator tries to match the prediction of the kernel t . The evaluator is also added to the objective of the reconstruction method, contributing as

$$\ell_{\text{E}}^{\text{R}}(\theta) = \frac{1}{m} \sum_{i=1}^m \sum_{k=1}^P |e_{\phi}(\hat{\mathbf{x}}_{\theta,i}, \omega)_k - 1|^2 \quad (7.14)$$

where we see that it acts as an adversary: the evaluator is trained to predict how closely a reconstructed $\hat{\mathbf{x}}$ matches a reference image \mathbf{x} , while the reconstructor aims at reconstructing images that look realistic to the evaluator. The updated objective of Equation 7.8 then becomes

$$\min_{\theta} \ell_{\text{R}}(\theta) = \min_{\theta} \ell_{\text{NLL}}(\theta) + \beta \ell_{\text{E}}^{\text{R}}(\theta) \quad (7.15)$$

The model is then trained by alternating updates between Equation 7.13 and Equation 7.15. Although the model is inspired by the adversarial training employed in GANs, significant differences remains, especially regarding the learning of a distribution in a GAN against point estimates in Zhang et al. (2019b), and the construction of an evaluator that matches a kernel rather than some more general distance such as a KL-divergence (Goodfellow et al., 2014) or a Wasserstein distance (Arjovsky et al., 2017).

7.4 Experiments

We first consider MRI adaptive sampling, and the move to image domain sampling on MNIST.

7.4.1 MRI adaptive sampling

Experimental setting

We carried out experiments on Fourier sampling with the same conditions as outlined in Section 6.3

Dataset. We used again the fastMRI (Zbontar et al., 2019) single-coil knee dataset for the experiments, using complex data and the same data normalization as previously. We cropped the data to 256×256 before resizing them to 128×128 . We used horizontal sampling masks, and experimented among other things with the mask distribution of Zhang et al. (2019b). This corresponds to the `c+rhz` setting in the previous chapter. Recall that for the mask distribution, the setting of Zhang et al. (2019b) meaning the model was trained with 10 center frequencies always selected and between 3 and 37 randomly selected elements. The data were postprocessed to have the magnitude of the ground truth in the range $[0, 1]$.

Reconstruction models. For this MRI experiment, we used two types of reconstruction models. The first one is the c-ResNet of Zhang et al. (2019b) that was trained to minimize the negative log-likelihood of Equation 7.8 along the adversarial contribution of the evaluator. The model was trained using Adam (Kingma and Ba, 2014) with $\beta = (0.5, 0.999)$, $lr = 6 \cdot 10^{-4}$ over 100 epochs. The learning rate was linearly decayed starting at 50 epochs until it reached 0 at the end of training.

The second type of model that we considered are different GAN versions. We observed that for Fourier-based conditional generation, the training of the Neural Conditioner (NC) (Belghazi et al., 2019) was generally unstable and generally did not lead to good performance, although it did very well in the image setting. We observed that the model proposed by Adler and Öktem (2018a) generally performed better than a standard conditional WGAN, and this is the one that we used in our experiments. We refer to it as the **AO** or **GAN** model. For the generator, we used the residual UNet (ResUNet) of Belghazi et al. (2019), and used also their ResNet-like architecture as discriminator. In the GAN case, the reconstruction was estimated by computing the empirical mean from a set of samples. In our experiments, we observed that merely 2 samples managed to provide an informative mean and variance, although averaging on more samples led to an increase in overall image quality.

For the WGAN we use Adam (Kingma and Ba, 2014), $\beta = (0, 0.9)$, TTUR (Heusel et al., 2017) with learning rates of $lr_G = 10^{-4}$ and $lr_D = 3 \cdot 10^{-4}$. The training is done over 100 epochs, and the learning rate was halved at 50 epochs. We used the gradient penalty (Gulrajani et al., 2017) with a weight of 10, following Adler and Öktem (2018a) along with a drift penalty with a weight of 10^{-3} . The output of the generator was followed by a data-consistency layer (Schlemper et al., 2018; Mardani et al., 2018), which has the role of replacing the generated data with the observed ones at the last layer, and a sigmoid to keep the range of the data in $[0, 1]$. Further details on the architecture and on the training are given in Appendix B.1.

Sampling methods. Similarly to what we did in the previous Chapter (Section 6.3), we evaluated the models on different policies, that we detail below.

- Random sampling (**Random**): Acquire a fixed proportion of low-frequency lines in Fourier and then randomly sample the remaining lines.
- **Evaluator**: Zhang et al. (2019b) proposed to train an evaluator that tries to estimate the current mean-squared error for each line in k-space.
- (Stochastic) Learning-based Compressive Sampling (**LBCS**) (Gözcü et al., 2018; Sanchez et al., 2020a)
- Policy Gradient model (**RL**): The greedy, adaptive deep reinforcement learning method of Bakker et al. (2020).

- **GAN (ours)**: We take samples from the conditional GAN posterior, transform them to Fourier and compute their empirical variance to guide the decision of what location to observe next.
- **MSE oracle**: An adaptive oracle acquiring at each step the line with the largest mean squared error.

For all the policies considered, we initialized the mask with 8 low-frequency lines, and then carried out sequential sampling until a rate of 50% sampling. Note however that the deep RL policy of Bakker et al. (2020) was trained only until 25% sampling rate for computational reasons, and to provide a fair comparison, the aggregated results will be computed up to 25% sampling.

Results

The results of our experiment are presented in Table 7.2 and Figure 7.1.

Similar trends can be observed on PSNR (provided in Appendix B.1) and SSIM, where the 1-step information policies (LBCS and RL) generally outperform the 0-step information provided by the Evaluator and the GAN. On the GAN model, the GAN policy largely outperform the Evaluator, and on the reconstructor, the opposite is true, although the gap is smaller. However, the trend is *not* the same if one goes until 50% sampling, as shown on Figure 7.1. Indeed, one sees that as the sampling horizon increases, the performance of the evaluator significantly worsens, whereas the GAN model matches the performance achieved by LBCS. Note that the Evaluator policy in this case is really the direct translation of the policy used in the case of Zhang et al. (2019b), but then reconstruction is carried out through the GAN. We see also that the RL policy is weaker on the reconstructor. This is due to the policy being trained on the GAN reconstruction model and then evaluated on the reconstructor. In addition, the results illustrate the inefficiency of sampling from the Gaussian NLL and transforming the samples into Fourier in order to sample: such a policy can perform even *worse* than random sampling.

We see also that while the MSE oracle provides an upper bound on the performance of our GAN policy, it is outperformed by LBCS. This further highlights the value of using 1-step information horizon over 0-step information horizon, as an adaptive 0-step oracle can be outperformed by a 1-step non-adaptive baseline.

We observe also interesting trade-offs. The evaluator shows a very good performance at low sampling rates ($< 10\%$) matching and even outperforming 1-step information methods, but its performance deteriorates at later stages. On the other hand, our GAN method is weaker at very early stages but stays much more consistent throughout the whole range of sampling rates.

We note also that using more than 2 samples can boost the overall performance of the GAN reconstruction. For instance, using 5 samples for the GAN reconstruction can boost the overall performance on the method: paired with the GAN policy, it reaches up to 31.43 dB and 0.759 SSIM, and on LBCS, it goes up to 32.10 dB and 0.773 PSNR. This amounts to an improvement of 0.8 to 1 dB and between 0.03 to 0.04 SSIM. However, averaging on increasingly many samples has a diminishing return, and in our experiments, a plateau is reached around 20 samples.

We also tried to perform sampling from the Gaussian NLL of Zhang et al. (2019b), and as expected, the policy performed just like random sampling, as the NLL does not capture the complexity of the posterior distribution.

Policy	Model			
	GAN		Reconstructor	
	PSNR	SSIM	PSNR	SSIM
Random	29.48	0.699	29.79	0.757
Gaussian NLL	29.48	0.699	29.69	0.747
Evaluator	30.03	0.710	30.77	0.765
GAN	30.67	0.726	30.48	0.762
LBCS	31.06	0.735	31.41	0.781
RL	30.88	0.731	30.68	0.766
MSE Oracle	30.98	0.734	31.23	0.778

Table 7.2 – Average PSNR and SSIM test set AUC (one AUC per image) on knee data in the $c+rh$ setting. The AUC is computed between the 8 initial lines and 25% sampling rate.

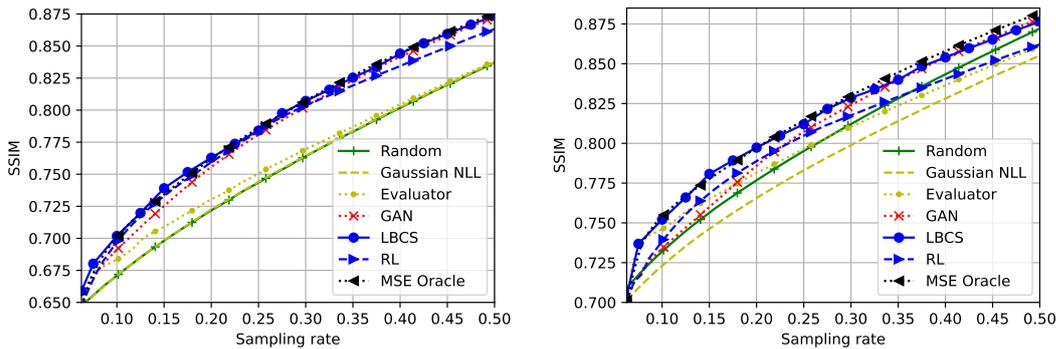


Figure 7.1 – SSIM plots of results shown in Table 7.2, showing the different performance of the policies. The left plot is the result of the evaluation on our GAN model, and the right plot is evaluated using the reconstructor of Zhang et al. (2019b). Although Table 7.2 computes the AUC until 25%, this table extends the evaluation until 50% sampling rate.

Policy distribution. Turning now to Figure 7.2, we can see how different the policy looks through the testing set for LBCS, the evaluator and our GAN. The probabilities are computed as average of the masks that are observed at sampling rate 25% of the

testing set, and we see very different behaviors. First, as LBCS is a fixed policy, it is clear that the same locations are acquired for all images, and as a result, the policy is either 0 or 1. Then the GAN policy is quite concentrated around low frequencies, and spans a range that is consistent with the locations acquired by LBCS. On the contrary, the evaluator tends to acquire a lot of high frequencies, and it is hard to discern much structure in the acquisition.

Note that the GAN policy might suffer from acquiring less often some very central frequencies. It might fail to do so due to the uncertainty on these locations being quite small, although the error remains non-negligible.

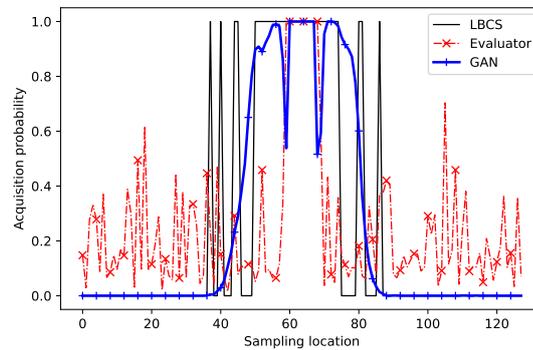


Figure 7.2 – Acquisition probability of sampling locations by LBCS, our method and the evaluator of Zhang et al. (2019b), evaluated on the testing set at 25% sampling rate, using the GAN model for reconstruction.

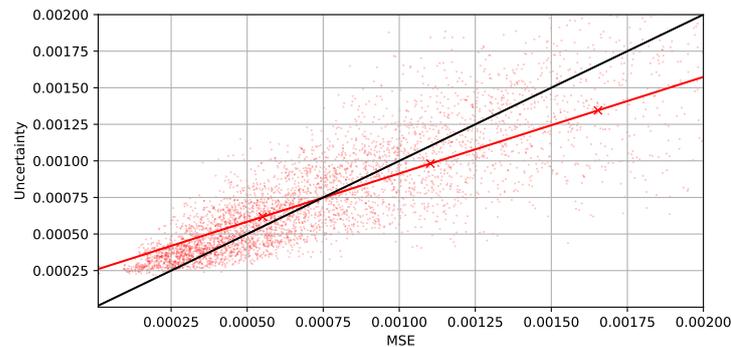


Figure 7.3 – Uncertainty (empirical variance) against MSE for the GAN model. Each point represents a given test sample at a given sampling rate. For clarity, not all data points are represented. The red line is the result of performing a linear regression on the data points, and has $R^2 = 0.705$. The black line corresponds to $\text{Uncertainty} = \text{MSE}$

Uncertainty calibration. We might also wonder how good is the uncertainty (aggregated empirical variance) at predicting the mean squared error. This is shown for our GAN model on Figure 7.3, where we see that overall estimation is quite good, but has some limitations still. At high errors, the model tends to be over-confident, and show a

Chapter 7. Uncertainty driven adaptive sampling via GANs

large variability in its uncertainty values, whereas at lower errors, it becomes slightly underconfident, although the uncertainty tends to have a lower variance.

However, note that this result is based on aggregating the uncertainty over the whole image, and might not reflect the ability of the model to predict good pixel-wise MSE. Nonetheless, the results suggest that the aggregated empirical variance of our model correlates well with MSE.

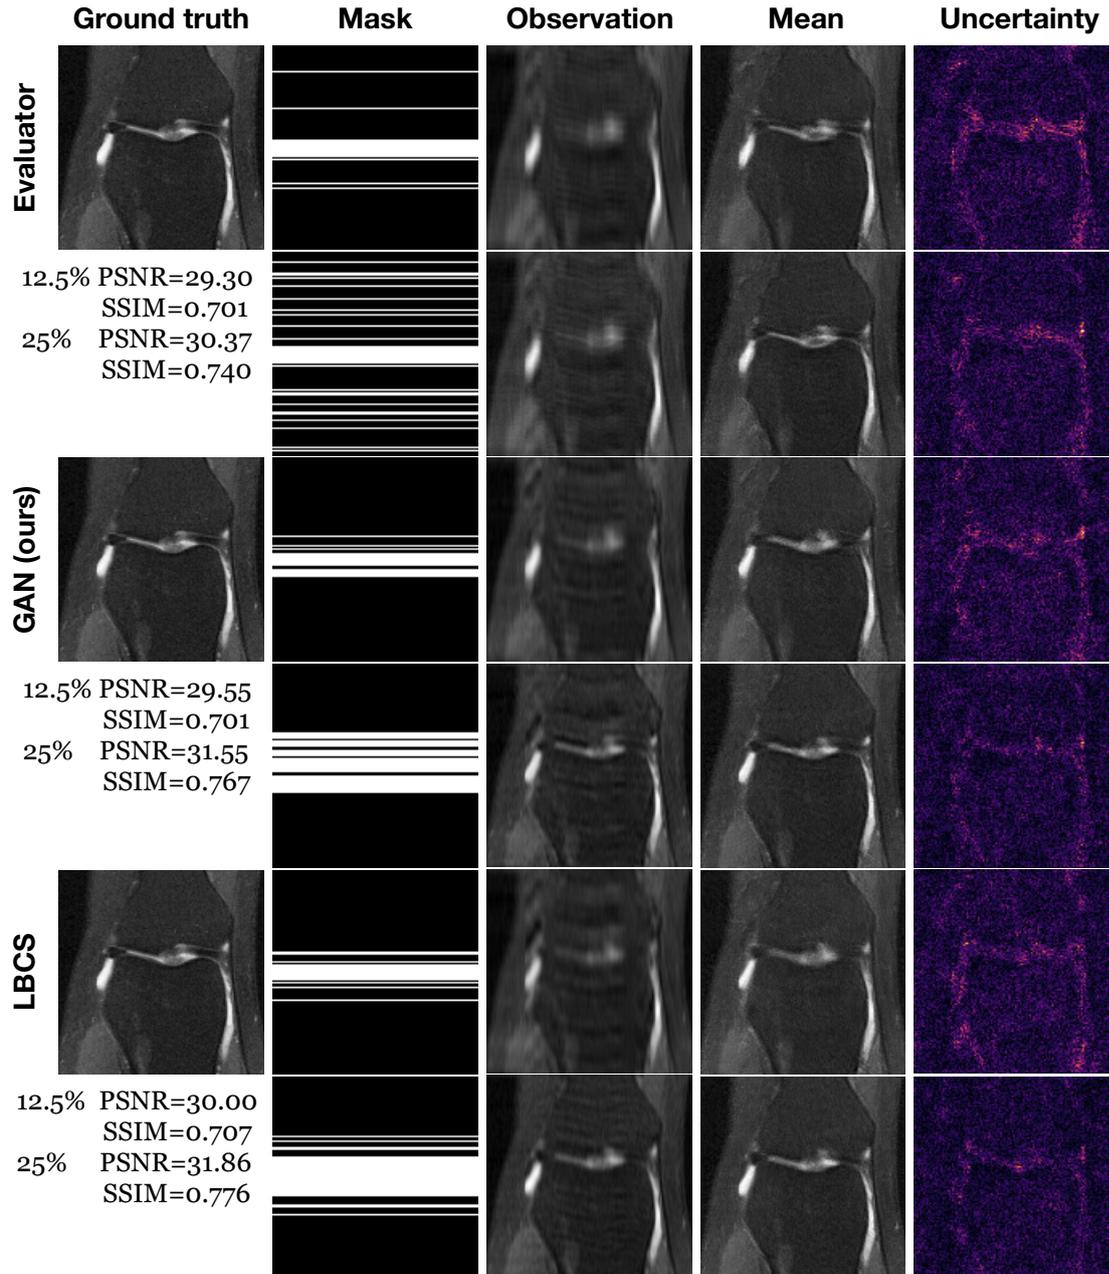


Figure 7.4 – Visual illustration of the evaluator, GAN and LBCS policies evaluated at 12.5% and 25% on a slice of knee data, using the GAN model for reconstruction.

7.4.2 Beyond Fourier sampling: image-domain conditional sampling

Although our method was motivated by Fourier sampling, the work of Belghazi et al. (2019) proposed a similar, GAN-based approach for inpainting tasks. This motivated us to consider the problem of adaptive sampling in image domain as well.

Experimental setting

We carried out our experiments on the MNIST dataset. The training set contains 60 000 handwritten digits of size 28×28 . We held out 10 000 digits for validation purposes. The test set consists of 10 000 digits. We zero-padded all images to 32×32 for convenience.

We evaluated the NC approach of Belghazi et al. (2019) using their discriminator architecture, as well as the ResUNet and cResNet as two generators, and we compared it to a conditional Wasserstein GAN.

Similarly to the MRI case, the models were trained with gradient penalty (Gulrajani et al., 2017) and the generators are followed with a data consistency layer as well as a sigmoid. All models were trained using Adam Kingma and Ba (2014) with learning rate 10^{-4} , $\beta = (0.5, 0.999)$ and weight decay 10^{-4} . All models were trained for 300 epochs. The images were undersampled using masks that were either drawn at random or drawn with a probability decaying away from the center.

The Neural Conditioner (NC) Belghazi et al. (2019) generator and discriminator were updated every other round. In the WGAN case, we optimized the generator once, followed by four updates of the discriminator, as done in Adler and Öktem (2018a).

Remark. Similarly to Belghazi et al. (2019), we found gradient regularization of the discriminator to be essential. However, we obtained very poor training results using the authors proposed gradient penalty, and observed that the gradient penalty of Gulrajani et al. (2017) led to a stable training.

Also, Belghazi et al. (2019) aim at controlling the encoder’s Lipschitz constant by using spectral-normalization on the encoding part of the generator. However, we did not observe any effect of doing this, and did not retain it in our model.

After training, we carried out several rounds of evaluation for different types of sampling as reported in Table 7.3. We found that 2 samples were sufficient for the adaptive sampling with GAS to yield good results. We also observed that more samples 20 to 100 made little improvement in the sampling quality or image quality, except that it suppressed the noise of the empirical posterior mean and variance a bit more efficiently.

Regarding the stochastic LBCS approach (Sanchez et al., 2020a), there are two parameters to be adjusted, namely the number of training samples to be used at each stage of the

greedy mask optimization, and the size of candidate locations to be used. We set these to 64 and 256 respectively. We used the negative MSE as the performance metric.

In this experiment, as we have smaller images in image domain, we performed pixel-wise masking rather than line-wise masking like in MRI. This greatly increases the number of candidate locations that the policy can sample. We used for training a mixture of random masks and masks that sample more frequently locations near the center of the image.

Note also that the approach of Zhang et al. (2019b) is *not* applicable here, as it is not clear how to train their evaluator in image domain, due to their introduction of spectral maps. Also, even if it were adapted, their approach scales very poorly to pixel-wise sampling, as this would require the computation of a spectral map for each individual pixel, all of which are fed at once as input of the evaluator.

Results

Our main results are presented on Table 7.3, which provides the AUC for PSNR, SSIM, as well as the downstream classification accuracy. This was computed using a simple convolutional network that was pretrained on MNIST without any undersampling. A visual assessment of different policies is shown on Figure 7.5.

Focusing first on Table 7.3, we see a very different picture than on the MRI case (Table 7.2). Here, our 0-step adaptive GAN policy very largely outperforms the 1-step fixed LBCS policy across all metrics. This highlights a fundamentally different paradigm between image domain and Fourier domain sampling. In image domain, it seems that adaptivity matters much more than in Fourier. This makes intuitive sense, and is well illustrated in the case of MNIST, where the sampling pattern focuses on areas that are ambiguous, which can be very different for different images: this is illustrated more clearly on Figure B.2 in Appendix B.2.3. This is also consistent with results on CIFAR10 shown in Appendix B.2.4.

It is surprising to see that, although the different generator architectures have roughly the same number of parameters, they can lead to very different performance when used as policies. The ResUNet performs significantly better with the NC and the cResNet shows a much better performance on the WGAN. However, there is not a significant difference on the performance with the LBCS policy. Although this is beyond the scope of this work, it would be interesting to form a better understanding of why such a discrepancy occurs.

Figure 7.5 features also an impressive ability of conditional GANs: although the model was only trained on pairs of data (\mathbf{x}, \mathbf{y}) from a *given* mode of the distribution, the model successfully produces varied conditional outcomes $\{\mathbf{x}_i\}$ that span *multiple modes*

Policy	Model	Architecture	PSNR	SSIM	Accuracy
Random	NC	ResUNet	17.10	0.76	0.67
	WGAN	c-ResNet	16.41	0.76	0.67
LBCS	NC	ResUNet	21.39	0.90	0.92
	NC	cResNet	20.33	0.88	0.93
	WGAN	ResUNet	20.29	0.89	0.90
	WGAN	cResNet	20.48	0.89	0.92
GAN (Ours)	NC	ResUNet	40.15	0.94	0.91
	NC	cResNet	31.47	0.93	0.93
	WGAN	ResUNet	31.73	0.92	0.88
	WGAN	cResNet	35.81	0.96	0.95

Table 7.3 – Average test set AUC (one AUC per image) on MNIST, in image domain.

of the conditional distribution. Figure 7.5 also shows an issue that our GAN policy can encounter: sometimes, the model can be confidently wrong: at 0.5% percent, the model thinks the digit should be a seven, which is incorrect³. However, adding more samples allows it to converge to the correct digit. Note that in this case, it shows a level of uncertainty that is generally low, as opposed to the random policy (VDS) and LBCS, which feature a higher level of uncertainty.

7.5 Related works

In this work, we have investigated how GANs can be used for inverse problems to perform reconstruction, uncertainty quantification and adaptive sampling. A flurry of different approaches have been proposed to tackle these problems in the literature, and we briefly discuss a few particularly relevant works now.

Although we have used GANs as the generative model in this work, many other approaches have been explored in the context of inverse problems. Variational Autoencoders (VAEs) (Rezende et al., 2014; Kingma and Welling, 2013) have been applied to inverse problems (Tonolini et al., 2019; Zhang et al., 2021) as well as the task of learning conditional distributions (Ivanov et al., 2019). But other approaches such as autoregressive models (Oord et al., 2016; Van Oord et al., 2016) or conditional normalizing flows (Papamakarios et al., 2019; Sun and Bouman, 2020; Denker et al., 2021) have also been explored. Lately, score-based approaches, based on the idea of learning to sample by learning the score of the distribution, have gained popularity (Ramzi et al., 2020; Song et al., 2021).

Although all of these methods can provide reconstruction and uncertainty quantification,

³For other generative models such as variational autoencoders, it has been shown that these models can frequently be overconfident (Nalisnick et al., 2018). Although evaluating the likelihood in GANs is an open problem, it is not unlikely that a similar behavior should be found in GANs.

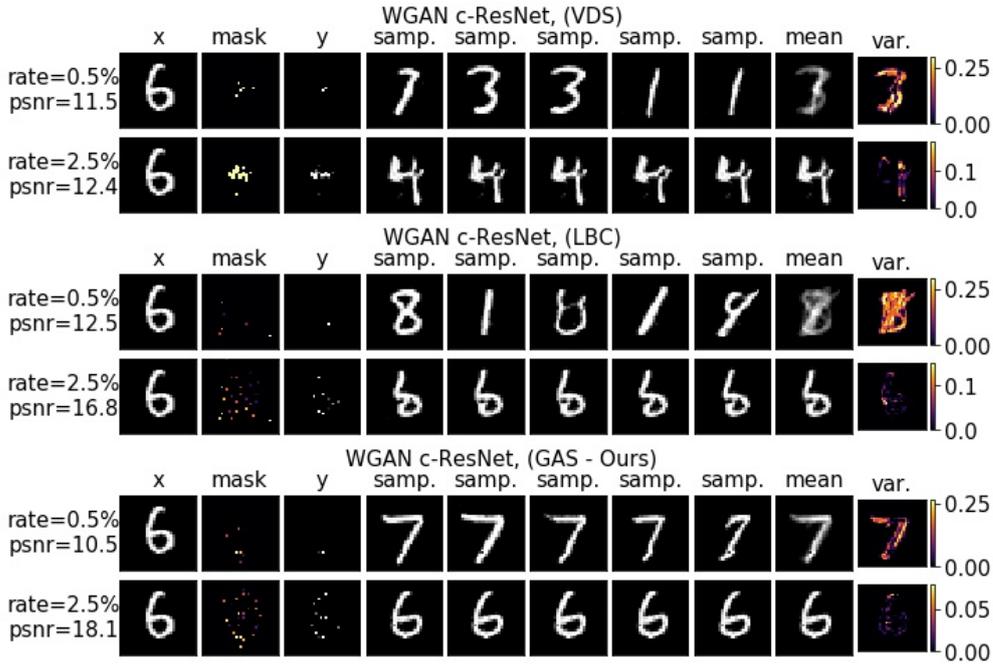


Figure 7.5 – Image domain illustration of the WGAN cResNet sampling using three different methods. The columns contain respectively: the ground truth image, the mask (with brighter color meaning that the location has been selected recently), the observation, several samples and the empirical posterior mean and variances. The first row shows the result obtained with 0.5% and 2.5% sampling respectively. Here, VDS is our random policy, LBC is LBCS and GAS stands for Generative Adaptive Sampling and is our GAN-based policy.

autoregressive models and score-based distribution suffer from slow inference times, limiting their potential use in a setting of adaptive acquisition in MRI, where the inference should take at most the same time as a readout. This can in principle be achieved by our model, which has an inference time in the order of milliseconds.

In the vein of the Neural Conditioner (NC) (Belghazi et al., 2019), we mention the works of Douglas et al. (2017); Gautam et al. (2020) and a very recent energy based approach (Strauss and Oliva, 2021). However, these approaches focus on missing data imputation or inpainting, and are not designed for settings where the acquisition does not happen in image domain. Other approaches train a conditional embedder produce a noise distribution which will subsequently enable conditional sampling (Whang et al., 2020), or motivate a conditional generator based on approximately triangular transport maps (Kovachki et al., 2020).

Generative adversarial networks have also been used in the context of inverse problems, and several works used unconditional GANs as priors (Bora et al., 2017; Anirudh et al., 2018; Latorre et al., 2019; Narnhofer et al., 2019; Jalali and Yuan, 2019; Hussein et al., 2020) and subsequently performed optimization in the latent space. Another approach has been to directly train GANs for reconstruction along additional losses in order to

increase the quality of the reconstruction (Yang et al., 2018). In such settings, the noise channel of the generator is often discarded at test time as the aim is to obtain the best image quality.

There has also been some recent interest into quantifying the uncertainty in the reconstruction. This notably motivated the negative log-likelihood in Zhang et al. (2019b), where the level of uncertainty can be used as a criterion to decide when sampling can be stopped. A recent work, based on Stein’s unbiased risk estimator, studied the calibration of the uncertainty for variational autoencoders (Edupuganti et al., 2020). Score-matching is also very promising for UQ, although it suffers from long inference times (Ramzi et al., 2020; Song et al., 2021).

As discussed, Zhang et al. (2019b) is the work more conceptually similar to ours. Their model aims at doing reconstruction and uncertainty quantification, while also providing adaptive sampling through their evaluator. It is also trained in an adversarial fashion, although it is not a GAN. However, their approach is *not* generative, and only yields point estimates of the mean and the variance by minimizing a multivariate Gaussian negative log likelihood. Their approach supposes the pixel-wise uncertainty to be independent random variables normally distributed around the mean, an assumption not required in Adler and Öktem (2018a); Belghazi et al. (2019). This assumption also implies that sampling from this distribution to estimate the Fourier variance is not possible, as this would require access to the full covariance matrix, and our results confirm this.

We note also that the optimization of the sampling pattern is not restricted to MRI, but has also been applied on the standard compressed sensing setting, notably motivated by early work on Bayesian Compressive Sensing (Ji et al., 2008; Wu et al., 2019b,a).

Finally, our methodology relies on the idea of sequentially acquiring the most uncertain measurement at each round, a concept which has first been proposed and leveraged in the context of optimal experimental design (Lindley, 1956; Chaloner and Verdinelli, 1995; Vanlier et al., 2012). While until recently, most approaches in this realm relied on a statistical prior model as well as expensive MCMC sampling (Vanlier et al., 2012; Pauwels et al., 2014), but variational approaches have also been explored (Foster et al., 2019). However, most of these approaches do not scale up to very high-dimensional problems. On the contrary, GANs are naturally fit to deal with high-dimensional problems, and provide with a natural way of empirically quantifying the uncertainty in this acquisition.

7.6 Discussion

Our results show the potential of conditional GANs in tackling inverse problems, as they offer a good reconstruction, uncertainty quantification and adaptive sampling at once, by only training in a standard adversarial fashion. Nonetheless, several points are worth

discussing

Generically trained generative reconstruction method. The current generative model was trained in a generic fashion and *not* specifically to optimize the quality of masks designed through it. The ability for designing masks stems purely from training it as a rigorous Bayesian modeling of the continuum of inverse problems. This allows the posterior to be conditioned on incrementally collected information in a closed loop fashion. However, one could imagine ways to improve the quality of the policy, by explicitly training the model to not only provide good samples, but simulate a *rollout* of samples. The discriminator could be then trained to distinguish not between samples but *sequences* of samples.

Alternatively, our method could easily be incorporated in a reinforcement learning-based framework aimed at jointly training reconstruction and sampling such as Jin et al. (2019). This could give the best of both worlds, giving principled uncertainty estimates to the RL sampler, moving beyond greedy sampling and possibly speeding up the training of the reconstruction method by focusing on regions with less reliable variance estimates instead of using masks sampled from parametric distributions as in this work.

Limitations of the model. While our approach and the one of Belghazi et al. (2019) are successfully producing conditioned samples on unseen masks shapes, we observe that they nonetheless depend on the type of masks used throughout training. For instance, our method, being trained on pixel-wise undersampling for MNIST, would not be as good in the inpainting of large holes. Hand-crafting the type of masks that we want to condition on in the future remains an important step towards obtaining good generalization performance, and further work needs to be done to have a method able to truly condition on arbitrarily structured observations.

We also observed an expected trade-off in performance between models trained on a single type of masks and a single sampling rate and models trained on varying types of masks and sampling rates. In this sense, adaptive sampling comes at a cost: a small drop in performance compared to a method *specialized* on a more restricted setting. Quantifying precisely this trade-off is beyond the scope of this work and should be further studied.

Impact of data consistency. We found the data-consistency layer to be a fundamental piece in all settings, as it forces the model to incorporate observed data into the output image, and the model is then much less prone to hallucinating unlikely features as more observations are being acquired. In addition, this ensures the fact that the input will never be discarded by the generator, and prevents the model from collapsing into a mode that is inconsistent with observations.

Limitations of the posterior estimation. We leveraged the approach of (Adler and Öktem, 2018a), which is the first of its kind to allow to construct a posterior estimator from samples of the joint distribution. While it works well empirically, the authors did

not provide any analysis or guarantees on how well the generator captures the tails of the posterior distribution. Our observations suggest that unusual images, that far away from the mean of the learned distribution, might not be accurately captured, i.e. the estimated variance is lower than expected. This could be due to the limited training data available, but might also be an artifact of the cWGAN approach which tends to struggle with capturing weaker modes of their distributions. Specifically, the problematic examples might be an indication that while the loss shown in Equation 7.6 avoids mode collapse, there might still be some modes that are underrepresented.

Adaptivity in image domain. Contrarily to our conclusions in Chapter 6 and Section 7.4.1, adaptive sampling seems to be greatly beneficial in image domain sampling. In particular, in a simple task such as MNIST, one can clearly see the benefit of adaptive policies, that tailor the sampling pattern to the relevant locations of the image (in the case of MNIST, the digit itself). This is also concurrent with our results on CIFAR10 (Appendix B.2.4). This raises the question of predicting *when* adaptive sampling would be beneficial over fixed sampling. Does it depend on the structure of the sampling domain itself, or would the distribution of images considered impact this as well? Does the adaptivity of the model depend on the distribution of masks used for training? All of these questions should be further investigated.

7.7 Conclusion

We presented a novel approach for reconstruction and adaptive sampling, rooted in Bayesian modeling. We showed that, in line with recent advances (Belghazi et al., 2019; Adler and Öktem, 2018a), GANs perform well at this task and yield useful posteriors for Fourier domain reconstruction as well as image domain, allowing both high quality reconstructions and derivation of a sampling heuristic. This natural sampling heuristic achieved surprisingly strong sampling results, without ever training to act as a policy, and is competitive with the state-of-the-art on the fronts of reconstruction, uncertainty quantification and adaptive sampling.

Bibliographic note

Zhaodong Sun provided initial results in the work. Igor Krawczuk helped with the code, implemented and provided the results for the methods of Zhang et al. (2019b) and Bakker et al. (2020), and carried out the CIFAR10 experiments. Volkan Cevher provided the initial motivation to look into the Deep Bayesian Inversion approach of Adler and Öktem (2018a) which started this project.

8 Conclusion and future works

Summary of the thesis

In this thesis, we studied the problem of optimizing sampling patterns for Cartesian MRI. More specifically, we considered the problem of optimizing sampling patterns in a sequential context, using learning-based, data-driven methods. Our work consists of both methodological novelties and insights into what impacts the performance of models, what design choices matter most and *why* some methods generally perform worse than others.

In **Chapter 4**, we brought a taxonomy of approaches that have been taken to tackle the problem of optimizing sampling patterns, which helped shed light on the development of approaches to optimize sampling for Cartesian MRI. More concretely, we proposed that similarly to reconstruction methods, the field of mask optimization has undergone a shift from *model-based* to *learning-based* approaches (Ravishankar et al., 2019; Doneva, 2020), but that it has also shifted from *model-driven* towards *data-driven* criteria to evaluate the quality of sampling masks. It was naturally observed that moving toward learning-based, data-driven approaches enables the model to perform better (Chauffert et al., 2013; Zijlstra et al., 2016), a trend that was also consistently confirmed throughout our work. We also observed that within the realm of learning-based, data-driven methods, several trends arose, where researchers tackled the problem of learning to design masks as either a fixed, static problem or a sequential, dynamic approach. Following the sequential approach to mask design, our work focused on studying the problem of optimizing sampling on pre-trained reconstruction models. This constitutes the precise framework of our contributions.

Within this context, in **Chapter 5**, we first investigated algorithms to drastically improve the scaling of LBCS (Gözcü et al., 2018). We showed that ideas from submodular optimization (Krause and Golovin, 2014) can be leveraged to scale up LBCS by several orders of magnitude, and in some cases without degrading its performance. Our algorithms,

Conclusion and future works

sLBCS and ILBCS, have been shown to provide strong mask designs for a range of clinically relevant MRI settings, including multicoil dynamic MRI and 3D MRI. These settings were previously not accessible to LBCS, due to the prohibitive cost of evaluating all candidate locations on the full dataset. sLBCS tackled this limitation by proposing a stochastic sampling of both candidate locations and data points, and achieved results on par with LBCS, and even sometimes outperforming the method. This method is easily implemented and broadly applicable, and we used it as a strong baseline throughout the rest of the thesis. ILBCS was proposed to address the case of multicoil and 3D MRI, where LBCS cannot be evaluated, due to its prohibitive computational cost. By using lazy evaluations, ILBCS can scale up to these problems. ILBCS presents very strong performances at high undersampling rates, but we hypothesized that its list of upper bounds becomes gradually worse as more sampling locations are added to the mask, making its performance gradually worse than baselines. However, several options were proposed to further improve its performance and scalability, and should be investigated as future work.

Our work illustrates the benefit of a flexible learning-based, data-driven approach to the problem of optimizing sampling. Moreover, by focusing on the problem of optimizing sampling patterns on pre-trained reconstruction models, sLBCS and ILBCS integrate a future-proof dimension, as they will remain readily available to any reconstruction method that is applicable to a sequence of sampling rates.

In **Chapter 6**, we aimed at resolving a seeming conflict between mask design methods that used deep reinforcement learning (Pineda et al., 2020; Bakker et al., 2020). These methods showed a potential to improve over sLBCS and ILBCS by integrating long-horizon sampling as well as patient-adaptive decisions. However, Pineda et al. (2020) and Bakker et al. (2020) reached seemingly contradictory conclusions. Pineda et al. (2020) suggested that long-term planning was the most important component, as their Dataset-Specific DDQN (non-adaptive) policy often matched the performance of their Subject-Specific DDQN (adaptive) policy. On the side of Bakker et al. (2020), it was found that a *greedily* trained adaptive deep RL policy could match the performance of policies trained with a long-horizon dependency. In other words, *does the ability to plan multiple steps ahead matter most, as suggested by Pineda et al. (2020), or is it rather the ability to adapt sampling to a given patient, as suggested by Bakker et al. (2020)?*

Our methodology was simply to consider a baseline policy that is greedily trained (no long-term planning) and non-adaptive, but that is trained to maximize the same reward function as considered by Pineda et al. (2020); Bakker et al. (2020). This baseline could then serve as the reference point from which one could observe what brings most performance improvement. However, it became quickly clear that our previously proposed algorithm, sLBCS, could exactly fulfill this role, and it was then natural to use it as the comparison baseline. Exhaustive experiments allowed us to bring elements of answers to this important question, albeit in a surprising way: it seems indeed that *neither long-term*

planning, nor adaptive sampling are able to bring a strong return on investment over sLBCS.

Exploring this question led us to consider the different processing choices that were done in the works of Pineda et al. (2020) and Bakker et al. (2020), such as the choices of architecture, how the models were pretrained, the metrics used for reporting and how results were aggregated and presented. Surprisingly again, we showed that *inconsistencies in reporting could easily lead to incoherent results*, and that subtle changes in the processing or the reporting pipeline could lead to cases where the SotA deep RL methods of Pineda et al. (2020); Bakker et al. (2020) were outperformed by sLBCS.

Moreover, the return on investment of moving from sLBCS to a deep RL policy model was at best marginal compared to other changes such as the architecture or the mask distribution used to pretrain it. Indeed, our results show very clearly that the benefit of using a deep RL policy over sLBCS is very limited compared to the improvements achieved by improving the reconstruction architecture or using a distribution of sampling masks tailored to the acquisition horizon considered. However, there is a clear benefit of moving from hand-crafted heuristics such as random sampling and low-pass acquisition policies to a *learning-based, data-driven* such as sLBCS.

These observations led us to wonder whether these results are due to current limitations of deep RL policies or if they describe some more fundamental limitations of the problem. We believe the second option to be more plausible, in light of more recent works such as Bakker et al. (2021), where in their case, the best performing adaptive policies *explicitly learn* to become non-adaptive. This also opens some further questions: could we design even simpler *learning-based, data-driven* approaches that retain the benefits of sLBCS, while being even easier to train? Can we also find new paradigms of acquisition that would allow to further improve over the *learning-based, data-driven* paradigm?

Finally, in **Chapter 7**, we proposed a novel approach to sampling, based on Bayesian modeling. Inspired by previous works to model inverse problems in a Bayesian fashion using GANs (Adler and Öktem, 2018a), we showed that in the context of MRI, this Bayesian modeling allows us to not only perform reconstruction and uncertainty quantification, but also provides a *natural* way to perform adaptive sampling. Indeed, we show that acquiring the location with the largest posterior variance provided a good, natural policy, without the model ever being trained explicitly to perform adaptive sampling. Our GAN-based approach was however outperformed by sLBCS and RL-based policies, but did outperform the approach of Zhang et al. (2019b). We postulated that the reason for this is that our GAN approach and the one of Zhang et al. (2019b) rely on 0-step information, i.e. base their decision on the current error, rather than receiving feedback about how their decision *actually* performed, which we describe as 1- or N -step information.

Conclusion and future works

Nonetheless, this last work provides a promising all-in-one approach to Bayesian modeling of inverse problems, and would deserve to be investigated in greater depth. GAN-based approach have generally achieved strong performances in MRI reconstruction (Chen et al., 2022), and seem like a promising approach to provide an integrated approach to MRI reconstruction, uncertainty quantification and sampling. But as our methodology is tied to learning the posterior rather than GANs themselves, we anticipate that novel contributions to Bayesian modeling would also benefit from these observations, further giving credit to learning the full posterior distribution rather than the pixel-wise mean and variance.

Insights

In addition to the methodological contributions, our work provides insights into how sampling policies can be designed efficiently. Figure 8.1 shows the connections that we established in this thesis. We used the common notions of fixed (non-adaptive) and adaptive policies and refined the distinction between greedy and long-horizon policies to include the *information horizon* accessible to the policy, distinguishing between 0-step and 1-step greedy policies. The first are driven by the error at the current reconstruction, whereas the second integrates feedback about the *actual* effect on reconstruction of adding a sampling location to the mask.

We found that this is the most significant distinction between the existing methods, as in Chapter 6, our results suggest that at least in their current state, neither long-horizon planning (N -steps information horizon) nor adaptivity allow significant and consistent improvements over (s)LBCS, a 1-step non-adaptive policy. The distinction between 0-step and 1-step methods was explored in Chapter 7, where we showed that even an adaptive oracle that acquires the line with the largest current mean-squared error is outperformed by a non-adaptive policy that leverages 1-step information.

Overall, our experiments suggest that the most important structure that a policy should include in order to achieve near state-of-the-art performance is 1-step information, i.e. taking into account the *real* impact on performance of acquiring a new location. Further refinements, at least in the current state, seem to bring at most very modest performance improvements, challenging the common trend to design ever-more complex models.

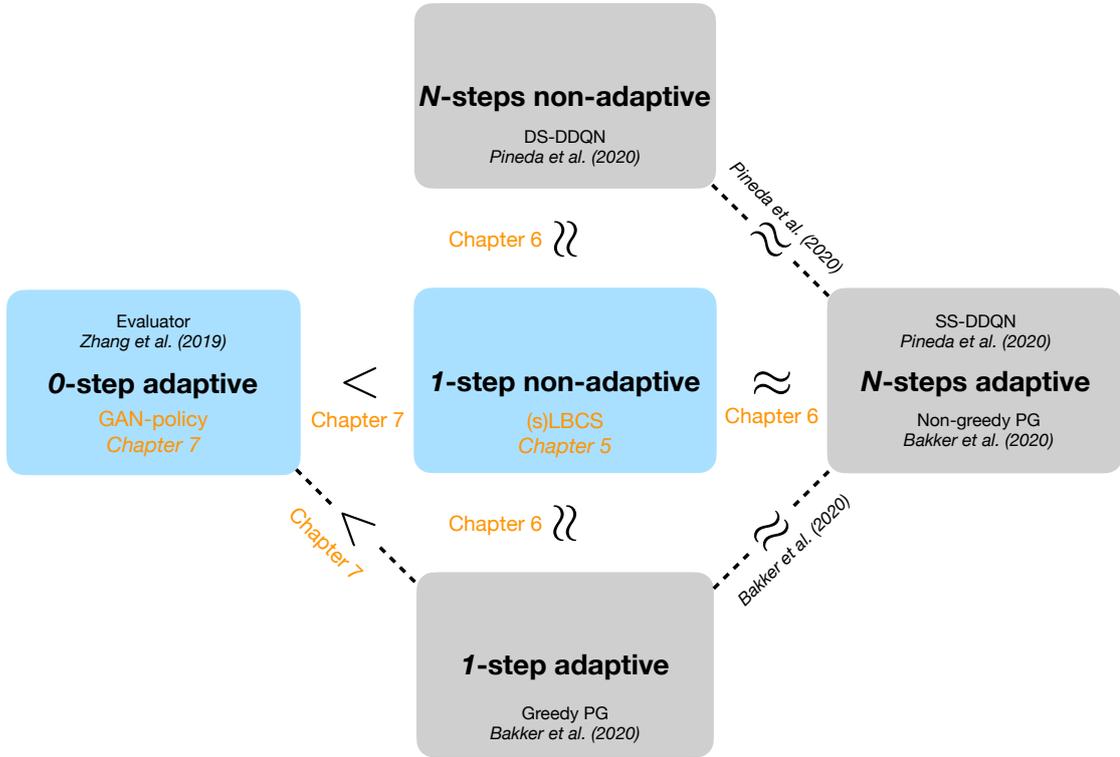


Figure 8.1 – Summary of the contributions of this thesis relating to the design of a mask optimization algorithm.

Future works

Nevertheless, even if our thesis brings methodological contributions and insights to the problem of learning sampling masks for Cartesian MRI, many important questions still remain open.

This thesis does *not* provide a definitive answer to whether the low return on investment of deep RL methods applied to MRI results from current limitation or if it arises from an intrinsic structure of the problem. It is not impossible that further advances in deep learning would allow to efficiently train complex models, and to achieve significant performance gain over methods such as (s)LBCS. Similarly, although the results of this thesis provide evidence toward the validity of the conclusions of Figure 8.1, it is not impossible either that a deeper connection would be able to explain these observations more adequately. As a result, further work, both at the theoretical and practical levels would be necessary in order to understand what really matters in sampling mask optimization.

Moreover, our results focus on *pre-trained reconstruction models* for sequential acquisitions, and it is not clear how *jointly* training the reconstruction and the policy models might affect the picture given in Figure 8.1. Indeed, joint training is challenging as it adds

Conclusion and future works

variance to the training procedure, and it is difficult to estimate its impact. Although a few works considered this problem (Jin et al., 2019; Van Gorp et al., 2021; Yin et al., 2021), it remains to be understood whether joint training mainly results in stronger reconstruction models by enabling them to focus on a more restricted set of masks, in stronger policies by co-designing them with the reconstruction model, or if both, or none, of these conclusions hold. Answering these questions would be an important step towards designing end-to-end, data-driven MRI pipelines.

A crucial interrogation relates however to the clinical value of optimized sampling masks for MRI sampling. It is known that there are some significant artifacts introduced by prospective undersampling, but their effect has not been sufficiently studied in the case of deep learning-based MRI reconstruction (Yu et al., 2022). However, there have not been any studies that consider the distortion of a prospective acquisition when using an optimized sampling. Are there ways to design masks that are robust to the distortions of prospective sampling and can retain good performance? Although robustness is a growing concern in MRI reconstruction (Antun et al., 2020; Johnson et al., 2021; Darestani et al., 2021), most works only consider retrospective undersampling.

Additionally, the world of MR imaging is not restricted to Cartesian imaging, and many exciting recent works also aim at optimizing sampling trajectories for non-Cartesian acquisitions (Lazarus et al., 2019; Weiss et al., 2019; Wang et al., 2021; Chaithya and Ciuciu, 2022). Non-Cartesian acquisitions can be in practice much faster than Cartesian trajectories, and although they suffer from different limitations, they are definitely a topic that would be worth further investigating.

Beyond these more conceptual directions, many technical questions need also to be addressed for the field of MRI sampling optimization to really mature into a clinically relevant application. We should clarify what really matters at different levels, first in terms of metrics: what kind of metrics do we want to use? Where does the reconstruction quality matter? In terms of the problems considered themselves: should we care about sequential reconstruction or optimize a fixed mask? Do we want to jointly train reconstruction methods and sampling policies? How can we ensure that a trained policy would be robust in a prospective acquisition?

These questions leave many avenues of research open, and promises exciting discoveries in the years to come.

A Appendix for Chapter 6

A.1 Implementation details

Pretraining regimes. The deep reconstructors used in Pineda et al. (2020); Bakker et al. (2020) are pre-trained by randomly sampling a mask from a set of distributions with different parameters. Bakker et al. (2020) uses a discrete set of sampling rates at (25%,25%,25%, 16.7%,16.7%, and 12.5%) with (25%,16.7%,12.5%, 16.7%,12.5%,12.5%) of the total selected frequencies being allocated to center frequencies and the rest sampled uniformly, while Zhang et al. (2019b) uses a distribution where 10 center frequencies are always acquired, and between 0 and 38 additional frequency lines are acquired following a uniform distribution (as a reminder, the total number of possible sampling locations is 128 lines or columns).

All UNet models of the ablations of the Bakker et al. (2020) setting were trained following the hyperparameters described in the paper, i.e. using Adam (Kingma and Ba, 2014) and training to 50 epochs, otherwise using the hyperparameters from Zbontar et al. (2019) (except for the $10\times$ learning rate drop 40 epochs, we instead kept the initial 10^{-3} rate throughout). We used a batch size of 32.

For the cResNet models used in the ablations we used Adam as an optimizer with learning rate 10^{-3} , betas of (0.9, 999), trained for 50 epochs and a batch size of 32. The architecture is the same as in Zhang et al. (2019b) except only outputting the mean of the image (no uncertainty channel), and using 72 channels for the 3 residual blocks, with 18, 36, 72 channels in the encoder and 72, 36, 18 channels in the decoder. In the comparison with Pineda et al. (2020), we used the checkpoint provided at <https://facebookresearch.github.io/active-mri-acquisition/misc.html> by the authors.

Computational hardware. We performed all of your experiments on a DGX-2 server using A100 GPUs. On this machine the reconstruction model of Pineda et al. (2020) fits into GPU memory with an effective batchsize of 50 (we use subbatching to enable

Appendix A. Appendix for Chapter 6

arbitrary batch sizes) and it took ~ 3 hours to train each of the reported LBCS masks. Meanwhile, masks for the comparisons of Bakker et al. (2020) could be trained in ~ 20 minutes, while training the RL policy took on the order of days as reported in the authors original paper. While one might be able to obtain a certain speedup using a more optimized or parallel RL algorithm, a key bottleneck is the sequential nature of the optimization.

A.2 Further experiments and results on Bakker’s setting

A.2.1 Data processing differences

In section 6.3, we mentioned that we slightly deviated from the setting of Bakker et al. (2020). We now detail these changes and show that they do not change the *relative* performance of the different methods.

1. **Train-test split:** we used a different train-test split than Bakker et al. (2020), as we randomly split 10% of the training set as a test set, and used the fastMRI validation set for test. We used the 50% more central slices, resulting in 15599 training slices, 1743 validation slices and 3564 test slices.
2. **Complex data undersampling:** Bakker et al. (2020) use magnitude ground truth images as the reference that is undersampled. We discussed in Section 6.2 the issue with this approach. We chose to use complex preprocessing of data, followed by taking magnitude of the observation obtained *after* undersampling.
3. **Data range:** PSNR and SSIM need to be provided with a maximal data range in their computation: Bakker et al. (2020) used the maximal intensity in the ground truth *volume*, while we used the maximal intensity of each ground truth *slice* or image.
4. **Data standardization:** Bakker et al. (2020) used a unique data standardization, where observations and ground truth data are standardized using their respective statistics, but denormalized, after reconstruction, using only the ground truth statistics. While this ensures a more closely matching data range, this introduces a mismatching data range that biases models leveraging data consistency, used in most state-of-the-art models. We performed *matched* data normalization and denormalization using the observation statistics for the observation, and the ground truth ones for the ground truth.

The third and fourth changes are related to postprocessing, and do *not* require retraining a model. We see on the three first columns of Table A.1 that these changes do not alter the relative ordering of the methods. However, in absolute numbers, the worse performing method in the initial setting outperforms the best method the best performing one after these changes (highlighted in italic on Table A.1). While comparing these numbers would be a mistake, this highlights the impact of subtle postprocessing changes and the need for care in comparing methods, especially across different papers.

Ablation over Magnitude vs Complex data. We first illustrate the impact of doing the appropriate preprocessing by evaluating the model of Bakker et al. (2020) using their data processing pipeline against a pipeline with complex preprocessing. The SSIM performance AUC is reported in Table A.1.

Appendix A. Appendix for Chapter 6

This would correspond to a simulation of what would happen if the model of Bakker et al. (2020) was to be used in deployment, as one would receive undersampled complex observations that are only then transformed to magnitude images. While this does not induce any *reversal* in this case, we see that in the matched setting, the gap between LBCS and RL shrinks. We see also again that training and evaluating on the matched setting leads to larger performance improvements than the ones obtained by training sophisticated policies.

Policy	Bakker's setting	Individual data range	Matched standardization	Mismatched setting	Complex preprocessing (matched - ours)
Random (SH)	0.6251	0.5964	0.5798	0.5595	0.5608
LtH (SH)	0.6073	0.5762	0.5650	0.5617	0.5719
LBCS (SH)	0.6375	0.6099	0.5928	0.5684	0.5732
RL (SH)	0.6388	0.6112	0.5941	0.5698	0.5738
NA Oracle (SH)	0.6383	0.6107	0.5937	0.5717	0.5738

Table A.1 – AUC on the test set, using SSIM, on various data processing on Bakker’s knee setting. The mismatched setting refers to the evaluation of the model trained in the setting of Bakker et al. (2020), but evaluated complex data preprocessing instead of magnitude data preprocessing. This makes the reconstruction and policy models to be out of the distribution. A model trained on the complex preprocessing is reported in the complex preprocessing (matched - ours) column.

A.2.2 Exact reproduction of Bakker et al.

In our experiments, we did not use the exact same train-val-test split and dataset undersampling as Bakker et al. (2020). We re-evaluated our model in the exact setting of Bakker et al. (2020), as illustrated on the first column of Table A.1. On Table A.2, we present the results at 25% sampling against the ones reported in the paper of Bakker et al. (2020).

Policy (SH)	Our code	Bakker	Difference
Random	0.6974	0.6948	+0.0026
Equi (one)	0.7078	0.7049	+0.0029
Equi (two)	0.7090	0.7064	+0.0026
NA Oracle	0.7235	0.7213	+0.0022
LBCS	0.7222	–	–
RL (Greedy)	0.7250	0.7230	+0.0020
RL $\gamma = 0.9$	0.7250	0.7232	+0.0018

Table A.2 – SSIM at 25% on the test set used by Bakker et al. (2020) and subsampling the dataset by a factor 2, in order to replicate their setting as close as possible. We compare the numbers that we obtain to the ones reported in the article of Bakker et al. (2020).

We see that the performance difference is in the same order of 0.0018 to 0.0029, and that the relative ordering of methods is preserved. We conclude from these results that

A.2. Further experiments and results on Bakker’s setting

our results are overall consistent with the results displayed in the paper of Bakker et al. (2020), although this is not an *exact* replication.

A.2.3 Further results in the Bakker experiment

We provide the results of our full ablation study on the components described in section 6.3. We report observations on the *cvb*, *cvz*, *c+rvz* and *c+rhz* settings. The summary of the full trajectory with an AUC is presented on Table A.3 and the performance at the end of sampling is shown in Table A.4.

It is difficult to consistently establish a trend for when reversals will happen, but several important observations can be made from the results.

First of all, by comparing Tables A.3 and Table A.4, we see that reversals can happen by considering one way of reporting over another (see *c+rvz* cResNet in both tables). This highlights again the impact of the way results are reported.

The effect of the masks used for pretraining the reconstructor is also interesting. Comparing the *cvb* and *cvz* results, one can consistently see the following trend. In the short horizon setting (LHS of the tables), using *cvb* leads to consistent improvements over *cvz*. In the long horizon setting, the opposite is true. Recall that the *b* masks are discretely distributed from sampling rates 12.5% to 25%, which matches the short horizon experiment range. The *z* masks span a continuous range, from roughly 7% to 37.5%, whereas the long horizon experiment spans a range from 3% to 25% sampling rate. These results suggest that matching the pretraining regime to the regime on which evaluation will be carried out has a significant influence on the performance of the sampling policy, an observation which has, to our knowledge, never been quantified before.

PSNR evaluation. We also provide a PSNR evaluation for the models trained on SSIM in Tables A.5 and A.6. There does not seem a clear trend or correlation between the policy’s performance on SSIM and PSNR. We observe however the same kind of dynamics in the settings, where cropped + resized data naturally have a higher PSNR than cropped ones, and cResNet improves the reconstruction quality of UNet.

Appendix A. Appendix for Chapter 6

Policy	Short Horizon				Long Horizon			
	cvb		c+rhz		cvb		c+rhz	
	UNet	CResNet	UNet	CResNet	UNet	CResNet	UNet	CResNet
Random	0.5608	0.5659	0.7292	0.7338	0.4348	0.4432	0.5860	0.5885
LtH	0.5719	0.5764	0.7309	0.7412	0.4849	0.5107	0.6678	0.6902
LBCS	0.5732	0.5828	0.7430	0.7526	0.5134	0.5243	0.7035	0.7096
RL	0.5738	0.5830	0.7430	0.7524	0.5142	0.5242	0.7035	0.7111
NA Oracle	0.5738	0.5832	0.7430	0.7527	0.5140	0.5247	0.7038	0.7099

Policy	Short Horizon				Long Horizon			
	cvz		c+rvz		cvz		c+rvz	
	UNet	CResNet	UNet	CResNet	UNet	CResNet	UNet	CResNet
Random	0.5580	0.5661	0.6851	0.6984	0.4483	0.4621	0.5252	0.5276
LtH	0.5663	0.5759	0.6739	0.7028	0.5122	0.5218	0.6292	0.6444
LBCS	0.5712	0.5822	0.7034	0.7235	0.5174	0.5328	0.6568	0.6712
RL	0.5717	0.5829	0.7042	0.7239	0.5183	0.5334	0.6582	0.6733
NA Oracle	0.5718	0.5835	0.7038	0.7236	0.5180	0.5336	0.6568	0.6717

Table A.3 – AUC on the test set, using SSIM, for the full ablation study using the model of Bakker et al. (2020). The short and long horizon results are not comparable with each other, as AUCs are integrated on the whole range of sampling rates. The top right part of the table (long horizon) replicated the results of table 6.3, excluding the standard deviation for legibility. The rest of the ablation were not averaged on several seeds due for computational reasons. Recall that *cvb* stands for cropped, vertical lines, Bakker-like mask distribution, *c+rhz* stands for cropped then resized, horizontal lines and Zhang-like masks, *cvz* stands for cropped, vertical lines and Zhang-like masks and *c+rvz* stands for cropped then resized, vertical lines and Zhang-like masks.

Policy	Short Horizon				Long Horizon			
	cvb		c+rhz		cvb		c+rhz	
	UNet	CResNet	UNet	CResNet	UNet	CResNet	UNet	CResNet
Random	0.607	0.6141	0.7565	0.7667	0.5249	0.5432	0.6567	0.6567
LtH	0.6267	0.6313	0.7602	0.7786	0.5832	0.6197	0.7325	0.7714
LBCS	0.6288	0.6413	0.7751	0.7923	0.6294	0.6417	0.7768	0.7941
RL	0.6298	0.6417	0.7745	0.7921	0.6298	0.6415	0.7761	0.7935
NA Oracle	0.6301	0.6421	0.7751	0.7926	0.6301	0.6428	0.7771	0.7942

Policy	Short Horizon				Long Horizon			
	cvz		c+rvz		cvz		c+rvz	
	UNet	CResNet	UNet	CResNet	UNet	CResNet	UNet	CResNet
Random	0.6053	0.6148	0.7201	0.7401	0.5502	0.5693	0.6191	0.6359
LtH	0.619	0.6309	0.7052	0.7487	0.6190	0.6309	0.7052	0.7487
LBCS	0.6282	0.6413	0.7411	0.7710	0.6293	0.6451	0.7453	0.7746
RL	0.6288	0.6416	0.7417	0.7714	0.6304	0.6452	0.7470	0.7738
NA Oracle	0.6292	0.643	0.741	0.7717	0.6304	0.6464	0.7457	0.7755

Table A.4 – SSIM at 25% sampling rate, using SSIM, with the model of Bakker et al. (2020). This is the counterpart of the results shown in Table A.3. Here, the results across short and long horizon are comparable.

A.2. Further experiments and results on Bakker’s setting

Policy	Short Horizon				Long Horizon			
	cvb		c+rhz		cvb		c+rhz	
	UNet	CResNet	UNet	CResNet	UNet	CResNet	UNet	CResNet
Random	24.206	24.180	27.641	28.365	20.802	20.836	22.375	22.371
LtH	24.486	24.420	27.628	28.640	23.570	23.578	26.607	27.317
LBCS	24.493	24.517	28.254	29.154	23.574	23.635	27.220	27.802
RL	24.466	24.506	28.247	29.147	23.585	23.646	27.210	27.757
NA Oracle	24.510	24.521	28.257	29.152	23.591	23.639	27.222	27.757

Policy	Short Horizon				Long Horizon			
	cvz		c+rvz		cvz		c+rvz	
	UNet	CResNet	UNet	CResNet	UNet	CResNet	UNet	CResNet
Random	23.978	24.184	25.845	26.949	20.967	21.437	20.428	20.625
LtH	24.108	24.396	25.641	27.048	23.441	23.703	24.487	25.510
LBCS	24.269	24.490	26.515	27.870	23.600	23.846	25.307	26.316
RL	24.226	24.498	26.530	27.861	23.581	23.858	25.240	26.246
NA Oracle	24.271	24.513	26.528	27.856	23.603	23.848	25.309	26.250

Table A.5 – AUC on the test set, using PSNR, for the full ablation study using the model of Bakker et al. (2020). The short and long horizon results are not comparable with each other, as AUCs are integrated on the whole range of sampling rates. The rest ablation were not averaged on several seeds for computational reasons. Recall that *cvb* stands for cropped, vertical lines, Bakker-like mask distribution, *c+rhz* stands for cropped then resized, horizontal lines and Zhang-like masks, *cvz* stands for cropped, vertical lines and Zhang-like masks and *c+rvz* stands for cropped then resized, vertical lines and Zhang-like masks.

Policy	Short Horizon				Long Horizon			
	cvb		c+rhz		cvb		c+rhz	
	UNet	CResNet	UNet	CResNet	UNet	CResNet	UNet	CResNet
Random	24.583	24.604	28.241	29.113	21.520	21.684	23.547	23.843
LtH	25.056	24.984	28.176	29.567	25.056	24.984	28.176	29.567
LBCS	25.068	25.117	29.088	30.214	25.043	25.077	29.226	30.234
RL	25.032	25.108	29.050	30.218	25.071	25.101	29.194	30.188
NA Oracle	25.095	25.120	29.090	30.233	25.111	25.111	29.211	30.229

Policy	Short Horizon				Long Horizon			
	cvz		c+rvz		cvz		c+rvz	
	UNet	CResNet	UNet	CResNet	UNet	CResNet	UNet	CResNet
Random	24.369	24.623	26.657	27.957	21.791	22.497	21.854	22.506
LtH	24.551	24.961	26.223	28.143	24.551	24.961	26.223	28.143
LBCS	24.817	25.089	27.439	29.069	24.850	25.084	27.518	29.145
RL	24.765	25.102	27.451	29.069	25.102	24.807	27.583	29.041
NA Oracle	24.813	25.123	27.429	29.043	24.837	25.101	27.446	29.072

Table A.6 – PSNR at 25% sampling rate, using PSNR, with the model of Bakker et al. (2020). This is the counterpart of the results shown in Table A.5, where the acronyms used are explicited. Here, the results across short and long horizon are comparable.

A.3 LBCS vs long-range adaptivity

Using the setting of Pineda et al. (2020), we studied qualitatively how the different policies yield different mask designs. As can be seen in fig. A.2a., d. and e, LBCS puts more emphasis on the center frequencies, but acquires similar sections of k-space as the SS-DDQN. It also creates a more symmetric mask, which is in line with Pineda et al. (2020) observations that SSIM creates more asymmetric masks. More interestingly, as can be seen in fig. A.2c., most of the variation is concentrated at the early sampling rates (left of the plot) with the std and especially the coefficient of variation ($\frac{\sigma}{\mu}$) decaying towards zero in most locations. This implies that while SS-DDQN is indeed adapting to each image individually, this mainly affects ordering early on and after 20 -40 samples (6 – 12% sampling rate) LBCS starts to catch up. This is also supported by two observations:

1. LBCS underperforms SSDQN by a larger margin than on the full evaluation, which we interpret as the mask being good *on average* and with a smaller sample there is a higher chance of individual suboptimality.
2. the per-image difference between LBCS and SSDQN grows in favor of SS-DDQN until about 40 samples where LBCS starts to slowly recover. The distribution of the difference however becomes much wider, implying there are images where LBCS performs wildly different from SS-DDQN.

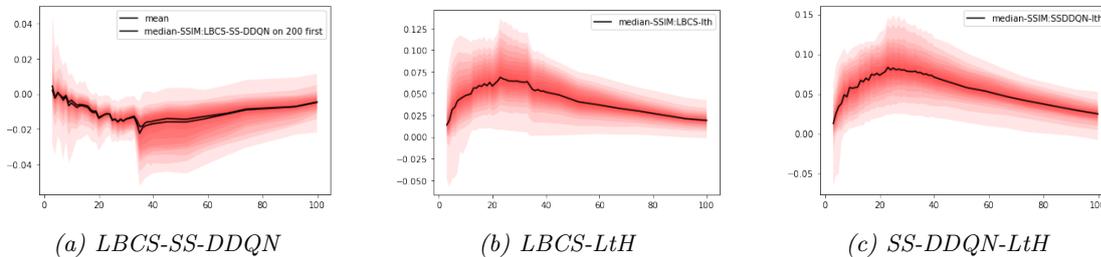


Figure A.1 – Per Image distribution of SSIM differences between sampling methods across the sampling process. Each shade corresponds to a 10 percent region, with the lightest shade indicating max and min regions.

A.3. LBCS vs long-range adaptivity

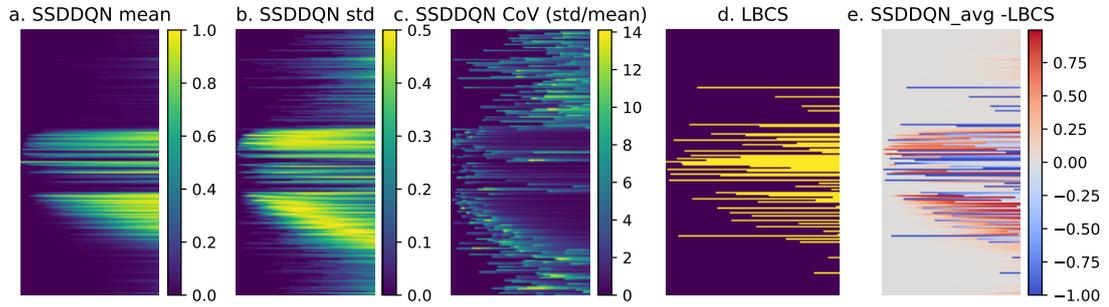
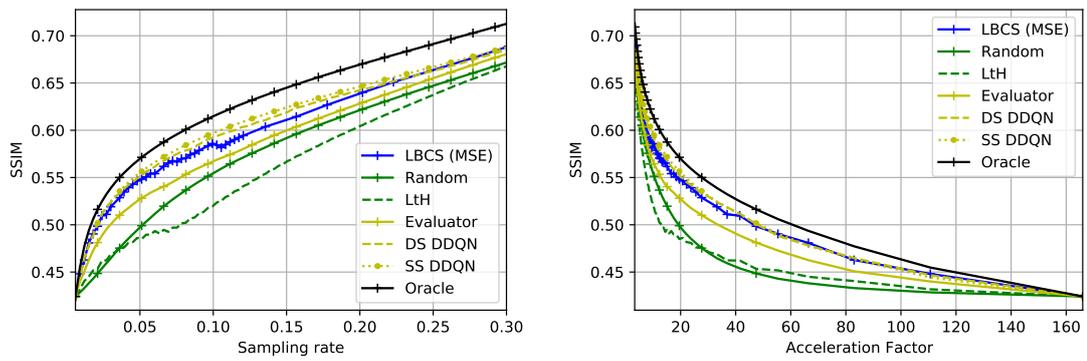


Figure A.2 – SS-DDQN variability and comparison with the LBCS mask. CoV stands for Coefficient of Variation here.



(a) SSIM vs sampling rate

(b) SSIM vs acceleration rate

Figure A.3 – Two ways to report the same result, SSIM version of fig. 6.3

A.4 LBCS complexity and parameters

Table A.8 summarizes the hyperparameters for sLBCS used and Table A.7 gives an overview of the computational complexity.

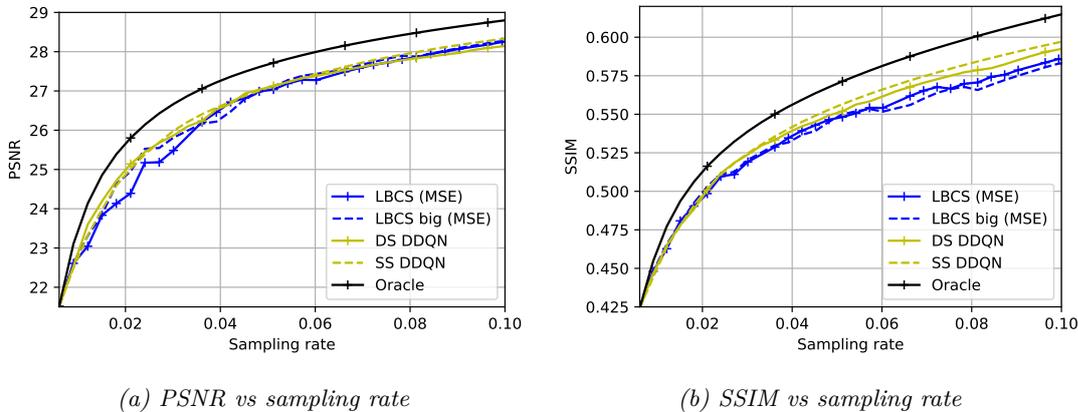
As we previously discussed, sLBCS is almost completely parallelizable, which leads to the stark runtime differences noted in Appendix A.1. As can be seen when comparing Figure 6.3 and Figure A.4, once a certain data batch size l and sampling candidate set cardinality k is reached, LBCS performance saturates, although it continues to benefit especially in the low sampling regime. This means that one could very feasibly reduce l for the comparison with Bakker et al. (2020), we attempted to use a larger batchsize to explore what would be the result of scaling sLBCS up.

Method	Forward	Backward	Total
(Bakker et al., 2020)	$q(8)E(50)B(16)H(n_r + n_p) = 6'400H(n_r + n_p)$	$6400Hn_p$	$6'400H(n_r + 2n_p)$
LBCS for Bakker et al. (2020)	$n(H)k(128)l(256)n_r = 32'768Hn_r$	0	$32'768Hn_r$
(Pineda et al., 2020)	$5e6(n_r + n_p)$	$5e6(n_p)$	$5e6(n_r + 2n_p)$
LBCS for Pineda et al. (2020)	$n(100)k(64)l(20)n_r = 128e3n_r$	0	$128e3n_r$

Table A.7 – Approximate computational cost of the compared methods. Note that at test time, LBCS is basically free while the RL policies will still need to be deployed. n_r, n_p are the parameter counts of the reconstruction and sampling policies respectively

Setting	Num. lines $ S $	Max. cardinality N	Candidate set size k	Data batch size l
Short horizon (Bakker et al., 2020)	128	16	128	256
Long horizon (Bakker et al., 2020)	128	28	128	256
(Pineda et al., 2020)	332	100	64	20
(Pineda et al., 2020) "big"	332	48	200	256

Table A.8 – The hyperparameters used for the stochastic LBCS masks throughout comparisons



(a) PSNR vs sampling rate

(b) SSIM vs sampling rate

Figure A.4 – PSNR and SSIM of LBCS with larger batch size (LBCS - big) trained on MSE, zoomed to the region where the smaller batch size LBCS underperformed. Note that it fully matches or outperforms both versions of DDQN on PSNR.

A.5 Visual comparisons

In this section, we present a visual comparison over a selected set of models, policies, settings and images. We present a visual evaluation of the models and policies at sampling rates 25% and 12.5% on Figures A.5 and A.6. In addition, we present a more exhaustive set of reconstruction, at various sampling rates, on Figure A.7, where we display both types of sequences that were used to generate the data, namely proton density (PD) and proton density, fat saturated (PDFS) of images (Zbontar et al., 2019).

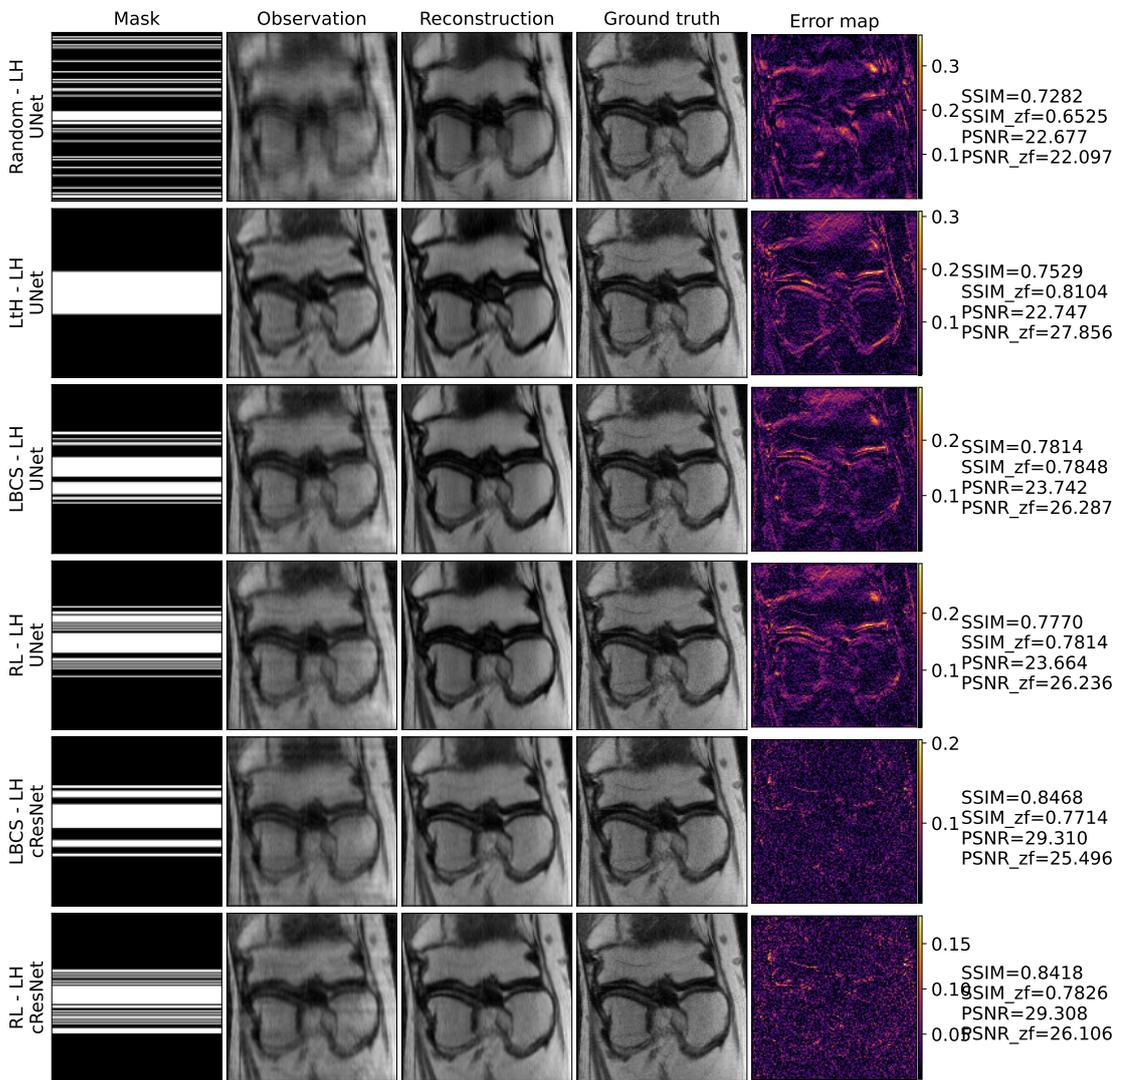


Figure A.5 – Visualization of masks, observations, reconstructions and ground truths and error maps ($|reconstruction - ground\ truth|$) at 25% sampling rate, for different policies (Random, LH, LBCS, RL). The data are processed according to the `c+rhz` setting, i.e. cropped then resized images, horizontal undersampling and Zhang-type distribution of masks. The SSIM and PSNR values are given on the right, and here, `zf` refers to zero-filled, and is the SSIM/PSNR taken between the observation and the ground truth.

Appendix A. Appendix for Chapter 6

Focusing first on Figures A.5 and A.6, we can observe that random and low-to-high sampling lie significantly behind the performance of LBCS and RL. This is the reason why they are not included in the rest of the figures. We can see on the observation that random sampling tends to miss important structures, and results in a severely aliased observation. On the contrary LtH, that focuses on low frequencies, obtains a good quality observation, but fails to yield an improvement after reconstruction, and generally loses out on higher frequency details, yielding a poorer performance especially on edges. In the rest of the comparison, it is hard to notice any significant visual difference between the images obtained by LBCS and the RL method of Bakker et al. (2020). It is however clear that in the `c+rhz` setting, the `cResNet` yields a significantly better and sharper reconstruction at 25% sampling, compared to the UNet. This is also confirmed by the results in Table A.4. On this particular image, this trend is also observed at 12.5% sampling rate on Figure A.6, but this is *not* a consistent trend, as this is not highlighted by the AUC computation of Table A.3.

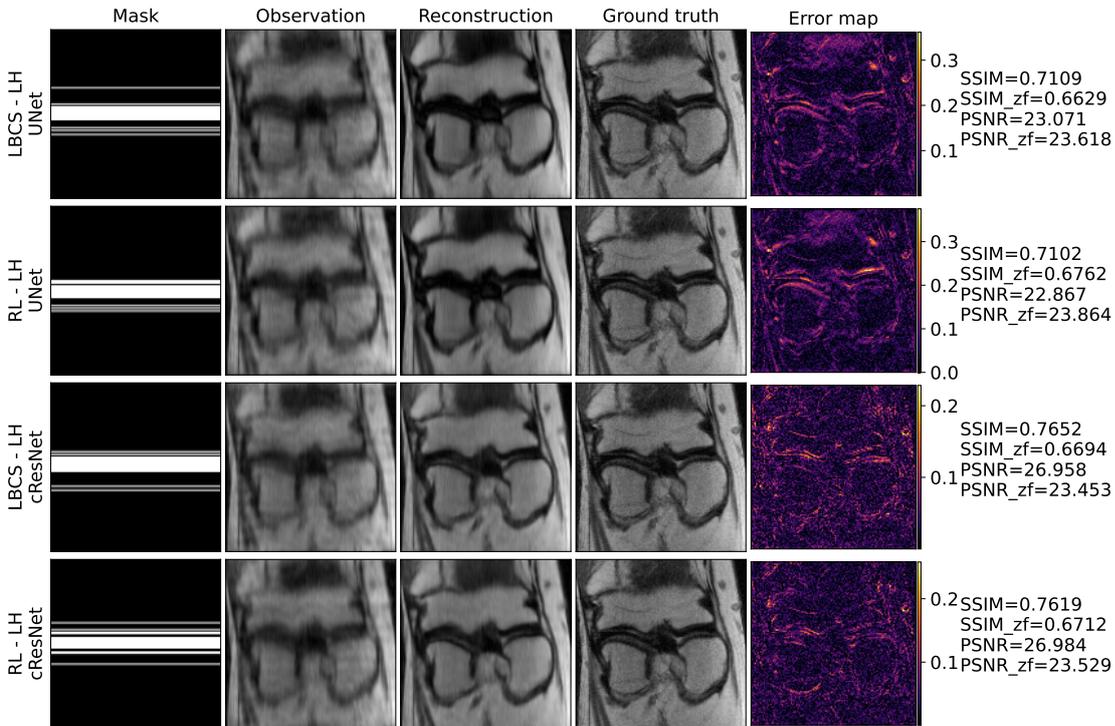


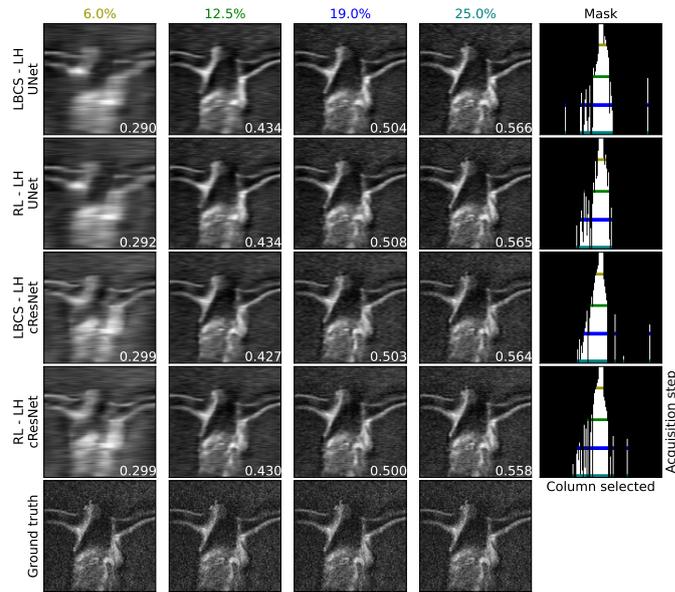
Figure A.6 – Additional visualization for the image displayed in Figure A.5. The results feature 12.5% sampling rate.

Turning now to Figure A.7, it is interesting first to discuss the policies obtained by LBCS and RL respectively. The LBCS policy is fixed for the reconstruction algorithm, so the first and third rows of Figures A.7a and A.7b will each feature the same policy. The adaptive RL policy, on the second and fourth rows seems to have central backbone of common frequency, but varies more around higher frequencies, as observed in the Figure 5 of the appendix of Bakker et al. (2020). However, in both cases, these differences have very little quantitative impact, and the same is true visually: there is no clear visual

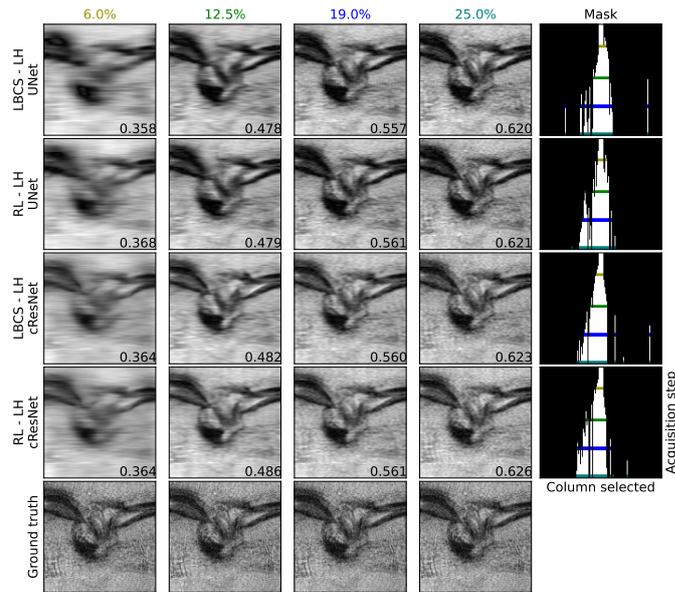
difference between the reconstructed images at either stage.

While this is not an exhaustive visual investigation, and does not directly assess the suitability of the different policies for various downstream tasks, this lack of visual difference could suggest that RL might not bring a significant improvement over simpler techniques such as LBCS in such cases. However, the question remains open in the case where the sampling policy would be tailored directly for the downstream task, rather than optimized for reconstruction quality, but this falls beyond the scope of the current work.

Appendix A. Appendix for Chapter 6



(a)



(b)

Figure A.7 – Visualization of reconstructed images at different sampling rates (6%, 12.5%, 19% and 25%) for two sampling policies (LBCS and RL) and two reconstruction algorithms (UNet and cResNet). The data are processed according to the *cvb* setting, i.e. cropped images, vertical undersampling and Bakker-type distribution of masks. The last row shows the ground truth (repeated), and each reconstruction has the corresponding SSIM displayed at the bottom right of the image. The rightmost column display the columns acquired during sampling (in white) as a function of the acquisition steps: starting on top with only center frequencies and progressively adding more and more lines to the sampling mask. The top plot (a) displays a proton density, fat saturated (PDFS) image, and the bottom plot displays a proton density (PD) image (Zbontar et al., 2019).

A.6 Detail on the practical recommendations

In this appendix, we discuss the recommendations issued in the conclusion, and provide the supporting evidence from our results.

Focus on improvements in the reconstructor architecture, mask distribution and algorithms used for training the reconstructor. Our results and ablations on the setting of Bakker et al. (2020) consistently show that the improvements obtained by changing the reconstructor architecture or the mask distribution are orders of magnitude more impactful than moving from LBCS to RL. This is supported the results of both Section 6.3.2 and Appendix A.2.3, where we see that moving from a UNet to a cResNet typically brings a significantly larger improvement (typically 0.01 SSIM) than what RL brings over LBCS in the best case (at most 0.0015 SSIM). This trend is also seen in Appendix A.2.3 when moving from *cvb* in the short horizon setting to *cvz* in the long horizon setting. We refer the reader to this section for more details.

Compare against strong baselines, such as LBCS. This point is established throughout our paper, where all results illustrate that, at best, RL methods bring moderate improvement over LBCS. This improvement often comes at the cost of prohibitive computational expense, even on higher end DGX-2 servers¹, while training LBCS required at most a couple of hours (cf. Appendix A.1).

Show sampling curves and use AUC to aggregate your results instead of performance at the final sampling rate. There is no consensus on how results should be aggregated from sampling curves. Bakker et al. (2020); Van Gorp et al. (2021) reported performance at the end of the acquisition, and Pineda et al. (2020) reported AUC curves computed on the acceleration factor. Gözcü et al. (2018); Yin et al. (2021) reported performance at selected sampling rates, and for other works such as Jin et al. (2019), it is not clear how results were aggregated.

We believe that reporting the AUC on sampling rates, computed on the whole range of acquisition steps allows to most meaningfully quantify the performance of a policy on its whole trajectory. It does not require to select a sampling rate at which the result should be evaluated, and using sampling rates as opposed to acceleration factor allows to equally weight the contribution of each acquired line.

Table 6.4 and Figure 6.3 compellingly illustrate that reporting the policy at a given sampling rate is not representative of its performance throughout the acquisition procedure.

Be mindful about preprocessing settings when evaluating a policy model. We recommend using the cropped+vertical setting with the data normalization implemented by Zbontar et al. (2019). We discussed in Section 6.2 that Bakker

¹The authors of (Pineda et al., 2020) confirmed to us that it took more than 20 days to train their model.

Appendix A. Appendix for Chapter 6

et al. (2020) used a data normalization that is, among other things, incompatible with data consistency, a commonly used building block for cascading networks (Schlemper et al., 2018; Zhang et al., 2019b). This can be prevented by using a normalization based on observation rather than ground truth statistics, as implemented in Zbontar et al. (2019).

Regarding the experimental setting, vertical masks have ubiquitously used on the fastMRI dataset (Zbontar et al., 2019; Huijben et al., 2020; Bakker et al., 2020; Pineda et al., 2020) and cropping has been the most common preprocessing to alleviate the computational burden of the large images (640×368) used in the dataset (Bakker et al., 2020; Huijben et al., 2020; Yin et al., 2021). Evaluating models on cropped data with vertical masks will then facilitate reproducibility among different works. We would additionally recommend to researchers to evaluate their models on the cropped+resized in addition to the cropped only setting, as the images display a significantly different field of view (cf. Figures A.5 and A.7).

B Appendix for Chapter 7

B.1 Implementation details

B.1.1 Knee experiments

The GAN experiment was carried out using a residual UNet (ResUNet) Belghazi et al. (2019), using the architecture described in Table B.2 below. The discriminator architecture is detailed in Table B.1. They consisted of respectively 1.2 million and 19.5 million parameters.

We also used the cascade of residual networks (cResNet) Zhang et al. (2019b) as our reconstruction baselines, along with their evaluator. One block of the architecture is given in Table B.3. The network was composed of 3 such blocks. The evaluator architecture is detailed in B.4. They amount to 29.5 million and 97.7 million parameters respectively. These numbers were obtained from the usage of Zhang et al. (2019b) by Pineda et al. (2020).

Type	Output size
Input	$6 \times 128 \times 128$
ResBlock ₁ (3, 2, 1)	$64 \times 64 \times 64$
ResBlock ₂ (3, 2, 1)	$128 \times 32 \times 32$
ResBlock ₃ (3, 2, 1)	$256 \times 16 \times 16$
ResBlock ₃ (3, 2, 1)	$512 \times 8 \times 8$
ResBlock ₃ (3, 2, 1)	$1024 \times 4 \times 4$
Global average pooling	$1024 \times 1 \times 1$
Conv(1, 1, 0)	$1 \times 1 \times 1$

Table B.1 – Discriminator ResNet. This was inspired by the discriminator of Belghazi et al. (2019). The model used ReLU activations, and consisted of 19.5 million parameters.

Appendix B. Appendix for Chapter 7

	Type	Output size	Comments
Input	–	$3 \times 128 \times 128$	
	ResBlock ₁ (3, 1, 1)	$64 \times 128 \times 128$	
Encoder	Avg Pool + ResBlock ₂ (3, 1, 1)	$64 \times 64 \times 64$	
	Avg Pool + ResBlock ₃ (3, 1, 1)	$64 \times 32 \times 32$	
	Avg Pool + ResBlock ₄ (3, 1, 1)	$64 \times 16 \times 16$	
	Avg Pool + ResBlock ₅ (3, 1, 1)	$64 \times 8 \times 8$	
Bottleneck	Avg Pool + ResBlock ₆ (3, 1, 1)	$64 \times 4 \times 4$	
	Upsample ₁ + ResBlock ₇ (3, 1, 1)	$64 \times 8 \times 8$	Cat[Upsample ₁ , ResBlock ₅]
Decoder	Upsample ₂ + ResBlock ₈ (3, 1, 1)	$64 \times 16 \times 16$	Cat[Upsample ₂ , ResBlock ₄]
	Upsample ₃ + ResBlock ₉ (3, 1, 1)	$64 \times 32 \times 32$	Cat[TrConv ₃ , ResBlock ₃]
	Upsample ₄ + ResBlock ₁₀ (3, 1, 1)	$64 \times 64 \times 64$	Cat[TrConv ₄ , ResBlock ₂]
	Upsample ₅ + ResBlock ₁₁ (3, 1, 1)	$64 \times 128 \times 128$	Cat[TrConv ₅ , ResBlock ₁]
	Conv(1, 1, 0) + Sigmoid + DC	$2 \times 128 \times 128$	

Table B.2 – ResUNet architecture used in the experiments, inspired by the architecture in Belghazi et al. (2019). The activation used throughout is a LeakyReLU with slope 0.2. We had $n_{channels,in}$ containing 2 channels for the observation, plus 1 channel for the noise. ResBlock(3, 1, 1) denotes a residual block with a 3×3 kernel size, stride 1 and padding 1. This results in a model with 1.2 million parameters. Bilinear upsampling is used.

	Type	Output size	Comments
Input	–	$2 \times 128 \times 128$	
	Conv ₁ (3, 2, 1)	$32 \times 64 \times 64$	
Encoder	Conv ₂ (3, 2, 1)	$64 \times 32 \times 32$	
	Conv ₃ (3, 2, 1)	$128 \times 16 \times 16$	
Bottleneck	ResBlock ₁ (3, 1, 1)	$128 \times 16 \times 16$	Skip-add from module $i - 1$
	ResBlock _{i} (3, 1, 1)	$128 \times 16 \times 16$	i blocks (total of 6 ResBlocks)
	ResBlock ₆ (3, 1, 1)	$128 \times 16 \times 16$	Skip to module $i + 1$
Decoder	TrConv ₁ (4, 2, 1)	$64 \times 32 \times 32$	
	TrConv ₂ (4, 2, 1)	$32 \times 64 \times 64$	
	TrConv ₃ (4, 2, 1)	$16 \times 128 \times 128$	
	Conv(1, 1, 0) + (Sigmoid) + DC	$3 \times 128 \times 128$	Sigmoid only at the last block

Table B.3 – One block of cResNet, used in Zhang et al. (2019b). The model used ReLU activations along with instance normalization. In the paper, we used 3 blocks with 6 residual blocks in the Bottleneck layer. We had $n_{channels,in}$ containing 2 channel for the observation. $n_{channels,out} = 3$: 2 channels for the reconstruction, and 1 channel for the variance. The entire model consisted of three such blocks, for a total of 29.5 million parameters.

	Type	Output size	Comments
Input spectral maps	–	$128 \times 128 \times 128$	
Input mask embedding	Conv(1, 1, 1)	$6 \times 1 \times 1$	Replicate & concatenate
Input tensor	–	$134 \times 128 \times 128$	
Evaluator	Conv(4, 2, 1) ₁	$256 \times 64 \times 64$	
	Conv(4, 2, 1) ₂	$512 \times 32 \times 32$	
	Conv(4, 2, 1) ₃	$1024 \times 16 \times 16$	
	Conv(4, 2, 1) ₄	$2048 \times 8 \times 8$	
	GAP	$1024 \times 1 \times 1$	
Output vector	Conv(1, 1, 0)	$128 \times 1 \times 1$	

Table B.4 – Evaluator architecture (Zhang et al., 2019b). First, spectral maps are created and concatenated with an embedded mask. Then, the data is processed through convolutional layers, using LeakyReLU(0.2) and instance normalization. This results in a total of 97.7 million parameters.

B.2 Extended results

We first provide additional knee results in Section B.2.1, then evaluate our GAN model on a small MRI dataset of brain images in Section B.2.2 and finish with additional MNIST evaluations in Section B.2.3.

B.2.1 Additional MRI knee results

In Figure B.1, we find the PSNR counterpart of the SSIM plot of Figure 7.1. It is interesting to see that there are noticeably different behaviors between the two models, and the results are overall consistent with the trends seen on the SSIM plot of Figure 7.1. Note also that the scales are different in each plot.

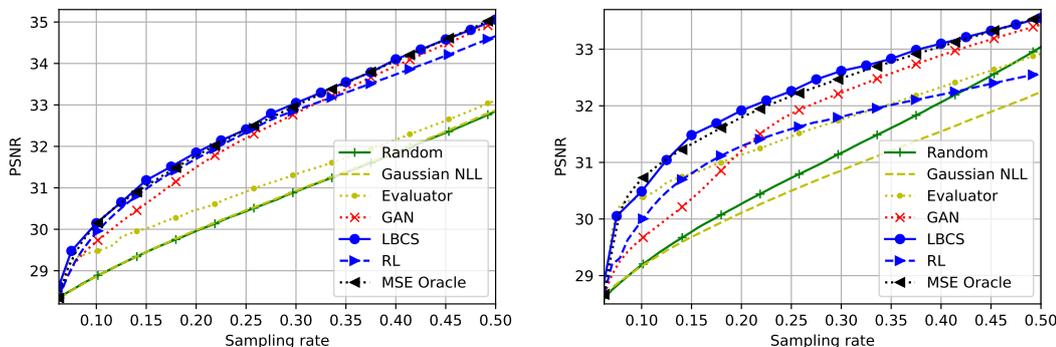


Figure B.1 – PSNR plots of results shown in Table 7.2, showing the different performance of the sampling methods. The left plot is the result of the evaluation on our GAN model, and the right plot is evaluated using the reconstructor of Zhang et al. (2019b). Although Table 7.2 computes the AUC until 25%, this table extends the evaluation until 50% sampling rate.

B.2.2 Brain experiment

Experimental setting

The dataset used in the first three experiments (subsections) below consists of a proprietary dataset of 2D T1-weighted brain scans. In our experiments, we use 100 slices of sizes 256×256 from five such subjects, 20 per subject. Three subjects (60 slices) were used for training the network, two subjects (30 slices) for testing. The data were then massively augmented with both rigid transformations and elastic deformations to counter overfitting as our dataset is very small, following the recommendations of Ronneberger et al. (2015); Schlemper et al. (2017). Namely, we apply both rigid transformations and elastic deformations. Specifically, at training time, each image was dynamically augmented with a randomly applied translation of ± 6 pixels along x - and y -axes, rotations of $[0, 2\pi)$, reflection along the x -axis with 50% probability. We also apply elastic deformations using the implementation in (Simard et al., 2003) with $\alpha \in [0, 40]$ and $\sigma \in [5, 8]$.

Appendix B. Appendix for Chapter 7

Data from individual coils was processed via a complex linear combination, where coil sensitivities were estimated from an 8×8 central calibration region of Fourier space (Bydder et al., 2002). The acquisition used a field of view (FOV) of 220×220 mm and a resolution of 0.9×0.7 mm. The slice thickness was 4.0mm. The imaging protocol comprised a flip angle of 70 deg, a TR/TE of 250.0/2.46 ms, with a scan time of 2 minutes and 10 seconds.

The training conditions for the Brain experiment were the same as the Knee experiment.

Results

We observe on Tables B.5 and B.6 results that are consistent with the ones on the knee dataset: LBCS outperforms both the Evaluator and the GAN, but it is interesting to observe that, in this setting, the GAN performance is quite close to the one of LBCS, especially in terms of SSIM. Note that in this context, the policy of Bakker et al. (2020) tends to perform worse than previously.

Policy	Model	
	GAN	Recon
Gaussian NLL	-	25.37
Evaluator	32.19	29.95
GAN	33.77	31.18
LBCS	34.03	32.03
RL	31.45	28.23

Table B.5 – Average PSNR test set AUC (one AUC per image) for the brain experiment.

Policy	Model	
	GAN	Recon
Gaussian NLL	-	0.72
Evaluator	0.83	0.85
GAN	0.88	0.87
LBCS	0.88	0.84
RL	0.80	0.80

Table B.6 – Average SSIM set AUC (one AUC per image) for the brain experiment.

B.2.3 Additional MNIST results

In Figure B.2, we provide additional reconstruction at different sampling rates for the different policies considered. Here GAS refers to generative adaptive sampling and corresponds to our GAN policy. We also provide the PSNR, SSIM and Accuracy against sampling rate on Figure B.3, which are the counterpart to Table 7.3. We see in particular that the strong performance of the GAN policy is due to the PSNR and SSIM reaching very high values around 5% sampling, where the sampling pattern has really acquired almost all pixels relevant to the digit. As there is almost no background, the error decreases very quickly. We see nonetheless that in every case, the adaptive policy manages to bring a noticeable improvement over LBCS, which was not the case in Fourier image.

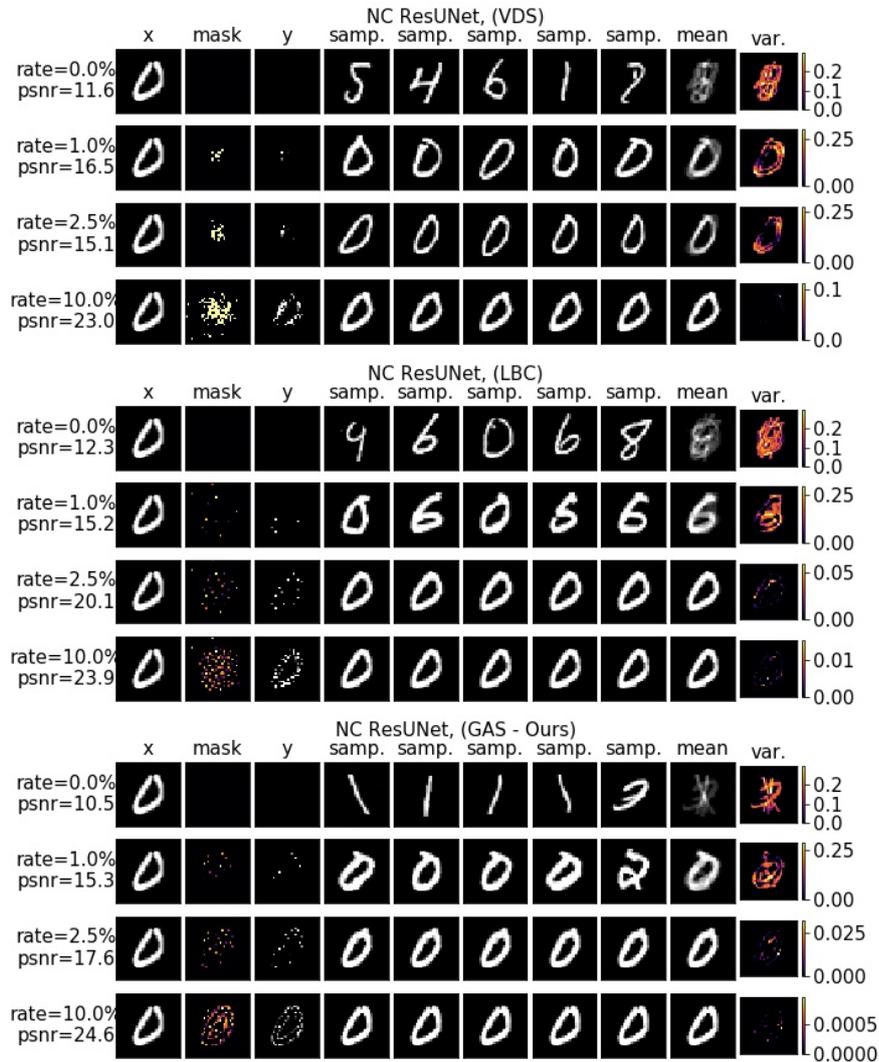


Figure B.2 – Additional Image domain results for the MNIST experiment, featuring different sampling rates, using a NC ResUNet model for conditional sampling.

Appendix B. Appendix for Chapter 7

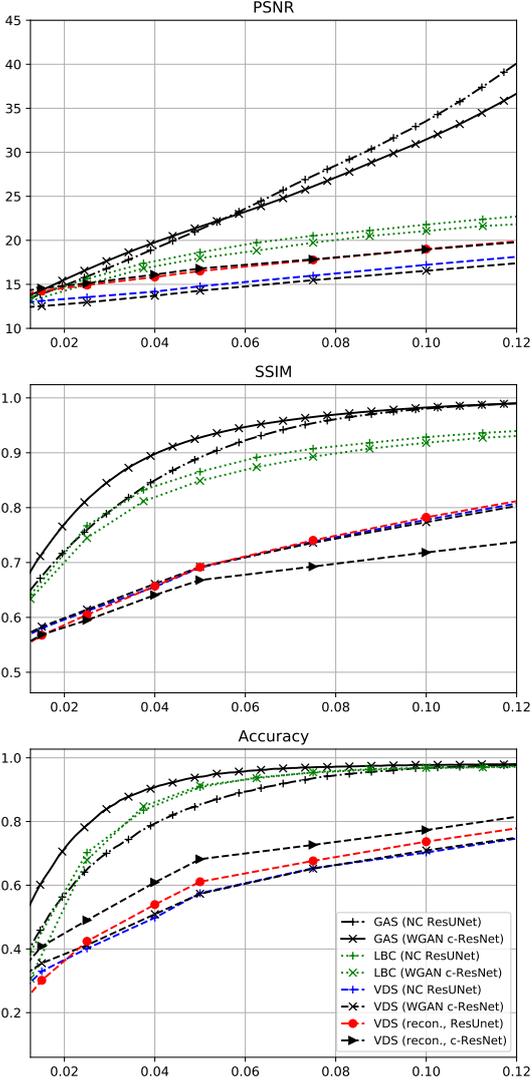


Figure B.3 – Image domain plots for the MNIST experiment, showing PSNR, SSIM and Accuracy for some selected models. These results are presented compactly in Table 7.3.

B.2.4 CIFAR10 experiment

Finally, we provide results on CIFAR10, for which the pixelwise subsampling task is much more challenging. We chose to use grayscale versions of the images.

Experiment setting. For CIFAR10, we modified the training procedures slightly. For the WGAN we use ExtraAdam (Gidel et al., 2018), TTUR (Heusel et al., 2017) with learning rates of $lr_G = 10^{-4}$ and $lr_D = 3 \cdot 10^{-4}$, one sided gradient penalty (Petzka et al., 2018) and betas of (0, 0.9). Number of epochs was set to 200 with the learning rates being halved every 50 steps. Batchsize was set to 128 and we used 8 GPUs in total, yielding an effective batch size of 1024. For the NC we retain batch size 128 and the Adam settings, simply moving to 8 GPUs and increasing training time to 200 epochs.

Results. Figure B.4 and fig. B.5 showcase GAS results using a WGAN and NC respectively. We evaluate VDS (variable-density sampling) and GAS (generative adaptive sampling, GAN policy) on a subset of the test set of CIFAR10 composed of 1000 images and report the results on Table B.7. The results are detailed in the plots of Figure B.6.

Algorithm	Model	Architecture	PSNR [dB] (\uparrow)	SSIM (\uparrow)	Accuracy [%] (\uparrow)
VDS	NC	ResUNet	9.32 ± 0.70	0.11 ± 0.03	0.10 ± 0.17
	WGAN	ResUNet	8.72 ± 0.62	0.09 ± 0.03	0.10 ± 0.17
VDS	Recon.	ResUNet	6.73 ± 0.51	0.09 ± 0.02	0.11 ± 0.17
	Recon.	c-ResNet	11.79 ± 0.72	0.13 ± 0.03	0.10 ± 0.16
GAS (Ours)	NC	ResUNet	14.36 ± 1.89	0.55 ± 0.08	0.27 ± 0.22
	WGAN	ResUNet	15.26 ± 2.34	0.57 ± 0.09	0.21 ± 0.18

Table B.7 – Average test set AUC (one AUC per image) with standard deviation on CIFAR10, in image domain. This was computed on a subset of the test set.

Appendix B. Appendix for Chapter 7

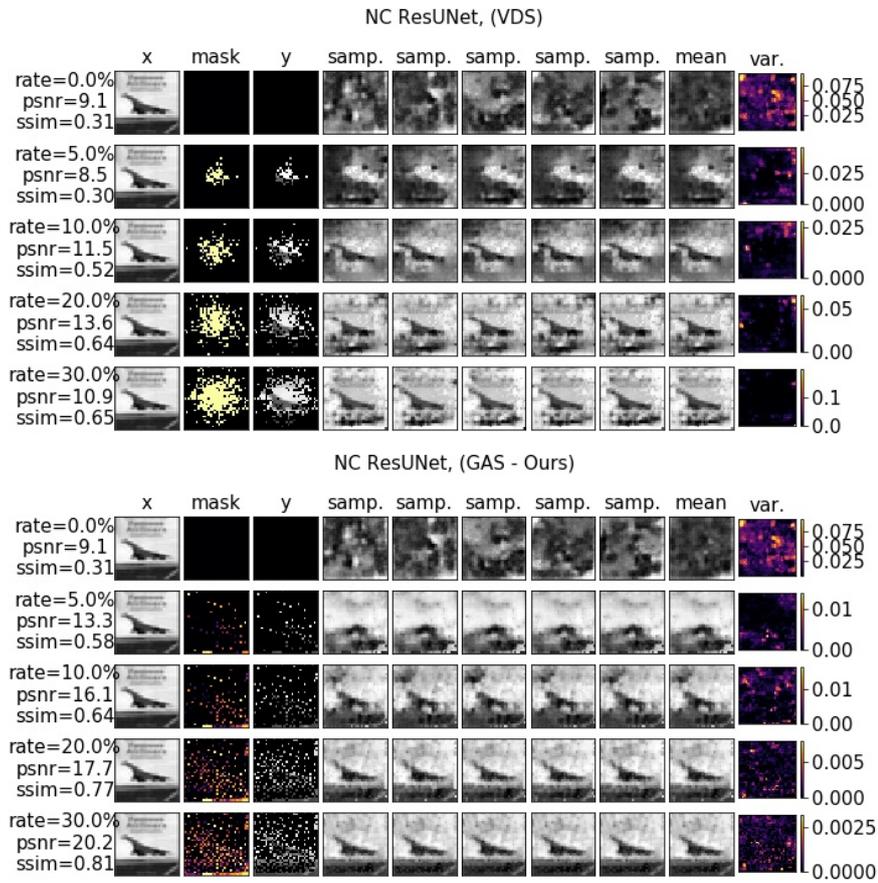


Figure B.4 – CIFAR10 results, featuring different sampling rates, using a NC ResUNet model for conditional sampling. Observe that while the model saw only masks going from 0.5% to 20% during training, it generalizes well up to 30%, gaining 2.5dB in the case of GAS. GAS stands for Generative Adaptive Sampling (GAN policy), and VDS refers to variable-density sampling policies.

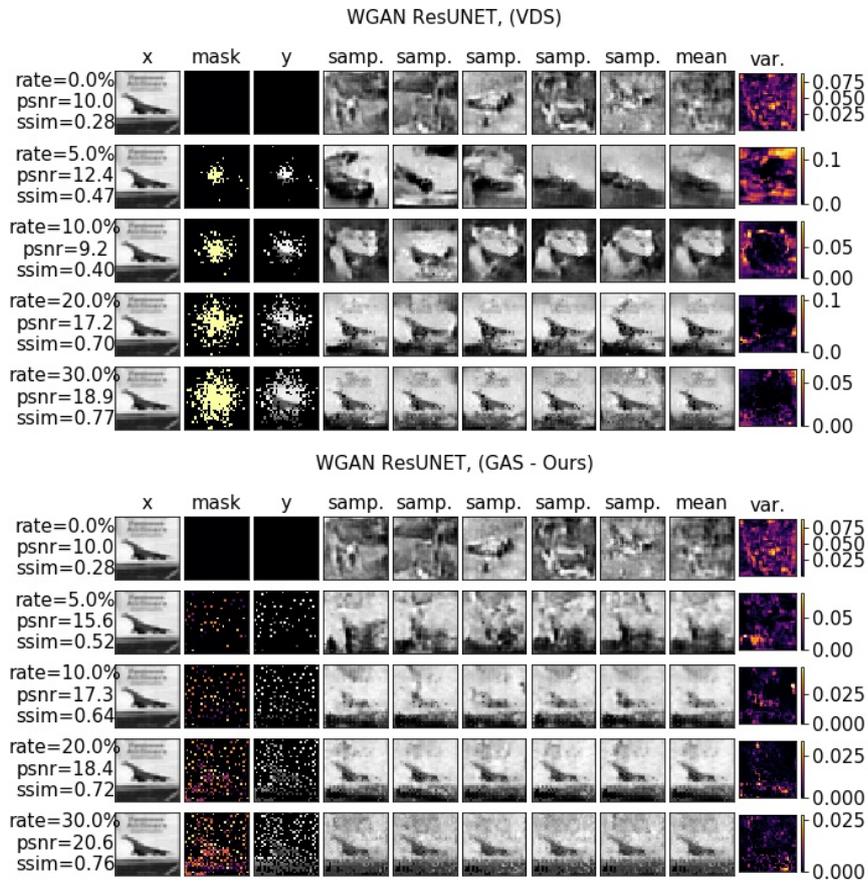


Figure B.5 – CIFAR10 results, featuring different sampling rates, using a WGAN ResUNET model for conditional sampling. Observe that while the model saw only masks going from 0.5% to 20% during training, it generalizes well up to 30%. GAS stands for Generative Adaptive Sampling (GAN policy), and VDS refers to variable-density sampling policies.

Appendix B. Appendix for Chapter 7

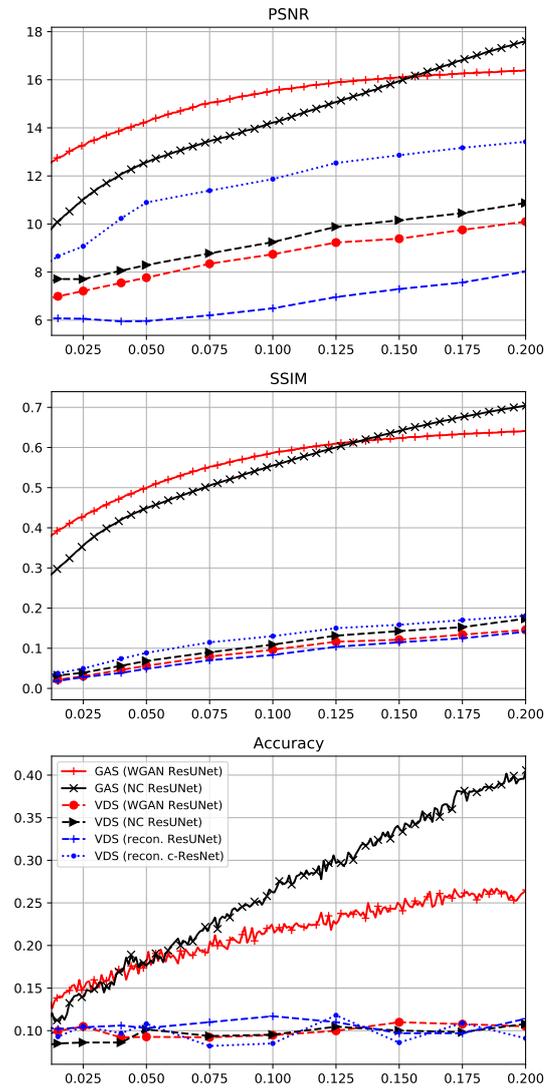


Figure B.6 – Image domain plots, showing PSNR, SSIM and Accuracy for some selected models. GAS stands for Generative Adaptive Sampling (GAN policy), and VDS refers to variable-density sampling policies.

Bibliography

- Adcock, B., Hansen, A. C., Poon, C., and Roman, B. (2017). Breaking the coherence barrier: A new theory for compressed sensing. In *Forum of Mathematics, Sigma*, volume 5. Cambridge University Press.
- Adcock, B., Hansen, A. C., and Roman, B. (2015). The quest for optimal sampling: Computationally efficient, structure-exploiting measurements for compressed sensing. In *Compressed Sensing and its Applications*, pages 143–167. Springer.
- Adler, J. and Öktem, O. (2018a). Deep bayesian inversion. *arXiv preprint arXiv:1811.05910*.
- Adler, J. and Öktem, O. (2018b). Learned primal-dual reconstruction. *IEEE transactions on medical imaging*, 37(6):1322–1332.
- Aggarwal, H. K. and Jacob, M. (2020). J-modl: Joint model-based deep learning for optimized sampling and reconstruction. *IEEE Journal of Selected Topics in Signal Processing*, 14(6):1151–1162.
- Anirudh, R., Thiagarajan, J. J., Kailkhura, B., and Bremer, T. (2018). An unsupervised approach to solving inverse problems using generative adversarial networks. *arXiv preprint arXiv:1805.07281*.
- Antun, V., Renna, F., Poon, C., Adcock, B., and Hansen, A. C. (2020). On instabilities of deep learning in image reconstruction and the potential costs of ai. *Proceedings of the National Academy of Sciences*, 117(48):30088–30095.
- Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein gan. *arXiv preprint arXiv:1701.07875*.
- Bahadir, C. D., Dalca, A. V., and Sabuncu, M. R. (2019). Learning-based optimization of the under-sampling pattern in MRI. In *International Conference on Information Processing in Medical Imaging*, pages 780–792. Springer.
- Bakker, T., Muckley, M. J., Romero-Soriano, A., Drozdal, M., and Pineda, L. (2021). On learning adaptive acquisition policies for undersampled multi-coil mri reconstruction.

Bibliography

- Bakker, T., van Hoof, H., and Welling, M. (2020). Experimental design for mri by greedy policy search. *Advances in Neural Information Processing Systems*, 33.
- Baldassarre, L., Li, Y.-H., Scarlett, J., Gözcü, B., Bogunovic, I., and Cevher, V. (2016). Learning-based compressive subsampling. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):809–822.
- Becker, S., Bobin, J., and Candès, E. J. (2011). NESTA: A fast and accurate first-order method for sparse recovery. *SIAM Journal on Imaging Sciences*, 4(1):1–39.
- Belghazi, M., Oquab, M., and Lopez-Paz, D. (2019). Learning about an exponential amount of conditional distributions. In *Advances in Neural Information Processing Systems 32*.
- Blaimer, M., Breuer, F., Mueller, M., Heidemann, R. M., Griswold, M. A., and Jakob, P. M. (2004). SMASH, SENSE, PILS, GRAPPA: how to choose the optimal method. *Topics in Magnetic Resonance Imaging*, 15(4):223–236.
- Block, K. T., Uecker, M., and Frahm, J. (2007). Undersampled radial mri with multiple coils. iterative image reconstruction using a total variation constraint. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 57(6):1086–1098.
- Bora, A., Jalal, A., Price, E., and Dimakis, A. G. (2017). Compressed sensing using generative models. *arXiv preprint arXiv:1703.03208*.
- Boyer, C., Chauffert, N., Ciuciu, P., Kahn, J., and Weiss, P. (2016). On the generation of sampling schemes for magnetic resonance imaging. *SIAM Journal on Imaging Sciences*, 9(4):2039–2072.
- Breuer, F. A., Blaimer, M., Mueller, M. F., Seiberlich, N., Heidemann, R. M., Griswold, M. A., and Jakob, P. M. (2006). Controlled aliasing in volumetric parallel imaging (2D CAIPIRINHA). *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 55(3):549–556.
- Bridson, R. (2007). Fast poisson disk sampling in arbitrary dimensions. In *SIGGRAPH sketches*, page 22.
- Brown, R. W., Cheng, Y.-C. N., Haacke, E. M., Thompson, M. R., and Venkatesan, R. (2014). *Magnetic resonance imaging: physical principles and sequence design*. John Wiley & Sons.
- Bydder, M., Larkman, D. J., and Hajnal, J. V. (2002). Combination of signals from array coils using image-based estimation of coil sensitivity profiles. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 47(3):539–548.

- Caballero, J., Price, A. N., Rueckert, D., and Hajnal, J. V. (2014). Dictionary learning and time sparsity for dynamic MR data reconstruction. *IEEE Transactions on Medical Imaging*, 33(4):979–994.
- Candes, E. and Romberg, J. (2007). Sparsity and incoherence in compressive sampling. *Inverse problems*, 23(3):969.
- Candès, E. J., Romberg, J., and Tao, T. (2006). Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Trans, on Inf. Theory*, 52(2):489–509.
- Candès, E. J. and Wakin, M. B. (2008). An introduction to compressive sampling. *IEEE signal processing magazine*, 25(2):21–30.
- Chaâri, L., Pesquet, J.-C., Benazza-Benyahia, A., and Ciuciu, P. (2011). A wavelet-based regularized reconstruction algorithm for SENSE parallel MRI with applications to neuroimaging. *Medical image analysis*, 15(2):185–201.
- Chaithya, G. and Ciuciu, P. (2022). Benchmarking learned non-cartesian k-space trajectories and reconstruction networks.
- Chaithya, G., Ramzi, Z., and Ciuciu, P. (2021). Learning the sampling density in 2d sparkling mri acquisition for optimized image reconstruction. In *2021 29th European Signal Processing Conference (EUSIPCO)*, pages 960–964. IEEE.
- Chaithya, G., Ramzi, Z., and Ciuciu, P. (2022). Hybrid learning of non-cartesian k-space trajectory and mr image reconstruction networks. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE.
- Chaloner, K. and Verdinelli, I. (1995). Bayesian experimental design: A review. *Statistical Science*, pages 273–304.
- Chauffert, N., Ciuciu, P., and Weiss, P. (2013). Variable density compressed sensing in MRI. theoretical vs heuristic sampling strategies. In *2013 IEEE 10th International Symposium on Biomedical Imaging*, pages 298–301. IEEE.
- Chen, Y., Schönlieb, C.-B., Liò, P., Leiner, T., Dragotti, P. L., Wang, G., Rueckert, D., Firmin, D., and Yang, G. (2022). Ai-based reconstruction for fast mri—a systematic review and meta-analysis. *Proceedings of the IEEE*, 110(2):224–245.
- Combettes, P. L. and Pesquet, J.-C. (2011). Proximal splitting methods in signal processing. In *Fixed-point algorithms for inverse problems in science and engineering*, pages 185–212. Springer.
- Coppo, S., Piccini, D., Bonanno, G., Chaptinel, J., Vincenti, G., Feliciano, H., Van Heeswijk, R. B., Schwitter, J., and Stuber, M. (2015). Free-running 4d whole-heart self-navigated golden angle mri: initial results. *Magnetic resonance in medicine*, 74(5):1306–1316.

Bibliography

- Coulom, R. (2006). Efficient selectivity and backup operators in monte-carlo tree search. In *International conference on computers and games*, pages 72–83. Springer.
- Darestani, M. Z., Chaudhari, A., and Heckel, R. (2021). Measuring robustness in deep learning based compressive sensing. *arXiv preprint arXiv:2102.06103*.
- Darestani, M. Z. and Heckel, R. (2021). Accelerated mri with un-trained neural networks. *IEEE Transactions on Computational Imaging*, 7:724–733.
- Daubechies, I., Defrise, M., and De Mol, C. (2004). An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 57(11):1413–1457.
- Denker, A., Schmidt, M., Leuschner, J., and Maass, P. (2021). Conditional invertible neural networks for medical imaging. *Journal of Imaging*, 7(11):243.
- Deshmane, A., Gulani, V., Griswold, M. A., and Seiberlich, N. (2012). Parallel MR imaging. *Journal of Magnetic Resonance Imaging*, 36(1):55–72.
- Doneva, M. (2020). Mathematical models for magnetic resonance imaging reconstruction: An overview of the approaches, problems, and future research areas. *IEEE Signal Processing Magazine*, 37(1):24–32.
- Donoho, D. L. (2006). Compressed sensing. *IEEE transactions on Information Theory*, 52(4):1289–1306.
- Douglas, L., Zarov, I., Gourgoulis, K., Lucas, C., Hart, C., Baker, A., Sahani, M., Perov, Y., and Johri, S. (2017). A universal marginalizer for amortized inference in generative models. *arXiv preprint arXiv:1711.00695*.
- Echternach, M., Sundberg, J., Arndt, S., Markl, M., Schumacher, M., and Richter, B. (2010). Vocal tract in female registers—a dynamic real-time mri study. *Journal of Voice*, 24(2):133–139.
- Edupuganti, V., Mardani, M., Vasanawala, S., and Pauly, J. (2020). Uncertainty quantification in deep mri reconstruction. *IEEE Transactions on Medical Imaging*, 40(1):239–250.
- El Gueddari, L., Ciuciu, P., Chouzenoux, E., Vignaud, A., and Pesquet, J.-C. (2019). Calibrationless oscar-based image reconstruction in compressed sensing parallel mri. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pages 1532–1536. IEEE.
- Engstrom, L., Ilyas, A., Santurkar, S., Tsipras, D., Janoos, F., Rudolph, L., and Madry, A. (2020). Implementation matters in deep rl: A case study on ppo and trpo. In *International Conference on Learning Representations*.

- Feng, L., Axel, L., Chandarana, H., Block, K. T., Sodickson, D. K., and Otazo, R. (2016a). Xd-grasp: Golden-angle radial MRI with reconstruction of extra motion-state dimensions using compressed sensing. *Magnetic resonance in medicine*, 75(2):775–788.
- Feng, L., Benkert, T., Block, K. T., Sodickson, D. K., Otazo, R., and Chandarana, H. (2016b). Compressed sensing for body MRI. *Journal of Magnetic Resonance Imaging*.
- Feng, L., Grimm, R., Block, K. T., Chandarana, H., Kim, S., Xu, J., Axel, L., Sodickson, D. K., and Otazo, R. (2014). Golden-angle radial sparse parallel MRI: Combination of compressed sensing, parallel imaging, and golden-angle radial sampling for fast and flexible dynamic volumetric MRI. *Magnetic Resonance in Medicine*, 72(3):707–717.
- Feng, L., Srichai, M. B., Lim, R. P., Harrison, A., King, W., Adluru, G., Dibella, E. V., Sodickson, D. K., Otazo, R., and Kim, D. (2013). Highly accelerated real-time cardiac cine MRI using $k-t$ SPARSE-SENSE. *Magnetic Resonance in Medicine*, 70(1):64–74.
- Fessler, J. A. and Sutton, B. P. (2003). Nonuniform fast fourier transforms using min-max interpolation. *IEEE Transactions on Signal Processing*, 51(2):560–574.
- Foster, A., Jankowiak, M., Bingham, E., Horsfall, P., Teh, Y. W., Rainforth, T., and Goodman, N. (2019). Variational bayesian optimal experimental design. In *Advances in Neural Information Processing Systems*, pages 14036–14047.
- Frahm, J., Haase, A., and Matthaei, D. (1986). Rapid NMR imaging of dynamic processes using the FLASH technique. *Magnetic Resonance in Medicine*, 3(2):321–327.
- Fritz, J., Raithel, E., Thawait, G. K., Gilson, W., and Papp, D. F. (2016). Six-fold acceleration of high-spatial resolution 3d space mri of the knee through incoherent k-space undersampling and iterative reconstruction—first experience. *Investigative radiology*, 51(6):400–409.
- Gamper, U., Boesiger, P., and Kozerke, S. (2008). Compressed sensing in dynamic MRI. *Magnetic Resonance in Medicine*, 59(2):365–373.
- Gautam, D., Lomeli, M., Gourgoulias, K., Thompson, D. H., and Johri, S. (2020). Masking schemes for universal marginalisers. *arXiv preprint arXiv:2001.05895*.
- Gidel, G., Berard, H., Vignoud, G., Vincent, P., and Lacoste-Julien, S. (2018). A variational inequality perspective on generative adversarial networks. *arXiv preprint arXiv:1802.10551*.
- Glockner, J. F., Hu, H. H., Stanley, D. W., Angelos, L., and King, K. (2005). Parallel MR imaging: a user’s guide. *Radiographics*, 25(5):1279–1297.
- Goodfellow, I. (2016). Nips 2016 tutorial: Generative adversarial networks. *arXiv preprint arXiv:1701.00160*.

Bibliography

- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680.
- Gözcü, B., Mahabadi, R. K., Li, Y.-H., Ilıcak, E., Çukur, T., Scarlett, J., and Cevher, V. (2018). Learning-based compressive MRI. *IEEE Transactions on Medical Imaging*.
- Gözcü, B., Sanchez, T., and Cevher, V. (2019). Rethinking sampling in parallel MRI: A data-driven approach. In *27th European Signal Processing Conference (EUSIPCO)*.
- Griswold, M. A., Jakob, P. M., Heidemann, R. M., Nittka, M., Jellus, V., Wang, J., Kiefer, B., and Haase, A. (2002a). Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magnetic Resonance in Medicine*, 47(6):1202–1210.
- Griswold, M. A., Walsh, D., Heidemann, R. M., Haase, A., and Jakob, P. M. (2002b). The use of an adaptive reconstruction for array coil sensitivity mapping and intensity normalization. In *International Society for Magnetic Resonance in Medicine (ISMRM), Proceedings of the 10th Scientific Meeting*, page 2410.
- Gronauer, S., Gottwald, M., and Diepold, K. (2021). The successful ingredients of policy gradient algorithms. In Zhou, Z.-H., editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 2455–2461. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. (2017). Improved training of wasserstein gans. In *Advances in neural information processing systems*, pages 5767–5777.
- Haldar, J. P. and Kim, D. (2019). Oedipus: An experiment design framework for sparsity-constrained mri. *IEEE transactions on medical imaging*.
- Hammernik, K., Klatzer, T., Kobler, E., Recht, M. P., Sodickson, D. K., Pock, T., and Knoll, F. (2018). Learning a variational network for reconstruction of accelerated mri data. *Magnetic resonance in medicine*, 79(6):3055–3071.
- Hasselt, H. (2010). Double q-learning. *Advances in neural information processing systems*, 23:2613–2621.
- Hassidim, A. and Singer, Y. (2017). Submodular optimization under noise. In *Conference on Learning Theory*, pages 1069–1122. PMLR.
- Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., and Meger, D. (2018). Deep reinforcement learning that matters. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Hennig, J., Nauerth, A., and Friedburg, H. (1986). Rare imaging: a fast imaging method for clinical mr. *Magnetic resonance in medicine*, 3(6):823–833.

- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. (2017). Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in neural information processing systems*, pages 6626–6637.
- Huijben, I. A., Veeling, B. S., and van Sloun, R. J. (2020). Deep probabilistic subsampling for task-adaptive compressed sensing. In *International Conference on Learning Representations*.
- Huijben, I. A. M., Veeling, B. S., and van Sloun, R. J. G. (2020). Learning sampling and model-based signal recovery for compressed sensing mri. In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8906–8910.
- Hussein, S. A., Tirer, T., and Giryes, R. (2020). Image-adaptive gan based reconstruction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3121–3129.
- Isola, P., Zhu, J.-Y., Zhou, T., and Efros, A. A. (2017). Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134.
- Ivanov, O., Figurnov, M., and Vetrov, D. (2019). Variational autoencoder with arbitrary conditioning. In *International Conference on Learning Representations*.
- Jackson, J. I., Meyer, C. H., Nishimura, D. G., and Macovski, A. (1991). Selection of a convolution function for fourier inversion using gridding (computerised tomography application). *IEEE Transactions on Medical Imaging*, 10(3):473–478.
- Jalal, A., Arvinte, M., Daras, G., Price, E., Dimakis, A. G., and Tamir, J. I. (2021). Robust compressed sensing mri with deep generative priors. *arXiv preprint arXiv:2108.01368*.
- Jalali, S. and Yuan, X. (2019). Solving linear inverse problems using generative models. In *2019 IEEE International Symposium on Information Theory (ISIT)*, pages 512–516. IEEE.
- Ji, S., Xue, Y., and Carin, L. (2008). Bayesian compressive sensing. *IEEE Transactions on signal processing*, 56(6):2346–2356.
- Jin, K. H., Lee, D., and Ye, J. C. (2016). A general framework for compressed sensing and parallel MRI using annihilating filter based low-rank Hankel matrix. *IEEE Transactions on Computational Imaging*, 2(4):480–495.
- Jin, K. H., McCann, M. T., Froustey, E., and Unser, M. (2017). Deep convolutional neural network for inverse problems in imaging. *IEEE Transactions on Image Processing*, 26(9):4509–4522.
- Jin, K. H., Unser, M., and Yi, K. M. (2019). Self-supervised deep active accelerated MRI. *arXiv preprint arXiv:1901.04547*.

Bibliography

- Johnson, P. M., Jeong, G., Hammernik, K., Schlemper, J., Qin, C., Duan, J., Rueckert, D., Lee, J., Pezzotti, N., Weerd, E. D., et al. (2021). Evaluation of the robustness of learned mr image reconstruction to systematic deviations between training and test data for the models from the fastmri challenge. In *International Workshop on Machine Learning for Medical Image Reconstruction*, pages 25–34. Springer.
- Jung, H., Sung, K., Nayak, K. S., Kim, E. Y., and Ye, J. C. (2009). k-t FOCUSS: A general compressed sensing framework for high resolution dynamic MRI. *Magn. Reson. Med.*, 61(1):103–116.
- Jung, H., Yoo, J., and Ye, J. C. (2007). Generalized kt blast and kt sense using focuss. In *2007 4th IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 145–148. IEEE.
- Karras, T., Aila, T., Laine, S., and Lehtinen, J. (2017). Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*.
- Karras, T., Laine, S., and Aila, T. (2019). A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410.
- Kazemian, S., Baur, C., Kuijper, A., van Ginneken, B., Navab, N., Albarqouni, S., and Mukhopadhyay, A. (2020). Gans for medical image analysis. *Artificial Intelligence in Medicine*, 109:101938.
- Kendall, A. and Gal, Y. (2017). What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in neural information processing systems*, pages 5574–5584.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kingma, D. P. and Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Knoll, F., Bredies, K., Pock, T., and Stollberger, R. (2011a). Second order total generalized variation (tgv) for mri. *Magnetic resonance in medicine*, 65(2):480–491.
- Knoll, F., Clason, C., Diwoy, C., and Stollberger, R. (2011b). Adapted random sampling patterns for accelerated MRI. *Magnetic resonance materials in physics, biology and medicine*, 24(1):43–50.
- Kovachki, N., Baptista, R., Hosseini, B., and Marzouk, Y. (2020). Conditional sampling with monotone gans. *arXiv preprint arXiv:2006.06755*.
- Krahmer, F. and Ward, R. (2013). Stable and robust sampling strategies for compressive imaging. *IEEE transactions on image processing*, 23(2):612–622.

- Krause, A. and Golovin, D. (2014). Submodular function maximization.
- Kutyniok, G., Lim, W.-Q., and Reisenhofer, R. (2016). Shearlab 3D: Faithful digital shearlet transforms based on compactly supported shearlets. *ACM Trans. Math. Softw.*, 42(1):5.
- Latorre, F., Eftekhari, A., and Cevher, V. (2019). Fast and provable ADMM for learning with generative priors. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché Buc, F., Fox, E., and Garnett, R., editors, *Advances in Neural Information Processing Systems 32*, pages 12027–12039. Curran Associates, Inc.
- Lauterbur, P. C. (1973). Image formation by induced local interactions: examples employing nuclear magnetic resonance. *nature*, 242(5394):190–191.
- Lazarus, C., Weiss, P., Chauffert, N., Mauconduit, F., El Gueddari, L., Destrieux, C., Zemmoura, I., Vignaud, A., and Ciuciu, P. (2019). Sparkling: variable-density k-space filling curves for accelerated t2*-weighted mri. *Magnetic resonance in medicine*, 81(6):3643–3661.
- Li, S., Zhu, Y., Xie, Y., and Gao, S. (2018). Dynamic magnetic resonance imaging method based on golden-ratio cartesian sampling and compressed sensing. *PloS one*, 13(1):e0191569.
- Liang, D., Cheng, J., Ke, Z., and Ying, L. (2020). Deep magnetic resonance image reconstruction: Inverse problems meet neural networks. *IEEE Signal Processing Magazine*, 37(1):141–151.
- Lindley, D. V. (1956). On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, pages 986–1005.
- Lingala, S. G., Hu, Y., DiBella, E., and Jacob, M. (2011). Accelerated dynamic MRI exploiting sparsity and low-rank structure: k-t SLR. *IEEE Transactions on Medical Imaging*, 30(5):1042–1054.
- Liu, D.-d., Liang, D., Liu, X., and Zhang, Y.-t. (2012). Under-sampling trajectory design for compressed sensing MRI. In *Engineering in Medicine and Biology Society (EMBC), 2012 Annual International Conference of the IEEE*, pages 73–76. IEEE.
- Lustig, M., Donoho, D., and Pauly, J. M. (2007). Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine*, 58(6):1182–1195.
- Lustig, M., Donoho, D. L., Santos, J. M., and Pauly, J. M. (2008). Compressed sensing MRI. *IEEE signal processing magazine*, 25(2):72–82.
- Maggioni, M., Boracchi, G., Foi, A., and Egiazarian, K. (2012). Video denoising, de-blocking, and enhancement through separable 4-D nonlocal spatiotemporal transforms. *IEEE Transactions on image processing*, 21(9):3952–3966.

Bibliography

- Majumdar, A. and Ward, R. K. (2012). Calibration-less multi-coil mr image reconstruction. *Magnetic resonance imaging*, 30(7):1032–1045.
- Mardani, M., Sun, Q., Donoho, D., Pappas, V., Monajemi, H., Vasanawala, S., and Pauly, J. (2018). Neural proximal gradient descent for compressive imaging. In *Advances in Neural Information Processing Systems*, pages 9573–9583.
- Marseille, G., De Beer, R., Fuderer, M., Mehlkopf, A., and Van Ormondt, D. (1996). Nonuniform phase-encode distributions for mri scan time reduction. *Journal of Magnetic Resonance, Series B*, 1(111):70–75.
- Meyer, C. H., Hu, B. S., Nishimura, D. G., and Macovski, A. (1992). Fast spiral coronary artery imaging. *Magnetic resonance in medicine*, 28(2):202–213.
- Minoux, M. (1978). Accelerated greedy algorithms for maximizing submodular set functions. In *Optimization techniques*, pages 234–243. Springer.
- Mirza, M. and Osindero, S. (2014). Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*.
- Mirzasoleiman, B., Badanidiyuru, A., Karbasi, A., Vondrák, J., and Krause, A. (2015). Lazier than lazy greedy. In *AAAI*, pages 1812–1818.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533.
- Monahan, G. E. (1982). State of the art—a survey of partially observable markov decision processes: theory, models, and algorithms. *Management science*, 28(1):1–16.
- Muckley, M. J., Riemenschneider, B., Radmanesh, A., Kim, S., Jeong, G., Ko, J., Jun, Y., Shin, H., Hwang, D., Mostapha, M., Arberet, S., Nickel, D., Ramzi, Z., Ciuciu, P., Starck, J.-L., Teuwen, J., Karkalousos, D., Zhang, C., Sriram, A., Huang, Z., Yakubova, N., Lui, Y. W., and Knoll, F. (2021). Results of the 2020 fastmri challenge for machine learning mr image reconstruction. *IEEE Transactions on Medical Imaging*, 40(9):2306–2317.
- Nalisnick, E., Matsukawa, A., Teh, Y. W., Gorur, D., and Lakshminarayanan, B. (2018). Do deep generative models know what they don’t know? *arXiv preprint arXiv:1810.09136*.
- Narnhofer, D., Hammernik, K., Knoll, F., and Pock, T. (2019). Inverse gans for accelerated mri reconstruction. In *Wavelets and Sparsity XVIII*, volume 11138, page 111381A. International Society for Optics and Photonics.
- Nemhauser, G. L. and Wolsey, L. A. (1978). Best algorithms for approximating the maximum of a submodular set function. *Mathematics of operations research*, 3(3):177–188.

- Nemhauser, G. L., Wolsey, L. A., and Fisher, M. L. (1978). An analysis of approximations for maximizing submodular set functions I. *Mathematical programming*, 14(1):265–294.
- Oord, A. v. d., Kalchbrenner, N., Vinyals, O., Espeholt, L., Graves, A., and Kavukcuoglu, K. (2016). Conditional image generation with pixelcnn decoders. *arXiv preprint arXiv:1606.05328*.
- O’sullivan, J. (1985). A fast sinc function gridding algorithm for fourier inversion in computer tomography. *IEEE Transactions on Medical Imaging*, 4(4):200–207.
- Otazo, R., Kim, D., Axel, L., and Sodickson, D. K. (2010). Combination of compressed sensing and parallel imaging for highly accelerated first-pass cardiac perfusion MRI. *Magnetic Resonance in Medicine*, 64(3):767–776.
- Pan, Z., Yu, W., Yi, X., Khan, A., Yuan, F., and Zheng, Y. (2019). Recent progress on generative adversarial networks (gans): A survey. *IEEE Access*, 7:36322–36333.
- Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., and Lakshminarayanan, B. (2019). Normalizing flows for probabilistic modeling and inference. *arXiv preprint arXiv:1912.02762*.
- Patel, D. V. and Oberai, A. A. (2021). Gan-based priors for quantifying uncertainty in supervised learning. *SIAM/ASA Journal on Uncertainty Quantification*, 9(3):1314–1343.
- Pauly, J. (2007). Non-cartesian reconstruction. URL: [http://www.stanford.edu/class/ee369c/notes/non cart rec](http://www.stanford.edu/class/ee369c/notes/non_cart_rec), 7.
- Pauwels, E., Lajaunie, C., and Vert, J.-P. (2014). A bayesian active learning strategy for sequential experimental design in systems biology. *BMC systems biology*, 8(1):102.
- Petzka, H., Fischer, A., and Lukovnikov, D. (2018). On the regularization of wasserstein GANs. In *International Conference on Learning Representations*.
- Pineda, L., Basu, S., Romero, A., Calandra, R., and Drozdal, M. (2020). Active mr k-space sampling with reinforcement learning. *arXiv preprint arXiv:2007.10469*.
- Pipe, J. G. and Menon, P. (1999). Sampling density compensation in mri: rationale and an iterative numerical solution. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 41(1):179–186.
- Pruessmann, K. P., Weiger, M., Scheidegger, M. B., Boesiger, P., et al. (1999). SENSE: sensitivity encoding for fast MRI. *Magnetic Resonance in Medicine*, 42(5):952–962.
- Puy, G., Vandergheynst, P., and Wiaux, Y. (2011). On variable density compressive sampling. *IEEE signal processing letters*, 18(10):595–598.

Bibliography

- Ramzi, Z., Chaithya, G., Starck, J.-L., and Ciuciu, P. (2022). Nc-pdnet: a density-compensated unrolled network for 2d and 3d non-cartesian mri reconstruction. *IEEE Transactions on Medical Imaging*.
- Ramzi, Z., Remy, B., Lanusse, F., Starck, J.-L., and Ciuciu, P. (2020). Denoising score-matching for uncertainty quantification in inverse problems.
- Rauhut, H. (2010). Compressive sensing and structured random matrices. In *Theoretical foundations and numerical methods for sparse recovery*, pages 1–92. de Gruyter.
- Ravishankar, S. and Bresler, Y. (2011a). Adaptive sampling design for compressed sensing MRI. In *Engineering in Medicine and Biology Society, EMBC, 2011 Annual International Conference of the IEEE*, pages 3751–3755. IEEE.
- Ravishankar, S. and Bresler, Y. (2011b). MR image reconstruction from highly under-sampled k-space data by dictionary learning. *IEEE Transactions on Medical Imaging*, 30(5):1028–1041.
- Ravishankar, S., Ye, J. C., and Fessler, J. A. (2019). Image reconstruction: From sparsity to data-adaptive methods and machine learning. *Proceedings of the IEEE*, 108(1):86–109.
- Rezende, D. J., Mohamed, S., and Wierstra, D. (2014). Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*.
- Roemer, P. B., Edelstein, W. A., Hayes, C. E., Souza, S. P., and Mueller, O. (1990). The NMR phased array. *Magnetic Resonance in Medicine*, 16(2):192–225.
- Roman, B., Hansen, A., and Adcock, B. (2014). On asymptotic structure in compressed sensing. *arXiv preprint arXiv:1406.4178*.
- Ronneberger, O., Fischer, P., and Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. (2016). Improved techniques for training gans. *Advances in neural information processing systems*, 29.
- Sanchez, T. (2018). Learning-based non-cartesian compressed sensing for dynamic MRI. Master Thesis.
- Sanchez, T., Gözcü, B., van Heeswijk, R. B., Eftekhari, A., Ilıcak, E., Çukur, T., and Cevher, V. (2020a). Scalable learning-based sampling optimization for compressive dynamic MRI. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8584–8588.

- Sanchez, T., Krawczuk, I., Sun, Z., and Cevher, V. (2020b). Uncertainty-driven adaptive sampling via GANs. In *NeurIPS 2020 Workshop on Deep Learning and Inverse Problems*. Available at <https://openreview.net/pdf?id=1WLYCQmtvW>.
- Sawyer, A. M., Lustig, M., Alley, M., Uecker, P., Virtue, P., Lai, P., and Vasanawala, S. (2013). Creation of fully sampled mr data repository for compressed sensing of the knee. In *Proceedings of the 22nd Annual Meeting for Section for Magnetic Resonance Technologists, Salt Lake City, Utah, USA*.
- Schlemper, J., Caballero, J., Hajnal, J. V., Price, A., and Rueckert, D. (2017). A deep cascade of convolutional neural networks for mr image reconstruction. In *International Conference on Information Processing in Medical Imaging*, pages 647–658. Springer.
- Schlemper, J., Caballero, J., Hajnal, J. V., Price, A. N., and Rueckert, D. (2018). A deep cascade of convolutional neural networks for dynamic MR image reconstruction. *IEEE Transactions on Medical Imaging*, 37(2):491–503.
- Seeger, M., Nickisch, H., Pohmann, R., and Schölkopf, B. (2010). Optimization of k-space trajectories for compressed sensing by bayesian experimental design. *Magn. Reson. Med.*, 63(1):116–126.
- Seeger, M., Steinke, F., and Tsuda, K. (2007). Bayesian inference and optimal design in the sparse linear model. In *Artificial Intelligence and Statistics*, pages 444–451. PMLR.
- Seeger, M. W. and Nickisch, H. (2008). Compressed sensing and bayesian experimental design. In *Proceedings of the 25th international conference on Machine learning*, pages 912–919.
- Sherry, F., Benning, M., Reyes, J. C. D. l., Graves, M. J., Maierhofer, G., Williams, G., Schönlieb, C.-B., and Ehrhardt, M. J. (2019). Learning the sampling pattern for MRI. *arXiv preprint arXiv:1906.08754*.
- Shimron, E., Tamir, J. I., Wang, K., and Lustig, M. (2021). Subtle inverse crimes: Naively training machine learning algorithms could lead to overly-optimistic results. *arXiv preprint arXiv:2109.08237*.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. (2017). Mastering the game of go without human knowledge. *nature*, 550(7676):354–359.
- Simard, P. Y., Steinkraus, D., Platt, J. C., et al. (2003). Best practices for convolutional neural networks applied to visual document analysis. *Icdar*, 3(2003).
- Singla, A., Tschitschek, S., and Krause, A. (2016). Noisy submodular maximization via adaptive sampling with applications to crowdsourced image collection summarization. In *Thirtieth AAAI Conference on Artificial Intelligence*.

Bibliography

- Sodickson, D. K. and Manning, W. J. (1997). Simultaneous acquisition of spatial harmonics (smash): fast imaging with radiofrequency coil arrays. *Magnetic Resonance in Medicine*, 38(4):591–603.
- Song, Y., Shen, L., Xing, L., and Ermon, S. (2021). Solving inverse problems in medical imaging with score-based generative models. *arXiv preprint arXiv:2111.08005*.
- Strauss, R. R. and Oliva, J. B. (2021). Arbitrary conditional distributions with energy. *arXiv preprint arXiv:2102.04426*.
- Sun, H. and Bouman, K. L. (2020). Deep probabilistic imaging: Uncertainty quantification and multi-modal solution characterization for computational imaging. *arXiv preprint arXiv:2010.14462*, 9.
- Sun, J., Li, H., Xu, Z., et al. (2016). Deep admn-net for compressive sensing mri. In *Advances in neural information processing systems*, pages 10–18.
- Sun, Z., Latorre, F., Sanchez, T., and Cevher, V. (2021). A plug-and-play deep image prior. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8103–8107. IEEE.
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Tonolini, F., Lyons, A., Caramazza, P., Faccio, D., and Murray-Smith, R. (2019). Variational inference for computational imaging inverse problems. *arXiv preprint arXiv:1904.06264*.
- Trzasko, J. D. and Manduca, A. (2011). Calibrationless parallel mri using clear. In *2011 conference record of the forty fifth Asilomar conference on signals, systems and computers (ASILOMAR)*, pages 75–79. IEEE.
- Tsai, C.-M. and Nishimura, D. G. (2000). Reduced aliasing artifacts using variable-density k-space sampling trajectories. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 43(3):452–458.
- Tsao, J., Boesiger, P., and Pruessmann, K. P. (2003). k-t BLAST and k-t SENSE: Dynamic MRI with high frame rate exploiting spatiotemporal correlations. *Magnetic Resonance in Medicine*, 50(5):1031–1042.
- Uecker, M., Lai, P., Murphy, M. J., Virtue, P., Elad, M., Pauly, J. M., Vasanawala, S. S., and Lustig, M. (2014). ESPIRiT: an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA. *Magn. Reson. Med.*, 71(3):990–1001.
- Uecker, M., Ong, F., Tamir, J. I., Bahri, D., Virtue, P., Cheng, J. Y., Zhang, T., and Lustig, M. (2015). Berkeley advanced reconstruction toolbox. In *Proc. Intl. Soc. Mag. Reson. Med.*, volume 23, page 2486.

- Uecker, M., Zhang, S., Voit, D., Karaus, A., Merboldt, K.-D., and Frahm, J. (2010). Real-time MRI at a resolution of 20 ms. *NMR in Biomedicine*, 23(8):986–994.
- Ulyanov, D., Vedaldi, A., and Lempitsky, V. (2018). Deep image prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9446–9454.
- Van Gorp, H., Huijben, I., Veeling, B. S., Pezzotti, N., and Van Sloun, R. J. (2021). Active deep probabilistic subsampling. In *International Conference on Machine Learning*, pages 10509–10518. PMLR.
- Van Hasselt, H., Guez, A., and Silver, D. (2016). Deep reinforcement learning with double q-learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30.
- Van Oord, A., Kalchbrenner, N., and Kavukcuoglu, K. (2016). Pixel recurrent neural networks. In *International Conference on Machine Learning*, pages 1747–1756. PMLR.
- Vanlier, J., Tiemann, C. A., Hilbers, P. A., and van Riel, N. A. (2012). A bayesian approach to targeted experiment design. *Bioinformatics*, 28(8):1136–1142.
- Vannesjo, S. J., Haeberlin, M., Kasper, L., Pavan, M., Wilm, B. J., Barmet, C., and Pruessmann, K. P. (2013). Gradient system characterization by impulse response measurements with a dynamic field camera. *Magnetic resonance in medicine*, 69(2):583–593.
- Vasanawala, S., Murphy, M., Alley, M. T., Lai, P., Keutzer, K., Pauly, J. M., and Lustig, M. (2011). Practical parallel imaging compressed sensing MRI: Summary of two years of experience in accelerating body MRI of pediatric patients. In *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, pages 1039–1043. IEEE.
- Vellagoundar, J. and Machireddy, R. R. (2015). A robust adaptive sampling method for faster acquisition of MR images. *Magnetic resonance imaging*, 33(5):635–643.
- Villani, C. (2008). *Optimal transport: old and new*, volume 338. Springer Science & Business Media.
- Walsh, D. O., Gmitro, A. F., and Marcellin, M. W. (2000). Adaptive reconstruction of phased array mr imagery. *Magnetic Resonance in Medicine*, 43(5):682–690.
- Wang, G., Luo, T., Nielsen, J.-F., Noll, D. C., and Fessler, J. A. (2021). B-spline parameterized joint optimization of reconstruction and k-space trajectories (bjork) for accelerated 2d mri. *arXiv preprint arXiv:2101.11369*.
- Wang, H., Liang, D., and Ying, L. (2009). Pseudo 2d random sampling for compressed sensing mri. In *2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 2672–2675. IEEE.

Bibliography

- Wang, Z. and Arce, G. R. (2009). Variable density compressed image sampling. *IEEE Transactions on image processing*, 19(1):264–270.
- Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. (2004). Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612.
- Weiss, T., Senouf, O., Vedula, S., Michailovich, O., Zibulevsky, M., and Bronstein, A. (2019). Pilot: Physics-informed learned optimal trajectories for accelerated mri. *arXiv preprint arXiv:1909.05773*.
- Weiss, T., Vedula, S., Senouf, O., Michailovich, O., Zibulevsky, M., and Bronstein, A. (2020). Joint learning of cartesian under sampling and reconstruction for accelerated mri. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8653–8657. IEEE.
- Weng, L. (2018). A (long) peek into reinforcement learning. <https://lilianweng.github.io/lil-log/2018/02/19/a-long-peek-into-reinforcement-learning.html>.
- Whang, J., Lindgren, E. M., and Dimakis, A. G. (2020). Approximate probabilistic inference with composed flows. *arXiv preprint arXiv:2002.11743*.
- Winkelmann, S., Schaeffter, T., Koehler, T., Eggers, H., and Doessel, O. (2007). An optimal radial profile order based on the golden ratio for time-resolved MRI. *IEEE Transactions on Medical Imaging*, 26(1):68–76.
- Wright, G. (1997). Magnetic resonance imaging. *IEEE Signal Processing Magazine*, 14(1):56–66.
- Wu, S., Dimakis, A., Sanghavi, S., Yu, F., Holtmann-Rice, D., Storcheus, D., Rostamizadeh, A., and Kumar, S. (2019a). Learning a compressed sensing measurement matrix via gradient unrolling. In Chaudhuri, K. and Salakhutdinov, R., editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6828–6839, Long Beach, California, USA. PMLR.
- Wu, Y., Rosca, M., and Lillcrap, T. (2019b). Deep compressed sensing. In Chaudhuri, K. and Salakhutdinov, R., editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 6850–6860, Long Beach, California, USA. PMLR.
- Yaman, B., Hosseini, S. A. H., Moeller, S., Ellermann, J., Uğurbil, K., and Akçakaya, M. (2020). Self-supervised physics-based deep learning mri reconstruction without fully-sampled data. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pages 921–925. IEEE.

- Yang, G., Yu, S., Dong, H., Slabaugh, G., Dragotti, P. L., Ye, X., Liu, F., Arridge, S., Keegan, J., Guo, Y., and Firmin, D. (2018). Dagan: Deep de-aliasing generative adversarial networks for fast compressed sensing mri reconstruction. *IEEE Transactions on Medical Imaging*, 37(6):1310–1321.
- Yin, T., Wu, Z., Sun, H., Dalca, A. V., Yue, Y., and Bouman, K. L. (2021). End-to-end sequential sampling and reconstruction for mr imaging. *arXiv preprint arXiv:2105.06460*.
- Yoo, J., Jin, K. H., Gupta, H., Yerly, J., Stuber, M., and Unser, M. (2021). Time-dependent deep image prior for dynamic mri. *IEEE Transactions on Medical Imaging*.
- Yu, T., Hilbert, T., Piredda, G. F., Joseph, A., Bonanno, G., Zenkhri, S., Omoumi, P., Cuadra, M. B., Canales-Rodríguez, E. J., Kober, T., et al. (2022). Validation and generalizability of self-supervised image reconstruction methods for undersampled mri. *arXiv preprint arXiv:2201.12535*.
- Zbontar, J., Knoll, F., Sriram, A., Murrell, T., Huang, Z., Muckley, M. J., Defazio, A., Stern, R., Johnson, P., Bruno, M., Parente, M., Geras, K. J., Katsnelson, J., Chandarana, H., Zhang, Z., Drozdal, M., Romero, A., Rabbat, M., Vincent, P., Yakubova, N., Pinkerton, J., Wang, D., Owens, E., Zitnick, C. L., Recht, M. P., Sodickson, D. K., and Lui, Y. W. (2019). fastMRI: An Open Dataset and Benchmarks for Accelerated MRI.
- Zhang, C., Barbano, R., and Jin, B. (2021). Conditional variational autoencoder for learned image reconstruction. *Computation*, 9(11):114.
- Zhang, H., Dauphin, Y. N., and Ma, T. (2019a). Fixup Initialization: Residual learning without normalization. In *International Conference on Learning Representations*.
- Zhang, Y., Peterson, B. S., Ji, G., and Dong, Z. (2014). Energy preserved sampling for compressed sensing MRI. *Computational and mathematical methods in medicine*, 2014.
- Zhang, Z., Romero, A., Muckley, M. J., Vincent, P., Yang, L., and Drozdal, M. (2019b). Reducing uncertainty in undersampled MRI reconstruction with active acquisition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2049–2058.
- Zhao, H., Gallo, O., Frosio, I., and Kautz, J. (2016). Loss functions for image restoration with neural networks. *IEEE Transactions on computational imaging*, 3(1):47–57.
- Zhu, J.-Y., Zhang, R., Pathak, D., Darrell, T., Efros, A. A., Wang, O., and Shechtman, E. (2017). Multimodal image-to-image translation by enforcing bi-cycle consistency. In *Advances in neural information processing systems*, pages 465–476.
- Zibetti, M. V., Herman, G. T., and Regatte, R. R. (2020). Fast data-driven learning of mri sampling pattern for large scale problems. *arXiv preprint arXiv:2011.02322*.

Bibliography

Zijlstra, F., Viergever, M. A., and Seevinck, P. R. (2016). Evaluation of variable density and data-driven k-space undersampling for compressed sensing magnetic resonance imaging. *Investigative radiology*, 51(6):410–419.

Selected Publications

Conference papers

Z. Sun, F. Latorre, T. Sanchez, and V. Cevher, "A plug-and-play deep image prior," in *ICASSP*, pp. 8103–8107, 2021.

T. Sanchez and I. Krawczuk et al., "Uncertainty-driven adaptive sampling via GANs," in *NeurIPS 2020 Workshop on Deep Learning and Inverse Problems*, 2020.

T. Sanchez et al., "Scalable learning-based sampling optimization for compressive dynamic MRI," in *ICASSP 2020*, pp. 8584–8588, 2020.

B. Gözcü, T. Sanchez, and V. Cevher, "Rethinking sampling in parallel MRI: A data-driven approach," in *27th European Signal Processing Conference*, 2019.

Pre-prints

T. Sanchez, I. Krawczuk, and V. Cevher, "On the benefits of deep RL in accelerated MRI sampling," 2021. Under review.

T. Sanchez, I. Krawczuk, Z. Sun, and V. Cevher, "Closed loop deep bayesian inversion: Uncertainty driven acquisition for fast MRI," 2019.