From Classical to Unsupervised-Deep-Learning Methods for Solving Inverse Problems in Imaging

Harshit Gupta

Thèse N° 7360 (septembre 2020)

Thèse présentée à la faculté des sciences et techniques de l'ingénieur
pour l'obtention du grade de docteur ès sciences
et acceptée sur proposition du jury

Prof. Dimitri Van De Ville, *président*
Prof. Michael Unser, *directeur de thèse*
Prof. François Fleuret, *rapporteur*
Prof. Ender Konukoglu, *rapporteur*
Dr. Sjors Scheres, *rapporteur*

École polytechnique fédérale de Lausanne—2020

Cover design by Annette Unser
Printing and binding by Repro-EPFL
Typeset with LaTeX

Available at **http://bigwww.epfl.ch/**

Abstract

In this thesis, we propose new algorithms to solve inverse problems in the context of biomedical images. Due to ill-posedness, solving these problems require some prior knowledge of the statistics of the underlying images. The traditional algorithms, in the field, assume prior knowledge related to smoothness or sparsity of these images. Recently, they have been outperformed by the second generation algorithms which harness the power of neural networks to learn required statistics from training data. Even more recently, last generation deep-learning-based methods have emerged which require neither training nor training data.

This thesis devises algorithms which progress through these generations. It extends these generations to novel formulations and applications while bringing more robustness. In parallel, it also progresses in terms of complexity, from proposing algorithms for problems with 1D data and an exact known forward model to the ones with 4D data and an unknown parametric forward model.

We introduce five main contributions. The last three of them propose deep-learning-based latest-generation algorithms that require no prior training.

1) We develop algorithms to solve the continuous-domain formulation of inverse problems with both classical Tikhonov and total-variation regularizations. We formalize the problems, characterize the solution set, and devise numerical approaches to find the solutions.

2) We propose an algorithm that improves upon end-to-end neural-network-based second generation algorithms. In our method, a neural network is first trained as a projector on a training set, and is then plugged in as a projector inside the projected gradient descent (PGD). Since the problem is nonconvex, we relax the PGD to ensure convergence to a local minimum under some constraints. This method outperforms all the previous generation algorithms for Computed Tomog-

raphy (CT).

3) We develop a novel time-dependent deep-image-prior algorithm for modalities that involve a temporal sequence of images. We parameterize them as the output of an untrained neural network fed with a sequence of latent variables. To impose temporal directionality, the latent variables are assumed to lie on a 1D manifold. The network is then tuned to minimize the data fidelity. We obtain state-of-the-art results in dynamic magnetic resonance imaging (MRI) and even recover intra-frame images.

4) We propose a novel reconstruction paradigm for cryo-electron-microscopy (CryoEM) called CryoGAN. Motivated by generative adversarial networks (GANs), we reconstruct a biomolecule's 3D structure such that its CryoEM measurements resemble the acquired data in a distributional sense. The algorithm is pose-or-likelihood-estimation-free, needs no ab initio, and is proven to have a theoretical guarantee of recovery of the true structure.

5) We extend CryoGAN to reconstruct continuously varying conformations of a structure from heterogeneous data. We parameterize the conformations as the output of a neural network fed with latent variables on a low-dimensional manifold. The method is shown to recover continuous protein conformations and their energy landscape.

Key words: inverse problems, deep learning, cryo electron-microscopy, continuous conformations, dynamic magnetic resonance, computed tomography, deep image prior, total variation

Résumé

Dans cette thèse, nous proposons de nouveaux algorithmes pour résoudre des problèmes inverses dans le cadre d'images biomédicales. En raison de la mauvaise pose, la résolution de ces problèmes nécessite une connaissance préalable des statistiques des images sous-jacentes. Les algorithmes traditionnels, sur le terrain, supposent des connaissances préalables liées à la fluidité ou à la rareté de ces images. Récemment, ils ont été dépassés par les algorithmes de deuxième génération qui exploitent la puissance des réseaux de neurones pour apprendre les statistiques requises à partir des données d'entraînement. Plus récemment encore, des méthodes basées sur l'apprentissage en profondeur de dernière génération sont apparues, qui ne nécessitent ni formation ni données de formation.

Cette thèse conçoit des algorithmes qui progressent à travers ces générations. Il étend ces générations à de nouvelles formulations et applications tout en apportant plus de robustesse. En parallèle, il progresse également en termes de complexité, de la proposition d'algorithmes pour des problèmes avec des données 1D et un modèle direct exact connu à ceux avec des données 4D et un modèle direct paramétrique inconnu.

Nous introduisons cinq contributions principales. Les trois derniers proposent des algorithmes de dernière génération basés sur le deep learning qui ne nécessitent aucune formation préalable.

1) Nous développons des algorithmes pour résoudre la formulation dans le domaine continu de problèmes inverses avec des régularisations Tikhonov classiques et à variation totale. Nous formalisons les problèmes, caractérisons l'ensemble de solutions et concevons des approches numériques pour trouver les solutions.

2) Nous proposons un algorithme qui améliore les algorithmes de deuxième génération basés sur un réseau de neurones de bout en bout. Dans notre méthode,

un réseau de neurones est d'abord formé en tant que projecteur sur un ensemble d'entraînement, puis branché en tant que projecteur à l'intérieur de la descente en gradient projeté (PGD). Comme le problème n'est pas convexe, nous assouplissons le DPI pour assurer la convergence vers un minimum local sous certaines contraintes. Cette méthode surpasse tous les algorithmes de génération précédente pour la tomodensitométrie (CT).

3) Nous développons un nouvel algorithme prioritaire d'image profonde dépendant du temps pour les modalités qui impliquent une séquence temporelle d'images. Nous les paramétrons comme la sortie d'un réseau neuronal non formé alimenté par une séquence de variables latentes. Pour imposer une directionnalité temporelle, les variables latentes sont supposées se trouver sur une variété 1D. Le réseau est ensuite réglé pour minimiser la fidélité des données. Nous obtenons des résultats de pointe en imagerie par résonance magnétique dynamique (IRM) et récupérons même des images intra-trame.

4) Nous proposons un nouveau paradigme de reconstruction pour la cryo microscopie électronique (CryoEM) appelé CryoGAN. Motivés par des réseaux contradictoires génératifs (GAN), nous reconstruisons la structure 3D d'une biomolécule de telle sorte que ses mesures CryoEM ressemblent aux données acquises dans un sens distributionnel. L'algorithme est sans estimation de vraisemblance ni de pose ou de vraisemblance, ne nécessite aucun ab initio et il est prouvé qu'il a une garantie théorique de récupération de la véritable structure.

5) Nous étendons CryoGAN pour reconstruire des conformations variant continuellement d'une structure à partir de données hétérogènes. Nous paramétrons les conformations comme la sortie d'un réseau neuronal alimenté avec des variables latentes sur une variété de faible dimension. Il est démontré que la méthode récupère les conformations protéiques continues et leur paysage énergétique.

Mots clés : problèmes inverses, deep learning, cryo-microscopie électronique, conformations continues, résonance magnétique dynamique, tomographie par computational, deep image prior, variation totale

Then even non-existence was not there, nor existence,
whence all creation had its origin,
the creator, whether he fashioned it or whether he did not,
the creator, who surveys it all from highest heaven,
he knows or maybe even he does not know.

-Nasadiya Sukt, Rigved (10:129), 1200 BC

Dedicated to Bittu (Parikshit) and my family.

Acknowledgement

This thesis is the result of support from my friends, family, and mentors who showed faith in me throughout this journey. Firstly, I thank Prof. Michael Unser for supervising this thesis. His intuition, passion, and curiosity helped me build my research instinct. I also thank him for trusting me during the tough times and encouraging me with both words and actions. In future, I would like to emulate the culture he has cultivated in the group.

I express sincere thanks to jury president Prof. Dimitri Van de Ville, the jury members, Prof. FranÃğois Fleuret, Prof. Ender Konukoglu, and Dr. Sjors Scheres for reviewing and accepting this thesis.

I would like to thank Dr. Michael T. McCann (Mike), the creative genius, who with his beautiful mind helped me solve many personal and research problems. I tried to learn his technique of brutally analyzing concepts and problems. This was immensely useful during the course of this thesis. I thank Dr. Soham Basu for all his advises and moral support. I thank Dr. Daniel Schmitter, Dr. Denis Fortun, and Luc Zheng for being there in the beginning years of my PhD.

I was fortunate to share my office with Thanh-an Pham from whom I learnt to have fun while carrying out a PhD. His humor has been a good company for the last three years. I am glad to have Shayan Aziznejad as my lab mate and a dear friend. Over the course of the thesis I learnt many life perspectives from him, shared numerous views, and had a lot of fun. I thank Pablo Garcia for all the fun coffee meetings and gaming sessions. I thank Dr. Kyong Hwan Jin for initiating my journey into deep learning.

I would like to thank Dr. Daniel Sage for his help with almost all the aspects of my research life at the lab and Dr. Philippe Thevenaz for all the help in writing the research papers. I thank Dr. Emrah Bostan and Dr. Pedram Pad for being really patient mentors and office mates during the initial part of my thesis. I thank Dr. Masih Nilchian for all the life wisdom he shared with me.

I thank Dr. Laurene Donati, Dr. Anais Bodoul, Thomas Debarre, Fangshu Yang, Pakshal Bohra, Dr. Quentin Denoyelle, and Dr. Jaejun Yoo for all the fun, humor, and discussions we had in the past few years. I thank all the past members of our lab Dr. Ferreol Soulez, Prof. Adrien Depeursinge, Dr. Emmanuel Soubies, Dr. Zsuzsanna PÃijspoki, Dr. Virginie Uhlmann, Prof. Arash Amini, Leello Dadi, and Carlos Garcia for the shared memories. I thank the recent members or affiliates of the lab Joaquim Campos, Thong Huy Phan, Alexis Goujon, Dr. Pol del Aguila Pla, and Yan Liu for bringing fresh perspectives.

I would like to specially thank the Indian community in Lausanne. I had fun of a lifetime with Ranjith, Sai, Harshal, Kunhal, Maneesha, Sanket, Sourabh, Sagar, Salil, Tejal, Anjali, Venkat, Anand, Rishikesh, Teju, Aparajitha, Yanisha, Sean, Kavitha, Nithin, Chethana, Shravan, Mohit, Murali, Amrita, and Mayank. I thank my friends Kaushal, Ravi, Gagandeep, Arpit, Pawan, Saurabh, Shashank, Puru, Lakshman, Nikit, Rupam, Rajan from IIT Guwahati and Kirti, Khushal, Arpit, Rishiraj, Rishabh, Rajwardhan, and Rahul from childhood for all their support.

Lastly, I thank my grandparents, parents, brother, and my entire long and wide family without whom I could not have completed this thesis.

Contents

Introduction

Imagine we have an old weighing machine and a sealed paper bag that contains two apples. We would like to know the weight of each apple. However, the bag must remain sealed. When weighed the total weight appears to be 501 grams. This problem of finding the individual weights of the apples, defined here as our signal, from an indirect hint about them, called measurement, can be regarded as an inverse problem. The measurement is obtained from the signal through a forward model. In this case, the forward model is the operation of summation.

This particular inverse problem is ill-posed since one cannot find the exact weights of the two apples based solely on this one measurement. This is because of two reasons. Firstly, there is an uncertainty on their exact total weight due to the fact that the old machine can have an error margin and that the bag has some unknown weight. These latter factors are called noise which corrupts the true measurement. Secondly, even if the true total weight was known, still there will be an infinite combination of weights that would sum up to this exact total.

To address this problem, we could take the advantage of prior knowledge about the apples. For example, if we know that in this particular season all the apples have almost the same weight, then we can estimate that each apple weighs around 250.5 grams. We could also factor in the weight of the bag and the error margin of the machine to get an even better estimate. If we know that both of these contribute together at most 1 gram of variation in the total weight, then we can claim that the weight of each apple lies now in an interval, somewhere between 250-251 grams.

The resolution of ill-posed inverse problems is fundamental to modern imaging [1] such as fluorescence microscopy, computational tomogaphy (CT), magnetic resonance imaging (MRI), cryo-electron microscopy (Cryo-EM), single-molecule localization microscopy (SMLM), *etc.*(see Figure 1). These techniques are essential in understanding the structure of internal organs, cells, biomolecules, *etc.* which is important in diagnosing diseases, drug-discovery, and to deepen the knowledge about life mechanisms. A shared trait of these techniques is that they do not give direct access to the desirable image (signal) but only provide information about it through an indirect physical mechanism (forward model). This information (measurement) then has to be processed adequately to obtain the required signal. Unlike the apple problem, the signal in these cases is numerically large and the forward model is much more complicated.

For example, in CT, x-rays propagate through the organ to be imaged [2] before being measured by a detector. The intensity of these rays is altered depending on the absorption properties of the organ tissues. This measured data thus contains the information about the organ absorption, albeit in a very convoluted form. The task then is to recover the image of the absorption index of the organ (signal) from the collected information (measurement). This is done by solving an inverse problem, which requires numerical *inversion* of the forward model associated with the device. This *inversion* is challenging since these problems are ill-posed. This ill-posedness stems from the following factors.

- The measurements are corrupted by noise. This can be thermal noise in the detector or intrinsic to the forward model or imaging setup.
- The measurements may not be fully informative because the forward model has an intrinsic null space.
- The measurements can be substantially fewer than the unknown variables in the signal. This means that there can be infinitely many signals that map to the same measurement. In fact, a reduction in the number of measurements is desirable. This is useful, for example, to reduce the radiation dose.

Therefore, solving these problems requires prior knowledge about the true signal. For example, it can be assumed that the CT image of the organ is piecewise-constant. The quality of reconstruction therefore substantially hinges on the ability to inject accurate prior knowledge during the reconstruction.

In all techniques (CT, MRI, Cryo-EM, SMLM, among others), the approaches to solve the underlying inverse problems have dramatically shifted from the classical-

---

[1] COVID-19 image has been obtained from the website of Centre of Disease Control.

methods introduced in the 1970s to the recently introduced deep-learning-based methods.

In the classical (first-generation methods), the reconstruction task is formulated as an optimization problem [3, 4]. Given a vector of measurements, the signal is reconstructed by solving this optimization problem. There are two components in this formulation. The first is called *data fidelity*. It ensures that the recovered signal is such that its simulated measurements are close to the acquired measurements. The second component is called *regularization*; it is required to counteract the ill-posedness. This component is based on the prior knowledge about the true signal. The two components ensure that the reconstructed signal is consistent with the measurements, while being compatible with the prior knowledge.

Depending on the type of prior knowledge used, the classical reconstruction methods can be broadly classified in two categories.

1. Tikhonov-Based: These methods use quadratic regularization (Tikhonov) which assumes that the true signal has low energy in some domain. The resulting optimization problem is easy to solve, often through a linear operation on the measurements. However, this prior results in reconstructed images being too smooth.

2. Sparsity-Based: Non-quadratic priors became popular since mid 2000s [5–9]. They assume images to be sparse in some domain and result in more realistic reconstruction since they are better than Tikhonov-based priors at capturing the statistics of the true images. However, the resulting optimization problem is harder to solve. The solution is obtained by a nonlinear iterative procedure that has necessitated the development of new specialized optimization routines [10–13].

The main shortcoming of classical methods is that their priors are hand-picked from a small set of well-behaved priors so that the resulting optimization problem can be solved via traditional optimization routines. This limits the type of prior knowledge that can be injected into the reconstruction.

In the first half of 2010s, deep learning [14] became an integral part of many computer-vision applications such as image segmentation, classification, and object detection [15]. Deep learning has then been utilized to solve inverse problems since 2015 and has resulted in state-of-the-art reconstructions [16–20]. These second-generational methods utilize a training set containing measurement-signal pairs.

Depending on the way the network is used in the reconstruction, this second generation of methods can itself be divided into two categories.

1. Direct Feedforward Reconstruction: In this category, the training set is first used to learn a neural network that maps the measurements to their corresponding training signal [17–20]. Once trained, the neural network is fed with the actual measurements and the output of the network then yields the reconstruction. These methods reconstruct images with unprecedented quality, owing to the capacity of neural networks to faithfully learn complicated mappings. However, these methods do not actively inject the information from the measurements and rely only on the training data to understand the inversion of the underlying physics-based forward model. This might reconstruct image which lack consistency with the actual measurements.

2. Iterative Reconstruction with Consistency Constraints: Many approaches have been proposed to enforce data consistency in the solution [21–25], including one of the contributions of this thesis [26]. In these approaches, reconstruction is performed iteratively with information from the forward model and the measurements being injected in conjunction with the ability of the network to reconstruct quality images. In essence, methods of this category combine the power of neural networks with the variational formulation of the classical methods. They produce better results and favor robustness when there is a mismatch between the training-data and the image to be reconstructed.

The main disadvantage of these supervised learning-based methods is that they require training-data.

Recently, there has been an emergence of third-generation methods that use deep learning without requiring training or training-data. The most prominent representatives of this generation are methods based on deep image priors [27], which use an untrained network to solve inverse problems. In this scheme, a fixed random input is fed to the network. The network is then optimized to ensure that its output is consistent with the acquired measurements. The success of this scheme is explained by the neural network architecture which imposes an implicit regularization that favours the reconstruction of natural-looking images. However, this deep image prior needs more theoretical and experimental analysis. This will be needed to understand its effect, applicability, and limits. Another category has also emerged. It is one of the contributions of this thesis; there, generative adversarial networks (GANs) is used for the reconstruction [28].

Main Contributions

This thesis brings five main contributions to the field of inverse problems in imaging. These contributions progress from the classical methods of the first generation to the deep-learning-based methods of the last generation. They extend these generations to novel formulations and applications, all the while bringing more robustness. In parallel, they also progress in terms of complexity, from algorithms for problems with 1D data and an exact known forward model to the problems with 4D data and an unknown parametric forward model. We summarize these contributions in Figure 2.

1. Continuous-Domain Extension of Classical Methods (Chapter 2)

In order to further our understanding of the classical methods of reconstruction, we formulate and solve 1D linear inverse problems in the continuous domain using Tikhonov-and-sparsity-based regularizations. Our object of recovery is a function resulting from the minimization of a convex objective functional composed of a data-fidelity term and regularization. For the latter, we consider and compare the continuous-domain Tikhonov and generalized total-variation (gTV) regularizations. Using representer theorems, we derive the parametric form of the solutions. This form is then used to discretize the problems and to find the numerical solutions. For the Tikhonov case, we obtain a smooth solution that lives in a fixed subspace determined by the forward model. In the gTV case, the solution is sparse and composed of a few functions that depend on the regularization operator. The number of these functions is upper-bounded by the number of measurements. These results are in resonance with the discrete counterparts of the two cases, ℓ_2-and-ℓ_1 regularizations. We illustrate these results through experiments in 1D. Moreover, for the scenario of multiple solutions in the gTV case, we devise a scheme to find the extreme points of the solution set which is theoretically ensured to be sparse.

Related Publication: H. Gupta, J. Fageot, M. Unser, "Continuous-domain solutions of linear inverse problems with Tikhonov *versus* generalized TV regularization," *IEEE Transactions on Signal Processing*, vol. 66(17), pp. 4670-84, July 2018.

2. Deep-Learning-Based PGD for Iterative Reconstruction (Chapter 3)

Many second-generation approaches learn a neural-network-based measurement-to-image regressor from the training data. This learnt regressor is then directly used to reconstruct the image from a given measurement. However, these approaches lack a feedback mechanism to enforce consistency between the reconstructed image and the actual measurements. Our proposal is a plug-and-play scheme where the projector in the projected gradient descent (PGD) is replaced by a trained neural network. This iterative approach alternatively enforces measurement consistency while a convolutional neural network (CNN) recursively projects the input to the space of desirable images. In order to ensure convergence, we propose a relaxed PGD and prove that, under certain conditions, the scheme converges to some local minimum of the non-convex optimization problem. Our experiments on sparse-view CT show that the scheme produces state-of-the-art results and brings more robustness with respect to noise and training-data mismatch. This method has become a fundamental contributor in the emergence of CNN-based iterative algorithms belonging to the second generation.

Related Publication: H. Gupta, K.H. Jin, H.Q. Nguyen, M.T. McCann, M. Unser, "CNN-based projected gradient descent for consistent CT image reconstruction," *IEEE Transactions on Medical Imaging*, vol. 37(6), pp. 1440-53, May 2018.

3. Time-Dependent Deep Image Prior (Chapter 4)

As part of our thesis, we participated in the development of a third-generation deep-learning-based method for dynamic MRI and other similar time-dependent inverse problems. Our method requires neither prior training nor additional data and extends deep image prior to temporal sequence. For this, we first assume a time-ordered sequence of latent variables that are forced to lie on a low-dimensional manifold. A neural network is then optimized to transform this sequence of latent variables into a sequence of images such that the data-fidelity with the sequence of acquired measurements is maximized. The semantic prior knowledge is injected by enforcing constraints on the latent variables sequence. For non-periodic data, we assume the latent variable sequence to lie on a line segment. For approximately periodic data like cardiac motion, we assume it to lie on a helix. We show that our scheme results in state-of-the-art performance for both synthetic and experimental data. Moreover, the interpolation of latent variables even lets us recover intra-frame images. To the best of our knowledge, this is the first deep-learning method applicable for dynamic MRI which requires neither prior training nor training-data.

Related Preprint: K.H. Jin*, H. Gupta*, J. Yerly, M. Stuber, M. Unser, “Time-dependent deep image prior for dynamic MRI,” *arXiv [eess.IV]*, October, 2019. *cofirst authors

4. CryoGAN: Cryo-EM Reconstruction using GAN Framework (Chapter 5)

In Cryo-EM, we obtain 2D noisy tomographic projection of the separate instance of the same but randomly oriented biomolecule. The task is to reconstruct the 3D structure from these projection data. The current methods are likelihood-based and, hence, require estimation of either pose or the conditional distribution over the space of poses for each projection, a computationally expensive routine. In our thesis, we introduce a novel paradigm that we name CryoGAN. It require neither pose estimation, *ab-initio*, additional training-data, nor prior training. We aim to reconstruct a biomolecule whose simulated Cryo-EM projections are distributionally similar to the acquired data. Similarly to the GAN framework, this is done by training a neural network called a discriminator which is optimized to distinguish the two distributions. We then learn a biomolecule such that the discriminator is unable to distinguish the two distributions. We mathematically prove that this quest for distribution matching results in the recovery of the true structure. Our scheme reaches 8.6 Å on realistic synthetic Cryo-EM data. Our results on experimental data are promising; the next step will be to improve the neural network architecture and to rely on multiresolution approaches to become competitive with the state-of-the-art in the field.

Related Preprints: H. Gupta*, M.T. McCann*, L. Donati, M. Unser. âĂIJCryo-GAN: A new reconstruction paradigm for single-particle Cryo-EM via deep adversarial learning,âĂİ *Biorxiv*, March 2020. *cofirst authors

5. Reconstructing Continuous Conformations in CryoEM using GANs (Chapter 6)

In the field of structural biology, the determination of the continuously varying conformations of a biomolecule is crucial in understanding its behaviour. The major challenge in Cryo-EM is to reconstruct these conformations of a biomolecule from their Cryo-EM data. To this end, we propose an extension of CryoGAN, termed Multi-CryoGAN. In addition to the advantages of CryoGAN, our method sidesteps the conformation estimation for each projection, a step necessary in the current methods implicitly or explicitly. We propose to parameterize the conformation manifold using a neural network that is driven by latent variables sampled from a distribution. The task then is to optimize the network such that the Cryo-EM projections of the generated conformations are distributionally similar to the acquired data. We provide a mathematical guarantee of recovery of the true conformation landscape. Our method can successfully reconstruct conformations of a heat-shock protein from their realistic synthetic data in both the continuous and discrete conformation cases. To the best of our knowledge, this is the first method that can recover continuous conformations of a biomolecule in a standalone manner (without the need of external routines).

Related Publication: H. Gupta, T.H. Phan, J. Yoo, M. Unser, âĂIJMulti-CryoGAN: Reconstruction of continuous conformations in Cryo-EM using Generative Adversarial Networks,âĂİ *In review.*

H. Gupta, T.H. Phan, J. Yoo, M. Unser, "Multi-CryoGAN: Reconstruction of continuous conformations in Cryo-EM using Generative Adversarial Networks," Proc. European Conference on Computer Vision Workshops (ECCVW 2020) (Online, August 23-28), in press.

Organization of the Thesis

Chapter 1 deals with the mathematical formulation of inverse problems and investigates the three generations of methods that solve these problems. Chapter 2 to Chapter 6 focus on each of the mentioned contributions. This is followed by conclusion and the outlook of this thesis.

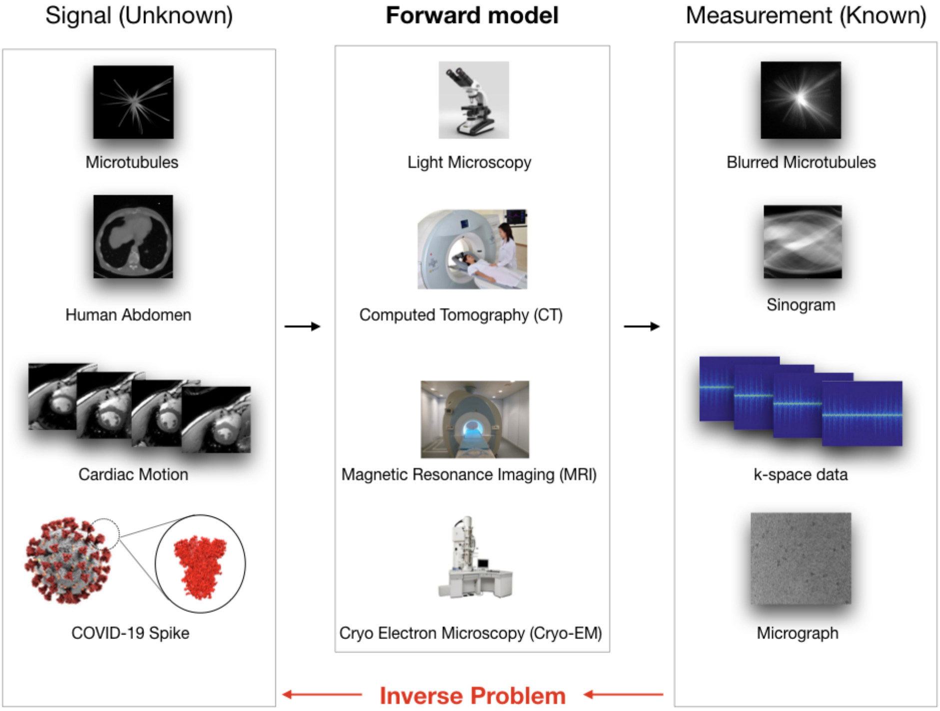

Figure 1: Inverse problems in biomedical imaging. The imaging device (forward model) is used to obtain the measurement from an unknown signal. This measurement is then used to reconstruct the signal. In conventional microscopy, the measurement is the blurred version of the signal. In CT, the scanner obtains the tomographic projection of the signal. An MRI scanner acquires the k-space data of the signal. [1]Cryo-EM provides the tomographic projection of separate instance of the same but randomly oriented biomolecule.

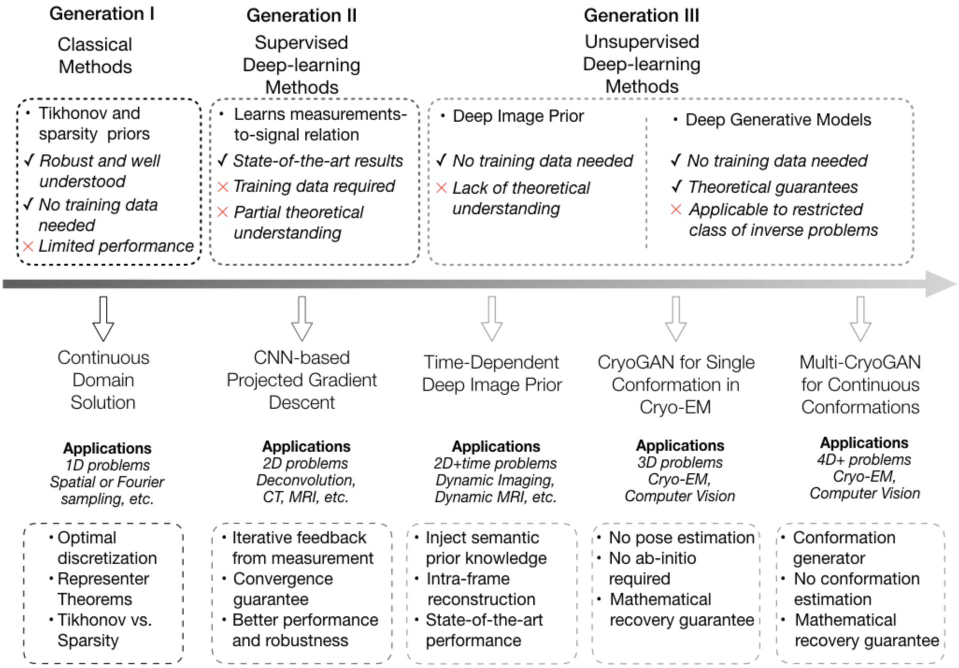

Figure 2: Contributions of this thesis in the context of the three generations of reconstruction methods for inverse problems.

Chapter 1

Linear Inverse Problems for Biomedical Imaging

[1] In this chapter, we will provide the mathematical formulation of inverse problems and the reconstruction methods.

1.1 Overview

Solving inverse problems implicitly requires inverting the forward model which incorporates the physical acquisition process by which the measurement is acquired from a signal. In the imaging applications dealt with in this thesis, the acquisition physics is such that the forward model is linear and hence can be represented by a linear operator $\mathbf{H} \in \mathbb{R}^M \rightarrow \mathbb{R}^N$. The measurement $\mathbf{y}$ is then given by

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n}, \tag{1.1}$$

where $\mathbf{x}$ is the space-domain signal[2] that we are interested in recovering and $\mathbf{n} \in \mathbb{R}^M$ is the noise intrinsic to the acquisition process. The noise can be assumed to be a random variable sampled from an appropriate probability distribution; for

---

[1] This chapter uses content from our work [26, 29].

[2] In this thesis, we use the term image and signal interchangeably.

example, a Gaussian distribution. Depending upon the imaging application, the signal can be an image, a movie, or a volume but here we describe it in the vectorized form. Moreover, the signal and the forward-model in reality are continuous-domain entities which need to be properly discretized in order to reach the formulation in (1.1). This aspect is discussed in detail in Chapter 2.

To find a suitable $\mathbf{x}$, the task is formulated as an optimization problem

$$\mathbf{x}^* = \arg\min_{\mathbf{x}\in\mathbb{R}^N} J(\mathbf{y}, \mathbf{x}), \tag{1.2}$$

where J is a suitable function on the space of signal and the given measurement. Hence, reconstruction procedure requires carefully choosing the function J and then finding the signal which minimizes it. Solving these problems when $M > N$ is easy and the algorithms that deal with this scenario are mature and efficient. In this thesis, we deal with the current trends in imaging. These trends may include significantly fewer measurements than the number of unknowns ($M \ll N$). For ex., this is useful in decreasing either the radiation dose in computed tomography (CT) or the scanning time in MRI. Moreover, the measurements are typically very noisy due to short integration times, which calls for some form of denoising.

This gives rise to an ill-posed problem in the sense that there may be an infinity of consistent signals (or images) that map to the same measurements $\mathbf{y}$. Thus, one challenge of the reconstruction algorithm is to select the best solution among a multitude of potential candidates. To counteract this ill-posedness, appropriate prior knowledge about the signal needs to be injected. The quality of the reconstruction method is therefore highly dependent on the data-fidelity and regularization used. The available reconstruction methods can be broadly arranged in three generations, which represent the continued efforts of the research community to address the aforementioned challenges.

1.2 Generation I - Classical Methods

In this generation the reconstruction task is formulated as solving

$$\mathbf{x}^* = \arg\min_{\mathbf{x}\in\mathbb{R}^N} (E(\mathbf{H}\mathbf{x}, \mathbf{y}) + \lambda R(\mathbf{x})), \tag{1.3}$$

where $E : \mathbb{R}^M \times \mathbb{R}^M \to \mathbb{R}^+$ is a data-fidelity term that favors solutions that are consistent with the measurements, $R : \mathbb{R}^N \to \mathbb{R}^+$ is a suitable regularizer that encodes prior knowledge about the image $\mathbf{x}$ to be reconstructed, and $\lambda \in \mathbb{R}^+$ is a tradeoff parameter. For example, E could be weighted least-squares $\|\mathbf{Hx} - \mathbf{y}\|^2$ and R is either a Tikhonov-or-sparsity-based regularization.

We discuss these two categories in detail.

1.2.1 Tikhonov-Prior-based Classical Methods

In these methods, the prior is expressed by a quadratic functional that is easy to differentiate. are simple, quadratic, and differentiable. For example $R(\mathbf{x}) = \|\mathbf{Lx}\|^2$ which yields,

$$\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathbb{R}^N} (E(\mathbf{Hx}, \mathbf{y}) + \lambda \|\mathbf{Lx}\|^2). \tag{1.4}$$

Due to the simplicity of the objective function, the reconstruction is easy to perform. In fact, for the case when $E(\mathbf{Hx}, \mathbf{y}) = \|\mathbf{Hx} - \mathbf{y}\|^2$ the solution is given by

$$\mathbf{x}^* = (\mathbf{H}^T\mathbf{H} + \lambda \mathbf{I})^{-1}\mathbf{H}^T\mathbf{y}. \tag{1.5}$$

Tikhonov-based classical methods are fast and provide excellent results when the number of measurements is large and the noise is small [30]. However, in case of extreme imaging, when M is smaller than N they reconstruct overly smooth image.

1.2.2 Sparsity-Prior-based Classical Methods

These methods avoid the shortcomings of the classical ones by assuming sparsity-based priors which are better at enforcing the true image statistics. For example, they could enforce the image to be sparse in some appropriate domain

$$\mathbf{x}^* = \arg \min_{\mathbf{x} \in \mathbb{R}^N} (E(\mathbf{Hx}, \mathbf{y}) + \lambda \|\mathbf{Lx}\|_1. \tag{1.6}$$

Here $\mathbf{L}$ could be a difference operator such that the regularization enforces image to be piecewise constant. Due to the non-differentiability of the regularization term, the solution to these problems cannot be found using simple gradient-based methods such as gradient descent. Instead the solutions are typically found iteratively by enforcing the data-fidelity and the regularization, alternatively. The latter is done using proximal operation. The overall skeleton of iterative algorithms can be summarised as follows

- *Data-fidelity update.* At iteration k, the estimate is updated to make it more consistent with its measurements. This is done by

  $$\tilde{\mathbf{x}}_k = \mathbf{x}_k - \gamma \nabla_{\mathbf{x}_k} E(\mathbf{H}\mathbf{x}_k, \mathbf{y}) \tag{1.7}$$

  where γ is an appropriate step size.

- *Regularization update.* In this step, the prior is enforced indirectly by solving

  $$\mathbf{x}_{k+1} = \arg\min_{\mathbf{x}} \|\tilde{\mathbf{x}}_k - \mathbf{x}\|_2^2 + \tau R(\mathbf{x}), \tag{1.8}$$

  where τ is an appropriate step-size. This step finds an image, which, on the one hand, minimizes the regularization cost and, on the other hand, maintains its similarity with the more data-consistent $\tilde{\mathbf{x}}_k$. Finding this solution is itself not easy since the non-differentiability of R is still an issue. However, this formulation decouples the data-fidelity term, which depends on the forward model, from the regularization term, rendering the latter task easier.

  Iterative schemes can be deployed to find the solution of (1.8). However, in some special regularization cases, an analytical form exist. For example, when $\mathbf{L} = \mathbf{I}$ such that $R(\mathbf{x}) = \|\mathbf{x}\|_1$, the solution of (1.8) is the soft-thresholding operation given by $\mathbf{x}^* = \text{sign}(\mathbf{x})(|\mathbf{x}| - \lambda)_+$. However, the solution for other $\mathbf{L}$ is generally not analytical and therefore requires additional computational overhead.

This alternate mechanism is the backbone of many schemes like Iterative Soft Thresholding Algorithm (ISTA) [31], Fast Iterative Soft Thresholding Algorithm (FISTA) [32], and Alternating Directions Methods of Multipliers (ADMM) [13]. In ISTA, the τ and γ are kept the same through the iterations. However, the rate of global convergence is sublinear, $O(1/k)$. In FISTA, these parameters are updated based on Nesterov acceleration. This brings faster speed of convergence than ISTA with convergence rate $O(1/k^2)$. Similar convergence rate is achieved in ADMM by decoupling the data-fidelity and regularization using auxiliary variables, and having a separate step to explicitly update these auxiliary variables.

Under the assumption that the functionals E and R are convex, the solution of (1.6) also satisfies

$$\mathbf{x}^* = \arg\min_{\mathbf{x}\in\mathcal{S}_R} E(\mathbf{H}\mathbf{x}, \mathbf{y}) \tag{1.9}$$

with $\mathcal{S}_R = \{\mathbf{x} \in \mathbb{R}^N : R(\mathbf{x}) \leq \tau\}$ for some unique τ that depends on the regularizaton parameter λ. Therefore, the solution has the best data fidelity among all images in the set $\mathcal{S}_R$ which is implicitly defined by R. This shows that the quality of the reconstruction depends heavily on the prior encoder R. As discussed, these priors are either handpicked (*e.g.*, total variation (TV) or the ℓ_1-norm of the wavelet coefficients of the image [5–9]) or learned through a dictionary [33–35]. However, in either case, they are restricted to well-behaved functionals that can be minimized via a convex routine [10–13]. This limits the type of prior knowledge that can be injected into the algorithm.

1.3 Generation II - Supervised-Deep-Learning-Based methods

Recently, researchers have obtained new state-of-the-art results for inverse problems with the help of the neural networks [16–20]. This has resulted in the emergence of new generation of deep-learning-based algorithms. A neural network CNN_ϕ is a function of the form

$$\mathit{CNN}_\phi = f_L(f_{L-1}(\dots f_1(\mathbf{x}))) \tag{1.10}$$

where $f_l = \rho(\mathbf{W}_l\mathbf{x}+\mathbf{a}_l)$ and $\phi = (\mathbf{W}_1, \mathbf{a}_1, \dots, \mathbf{W}_L, \mathbf{a}_L)$. There ρ is a pointwise nonlinear function and thus CNN_ϕ is composed of L layers each of which constitutes of a linear followed by a non-linear function. The parameters ϕ are learnable and thus, tune the function CNN_ϕ to behave in a desirable way. If the linear part is implemented by a convolution filter, then the network is called a convolutional neural network (CNN). These networks are more efficient for processing images since they are able to extract their spatial correlation in a shift invariant fashion.

This generation can be further divided into two categories.

1.3.1 Direct Feedforward Reconstruction

In these methods, it is assumed that we have access to training data

$$\{(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_Q, \mathbf{y}_Q)\}$$

which constitute pairs of signal and measurement. The network is first trained to map these measurement to their true corresponding image, and then is deployed on the new unseen measurement. This is given by

$$\text{Training}: \quad \phi^* = \arg\min \sum_{q=1}^{Q} \|CNN_{\phi}(\mathbf{y}^q) - \mathbf{x}^q\|^2, \tag{1.11}$$

$$\text{Reconstruction}: \quad \mathbf{x}^* = CNN_{\phi^*}(\mathbf{y}). \tag{1.12}$$

Often the measurement is followed by a fixed linear step like Back Projection (BP) or Filtered Back-Projection (FBP) before being fed to the CNN [17,18]. This helps in better convergence and to implicitly inject the physical knowledge during the learning.

Although it is now well documented that these methods have the capacity to outperform the classical algorithms, there is still limited theoretical understanding regarding them. Moreover, these methods are extremely sensitive to the mismatch between the statistics of the training and testing data. This is because the measurement information is injected into the reconstruction only once.

1.3.2 Iterative Reconstruction

In these methods, the advantages of classical and deep-learning methods are combined to obtain improved performance with robustness guarantees. Here the reconstruction is performed iteratively. Each step of iteration constitute a sub-step to maximize data-fidelity and a sub-step to clean the image using CNN. Our contribution in Chapter 3 discusses this framework in more depth.

1.4 Generation III - Unsupervised Deep-Learning-based methods

The second generation algorithms rely on the power of the neural networks to reconstruct state-of-the-art results but they require rich training data. The latter may not be available for many modalities. The latest generation algorithms have been devised to address this issue.

1.4.1 Deep Image Prior

In this algorithm [27], a neural network is optimized (not trained) to map a random vector $\mathbf{z}$ to an image under the constraint that the latter is the most consistent with the measurement $\mathbf{y}$. This is formulated by

$$\text{Optimization}: \quad \phi^* = \arg\min_{\phi} \|\mathbf{H}\{\mathit{CNN}_{\phi}(\mathbf{z})\} - \mathbf{y}\|^2, \tag{1.13}$$

$$\text{Solution}: \quad \mathbf{x}^* = \mathit{CNN}_{\phi^*}(\mathbf{z}), \tag{1.14}$$

where ϕ are the parameters of the CNN and $\mathbf{H}\mathit{CNN}_{\phi^*}(\mathbf{z}) = \mathbf{x}^*$ are the measurements associated with the reconstruction. CNNs have the ability to represent a large set of images. Therefore, they are susceptible to reconstruct undesirable images in (1.13), which minimizes the data fidelity. Yet, it has been observed that when the optimization is carried out using standard techniques the algorithm reconstructs desirable natural and biomedical images. In the case when the measurement is corrupted with noise, early stopping is required.

This bias or prior towards desirable images has been ascribed to the architecture of the network which not only lets it represent natural and biomedical images more easily than other images but even helps in reaching faster to these desirable images. In the literature this bias is called deep image prior.

1.4.2 Deep Generative Models

Methods like GAN [36] and Variational auto encoder (VAE) [37] have become quite popular in generative modelling. Until recently, their application in inverse problem was limited to learn priors from the training data. In this thesis, we show that there are certain modalities like Cryo-EM where deep generative models could play a central role in reconstruction without the need of training data. We shall delve deeper into this aspect in Chapter 5 and 6.

1.5 Summary

In this chapter, we give an overview of the reconstruction methods to solve linear inverse problems. We describe the optimization problem that are formulated for these methods and the routines deployed to solve them numerically. This will act as the background for to understand the next parts of our thesis.

Part I

First Generation

Chapter 2

Continuous-Domain Extension of Classical Methods

[1] In chapter 1, we formulated linear inverse problems. A subtle point to note in that formulation was that the signal, measurement, and the forward model were represented in discrete-domain (vectors). It is important to understand that the true signal and the forward model are continuous-domain entities. The measurements are generally discrete since they are obtained from a finite number of detectors. For example, MRI measurements are samples of the Fourier transform at a finite number of different frequencies of a continuous-domain signal. In this chapter, we further our understanding of classical methods by appropriate extension to solve continuous-domain linear inverse problems.

2.1 Overview

Although the signals and forward model that one encounters are in continuous-domain, they are discretized in order to numerically solve the inverse problems. This is done by choosing an arbitrary but suitable basis $\{\varphi_n\}$ and to represent the

---

[1] The content of this chapter are based on our work [29]

signal in continuous-domain as

$$f(x) = \sum_{n=1}^{N} \mathrm{f}_n \varphi_n(x), \tag{2.1}$$

where $\mathbf{f} = (\mathrm{f}_1, \dots, \mathrm{f}_N) \in \mathbb{R}^N$. Given the measurements $\mathbf{y} \in \mathbb{R}^M$, the task then is to find the expansion coefficients $\mathbf{f}$ by minimizing

$$\mathbf{f}^* = \arg\min_{\mathbf{f}\in\mathbb{R}^N} \left(\underbrace{\|\mathbf{y} - \mathbf{H}\mathbf{f}\|_2^2}_{\mathrm{I}} + \lambda \underbrace{\|\mathbf{L}\mathbf{f}\|_2^2}_{\mathrm{II}} \right), \tag{2.2}$$

where $\mathbf{H} \in \mathbb{R}^{M\times N}$ has elements $[\mathbf{H}]_{m,n} = \langle h_m, \varphi_n \rangle$. The analysis functions $\{h_m\}_{m=1}^M$ specify the forward model which encodes the physics of the measurement process. Term I in (2.2) is the data fidelity. It ensures that the recovered signal is close to the measurements. Term II is the regularization, which encodes the prior knowledge about the signal. The regularization is imposed on some transformed version of the signal coefficients using the matrix $\mathbf{L}$. Linear reconstruction algorithms [38,39] can be used to solve Problem (2.2). As discussed earlier, the notion that the real-world signals are sparse in some basis (e.g., wavelets) has resulted in better signal reconstruction in some fields. This prior is imposed by using the sparsity-promoting ℓ_1-regularization norm [40], [8] and results in the minimization problem

$$\mathbf{f}^* = \arg\min_{\mathbf{f}\in\mathbb{R}^N} \left(\|\mathbf{y} - \mathbf{H}\mathbf{f}\|_2^2 + \lambda \|\mathbf{L}\mathbf{f}\|_1 \right) \tag{2.3}$$

which can be efficiently solved using iterative algorithms [7], [41]. The solutions to (2.2), (2.3), and their variants with generalized data-fidelity terms are well known [42], [43], [44], [45].

While those discretization paradigms are well studied and used successfully in practice, it remains that the use of a prescribed basis $\{\varphi_n\}$, as in (2.1), is somewhat arbitrary.

In this chapter, we propose to bypass this limitation by reformulating and solving the linear inverse problem directly in the continuous domain. To that end, we impose the regularization in the continuous domain, too, and restate the reconstruction task as a functional minimization. We show that this new formulation leads to the identification of an optimal basis for the solution which then suggests a natural way to discretize the problem.

2.1.1 Contributions

Our contributions are two folds and are summarized as follows:

Theoretical.

- Given $\mathbf{y} \in \mathbb{R}^M$, we formalize 1D inverse problem in the continuous domain as
$$f_{\mathrm{R}}^* = \arg\min_{f \in \mathcal{X}} \underbrace{\left(\|\mathbf{y} - \mathrm{H}\{f\}\|_2^2 + \lambda \mathrm{R}(f)\right)}_{\mathrm{J}_{\mathrm{R}}(f|\mathbf{y})}, \tag{2.4}$$
where f is a function that belongs to a suitable function space $\mathcal{X}$. Similarly to the discrete regularization terms $\|\mathbf{Lf}\|_{\ell_2}^2$ and $\|\mathbf{Lf}\|_{\ell_1}$ in (2.2) and (2.3), we focus on their continuous-domain counterparts $\mathrm{R}(f) = \|\mathrm{L}f\|_{L_2}^2$ and $\mathrm{R}(f) = \|\mathrm{L}f\|_{\mathcal{M}}$, respectively. There, L and H are the continuous-domain versions of **L** and **H**, while $\|\mathrm{L}f\|_{\mathcal{M}}$ is the proper continuous-domain counterpart of the discrete ℓ_1 norm. We show that the effect of these regularizations is similar to the one of their discrete counterparts.

- We provide the parametric form of the solution (representer theorem) that minimizes $\mathrm{J}_{\mathrm{R}}(f|\mathbf{y})$ in (2.4) for the Tikhonov regularization $\mathrm{R}(f) = \|\mathrm{L}f\|_{L_2}^2$ and the generalized total-variation (gTV) regularization $\mathrm{R}(f) = \|\mathrm{L}f\|_{\mathcal{M}}$. Our results underline how the discrete regularization resonates with the continuous-domain one. The optimal solution for the Tikhonov case is smooth, while it is sparse for the gTV case. The optimal bases in the two cases are intimately connected to the operators L and H.

- We present theoretical results that are valid for any convex, coercive, and lower-semicontinuous data-fidelity term which is proper in the range of H. This includes the case when the data-fidelity term is $\|\mathbf{y} - \mathrm{H}\{f\}\|_2^2$. In this sense, for the gTV case our work extends the results in [46] which only deals with indicator function over a feasible convex-compact set as a data-fidelity term.

Algorithmic.

- We propose a discretization scheme to minimize $\mathrm{J}_{\mathrm{R}}(f|\mathbf{y})$ in the continuous domain. Even though the minimization of $\mathrm{J}_{\mathrm{R}}(f|\mathbf{y})$ is an infinite-dimensional problem, the knowledge of the optimal basis of the solution makes the problem finite-dimensional: it boils down to the search for a set of optimal expansion coefficients.

- We devise an algorithm to find a sparse solution when the gTV solution is non-unique. For this case, the optimization problem turns out to be a LASSO [43] minimization with non-unique solution. We introduce a combination of FISTA [32] and the simplex algorithm to find a sparse solution which we prove to be an extreme point of the solution set.

This chapter is organized as follows: In Sections 2.2 and 2.4, we present the formulation and the theoretical results of the inverse problem for the two regularization cases. In Section 2.5, we compare the solutions of the two cases. We present our numerical algorithm in Section 2.6 and illustrate its behavior with various examples in Section 2.7. The mathematical proofs of the main theorems are given in the appendices.

2.2 Continuous-Domain Formulation of Inverse Problems

In our formulation of a linear inverse problem, the signal f is a function of the continuous-domain variable $x \in \mathbb{R}$. The task is then to recover f from the vector of measurements $\mathbf{y} = \mathrm{H}\{f\} + \mathbf{n} \in \mathbb{R}^M$, where $\mathbf{n}$ is an unknown noise component that is typically assumed to be i.i.d. Gaussian.

In the customary discrete formulation, the basis of the recovered function is already chosen and, therefore, all that remains is to recover the expansion coefficients of the signal representation (2.1). In this scenario, one often includes matrices $\mathbf{H}$ and $\mathbf{L}$ that directly operate on these coefficients. However, for our continuous-domain formulation, the operations have to act directly on the function f. For this reason, we also need the continuous-domain counterparts of the measurement and regularization operators. The entities that enter our formulation are described next.

2.2.1 Measurement Operator

The system matrix $\mathbf{H}$ in (2.2) and (2.3) is henceforth replaced by the operator $\mathrm{H} : \mathcal{X} \to \mathbb{R}^M$ that maps the continuous-domain functions living in a Banach space $\mathcal{X}$ to the linear measurements $\mathbf{y} \in \mathbb{R}^N$. This operator is described as

$$\mathrm{H}\{f\} = (\langle h_1, f\rangle, \dots, \langle h_M, f\rangle) = (y_1, \dots, y_M) = \mathbf{y}, \tag{2.5}$$

where $\langle h, g\rangle = \int_{\mathbb{R}} h(x)g(x)\,\mathrm{d}x$, which in the case of generalized functions should be interpreted as the duality product. Furthermore, the map $h_m : f \mapsto \langle h_m, f\rangle$ is assumed to be continuous in $\mathcal{X} \to \mathbb{R}$. For example, the components of the measurement operator that samples a function at the locations $x_1, \dots, x_M$ are represented by $h_m = \delta(\cdot - x_m)$ such that $\langle \delta(\cdot - x_m), f\rangle = f(x_m)$. Similarly, Fourier measurements at pulsations $\omega_1, \dots, \omega_M$ are obtained by taking $h_m = \mathrm{e}^{-\mathrm{j}\omega_m \cdot}$.

2.2.2 Data-Fidelity Term

As extension of the conventional quadratic data-fidelity term $\|\mathbf{y}-\mathbf{Hf}\|_2^2$, we consider a general cost functional $E(\mathbf{y}, \cdot) : \mathbb{R}^M \to \mathbb{R}^+ \cup \{\infty\}$ with some assumptions (see Assumption 2 in Section 2.4), that measures the discrepancy between the given measurements $\mathbf{y}$ and the values $\mathrm{H}\{f\}$ predicted from the reconstruction. A relevant example is the weighted quadratic data-fidelity term, which is often used when the measurement noise is Gaussian with diagonal covariance. Similarly, we can use $\|\mathbf{y} - \mathrm{H}\{f\}\|_1$, for example, when the additive noise is Laplacian. Alternatively, when the measurements are noiseless, we use the indicator function

$$\mathcal{I}(\mathbf{y}, \mathrm{H}\{f\}) = \begin{cases} 0, & \mathbf{y} = \mathrm{H}\{f\} \\ \infty, & \mathbf{y} \neq \mathrm{H}\{f\}, \end{cases} \tag{2.6}$$

which imposes an exact fit.

2.2.3 Regularization Operator

Since the underlying signal is continuously defined, we need to replace the regularization matrix $\mathbf{L}$ in (2.2) and (2.3) by a regularization operator $\mathrm{L} : \mathcal{X} \to \mathcal{Y}$, where $\mathcal{X}$ and $\mathcal{Y}$ are appropriate (generalized) function spaces to be defined in Section 2.2.5. The typical example that we have in mind is the derivative operator $\mathrm{L} = \mathrm{D} = \frac{\mathrm{d}}{\mathrm{d}x}$.

The continuous-domain regularization is then imposed on $\mathrm{L}f$. We assume that the operator L is admissible in the sense of definition 2.2.1.

Definition 2.2.1. *The operator* $\mathrm{L} : \mathcal{X} \to \mathcal{Y}$ *is called* **spline-admissible** *if*

- *it is linear and shift-invariant;*
- *its null space* $\mathcal{N}_{\mathrm{L}} = \{p \in \mathcal{X} : \mathrm{L}p = 0\}$ *is finite-dimensional;*
- *it admits the Green's function* $\rho_{\mathrm{L}} : \mathbb{R} \to \mathbb{R}$ *with the property that* $\mathrm{L}\rho_{\mathrm{L}} = \delta$*.*

Given that $\widehat{\mathrm{L}}$ is the frequency response of L, the Green's function can be calculated through the inverse Fourier transform $\rho_{\mathrm{L}} = \mathrm{F}^{-1}\left\{\frac{1}{\widehat{\mathrm{L}}}\right\}$. For example, if $\mathrm{L} = \mathrm{D}$, then $\rho_{\mathrm{D}}(x) = \frac{1}{2}\mathrm{sign}(x)$. Here the Fourier transform, $\mathrm{F} : f \mapsto \mathrm{F}f = \int_{\mathbb{R}} f(x)\mathrm{e}^{-\mathrm{j}x(\cdot)}\,\mathrm{d}x$, is defined when the function is integrable and can be extended in the usual manner to $f \in \mathcal{S}'(\mathbb{R})$ where $\mathcal{S}'(\mathbb{R})$ is 'Schwartz' space of tempered distributions. In cases such as $\rho_{\mathrm{L}} = \mathrm{F}^{-1}\left\{\frac{1}{\widehat{\mathrm{L}}}\right\}$ when the argument is non-integrable, the definition should be seen in terms of generalized Fourier Transform [47, Defintion 8.9] which treats the argument as a distribution.

2.2.4 Regularization Norms

Since the optimization is done in the continuous domain, we also have to specify the proper counterparts of the ℓ_2 and ℓ_1 norms, as well as the corresponding vector spaces.

i. Quadratic (or Tikhonov) regularization: $\mathrm{R}_2(f) = \|\mathrm{L}f\|_{L_2}^2$, where

$$\|w\|_{L_2}^2 = \int_{\mathbb{R}} |w(x)|^2\,\mathrm{d}x. \tag{2.7}$$

ii. Generalized total variation: $\mathrm{R}_1(f) = \|\mathrm{L}f\|_{\mathcal{M}}$, where

$$\|w\|_{\mathcal{M}} = \sup_{\varphi \in \mathcal{S}(\mathbb{R}), \|\varphi\|_\infty = 1} \langle w, \varphi \rangle. \tag{2.8}$$

There $\mathcal{S}(\mathbb{R})$ is the 'Schwartz' space of smooth and rapidly decaying functions, which is also the dual of $\mathcal{S}'(\mathbb{R})$. Moreover, $\mathcal{M}=\{w \in \mathcal{S}'(\mathbb{R}) \mid \|w\|_{\mathcal{M}} < \infty\}$. In particular, when $w \in L_1 \subset \mathcal{M}$, we have that

$$\|w\|_{\mathcal{M}}=\int_{\mathbb{R}} |w(x)| \,\mathrm{d}x = \|w\|_{L_1}. \tag{2.9}$$

Yet, we note that $\mathcal{M}$ is slightly larger than L_1 since it also includes the Dirac distribution δ with $\|\delta\|_{\mathcal{M}} = 1$. The popular TV norm is recovered by taking $\|f\|_{\mathrm{TV}} = \|\mathrm{D}f\|_{\mathcal{M}}$ [46].

2.2.5 Search Space

The Euclidean search space $\mathbb{R}^N$ is replaced by spaces of functions, namely,

$$\mathcal{X}_2=\{f : \mathbb{R} \to \mathbb{R} \mid \|\mathrm{L}f\|_{L_2} < +\infty\}, \tag{2.10}$$
$$\mathcal{X}_1=\{f : \mathbb{R} \to \mathbb{R} \mid \|\mathrm{L}f\|_{\mathcal{M}} < +\infty\}. \tag{2.11}$$

In other words, our search (or native) space is the largest space over which the regularization is well defined. It turns out that $\mathcal{X}_2$ and $\mathcal{X}_1$ are Hilbert and Banach spaces, respectively. However, this is nontrivial to see since these spaces contain the null space which makes $\|\mathrm{L}f\|_{L_2}$ and $\|\mathrm{L}f\|_{\mathcal{M}}$ semi-norms. This null-space can be taken care off by using an appropriate inner-product $\langle\cdot,\cdot\rangle_{\mathcal{N}_{\mathrm{L}}}$ (norm $\|\cdot\|_{\mathcal{N}_{\mathrm{L}}}$, respectively) such that $\langle\cdot,\cdot\rangle_{\mathcal{X}_2} = \langle \mathrm{L}\cdot, \mathrm{L}\cdot\rangle + \langle\cdot,\cdot\rangle_{\mathcal{N}_{\mathrm{L}}}$ ($\|\cdot\|_{\mathcal{X}_1} = \|\mathrm{L}\cdot\|_{\mathcal{M}} + \|\cdot\|_{\mathcal{N}_{\mathrm{L}}}$, respectively) is the inner-product (norm, respectively) on $\mathcal{X}_2$ ($\mathcal{X}_1$, respectively). The structure of these spaces has been studied in [46] and is recalled in the Appendix A.1.6

As we shall see in Section 2.4, the solution of (2.4) will be composed of splines; therefore, we also review the definition of splines.

Definition 2.2.2 (Nonuniform L-spline). *A function $f : \mathbb{R} \to \mathbb{R}$ is called a nonuniform* L*-spline with spline knots $(x_1, \dots, x_K)$ and weights $(a_1, \dots, a_K)$ if*

$$\mathrm{L}f = \sum_{k=1}^{K} a_k \delta(\cdot - x_k). \tag{2.12}$$

By solving the differential equation in (2.12), we find that the generic form of the nonuniform spline f is

$$f = p_0 + \sum_{k=1}^{K} a_k \rho_{\mathrm{L}}(\cdot - x_k), \tag{2.13}$$

where $p_0 \in \mathcal{N}_{\mathrm{L}}$.

2.3 Related Work

The use of $\mathrm{R}(f) = \|\mathrm{L}f\|_{L_2}^2$ goes back to Tikhonov's theory of regularization [38] and to kernel methods in machine learning [48]. In the learning community, representer theorems (RT) as in [49], [50] use the theory of reproducing-kernel Hilbert spaces (RKHS) to state the solution of the problem for the restricted case where the measurements are samples of the function. For the generalized-measurement case, there are also tight connections between these techniques and variational splines and radial-basis functions [51], [52], [47]. These representer theorems, however, either have restrictions on the empirical risk functional or on the class of measurement operators.

Specific spline-based methods with quadratic regularization have been developed for inverse problems. In particular, [53], [54] used variational calculus. Here, we strengthen these results by proving the uniqueness and existence of the solution of (2.4) for $\mathrm{R}(f) = \|\mathrm{L}f\|_{L_2}^2$. We revisit the derivation of the result using the theory of RKHS.

Among more recent non-quadratic techniques, the most popular ones rely on (TV) regularization which was introduced as a noise-removal technique in [55] and is widely used in computational imaging and compressed sensing, although always in discrete settings. Splines as solutions of TV problems for restricted scenarios have been discussed in [56]. More recently, a RT for the continuous-domain $\mathrm{R}(f) = \|\mathrm{L}f\|_{\mathcal{M}}$ in a general setting has been established in [46], extending the seminal work of Fisher and Jerome [57]. The solution has been shown to be composed of splines that are directly linked to the differential operator L. Other recent contributions on inverse problems in the space of measures include [58–62]. In particular, in this chapter, we extend the result of [46] to an unconstrained version of the problem. The unconstrained formulation is useful in devising numerical algorithms which are one of the main contributions of our work. In addition our results are valid for a much larger set of data-fidelity terms than [46]. This is useful in practical scenarios where one may use data-fidelity terms depending on factors like distribution of noise, *etc.*.

2.4 Theoretical Results

To state our theorems, we need some technical assumptions.

Assumption 1. *Let the search space $\mathcal{X}$ and the regularization space $\mathcal{Y}$ be Banach spaces such that the following holds.*

i. The functionals h_m for $m \in \{1, \dots, M\}$ are linear continuous over $\mathcal{X}$ and the vector-valued functional $\mathrm{H} : \mathcal{X} \to \mathbb{R}^M$ gives the linear measurements $f \mapsto \mathrm{H}\{f\} = (\langle h_1, f\rangle, \dots, \langle h_M, f\rangle)$.

ii. The regularization operator $\mathrm{L} : \mathcal{X} \to \mathcal{Y}$ is spline-admissible. *Its finite-dimensional null space $\mathcal{N}_{\mathrm{L}}$ has the basis $\boldsymbol{p} = (p_1, \dots, p_{N_0})$.*

iii. The inverse problem is well posed over the null space. This means that, for any pair $p_1, p_2 \in \mathcal{N}_{\mathrm{L}}$, we have that

$$\mathrm{H}\{p_1\} = \mathrm{H}\{p_2\} \Leftrightarrow p_1 = p_2. \tag{2.14}$$

In other words, different null-space functions result in different measurements.

In particular, Condition iii) is equivalent to $\mathcal{N}_{\mathrm{L}} \cap \mathcal{N}_{\mathrm{H}} = \{0\}$, where $\mathcal{N}_{\mathrm{H}}$ is the null space of the vector-valued measurement functional. This property prevents from having a nonzero $f_0 \in \mathcal{N}_{\mathrm{L}} \cap \mathcal{N}_{\mathrm{H}}$ whose addition to any $f \in \mathcal{X}$ can neither be detected by the data-fidelity term nor by the regularization term. This is essential in ensuring the boundedness of the set of the minimizers.

Assumption 2. *For a given $\mathbf{y} \in \mathbb{R}^M$, the functional $E(\mathbf{y}, \cdot) : \mathbb{R}^M \to \mathbb{R}^+ \cup \{\infty\}$ is convex, coercive, and lower semi-continuous on the whole $\mathbb{R}^M$, and is proper (has finite value for at least one input) in the range of* H.

Assumption 2'. *For a given $\mathbf{y} \in \mathbb{R}^M$, the functional $E(\mathbf{y}, \cdot)$ satisfies Assumption 2 as well as one of the following.*

i. It is strictly convex; or

ii. it is an indicator function $I(\mathbf{y}, \cdot)$.

As we shall see later, stronger results can be derived for the $E(\mathbf{y}, \cdot)$ that satisfy Assumption 2'.

Two remarks are in order. Firstly, the condition of being proper in the range of H implies that there exists an $f \in \mathcal{X}$ such that $E(\mathbf{y}, \mathrm{H}\{f\})$ is finite. Secondly, when $E(\mathbf{y}, \cdot)$ is strictly convex or is such that its range does not include ∞, it is redundant to ensure that it is proper in the range of H.

We now state our two main results. Their proofs are given in Appendix A.1.1 and Appendix A.1.2, respectively.

2.4.1 Inverse Problem with Tikhonov/L2 Regularization

Theorem 2.4.1. *Let Assumptions 1 and 2 hold with the search space $\mathcal{X} = \mathcal{X}_2$ and regularization space $\mathcal{Y} = L_2$. Then, the set*

$$\mathcal{V}_2 = \arg\min_{f \in \mathcal{X}_2} \left(E(\mathbf{y}, \mathrm{H}\{f\}) + \lambda \|\mathrm{L}f\|_{L_2}^2 \right) \tag{2.15}$$

of minimizers is nonempty, convex, and such that any $f_2^* \in \mathcal{V}_2$ is of the form

$$f_2^*(x) = \sum_{m=1}^{M} a_m \varphi_m(x) + \sum_{n=1}^{N_0} b_n p_n(x), \tag{2.16}$$

where $\varphi_m = \mathrm{F}^{-1}\left\{\frac{\widehat{h}_m}{|\widehat{\mathrm{L}}|^2}\right\}$, and $\mathbf{a} = (a_1, \ldots, a_M)$, and $\mathbf{b} = (b_1, \ldots, b_{N_0})$ are expansion coefficients such that

$$\sum_{m=1}^{M} a_m \langle h_m, p_n \rangle = 0 \tag{2.17}$$

for all $n \in \{1, \ldots, N_0\}$. Moreover, if $E(\mathbf{y}, \cdot)$ satisfies Assumption 2' then the minimizer is unique (the set $\mathcal{V}_2$ is singleton).

2.4.2 Inverse Problem with gTV Regularization

Theorem 2.4.2. *Let Assumptions 1 and 2 hold for the search space* $\mathcal{X} = \mathcal{X}_1$ *and regularization space* $\mathcal{Y} = \mathcal{M}$*. Moreover, assume that* H *is weak*-continuous (see Appendix A.1.6)). Then, the set*

$$\mathcal{V}_1 = \arg \min_{f \in \mathcal{X}_1} \left(E(\mathbf{y}, \mathrm{H}\{f\}) + \lambda \|\mathrm{L}f\|_{\mathcal{M}} \right) \tag{2.18}$$

of minimizers is nonempty, closed-convex, weak-compact, and its extreme points are nonuniform* L*-splines of the form*

$$f_1^*(x) = \sum_{k=1}^{K} a_k \rho_{\mathrm{L}}(x - x_k) + \sum_{n=1}^{N_0} b_n p_n(x) \tag{2.19}$$

for some $K \leq (M - N_0)$*. The unknown knots* $(x_1, \dots, x_K)$*, and the expansion coefficients* $\mathbf{a} = (a_1, \dots, a_K)$ *and* $\mathbf{b} = (b_1, \dots, b_{N_0})$ *are the parameters of the solution with* $\|\mathrm{L}f_1^*\|_{\mathcal{M}} = \|\mathbf{a}\|_1$*. The solution set* $\mathcal{V}_1$ *is the closed-convex hull of these extreme points. Moreover, if Assumption 2' is satisfied then all the solutions have the same measurement; i.e.,* $\mathbf{y}_{\mathcal{V}_1} = \mathrm{H}\{f\}$, $\forall f \in \mathcal{V}_1$.

A sufficient condition for weak*-continuity of h_m is $\int_{\mathbb{R}} |h_m(x)|(1+|x|)^D \mathrm{d}x < \infty$ ([46, Theorem 6]), meaning that h_m should exhibit some minimal decay at infinity (see Appendix A.1.6). Here $D = \inf\{n \in \mathbb{N} : (\operatorname{ess\,sup}_{x \in \mathbb{R}} \rho_{\mathrm{L}}(1+|x|)^n) < +\infty\}$. The ideal sampling is feasible as well, provided that the ρ_{L} is continuous; a detailed proof of the weak*-continuity of $\delta(\cdot - x_n)$ for the case $\mathrm{L} = \mathrm{D}^2$ can be found in [63, Proposition 6].

We remark that [46, Theorem 2] is a special case of Theorem 4. The former states the same result as Theorem 4 for the minimization problem

$$\mathcal{V}_1 = \arg \min_{\mathrm{H}\{f\} \in \mathcal{C}} \|\mathrm{L}f\|_{\mathcal{M}}, \tag{2.20}$$

where $\mathcal{C}$ is feasible, convex, and compact. Feasibility of $\mathcal{C}$ means that the set $\mathcal{C}_{\mathcal{X}_1} = \{f \in \mathcal{X}_1 : \mathrm{H}f \in \mathcal{C}\}$ is nonempty. In our setting, problem (2.20) can be obtained by using an indicator function over the feasible set $\mathcal{C}$ as the data-fidelity term. However, Theorem 4 covers other more useful cases of E; for example, $\|\mathbf{y} - \mathrm{H}\{f\}\|_1$ and $\|\mathbf{y} - \mathrm{H}\{f\}\|_2^2$. Moreover, as discussed earlier, when data-fidelity terms are strictly convex or do not have ∞ in their range, they are proper in the range of H for any $\mathbf{y} \in \mathbb{R}^M$. This means that they do not require careful selection of $\mathcal{C}$ in order to satisfy the feasibility condition. This is helpful in directly devising and deploying algorithms to find the minimizers.

Also, fundamentally (2.20) only penalizes the regularization value, whereas Theorem 4 additionally penalizes a data-fidelity term that can recover more desirable solutions. In fact, Theorem 4 also covers cases such as

$$\mathcal{V}_1 = \arg \min_{\mathrm{H}\{f\} \in \mathcal{C}} E(\mathbf{y}, \mathrm{H}\{f\}) + \lambda \|\mathrm{L} f\|_{\mathcal{M}}, \tag{2.21}$$

which allow more control than (2.20) over the data-fidelity of the recovered solution.

2.4.3 Illustration with Ideal Sampling

Here, we discuss the regularized case where noisy data points $((x_1, y_1), \ldots, (x_M, y_M))$ are fitted by a function. The measurement functionals in this case are the shifted Dirac impulses $h_m = \delta(\cdot - x_m)$ whose Fourier transform is $\widehat{h}_m(\omega) = \mathrm{e}^{-\mathrm{j}\omega x_m}$. We choose $\mathrm{L} = \mathrm{D}^2$ and $E = \|\mathbf{y} - \mathrm{H}\{f\}\|_2^2$ which satisfies Assumption 2'.i). Here D^2 is the generalized second-order derivative. For the L_2 problem, we have that

$$f_2^* = \arg \min_{f \in \mathcal{X}_2} \left(\sum_{m=1}^{M} |y_m - f(x_m)|^2 + \lambda \|\mathrm{D}^2 f\|_{L_2}^2 \right). \tag{2.22}$$

As given in Theorem 2.4.1, f_2^* is unique and has the basis function $\varphi_m(x) = \mathrm{F}^{-1}\left\{\frac{\mathrm{e}^{-\mathrm{j}(\cdot)x_m}}{|-(\cdot)^2|^2}\right\}(x) = \frac{1}{12}|x - x_m|^3$. The resulting solution is piecewise linear. It can be expressed as

$$f_2^*(x) = b_1 + b_2 x + \sum_{m=1}^{M} \frac{1}{12} a_m |x - x_m|^3, \tag{2.23}$$

where $b_1 + b_2 x \in \mathcal{N}_{\mathrm{D}^2}$ is a linear function.
We contrast (2.22) with the gTV version

$$f_1^* = \arg \min_{f \in \mathcal{X}_1} \left(\sum_{m=1}^{M} |y_m - f(x_m)|^2 + \lambda \underbrace{\|\mathrm{D}^2 f\|_{\mathcal{M}}}_{\|\mathrm{D} f\|_{\mathrm{TV}}} \right). \tag{2.24}$$

In this scenario, the term $\|\mathrm{D}^2 f\|_{\mathcal{M}}$ is the total variation of the function $\mathrm{D}f$. It penalizes solutions whose slope varies too much from one point to the next.

The Green's function in this case is $\rho_{\mathrm{D}^2} = \frac{|x|}{2}$. Based on Theorem 2.4.2, any extreme point of (2.24) is of the form

$$f_1^*(x) = b_1 + b_2 x + \frac{1}{2}\sum_{k=1}^{K} a_k' |x - \tau_k|, \tag{2.25}$$

which is a piecewise linear function composed of a linear term b_1+b_2x and $K \leq (M-1)$ basis functions, $\{|x-\tau_k|\}_{k=1}^K$. The knots (or locations) $\{\tau_k\}_{k=1}^K$ are not fixed a priori and usually differ from the measurement points $\{x_m\}_{m=1}^M$.

The two solutions and their basis functions are illustrated in Figure 2.1 for specific data. This example demonstrates that the mere replacement of the L_2 penalty with the gTV norm has a fundamental effect on the solution: piecewise-cubic functions having knots at the sampling locations are replaced by piecewise-linear functions with a lesser number of adaptive knots. Moreover, in the gTV case, the regularization has been imposed on the generalized second-order derivative of the function $\left(\|\mathrm{D}^2 f\|_{\mathcal{M}}\right)$, which uncovers the innovations $\mathrm{D}^2 f_1^* = \sum\limits_{k=1}^{K} a_k' \delta(\cdot - \tau_k)$. By contrast, when $\mathrm{R}_2(f) = \|\mathrm{D}^2 f\|_{L_2}^2 = \langle \mathrm{D}^{2*}\mathrm{D}^2 f, f\rangle$, the recovered solution is such that $\mathrm{D}^{2*}\mathrm{D}^2 f_2^* = \sum_{m=1}^M a_m \delta(\cdot - x_m)$, where $\mathrm{D}^{2*} = \mathrm{D}^2$ is the adjoint operator of D^2. Thus, in both cases, the recovered functions are composed of the Green's function of the corresponding active operators: D^2 vs. $\mathrm{D}^{2*}\mathrm{D}^2 = \mathrm{D}^4$.

2.5 Comparison

We now discuss and contrast the results of Theorems 2.4.1 and 2.4.2. In either case, the solution is composed of a primary component and a null-space component whose regularization cost vanishes.

Nature of the Primary Component The solutions for the gTV regularization are composed of atoms within the infinitely large dictionary $\{\rho_{\mathrm{L}}(\cdot - \tau)\}, \forall \tau \in \mathbb{R}$, whose shapes depend only on L. In contrast, the L_2 solutions are composed of fixed atoms $\{\varphi_m\}_{m=1}^M$ whose shapes depend on both L and H. As the shape of the atoms of the gTV solutions does not depend on H, this makes it easier to inject

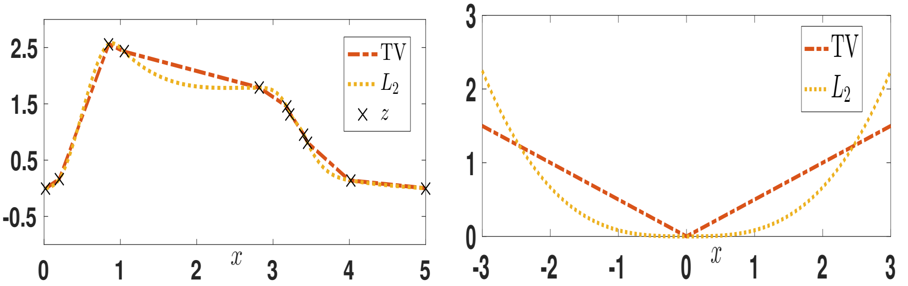

(a) $f_1^*(x)$ and $f_2^*(x)$. (b) $\rho_{\mathrm{D}^2}(x)$ and $\rho_{\mathrm{D}^2*\mathrm{D}^2}(x)$.

Figure 2.1: Reconstructions of a signal from nonuniform samples for $\mathrm{L} = \mathrm{D}^2$: (a) Tikhonov ($L_2$) vs. gTV solution, and (b) Corresponding basis functions ρ_{D^2} vs. $\rho_{\mathrm{D}^2*\mathrm{D}^2}$.

prior knowledge in that case. The weights and the location of the atoms of the gTV solution are adaptive and found through a data-dependent procedure which results in a sparse solution that turns out to be a nonuniform spline. By contrast, the L_2 solution lives in a fixed finite-dimensional space.

Null-Space Component. The second component in either solution belongs to the null space of the operator L. As its contribution to regularization vanishes, the solutions tend to have large null-space components in both instances.

Oscillations. The modulus of the Fourier transform of the basis function of the gTV case, $\left|\left\{\frac{1}{\widehat{\mathrm{L}}}\right\}\right|$ typically decays faster than that of the L_2 case, $\left|\left\{\frac{\widehat{h}_m}{|\widehat{\mathrm{L}}|^2}\right\}\right|$. Therefore, the gTV solution exhibits weaker Gibbs oscillations at edges.

Uniqueness of the Solution. Our hypotheses guarantee existence. Moreover, the minimizer of the L_2 problem is unique when Assumption 2' is true. By contrast, even for this special category of $E(\mathbf{y}, \cdot)$, the gTV problem can have infinitely many solutions, despite all having the same measurements. Remarkably, however, when the gTV solution is unique, it is guaranteed to be an L-spline.

Nature of the Regularized Function. One of the main differences between the reconstructions f_2^* and f_1^* is their sparsity. Indeed, $\mathrm{L}f_1^*$ uncovers Dirac impulses situated at $(M-1)$ locations for the gTV case, with $\mathrm{L}f_1^* = \sum_{m=1}^{M-1} a_m \delta(\cdot - \tau_m)$. In return, $\mathrm{L}f_2^*$ is a nonuniform L-spline convolved with the measurement functions, whose temporal support is not localized. This allows us to say that the gTV solution is sparser than the Tikhonov solution.

2.6 Discretization and Algorithms

We now lay down the discretization procedure that translates the continuous-domain optimization into a more tractable finite-dimensional problem. Theorems 2.4.1 and 2.4.2 imply that the infinite-dimensional solution lives in a finite-dimensional space that is characterized by the basis functions $\{\varphi_m\}_{m=1}^{M}$ for L_2 and $\{\rho_{\mathrm{L}}(\cdot-\tau_k)\}_{k=1}^{K}$ for gTV, in addition to $\{p_n\}_{n=1}^{N_0}$ as basis of the null space. Therefore, the solutions can be uniquely expressed with respect to the finite-dimensional parameter $\mathbf{a}\in\mathbb{R}^M$ or $\mathbf{a}\in\mathbb{R}^K$, respectively, and $\mathbf{b}\in\mathbb{R}^{N_0}$. Thus, the objective functional $\mathrm{J}_{\mathrm{R}_i}(f|\mathbf{y},\lambda)$, for a given $i\in\{1,2\}$, can be discretized to get the objective functional $\mathrm{J}_{\mathrm{R}_i}(\mathbf{a},\mathbf{b}|\mathbf{y},\lambda)$. Its minimization is done numerically, by expressing $\mathrm{H}\{f\}$ and $\|\mathrm{L}f\|_{L_2}^2$ or $\|\mathrm{L}f\|_{\mathcal{M}}$ in terms of $\mathbf{a}$ and $\mathbf{b}$. We discuss the strategy to achieve $\mathrm{J}_{\mathrm{R}_i}(\mathbf{a},\mathbf{b}|\mathbf{y},\lambda)$ and its minima for the two cases. From now onwards, we will use J_i for $\mathrm{J}_{\mathrm{R}_i}$ where $i\in\{1,2\}$.

2.6.1 Tikhonov Regularization

For the L_2 regularization, given $\lambda>0$, the solution

$$f_2^*=\arg\min_{f\in\mathcal{X}_2}\underbrace{\left(E(\mathbf{y},\mathrm{H}\{f\})+\lambda\|\mathrm{L}f\|_{L_2}^2\right)}_{\mathrm{J}_2(f|\mathbf{y},\lambda)} \tag{2.26}$$

can be expressed as

$$f_2^*=\sum_{m=1}^{M}a_m\varphi_m+\sum_{n=1}^{N_0}b_np_n. \tag{2.27}$$

Recall that $\mathrm{L}^*\mathrm{L}\varphi_m=h_m$, so that

$$\mathrm{L}^*\mathrm{L}f_2^*=\sum_{m=1}^{M}a_mh_m. \tag{2.28}$$

The corresponding $J_2(\mathbf{y}|\lambda, \mathbf{a}, \mathbf{b})$ is then found by expressing $\mathrm{H}\{f_2^*\}$ and $\|\mathrm{L}f_2^*\|_{L_2}^2$ in terms of $\mathbf{a}$ and $\mathbf{b}$. Due to the linearity of the model,

$$\begin{aligned}\mathrm{H}\{f_2^*\} &= \sum_{m=1}^{M} a_m \mathrm{H}\{\varphi_m\} + \sum_{n=1}^{N_0} b_n \mathrm{H}\{p_n\} \\ &= \mathbf{Va} + \mathbf{Wb}, \qquad (2.29)\end{aligned}$$

where $[\mathbf{V}]_{m,n} = \langle h_m, \varphi_n \rangle$ and $[\mathbf{W}]_{m,n} = \langle h_m, p_n \rangle$. Similarly,

$$\langle \mathrm{L}f_2^*, \mathrm{L}f_2^* \rangle = \langle \mathrm{L}^*\mathrm{L}f_2^*, f_2^* \rangle = \left\langle \sum_{m=1}^{M} a_m h_m, f_2^* \right\rangle \qquad (2.30)$$

$$= \mathbf{a}^T\mathbf{Va} + \mathbf{a}^T\mathbf{Wb} = \mathbf{a}^T\mathbf{Va}, \qquad (2.31)$$

where (2.30) uses (2.28) and where (2.31) uses the orthogonality property (2.17), which we can restate as $\mathbf{a}^T\mathbf{W} = \mathbf{0}$. By substituting these reduced forms in (2.26), the discretized problem becomes

$$f_2^* = \arg\min_{\mathbf{a},\mathbf{b}} \underbrace{\left(E(\mathbf{y}, \mathbf{Va} + \mathbf{Wb}) + \lambda \mathbf{a}^T\mathbf{Va}\right)}_{\mathrm{J}_2(\mathbf{a},\mathbf{b}|\mathbf{y},\lambda) = \mathrm{J}_2(f_2^*|\mathbf{y},\lambda)}. \qquad (2.32)$$

Due to Assumption 2, this problem is convex. If E is differentiable with respect to the parameters, the solution can be found by gradient descent.
When $E(\mathbf{y}, \mathrm{H}\{f\}) = \|\mathbf{y} - \mathrm{H}\{f\}\|_2^2$, the problem is reduced to

$$\arg\min_{\mathbf{a},\mathbf{b}} \underbrace{\left(\|\mathbf{y} - (\mathbf{Va} + \mathbf{Wb})\|_2^2 + \lambda \mathbf{a}^T\mathbf{Va}\right)}_{\mathrm{J}_2(\mathbf{a},\mathbf{b}|\mathbf{y},\lambda)} \qquad (2.33)$$

which is very similar to (2.2). This criterion is convex with respect to the coefficients $\mathbf{a}$ and $\mathbf{b}$. Enforcing that the gradient of J_2 vanishes with respect to $\mathbf{a}$ and $\mathbf{b}$ and setting the gradient to $\mathbf{0}$ then yields M linear equations with respect to the $M + N_0$ variables, while the orthogonality property (2.17) gives N_0 additional constraints. The combined equations correspond to the linear system

$$\begin{bmatrix} \mathbf{V} + \lambda\mathbf{I} & \mathbf{W} \\ \mathbf{W}^T & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix} = \begin{bmatrix} \mathbf{y} \\ \mathbf{0} \end{bmatrix}. \qquad (2.34)$$

The system matrix so obtained can be proven to be positive definite due to the property of Gram matrices generated in an RKHS and the admissibility condition of the measurement functional (Assumption 1). This ensures that the matrix is always invertible. The consequence is that the reconstructed signal can be obtained by solving a linear system of equation, for instance by QR decomposition or by simple matrix inversion. The derived solution is the same as the least-square solution in [54].

2.6.2 gTV Regularization

In the case of gTV regularization, the problem to solve is

$$f_1^* = \arg\min_{f\in\mathcal{X}_1} \underbrace{\left(E(\mathbf{y}, \mathrm{H}\{f\}) + \lambda\|\mathrm{L}f\|_{\mathcal{M}}\right)}_{\mathrm{J}_1(f|\mathbf{y},\lambda)}. \tag{2.35}$$

According to Theorem 2.4.2, an extreme-point solution of (2.35) is

$$f_1^*(x) = \sum_{k=1}^{K} a_k \rho_{\mathrm{L}}(x-\tau_k) + \sum_{n=1}^{N_0} b_n p_n(x) \tag{2.36}$$

and satisfies

$$\mathrm{L}f_1^* = w_1 = \sum_{k=1}^{K} a_k \delta(\cdot - \tau_k) \tag{2.37}$$

with $K \leq (M - N_0)$. Theorem 2.4.2 implies that we only have to recover a_k, τ_k, and the null-space component p to recover f_1^*.

In our experiments, we shall consider the case of measurement functionals whose support is limited to $[0, T]$. We therefore only reconstruct the restriction of the signal in this interval. Since we usually know neither K nor τ_k beforehand, our solution is to quantize the x-axis and look for τ_k in the range $[0, T]$ on a grid with $N \gg K$ points. We control the quantization error with the grid step $\Delta = T/N$.

The discretized problem is then to find $\mathbf{a} \in \mathbb{R}^N$ with fewer than $(M - N_0)$ nonzero coefficients and $\mathbf{b} \in \mathbb{R}^{N_0}$ such that

$$f_{1,\Delta}^*(x) = \sum_{n=0}^{N-1} a_n \rho_{\mathrm{L}}(x - n\Delta) + \sum_{n=1}^{N_0} b_n p_n(x) \tag{2.38}$$

with $K \leq (M - N_0) \ll N$ nonzero coefficients a_n, satisfies a computationally feasible variant of (2.35). In other words, we solve the restricted version of (2.35)

$$\min_{f \in \mathcal{X}_{1,\Delta}} \underbrace{\left(E(\mathbf{y}, \mathrm{H}\{f\}) + \lambda \|\mathrm{L}f\|_{\mathcal{M}}\right)}_{J_{1,\Delta}(\mathbf{y}|\lambda,f)}, \tag{2.39}$$

where

$$\mathcal{X}_{1,\Delta} = \left\{ \sum_{n=0}^{N-1} a_n \rho_{\mathrm{L}}(\cdot - n\Delta) + \sum_{n=1}^{N_0} b_n p_n \middle| (\mathbf{a}, \mathbf{b}) \in \mathbb{R}^{N+N_0} \right\}.$$

Similarly to the L_2 case, $\mathrm{J}_{1,\Delta}(\mathbf{a}, \mathbf{b}|\mathbf{y}, \lambda)$ is found by expressing $\mathrm{H}\{f^*_{1,\Delta}\}$ and $\|\mathrm{L}f^*_{1,\Delta}\|_{\mathcal{M}}$ in terms of $\mathbf{a}$ and $\mathbf{b}$. For this, we use the properties that $\mathrm{L}\rho_{\mathrm{L}} = \delta$, $\|\delta\|_{\mathrm{TV}} = 1$, and $\mathrm{L}p_n = 0$ for $n \in [1 \dots N_0]$. This results in

$$\mathrm{H}\{f^*_{1,\Delta}\} = \mathbf{Pa} + \mathbf{Qb}, \tag{2.40}$$

$$\|\mathrm{L}f^*_{1,\Delta}\|_{\mathcal{M}} = \|\mathbf{a}\|_1, \tag{2.41}$$

where $\mathbf{a} = (a_0, \dots, a_{N-1})$, $[\mathbf{P}]_{m,n} = \langle h_m, \rho_{\mathrm{L}}(\cdot - n\Delta)\rangle$ for $n \in [0 \dots N-1]$, $[\mathbf{Q}]_{m,n} = \langle h_m, p_n\rangle$ for $n \in [1 \dots N_0]$, $\|\mathbf{a}\|_1 = \sum_{n=1}^{N} |a_n|$, and where N is the initial number of Green's functions of our dictionary. The new discretized objective functional is

$$f^*_{1,\Delta} = \arg\min_{\mathbf{a},\mathbf{b}} \underbrace{\left(E(\mathbf{y}, (\mathbf{Pa} + \mathbf{Qb})) + \lambda \|\mathbf{a}\|_1\right)}_{\mathrm{J}_{1,\Delta}(\mathbf{a},\mathbf{b}|\mathbf{y},\lambda) = \mathrm{J}_{1,\Delta}(f^*_{1,\Delta}|\mathbf{y},\lambda)}. \tag{2.42}$$

Note that (2.42) is the exact discretization of the infinite-dimensional problem (2.39). However, additional theories, such as $\Gamma-$convergence [64–66], are needed to show that the recovered signal $f^*_{1,\Delta}$ converges (in the weak sense) to one of the solution of the original problem (2.35) when the discretization step Δ goes to 0 (or when N is large enough). We leave this analysis for the future work.

When E is differentiable with respect to the parameters, a minimum can be found by using proximal algorithms where the slope of $\|\mathbf{a}\|_1$ is defined by a Prox operator. We discuss the two special cases when E is either an indicator function or a quadratic data-fidelity term.

Exact Fit with $E = \mathcal{I}(\mathbf{y}, \mathrm{H}\{f\})$

To perfectly recover the measurements, we impose an infinite penalty when the recovered measurements differ from the given ones. In view of (2.40) and (2.41), this corresponds to solving

$$(\mathbf{a}^*, \mathbf{b}^*) = \arg\min_{\mathbf{a},\mathbf{b}} \|\mathbf{a}\|_1 \quad \text{subject to} \quad \mathbf{Pa} + \mathbf{Qb} = \mathbf{y}. \tag{2.43}$$

We then recast Problem (2.43) as the linear program

$$\begin{aligned}(\mathbf{a}^*, \mathbf{u}^*, \mathbf{b}^*) = \min_{\mathbf{a},\mathbf{u},\mathbf{b}} \sum_{n=1}^{N} u_n \text{ subject to } \mathbf{u} + \mathbf{a} &\geq \mathbf{0},\\ \mathbf{u} - \mathbf{a} &\geq \mathbf{0},\\ \mathbf{Pa} + \mathbf{Qb} &= \mathbf{y},\end{aligned} \tag{2.44}$$

where the inequality $\mathbf{x} \geq \mathbf{y}$ between any 2 vectors $\mathbf{x} \in \mathbb{R}^N$ and $\mathbf{y} \in \mathbb{R}^N$ means that $x_n \geq y_n$ for $n \in [1 \dots N]$. This linear program can be solved by a conventional simplex or a dual-simplex approach [67], [68].

Least Squares Fit with $E = \|\mathbf{y} - \mathrm{H}\{f\}\|_2^2$

When E is a quadratic data-fidelity term, the problem becomes

$$(\mathbf{a}^*, \mathbf{b}^*) = \arg\min_{\mathbf{a},\mathbf{b}} \left(\|\mathbf{y} - (\mathbf{Pa} + \mathbf{Qb})\|_2^2 + \lambda\|\mathbf{a}\|_1\right), \tag{2.45}$$

which is more suitable when the measurements are noisy. The discrete version (2.45) is similar to (2.3), the fundamental difference being in the nature of the underlying basis function.

The problem is converted into a LASSO formulation [43] by decoupling the computation of $\mathbf{a}^*$ and $\mathbf{b}^*$. Suppose that $\mathbf{a}^*$ is fixed, then $\mathbf{b}^*$ is found by differentiating (2.45) and equating the gradient to $\mathbf{0}$. This leads to

$$\mathbf{b}^* = \left(\mathbf{Q}^T\mathbf{Q}\right)^{-1} \mathbf{Q}^T(\mathbf{y} - \mathbf{Pa}^*). \tag{2.46}$$

Upon substitution in (2.45), we get that

$$\mathbf{a}^* = \arg\min_{\mathbf{a}} \left(\|\mathbf{Q}'\mathbf{y} - \mathbf{Q}'\mathbf{Pa}\|_2^2 + \lambda\|\mathbf{a}\|_1\right), \tag{2.47}$$

where $\mathbf{Q}' = \left(\mathbf{I} - \mathbf{Q}\left(\mathbf{Q}^T\mathbf{Q}\right)^{-1}\mathbf{Q}^T\right)$ and $\mathbf{I}$ is the $(M \times M)$ identity matrix. Problem (2.47) can be solved using a variety of optimization techniques such as interior-point methods or proximal-gradient methods, among others. We employ the popular iterative algorithm FISTA [32], which has an $\mathcal{O}(1/t^2)$ convergence rate with respect to its iteration number t. However, in our case, the system matrices are formed by the measurements of the shifted Green's function on a fine grid. This leads to high correlations among the columns and introduces two issues.

- If LASSO has multiple solutions, then FISTA can converge to a solution within the solution set, whose sparsity index is greater than M. A similar type of limitation has been discussed in [69].

- If LASSO has a unique solution, then the convergence to the exact solution can be slow. The convergence rate is inversely proportional to the Lipschitz constant of the gradient of a quadratic loss function $\left(\max \mathrm{Eig}\left(\mathbf{H}^T\mathbf{H}\right)\right)$, which is typically high for the system matrix obtained through our formulation.

We address these issues by using a combination of FISTA and simplex, governed by the following Lemma 2.6.1 and Theorem 2.6.2. The properties of the solution of the LASSO problem have been discussed in [70], [71], [72]. We quickly recall one of the main results from [70].

Lemma 2.6.1 ([70, Lemma 1 and 11]). *Let $\mathbf{y} \in \mathbb{R}^M$ and $\mathbf{H} \in \mathbb{R}^{M\times N}$, where $M < N$. Then, the solution set*

$$\alpha_\lambda = \left\{ \arg\min_{\mathbf{a}\in\mathbb{R}^N} \left(\|\mathbf{y} - \mathbf{H}\mathbf{a}\|_2^2 + \lambda\|\mathbf{a}\|_1\right)\right\} \tag{2.48}$$

has the same measurement $\mathbf{H}\mathbf{a}^* = \mathbf{y}_0$ for any $\mathbf{a}^* \in \alpha_\lambda$. Moreover, if the solution is not unique, then any two solutions $\mathbf{a}^{(1)}, \mathbf{a}^{(2)} \in \alpha_\lambda$ are such that their mth element satisfies $\left\{\mathrm{sign}\left(\mathbf{a}_m^{(1)}\right)\mathrm{sign}\left(\mathbf{a}_m^{(2)}\right) \geq 0\right\}$ for $m \in [1 \dots M]$. In other words, any two solutions have the same sign over their common support.

We use Lemma 2.6.1 to infer Theorem 2.6.2, whose proof is given in Appendix 2.6.2.

Theorem 2.6.2. *Let* $\mathbf{y} \in \mathbb{R}^M$ *and* $\mathbf{H} \in \mathbb{R}^{M\times N}$, *where* $M < N$. *Let* $\mathbf{y}_{0,\lambda} = \mathbf{H}\mathbf{a}^*, \forall \mathbf{a}^* \in \alpha_\lambda$, *be the measurement of the solution set* α_λ *of the LASSO formulation*

$$\mathbf{a}^* = \arg \min_{\mathbf{a}\in\mathbb{R}^N} \left(\|\mathbf{y} - \mathbf{H}\mathbf{a}\|_2^2 + \lambda\|\mathbf{a}\|_1\right). \tag{2.49}$$

Then, the solution $\mathbf{a}^*_{\mathrm{SLP}}$ *(obtained using the simplex algorithm) of the linear program corresponding to the problem*

$$\mathbf{a}^*_{\mathrm{SLP}} = \arg\min \|\mathbf{a}\|_1 \quad \textit{subject to} \quad \mathbf{H}\mathbf{a} = \mathbf{y}_{0,\lambda} \tag{2.50}$$

is an extreme point of α_λ. *Moreover,* $\|\mathbf{a}^*_{\mathrm{SLP}}\|_0 \le M$.

Theorem 2.6.2 helps us to find an extreme point of the solution set α_λ of a given LASSO problem in the case when its solution is non-unique. To that end, we first use FISTA to solve the LASSO problem until it converges to a solution $\mathbf{a}^*_{\mathrm{F}}$. By setting $\mathbf{y}_{0,\lambda} = \mathbf{H}\mathbf{a}^*_{\mathrm{F}}$, Lemma 2.6.1 then implies that $\mathbf{H}\mathbf{a} = \mathbf{y}_{0,\lambda}, \forall \mathbf{a} \in \alpha_\lambda$. We then run the simplex algorithm to find

$$\mathbf{a}^*_{\mathrm{SLP}} = \arg\min \|\mathbf{a}\|_1 \quad \text{subject to} \quad \mathbf{H}\mathbf{a} = \mathbf{H}\mathbf{a}_{\mathrm{F}},$$

which yields an extreme point of α_λ by Theorem 2.6.2.
An example where the LASSO problem has a non-unique solution is shown in Figure 2.2.b. In this case, FISTA converges to a non-sparse solution with $\|\mathbf{a}^*_{\mathrm{F}}\| > M$, shown as solid stems. This implies that it is not an extreme point of the solution set. The simplex algorithm is then deployed to minimize the ℓ_1 norm such that the measurement $\mathbf{y}_0 = \mathbf{H}\mathbf{a}^*_{\mathrm{F}}$ is preserved. The final solution shown as dashed stems is an extreme point with the desirable level of sparsity. The continuous-domain relation of this example is discussed later.
The solution of the continuous-domain formulation is a convex set whose extreme points are composed of at most M shifted Green's functions. To find the position of these Green's functions, we discretize the continuum into a fine grid and then run the proposed two-step algorithm. If the discretization is fine enough, then the continuous-domain function that corresponds to the extreme point of the LASSO formulation is a good proxy for the actual extreme point of the convex-set solution of the original continuous-domain problem. This makes the extreme-point solutions of the LASSO a natural choice among the solution set. For the case when there is a unique solution but the convergence is too slow owing to the high value of

the Lipschitz constant of the gradient of the quadratic loss, the simplex algorithm is used after the FISTA iterations are stopped using an appropriate convergence criterion. For FISTA, the convergence behavior is ruled by the number of iterations t as

$$F(\mathbf{a}_t) - F(\mathbf{a}^*) \leq \frac{C}{(t+1)^2}, \tag{2.51}$$

where F is the LASSO functional and

$$C = 2\|\mathbf{a}_0 - \mathbf{a}^*\|_2^2 \max \mathrm{Eig}\left(\mathbf{H}^T\mathbf{H}\right) \tag{2.52}$$

(see [32]). This implies that an ϵ neighborhood of the minima of the functional is obtained in at most $t = \sqrt{C/\epsilon}$ iterations. To ensure convergence, it is also advisable to rely on the modified version of FISTA proposed in [73].

However, there is no direct relation between the functional value and the sparsity index of the iterative solution. Using the simplex algorithm as the next step guarantees the upper bound M on the sparsity index of the solution. Also, $F(\mathbf{a}^*_{\mathrm{SLP}}) \leq F(\mathbf{a}^*_{\mathrm{F}})$. This implies that an ϵ-based convergence criterion, in addition to the sparsity-index-based criterion like $\mathbf{a}^*_{\mathrm{F}} \leq M$, can be used to stop FISTA. Then, the simplex scheme is deployed to find an extreme point of the solution set with a reduced sparsity index.

Note that when $E(\mathbf{y}, \cdot)$ is not strictly convex, the solution set can have non-unique measurements. In that case, it is still possible to further sparsify a recovered solution by using the discussed Simplex approach.

2.6.3 Alternative Grid-free Techniques

Our proposed method relies on a grid based discretizatoin of the infinite-dimensional problem. For the sake of completeness, we discuss here alternative techniques for reconstructing continuous-domain sparse signals which employ grid-free optimization. Although elegant, these techniques have a more restricted range of applicability. The Taut-string algorithm (see [74]) can fit L-splines for $\mathrm{L} = \mathrm{D}^n$ but is devised for ideal sampling only. In [59, 60, 69, 75–77] the dual problem is considered for the optimization with an added emphasis on recovering the ground-truth signal. These methods, however, only deal with $\mathrm{L} = \mathrm{Id}$ and limited measurement operators.

Recently, in [62], motivated from [46], results for more general L and H have been derived. There the optimization is carried out in two steps; firstly, a finite dimensional dual problem involving two infinite-dimensional convex-constraints-sets is solved; secondly, the support of this solution is identified which is finally used to solve a finite-dimensional primal problem. Remarkably, for some specific cases, solving each of these steps is feasible which results in an exact finite-dimensional formulation (see for example [62, Section 2.4.2 and 2.4.3]).

2.7 Illustrations

We discuss the results obtained for the cases when the measurements are random samples either of the signal itself or of its continuous-domain Fourier transform. The operators of interest are $\mathrm{L} = \mathrm{D}$ and $\mathrm{L} = \mathrm{D}^2$. The ground truth (GT) signal f_{GT} is solution of the stochastic differential equation $\mathrm{L}f_{\mathrm{GT}} = w_{\mathrm{GT}}$ [78] for the two cases when w_{GT} is

- **Impulsive Noise**. Here, the innovation w_{GT} is a compound-Poisson noise with Gaussian jumps, which corresponds to a sum of Dirac impulses whose amplitudes follow a Gaussian distribution. The corresponding process f_{GT} has then the particularity of being piecewise smooth [79]. This case is matched to the regularization operator $\|\mathrm{L}f\|_{\mathcal{M}}$ and is covered by Theorem 2.4.2 which states that the minima f_1^* for this regularization case is such that

$$w_1^* = \mathrm{L}f_1^* = \sum_{k=1}^{K} a_k \delta(\cdot - x_k), \tag{2.53}$$

  which is a form compatible with a realization of an impulsive white noise.

- **Gaussian White Noise**. This case is matched to the regularization operator $\|\mathrm{L}f\|_{L_2}$. Unlike the impulsive noise, $w_2^* = \mathrm{L}f_2^*$ is not localized to finite points and therefore is a better model for the realization of a Gaussian white noise.

In all experiments, we also constrain the test signals to be compactly supported. This can be achieved by putting linear constraints on the innovations of the signal. In Sections 2.7.1 and 2.7.3, we confirm experimentally that matched regularization recovers the test signals better than non-matched regularization. While reconstructing the Tikhonov and gTV solutions when the measurements are noisy, the parameter λ in (2.34) and (2.45) is tuned using a grid search to give the best recovered SNR.

2.7.1 Random Sampling

In this experiment, the measurement functionals are Dirac impulses with the random locations $\{x_m\}_{m=1}^{M}$. The regularization operator is $\mathrm{L} = \mathrm{D}^2$. It corresponds to $\rho_{\mathrm{D}^2}(x) = -\frac{1}{2}|x|$ and $\varphi_{\mathrm{D}^2}(x) = \left(\rho_{\mathrm{L}^*\mathrm{L}} * h_m\right)(x) = |x - x_m|^3/12$. The null space is $\mathcal{N}_{\mathrm{D}^2} = \mathrm{span}\{1, x\}$ for this operator. This means that the gTV-regularized solution is piecewise linear and that the L_2-regularized solution is piecewise cubic. We compare in Figures 2.3.a and 2.3.b the recovery from noiseless samples of a second-order process, referred to as ground truth (GT). It is composed of sparse (impulsive Poisson) and non-sparse (Gaussian) innovations, respectively [80]. The sparsity index—the number of impulses or non-zero elements—for the original sparse signal is 9. The solution for the gTV case is recovered with $\Delta = 0.05$ and $N = 200$. The sparsity index of the gTV solution for the sparse and Gaussian cases are 9 and 16, respectively. As expected, the recovery of the gTV-regularized reconstruction is better than that of the L_2-regularized solution when the signal is sparse. For the Gaussian case, the situation is reversed.

2.7.2 Multiple Solutions

We discuss the case when the gTV solution is non-unique. We show in Figure 2.2.a examples of solutions of the gTV-regularized random-sampling problem obtained using FISTA alone (f_{F}^*) and FISTA + simplex (linear programming, f_{SLP}^*). In this case, $M = 30$, $\mathrm{L} = \mathrm{D}^2$, and $\lambda = 0.182$. The continuous-domain functions f_{F}^* and f_{SLP}^* have basis functions whose coefficients are the (non-unique) solutions of a given LASSO problem, as shown in Figure 2.2.b. The ℓ_1 norms of the corresponding coefficients are the same. Also, it holds that

$$\|\mathrm{D}^2 f_{\mathrm{F}}^*\|_{\mathcal{M}} = \|\mathrm{D}^2 f_{\mathrm{SLP}}^*\|_{\mathcal{M}} = \|\mathrm{D} f_{\mathrm{F}}^*\|_{\mathrm{TV}} = \|\mathrm{D} f_{\mathrm{SLP}}^*\|_{\mathrm{TV}}, \tag{2.54}$$

which implies that the TV norm of the slope of f_{F}^* and f_{SLP}^* are the same. This is evident from Figure 2.2.c. The arc-length of the two curves are the same. The signal f_{SLP}^* is piecewise linear $(21 < M)$, carries a piecewise-constant slope, and is by definition, a non-uniform spline of degree 1. By contrast, f_{F}^* has many more knots and even sections whose slope appears to be piecewise-linear.

Theorem 2.4.2 asserts that the extreme points of the solution set of the gTV regularization need to have fewer than M knots. Remember that f_{SLP}^* is obtained by combining FISTA and simplex; this ensures that the basis coefficients of f_{SLP}^*

are the extreme points of the solution set of the corresponding LASSO problem (Theorem 2.6.2) and guarantees that the number of knots is smaller than M.

This example shows an intuitive relationship between the continuous-domain and the discrete-domain formulations of inverse problems with gTV and ℓ_1 regularization, respectively. The nature of the continuous-domain solution set and its extreme points resonates with its corresponding discretized version. In both cases, the solution set is convex and the extreme points are sparse.

No. of		D		D^2	
impulses	Sparsity	TV	L_2	TV	L_2
10	Strong	**19.60**	15.7	**52.08**	41.54
100	Medium	**16.58**	16.10	**41.91**	41.26
2000	Low	14.45	**16.14**	39.68	**41.40**
-	Gaussian	14.30	**16.32**	40.05	**41.23**

(a) Noiseless case.

No. of		D		D^2	
impulses	Sparsity	TV	L_2	TV	L_2
10	Strong	**17.06**	11.52	**25.55**	24.60
100	Medium	**13.24**	10.94	**24.44**	24.24
2000	Low	10.61	**11.13**	25.80	**26.19**
-	Gaussian	10.40	**11.10**	24.95	**25.48**

(b) Noisy case.

Table 2.1: Comparison of TV and L_2 recovery from their (top table) noiseless and (bottom table) noisy (with 40 dB SNR) random Fourier samples. The results have been averaged over 40 realizations.

2.7.3 Random Fourier Sampling

Let now the measurement functions be $h_m(x) = \mathrm{rect}\left(\frac{x}{T}\right) \mathrm{e}^{-\mathrm{j}\omega_m x}$, where T is the window size. The samples are thus random samples of the continuous-domain Fourier transform of a signal restricted to a window. For the regularization operator $\mathrm{L} = \mathrm{D}$, the Green's function is $\rho_{\mathrm{D}}(x) = \frac{1}{2}\mathrm{sign}(x)$ and the basis is $\varphi_{\mathrm{D},m}(x) =$

$\left(\frac{1}{2}|\cdot| * h_m\right)(x)$. Figure 2.4.a and 2.4.b correspond to a first-order process with sparse and Gaussian innovations, respectively. The grid step $\Delta = 0.05$, $M = 41$, and $N = 200$. The sparsity index of the gTV solution for the sparse and Gaussian cases is 36 and 39, respectively. For the original sparse signal (GT), it is 7. The oscillations of the solution in the L_2-regularized case are induced by the sinusoidal form of the the measurement functionals. This also makes the L_2 solution intrinsically smoother than its gTV counterpart. Also, the quality of the recovery depends on the frequency band used to sample.

In Figures 2.4.c and 2.4.d, we show the zoomed version of the recovered second-order process with sparse and Gaussian innovations, respectively. The grid step is $\Delta = 0.05$, $M = 41$ and $N = 200$. The operator $\mathrm{L} = \mathrm{D}^2$ is used for the regularization. This corresponds to $\rho_{\mathrm{D}^2}(x) = \frac{|x|}{2}$ and $\varphi_{\mathrm{D}^2,m}(x) = \left(\frac{1}{12}|\cdot|^3 * h_m\right)(x)$. The sparsity index of the gTV solution in the sparse and Gaussian cases is 10 and 36, respectively. For the original sparse signal (GT), it is 10. Once again, the recovery by gTV is better than by L_2 when the signal is sparse. In the Gaussian case, the L_2 solution is better.

The effect of sparsity on the recovery of signals from their noiseless and noisy (40 dB SNR) Fourier samples are shown in Table 1. The sample frequencies are kept the same for all the cases. Here, $M = 41$, $N = 200$, $T = 10$, and the grid step $\Delta = 0.05$. We observe that reconstruction performances for random processes based on impulsive noise are comparable to that of Gaussian processes when the number of impulses increases. This is reminiscent of the fact that generalized-Poisson processes with Gaussian jumps are converging in law to corresponding Gaussian processes [81].

2.8 Summary

In this chapter, we consider 1D linear inverse problems that are formulated in the continuous domain. The object of recovery is a function that is assumed to minimize a convex objective functional. The solutions are constrained by imposing a continuous-domain regularization. We derive the parametric form of the solution (representer theorems) for Tikhonov (quadratic) and generalized total-variation (gTV) regularizations. We show that, in both cases, the solutions are splines that are intimately related to the regularization operator. In the Tikhonov case, the solution is smooth and constrained to live in a fixed subspace that depends on the measurement operator. By contrast, the gTV regularization results in a sparse solution composed of only a few dictionary elements that are upper-bounded by the number of measurements and independent of the measurement operator. Our findings for the gTV regularization resonates with the minimization of the ℓ_1 norm, which is its discrete counterpart and also produces sparse solutions. The formulations and the results of this chapter are summarized in Figure 2.5. Finally, we find the experimental solutions for some measurement models in one dimension. We discuss the special case when the gTV regularization results in multiple solutions and devise an algorithm to find an extreme point of the solution set which is guaranteed to be sparse.

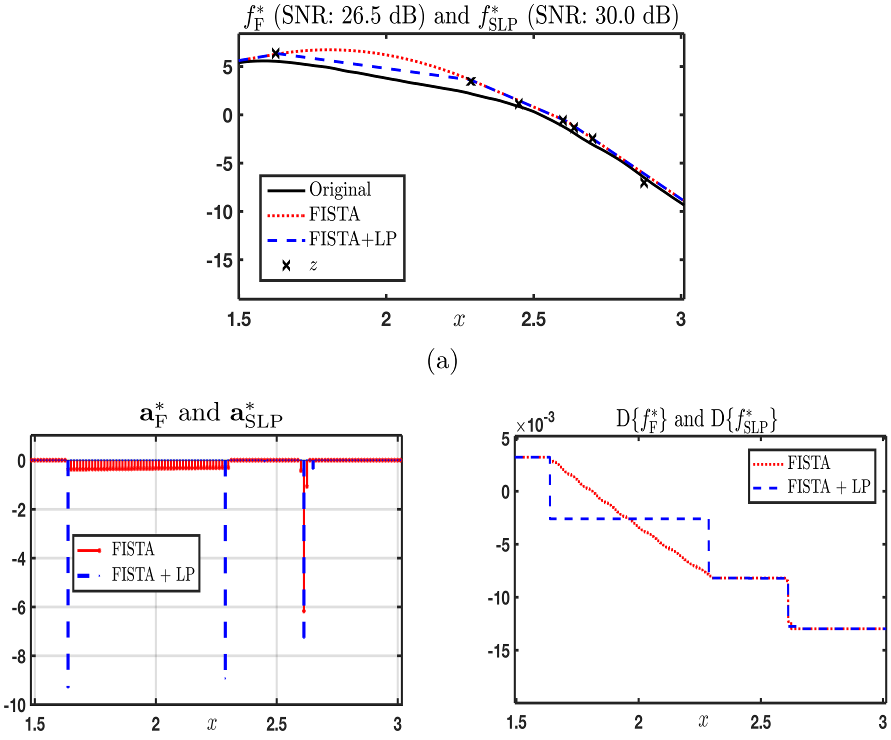

(a)

(b) (c)

Figure 2.2: Illustration of inability of FISTA to deliver a sparse solution : (a) comparison of solutions, f_{F}^* vs. f_{SLP}^* for continuous-domain gTV problem, (b) signal innovations with sparsity index 64 ($> M$) and 21 ($< M$), respectively, and (c) derivative of the two solutions. The two signal innovations in (b) are solutions of the same Lasso problem, but only $\mathbf{a}_{\mathrm{SLP}}^*$ is an extreme point of the solution set. The original signal is a second-order process ($\mathrm{L} = \mathrm{D}^2$) and the measurements are $M = 30$ nonuniform noisy samples (SNR $=$ 40 dB). The parameters are $\lambda = 0.182$, $N = 400$, and grid step $\Delta = \frac{1}{80}$.

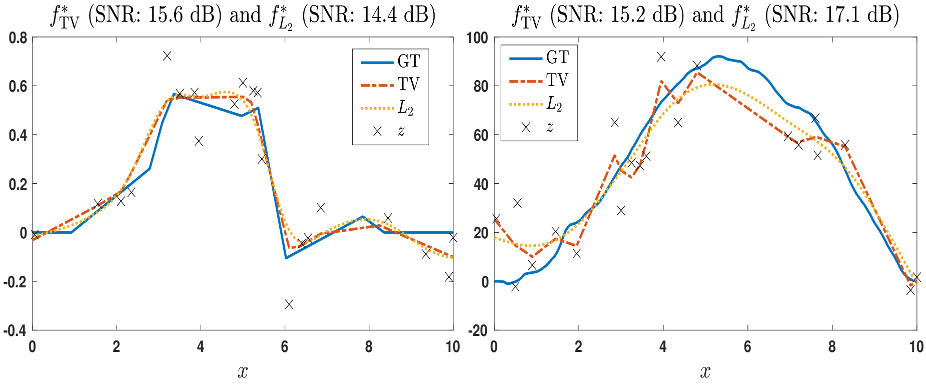

(a) Sparse Signal (b) Gaussian Signal

Figure 2.3: Recovery of sparse (a) and Gaussian (b) second-order processes (GT) using $\mathrm{L} = \mathrm{D}^2$ from their nonuniform samples corrupted with 40 dB measurement noise.

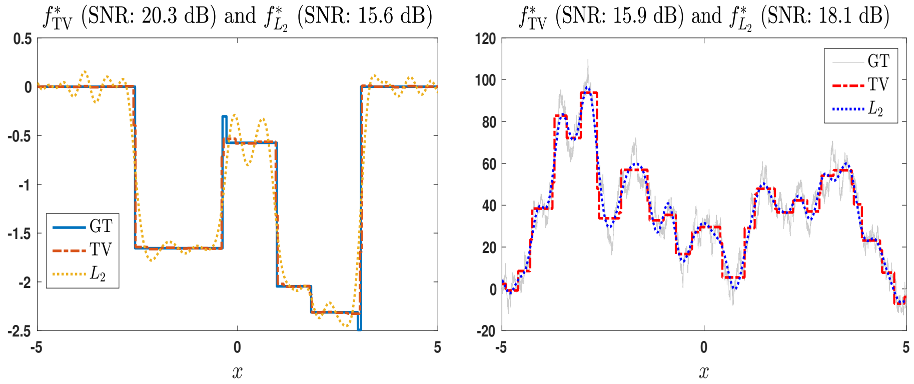

(a) Sparse Signal (b) Gaussian Signal

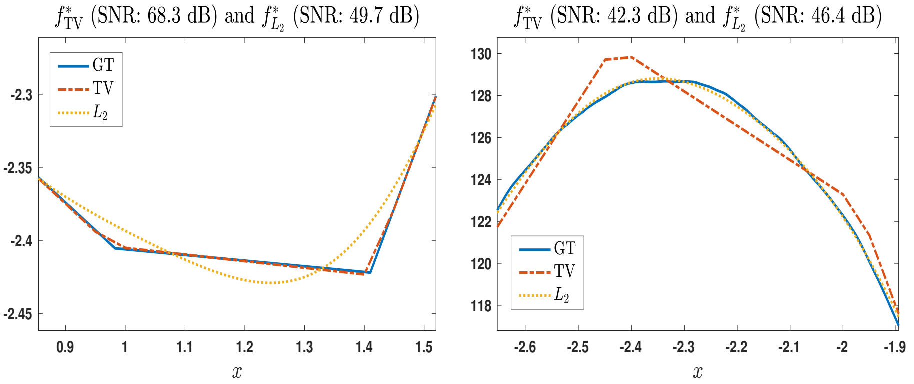

(c) Sparse Signal (d) Gaussian Signal

Figure 2.4: Recovery of first-order (first row) and second-order (second row) processes from their random noiseless Fourier samples. In all the cases, $M = 41$ and $N = 200$. In the interest of clarity, (c) and (d) contain the zoomed versions of the actual signals.

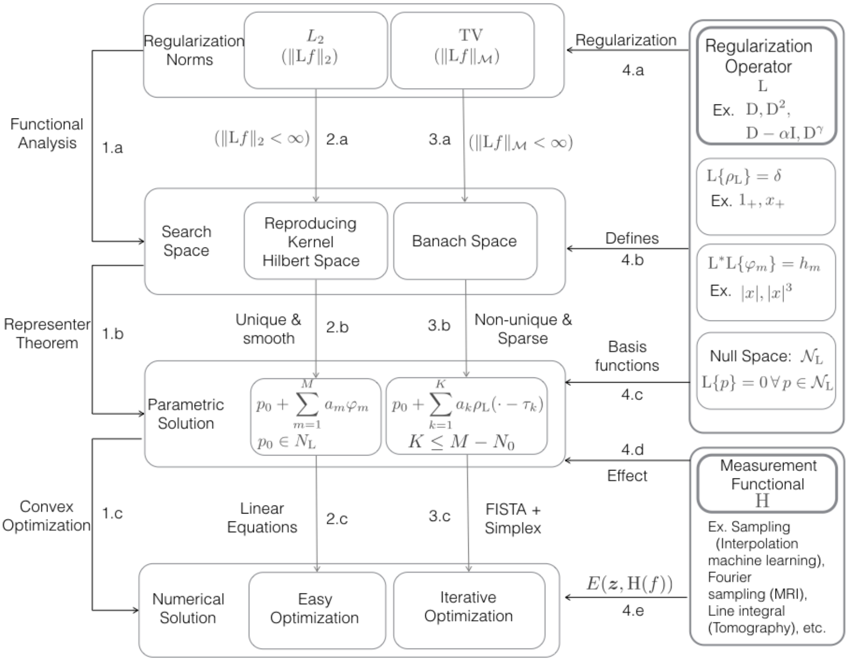

Figure 2.5: Summary of the whole scheme. The regularization operator with a given norm {4.a} defines the search space for the solution{1.a, 4.b}. Representer theorems then give the parametric representation of the solution {1.b}. The numerical solution is then recovered by optimizing over the parameters to minimize $\mathrm{J}_{\mathrm{R}}(f|\mathbf{y})$ {1.c}.

Part II

Second Generation

Chapter 3

Deep-Learning-based PGD for Iterative Reconstruction

[1]As discussed earlier, supervised deep-learning-based reconstruction methods outperform the classical methods. However, this comes at the cost of robustness. In this chapter, we introduce an iterative framework inspired from classical methods in order to bring more robustness and better performance to the deep-learning methods.

3.1 Overview

While medical imaging is a fairly mature area, there is recent evidence that it may still be possible to reduce the radiation dose and/or speedup the acquisition process without compromising image quality. This can be accomplished with the help of sophisticated reconstruction algorithms that incorporate some prior knowledge (*e.g.*, sparsity) on the class of underlying images [7]. The reconstruction task is usually formulated as an inverse problem where the image-formation physics are modeled by an operator $\mathbf{H} : \mathbb{R}^N \to \mathbb{R}^M$ (called the *forward model*). The measurement equation is $\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n} \in \mathbb{R}^M$, where $\mathbf{x} \in \mathbb{R}^N$ is the space-domain image that

---

[1] This content of this chapter are based on our work [26].

we are interested in recovering and $\mathbf{n} \in \mathbb{R}^M$ is the noise intrinsic to the acquisition process.

In the case of *extreme imaging*, the number of measurements is reduced as much as possible to decrease either the radiation dose in computed tomography (CT) or the scanning time in MRI. Moreover, the measurements are typically very noisy due to short integration times, which calls for some form of denoising. Indeed, there may be significantly fewer measurements than the number of unknowns ($M << N$). This gives rise to an ill-posed problem in the sense that there may be an infinity of consistent images that map to the same measurements $\mathbf{y}$. Thus, one challenge of the reconstruction algorithm is to select the best solution among a multitude of potential candidates.

Recently, a surge in using deep learning to solve inverse problems in imaging [16–20], has established new state-of-the-art results for tasks such as sparse-view CT reconstruction [17]. Rather than reconstructing the image from the measurements $\mathbf{y}$ directly, the most successful strategies have been to train the CNN as a regressor between a rough initial reconstruction $\mathbf{A}\mathbf{y}$, where $\mathbf{A} : \mathbb{R}^M \to \mathbb{R}^N$, and the final, desired reconstruction [17, 18]. This initial reconstruction could be obtained using classical algorithms (*e.g.*, FBP, BP) or by some other linear operation. Once the training is complete, the reconstruction for a new measurement $\mathbf{y}$ is given by $\mathbf{x}^* = \mathit{CNN}_{\boldsymbol{\theta}^*}(\mathbf{A}\mathbf{y})$, where $\mathit{CNN}_{\boldsymbol{\theta}} : \mathbb{R}^N \to \mathbb{R}^N$ denotes the CNN as a function and $\boldsymbol{\theta}^*$ denotes the internal parameters of the CNN after training. These schemes exploit the fact that the structure of images can be learned from representative examples. CNNs are favored because of the way they encode the data in their hidden layers. In this sense, a CNN can be seen as a good prior encoder.

Although the results reported so far are remarkable in terms of image quality, there is still some concern as to whether or not they can be trusted, especially in the context of diagnostic imaging. The main limitation of direct algorithms such as [17] is that they do not provide any guarantee on the worst-case performance. Moreover, even in the case of noiseless (or low-noise) measurements, there is no insurance that the reconstructed image is consistent with the measurements because, unlike for the iterative schemes, there is no feedback mechanism that imposes this consistency.

3.1.1 Contributions

In this work, we propose a simple yet effective iterative scheme (see Figure 3.1), which tries to incorporate the advantages of the existing algorithms and side-steps

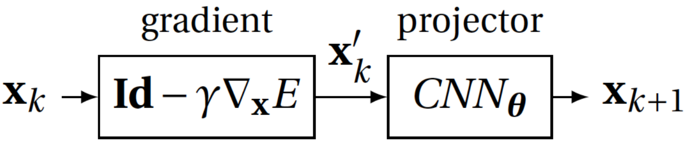

(a) Projected gradient descent

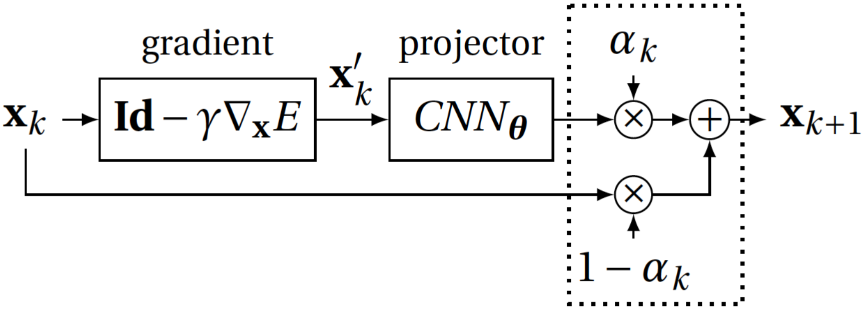

(b) Relaxed projected gradient descent

Figure 3.1: (a) Block diagram of projected gradient descent using a CNN as the projector and E as the data-fidelity term. The gradient step promotes consistency with the measurements and the projector forces the solution to belong to the set of desired solutions. If the CNN is only an approximate projector, the scheme may diverge. (b) Block diagram of the proposed relaxed projected gradient descent. The α_ks are updated in such a way that the algorithm always converges (see Algorithm 1 for more details).

their disadvantages. Specifically:

- We first propose to learn a CNN that acts as a projector onto a set $\mathcal{S}$ which can be intuitively thought of as the manifold of the data (*e.g.*, biomedical images). In this sense, our CNN encodes the prior knowledge of the data. Its purpose is to map an input image to an output image that is more similar to the training data.
- Given a measurement $\mathbf{y}$, we initialize our reconstruction using a classical algorithm.
- We then iteratively alternate between minimizing the data-fidelity term and projecting the result onto the set $\mathcal{S}$ by applying a suitable variant of the projected gradient descent (PGD) which ensures convergence.

Besides the design of the implementation, our contribution is in the proposal of the relaxed form of PGD that is guaranteed to converge and under certain conditions can also find a local minima of a nonconvex inverse problem. Moreover, as we shall see later, this method outperforms existing algorithms on low-dose x-ray CT reconstructions.

3.1.2 Related and Prior Work

Deep learning has already shown promising results in image denoising, superresolution, and deconvolution. Recently, it has also been used to solve inverse problems in imaging using limited data [17–20], and in compressed sensing [82]. However, as discussed earlier, these regression-based approaches lack a feedback mechanism that could be beneficial in solving inverse problems.

Another usage of deep learning is to complement iterative algorithms. This includes learning a CNN as an unrolled version of the iterative shrinkage-thresholding algorithm (ISTA) [83] or ADMM [84]. In [21], inverse problems involving non-linear forward models are solved by partially learning the gradient descent. In [85], the iterative algorithm is replaced by a recurrent neural network (RNN). Recently, in [86], a cascade of CNNs is used to reconstruct images. Within this cascade the data-fidelity is enforced at multiple steps. However, in all of these approaches the training is performed end-to-end, meaning that the network parameters are dependent on the iterative scheme chosen.

These approaches differ from plug-and-play ADMM [87–89], where an independent off-the-shelf denoiser or a trained operator is plugged into the iterative scheme of the alternating-direction method of multipliers (ADMM) [13]. ADMM is an iterative optimization technique that alternates between (i) a linear solver that reinforces consistency with respect to the measurements; and (ii) a nonlinear operation that re-injects the prior. The idea of plug-and-play ADMM is to replace (ii), which resembles denoising, with an off-the-shelf denoiser. Plug-and-play ADMM is more general than the optimization framework (1.6) but still lacks theoretical justifications. In fact, there is little understanding yet of the connection between the use of a given denoiser and the regularization it imposes (though this link has recently been explored in [90]).

In [91], a generative adversarial network (GAN) trained as a projector onto a set, has been used with the plug-and-play ADMM. Similarly, in [92], the inverse problem is solved over a set parameterised by a generative model. However, it requires a precise initialization of the parameters. In [93], similarly to us, the projector in PGD is replaced with a neural network. However, the scheme lacks convergence guarantee and a rigorous theoretical analysis.

Our scheme is similar in spirit to plug-and-play ADMM, but is simpler to analyze. Although our methodology is generic and can be applied in principle to any inverse problem, our experiments here involve sparse-view x-ray CT reconstruction. For a recent overview of the field, see [94]. Current approaches to sparse-view CT reconstruction follow the formulation (1.6), e.g., using a penalized weighted least-squares data term and sparsity-promoting regularizer [95], dictionary learning-based regularizer [96], or generalized total variation regularizer [97]. There are also prior works on the direct application of CNNs to CT reconstruction. These methods generally use the CNN to denoise the sinogram [98] or the reconstruction obtained from a standard technique [17, 99–101]; as such, they do not perform the reconstruction directly.

3.1.3 Roadmap

The chapter is organized as follows: In Section 3.2, we discuss the mathematical framework that motivates our approach and justify the use of a projector onto a set as an effective strategy to solve inverse problems. In Section 3.3, we present our algorithm, which is a relaxed version of PGD. It has been modified so as to converge in practical cases where the projection property is only approximate. We discuss in Section 3.4 a novel technique to train the CNN as a projector onto a set, especially when the training data is small. This is followed by experiments (Section 3.5), results and discussions (Section 3.6 and Section 3.7), and conclusions (Section 3.8).

3.2 Theoretical Framework

Our goal is to use a trained CNN iteratively inside PGD to solve an inverse problem. To understand why this scheme will be effective, we first analyze how using a projector onto a set, combined with gradient descent, can be helpful in solving inverse problems. Properties of PGD using an orthogonal projector onto a convex set are known [102]. Here, we extend these results for any projector onto a non-convex set. This extension is required because there is no guarantee that the set of desirable reconstruction images is convex. Proofs of all the results in this section can be found in the Appendix.

3.2.1 Notation

We consider the finite-dimensional Hilbert space $\mathbb{R}^N$ equipped with the scalar product $\langle \cdot\,,\,\cdot \rangle$ that induces the ℓ_2 norm $\|\cdot\|_2$. The spectral norm of the matrix $\mathbf{H}$, denoted by $\|\mathbf{H}\|_2$, is equal to its largest singular value. For $\mathbf{x} \in \mathbb{R}^N$ and $\varepsilon > 0$, we denote by $\mathcal{B}_\varepsilon(\mathbf{x})$ the ℓ_2-ball centered at $\mathbf{x}$ with radius ε, *i.e.*,

$$\mathcal{B}_\varepsilon(\mathbf{x}) = \left\{ \mathbf{z} \in \mathbb{R}^N : \|\mathbf{z} - \mathbf{x}\|_2 \leq \varepsilon \right\}.$$

The operator $T : \mathbb{R}^N \to \mathbb{R}^N$ is Lipschitz-continuous with constant L if

$$\|T(\mathbf{x}) - T(\mathbf{z})\|_2 \leq L \, \|\mathbf{x} - \mathbf{z}\|_2 , \quad \forall \mathbf{x}, \mathbf{z} \in \mathbb{R}^N.$$

It is contractive if it is Lipschitz-continuous with constant $L < 1$ and non-expansive if $L = 1$. A fixed point $\mathbf{x}^*$ of T (if any) satisfies $T(\mathbf{x}^*) = \mathbf{x}^*$.

Given the set $\mathcal{S} \subset \mathbb{R}^N$, the mapping $P_\mathcal{S} : \mathbb{R}^N \to \mathcal{S}$ is called a projector if it satisfies the idempotent property $P_\mathcal{S} P_\mathcal{S} = P_\mathcal{S}$. It is called an orthogonal projector if

$$P_\mathcal{S}(\mathbf{x}) = \inf_{\mathbf{z} \in \mathcal{S}} \|\mathbf{x} - \mathbf{z}\|_2 , \quad \forall \mathbf{x} \in \mathbb{R}^N.$$

3.2.2 Constrained Least Squares

Consider the problem of the reconstruction of the image $\mathbf{x} \in \mathbb{R}^N$ from its noisy measurements $\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n}$, where $\mathbf{H} \in \mathbb{R}^{M \times N}$ is the linear forward model and $\mathbf{n} \in \mathbb{R}^M$ is additive white Gaussian noise. The framework is also applicable to Poisson noise model-based CT via a suitable transformation, as shown in Appendix A.2.2.

Our reconstruction incorporates a strong form of prior knowledge about the original image: We assume that $\mathbf{x}$ must lie in a set $\mathcal{S} \subset \mathbb{R}^N$ that contains all objects of interest. The proposed way to make the reconstruction consistent with the measurements as well as with the prior knowledge is to solve the constrained least-squares problem

$$\min_{\mathbf{x} \in \mathcal{S}} \frac{1}{2} \|\mathbf{H}\mathbf{x} - \mathbf{y}\|_2^2 . \tag{3.1}$$

The condition $\mathbf{x} \in \mathcal{S}$ in (3.1) plays the role of a regularizer. If no two points in $\mathcal{S}$ have the same measurements and in case $\mathbf{y}$ is noiseless, then out of all the points in $\mathbb{R}^N$ that are consistent with the measurement $\mathbf{y}$, (3.1) selects a unique point $\mathbf{x}^* \in \mathcal{S}$. In this way, the ill-posedness of the inverse problem is bypassed. When the measurements are noisy, (3.1) returns a point $\mathbf{x}^* \in \mathcal{S}$ such that $\mathbf{y}^* = \mathbf{H}\mathbf{x}^*$ is as close as possible to $\mathbf{y}$. Thus, it also denoises the measurement, where the quantity $\mathbf{y}^*$ can be regarded as the denoised version of $\mathbf{y}$.

The point $\mathbf{x}^* \in \mathcal{S}$ is called a local minimizer of (3.1) if

$$\exists \varepsilon > 0 : \|\mathbf{H}\mathbf{x}^* - \mathbf{y}\|_2 \leq \|\mathbf{H}\mathbf{x} - \mathbf{y}\|_2 , \forall \mathbf{x} \in \mathcal{S} \cap \mathcal{B}_\varepsilon(\mathbf{x}^*).$$

3.2.3 Projected Gradient Descent

When $\mathcal{S}$ is a closed convex set, it is well known [102] that a solution of (3.1) can be found by PGD

$$\mathbf{x}_{k+1} = P_{\mathcal{S}}(\mathbf{x}_k - \gamma \mathbf{H}^{\mathrm{T}}\mathbf{H}\mathbf{x}_k + \gamma \mathbf{H}^{\mathrm{T}}\mathbf{y}), \tag{3.2}$$

where γ is a step size chosen such that $\gamma < 2/\left\|\mathbf{H}^{\mathrm{T}}\mathbf{H}\right\|_2$. This algorithm combines the orthogonal projection onto $\mathcal{S}$ with the gradient descent with respect to the quadratic objective function, also called the Landweber update [103]. PGD [104, Section 2.3] is a subclass of the forward-backward splitting [105, 106], which is known in the ℓ_1-minimization literature as iterative shrinkage/thresholding algorithms (ISTA) [10, 11, 107].

In our problem, $\mathcal{S}$ is presumably non-convex, but we propose to still use the update (3.2) with some projector $P_{\mathcal{S}}$ that may not be orthogonal. In the rest of this section, we provide sufficient conditions on the projector $P_{\mathcal{S}}$ (not on $\mathcal{S}$ itself) under which (3.2) leads to a local minimizer of (3.1). Similarly to the convex case, we characterize the local minimizers of (3.1) by the fixed points of the combined operator

$$G_\gamma(\mathbf{x}) = P_{\mathcal{S}}(\mathbf{x} - \gamma \mathbf{H}^{\mathrm{T}}\mathbf{H}\mathbf{x} + \gamma \mathbf{H}^{\mathrm{T}}\mathbf{y}) \tag{3.3}$$

and then show that some fixed point of that operator must be reached by the iteration $\mathbf{x}_{k+1} = G_\gamma(\mathbf{x}_k)$ as $k \to \infty$, regardless of the initial point $\mathbf{x}_0$. We first state a sufficient condition for each fixed point of G_γ to become a local minimizer of (3.1).

Proposition 3.2.1. *Let $\gamma > 0$ and $P_{\mathcal{S}}$ be such that, for all $\mathbf{x} \in \mathbb{R}^N$,*

$$\langle \mathbf{z} - P_{\mathcal{S}}\mathbf{x}\,,\, \mathbf{x} - P_{\mathcal{S}}\mathbf{x}\rangle \leq 0, \quad \forall \mathbf{z} \in \mathcal{S} \cap \mathcal{B}_\varepsilon(P_{\mathcal{S}}\mathbf{x}), \tag{3.4}$$

for some $\varepsilon > 0$. Then, any fixed point of the operator G_γ in (3.3) is a local minimizer of (3.1). Furthermore, if (3.4) is satisfied globally, in the sense that

$$\langle \mathbf{z} - P_{\mathcal{S}}\mathbf{x}\,,\, \mathbf{x} - P_{\mathcal{S}}\mathbf{x}\rangle \leq 0, \quad \forall \mathbf{x} \in \mathbb{R}^N, \mathbf{z} \in \mathcal{S}, \tag{3.5}$$

then any fixed point of G_γ is a solution of (3.1).

Two remarks are in order. First, (3.5) is a well-known property of orthogonal projections onto closed convex sets. It actually implies the convexity of $\mathcal{S}$ (see Proposition 3.2.2). Second, (3.4) is much more relaxed and easily achievable, for example, as stated in Proposition 3.2.3, by orthogonal projections onto unions of closed convex sets. (Special cases are unions of subspaces, which have found some applications in data modeling and clustering [108]).

Proposition 3.2.2. *If $P_{\mathcal{S}}$ is a projector onto $\mathcal{S} \subset \mathbb{R}^N$ that satisfies (3.5), then $\mathcal{S}$ must be convex.*

Proposition 3.2.3. *If $\mathcal{S}$ is a union of a finite number of closed convex sets in $\mathbb{R}^N$, then the orthogonal projector $P_{\mathcal{S}}$ onto $\mathcal{S}$ satisfies (3.4).*

Propositions 1-3 suggest that, when $\mathcal{S}$ is non-convex, the best we can hope for is to find a local minimizer of (3.1) through a fixed point of G_γ. Theorem 3.2.4 provides a sufficient condition for PGD to converge to a unique fixed point of G_γ.

Theorem 3.2.4. *Let* $\lambda_{\max}$ *and* $\lambda_{\min}$ *be the largest and smallest eigenvalues of* $\mathbf{H}^{\mathrm{T}}\mathbf{H}$*, respectively. If* $P_{\mathcal{S}}$ *satisfies* (3.4) *and is Lipschitz-continuous with constant* $L < (\lambda_{\max} + \lambda_{\min})/(\lambda_{\max} - \lambda_{\min})$*, then, for* $\gamma = 2/(\lambda_{\max} + \lambda_{\min})$*, the sequence* $\{\mathbf{x}_k\}$ *generated by* (3.2) *converges to a local minimizer of* (3.1)*, regardless of the initialization* $\mathbf{x}_0$*.*

It is important to note that the projector $P_{\mathcal{S}}$ can never be contractive since it preserves the distance between any two points on $\mathcal{S}$. Therefore, when $\mathbf{H}$ has a non-trivial null space, the condition $L < (\lambda_{\max} + \lambda_{\min})/(\lambda_{\max} - \lambda_{\min})$ of Theorem 3.2.4 is not feasible. The smallest possible Lipschitz constant of $P_{\mathcal{S}}$ is $L = 1$, which means that $P_{\mathcal{S}}$ is non-expansive. Even with this condition, it is not guaranteed that the combined operator F_γ has a fixed point. This limitation can be overcome when F_γ is assumed to have a nonempty set of fixed points. Indeed, we state in Theorem 3.2.5 that one of them must be reached by iterating the averaged operator $\alpha\,\mathrm{Id} + (1-\alpha)G_\gamma$, where $\alpha \in (0,1)$ and Id is the identity operator. We call this scheme averaged PGD (APGD).

Theorem 3.2.5. *Let* $\lambda_{\max}$ *be the largest eigenvalue of* $\mathbf{H}^{\mathrm{T}}\mathbf{H}$*. If* $P_{\mathcal{S}}$ *satisfies* (3.4) *and is a non-expansive operator such that* G_γ *in* (3.3) *has a fixed point for some* $\gamma < 2/\lambda_{\max}$*, then the sequence* $\{\mathbf{x}_k\}$ *generated by APGD, with*

$$\mathbf{x}_{k+1} = (1-\alpha)\mathbf{x}_k + \alpha G_\gamma(\mathbf{x}_k) \tag{3.6}$$

for any $\alpha \in (0,1)$*, converges to a local minimizer of* (3.1)*, regardless of the initialization* $\mathbf{x}_0$*.*

3.3 Relaxation with Guaranteed Convergence

Despite their elegance, Theorems 3.2.4 and 3.2.5 are not directly productive when we construct the projector $P_{\mathcal{S}}$ by training a CNN because it is unclear how to enforce the Lipschitz continuity of $P_{\mathcal{S}}$ on the CNN architecture. Without putting any constraints on the CNN, however, we can still achieve the convergence of the reconstruction sequence by modifying PGD as described in Algorithm 1; we name it relaxed projected gradient descent (RPGD). In Algorithm 1, the projector $P_{\mathcal{S}}$ is replaced by the general nonlinear operator F. We also introduce a sequence $\{c_k\}$ that governs the rate of convergence of the algorithm and a sequence $\{\alpha_k\}$ of relaxation parameters that evolves with the algorithm. The convergence of RPGD is guaranteed by Theorem 3.3.1. More importantly, if the nonlinear operator F is actually a projector and the relaxation parameters do not go all the way to 0, then RPGD converges to a meaningful point.

Algorithm 1 Relaxed projected gradient descent (RPGD)

Input: $\mathbf{H}$, $\mathbf{y}$, $\mathbf{A}$, nonlinear operator F, step size $\gamma > 0$, positive sequence $\{c_n\}_{n\geq 1}$, $\mathbf{x}_0 = \mathbf{A}\mathbf{y} \in \mathbb{R}^N$, $\alpha_0 \in (0, 1]$.
Output: reconstructions $\{\mathbf{x}_k\}$, relaxation parameters $\{\alpha_k\}$.

```
$k \leftarrow 0$
while not converged do
  $\mathbf{z}_k = F(\mathbf{x}_k - \gamma\mathbf{H}^{\mathrm{T}}\mathbf{H}\mathbf{x}_k + \gamma\mathbf{H}^{\mathrm{T}}\mathbf{y})$
  if $k \geq 1$ then
    if $\|\mathbf{z}_k - \mathbf{x}_k\|_2 > c_k \|\mathbf{z}_{k-1} - \mathbf{x}_{k-1}\|_2$ then
      $\alpha_k = c_k\|\mathbf{z}_{k-1} - \mathbf{x}_{k-1}\|_2 / \|\mathbf{z}_k - \mathbf{x}_k\|_2\, \alpha_{k-1}$
    else
      $\alpha_k = \alpha_{k-1}$
    end if
  end if
  $\mathbf{x}_{k+1} = (1 - \alpha_k)\mathbf{x}_k + \alpha_k\mathbf{z}_k$
  $k \leftarrow k + 1$
end while
```

Theorem 3.3.1. *Let the input sequence $\{c_k\}$ of Algorithm 1 be asymptotically upper-bounded by $C < 1$. Then, the following statements hold true for the reconstruction sequence $\{\mathbf{x}_k\}$:*

(i) $\mathbf{x}_k \to \mathbf{x}^*$ *as* $k \to \infty$, *for all choices of* F*;*

(ii) if F is continuous and the relaxation parameters $\{\alpha_k\}$ are lower-bounded by $\varepsilon > 0$, then $\mathbf{x}^*$ is a fixed point of

$$G_\gamma(\mathbf{x}) = F(\mathbf{x} - \gamma\mathbf{H}^{\mathrm{T}}\mathbf{H}\mathbf{x} + \gamma\mathbf{H}^{\mathrm{T}}\mathbf{y}); \tag{3.7}$$

Table 3.1: Convergence of different algorithms for different cases. Here G_γ and F are related by (3.7), and the terms global and local minima are in the context of the problem (3.1). Note that when F is CNN_θ, only RPGD offers convergence guarantee.

Algorithm	$\mathcal{S}$	F	Converges	Converged Point Property
PGD	Convex	$P_\mathcal{S}$	Always	Global Minima
	NonConvex	$P_\mathcal{S}$	Not always	Fixed Point (FP) of G_γ
	NonConvex	$P_\mathcal{S}$ with (3.4)	Not always	Local Minima
	NonConvex	CNN_θ	Not always	Unknown
APGD	NonConvex	$P_\mathcal{S}$ with (3.4) and $L = 1$	If G_γ has FP	Local Minima
	NonConvex	CNN_θ	Not always	Unknown
RPGD	NonConvex	$P_\mathcal{S}$ with (3.4)	Always	Local Minima if $\alpha_k > 0$
	NonConvex	CNN_θ	Always	FP of G_γ if $\alpha_k > 0$

(iii) if, in addition to (ii)*, F is indeed a projector onto $\mathcal{S}$ that satisfies* (3.4)*, then* $\mathbf{x}^*$ *is a local minimizer of* (3.1)*.*

We prove Theorem 3.3.1 in Appendix A.2.1. Note that the weakest statement here is (i); it guarantees that RPGD always converges, albeit not necessarily to a fixed point of G_γ. Moreover, the assumption about the continuity of F in (ii) is automatically satisfied when F is a CNN.

In summary, we have described three algorithms: PGD, APGD, and RPGD. PGD is a standard algorithm which, in the event of convergence, finds a local minima of (3.1); however, it does not always converge. APGD ensures convergence under the broader set of conditions given in Theorem 3.2.5; but, in order to have these properties, both PGD and APGD necessarily need a projector. While, we shall train our CNN to act like a projector, it may not exactly fulfill the required conditions. This is the motivation for RPGD, which, unlike PGD and APGD, is guaranteed to converge. It also retains the desirable properties of PGD and APGD: it finds a local minima of (3.1), given that the conditions (ii) and (iii) of Theorem 3.3.1 are satisfied. Note, however, that when the set $\mathcal{S}$ is nonconvex, this local minimum may not be a global minimum. The results of Section 3.2 and 3.3

are summarized in Table 3.1.

3.4 Training a CNN as a Projector

For any point $\mathbf{x} \in \mathcal{S}$, a projector onto $\mathcal{S}$ should satisfy $P_{\mathcal{S}}\mathbf{x} = \mathbf{x}$. Moreover, we want that

$$\mathbf{x} = P_{\mathcal{S}}(\tilde{\mathbf{x}}), \tag{3.8}$$

where $\tilde{\mathbf{x}}$ is any perturbed version of $\mathbf{x}$. Given the training set, $\{\mathbf{x}^1, \ldots, \mathbf{x}^Q\}$ of Q points drawn from the set $\mathcal{S}$, we generate the ensemble

$$\{\{\tilde{\mathbf{x}}^{1,1}, \ldots, \tilde{\mathbf{x}}^{Q,1}\}, \ldots, \{\tilde{\mathbf{x}}^{1,N} \ldots, \tilde{\mathbf{x}}^{Q,N}\}\}$$

of $N \times Q$ perturbed points and train the CNN by minimizing the loss function

$$J(\boldsymbol{\theta}) = \sum_{n=1}^{N} \underbrace{\sum_{q=1}^{Q} \|\mathbf{x}^q - CNN_{\boldsymbol{\theta}}(\tilde{\mathbf{x}}^{q,n})\|_2^2}_{J_n(\boldsymbol{\theta})}. \tag{3.9}$$

The optimization proceeds by stochastic gradient descent for T epochs, where an epoch is defined as one pass though the training data.

It remains to select the perturbations that generate the $\mathbf{x}^{q,n}$. Our goal here is to create a diverse set of perturbations so that the CNN does not overfit one specific type. In our experiments, while training for the tth epoch, we chose

$$\tilde{\mathbf{x}}^{q,1} = \mathbf{x}^q \tag{3.10}$$
$$\tilde{\mathbf{x}}^{q,2} = \mathbf{A}\mathbf{H}\mathbf{x}^q \tag{3.11}$$
$$\tilde{\mathbf{x}}^{q,3} = CNN_{\boldsymbol{\theta}_{t-1}}(\tilde{\mathbf{x}}^{q,2}), \tag{3.12}$$

where $\mathbf{A}$ is a classical linear reconstruction algorithm (FBP in our experiments), and $\boldsymbol{\theta}_t$ are the CNN parameters after t epochs. Equations (3.10), (3.11), and (3.12) correspond to no perturbation, a linear perturbation, and a dynamic nonlinear perturbation, respectively. We now comment on each perturbation in detail.

Keeping $\tilde{\mathbf{x}}^{q,1}$ in the training ensemble will train the CNN with the defining property of the projector: the projector maps a point in the set $\mathcal{S}$ onto itself. If the CNN were trained only with (3.10), it would be an autoencoder [14].

To understand the perturbation $\tilde{\mathbf{x}}^{q,2}$ in (3.11), recall that $\mathbf{AHx}^q$ is the classical linear reconstruction of $\mathbf{x}^q$ from its measurement $\mathbf{y} = \mathbf{Hx}^q$. Perturbation (3.11) is indeed useful because we initialize RPGD with $\mathbf{AHx}^q$. Using only (3.11) for training would return the same CNN as in [17].

The perturbation $\tilde{\mathbf{x}}^{q,3}$ in (3.12) is the output of the CNN whose parameters $\boldsymbol{\theta}_t$ change with every epoch t; thus, it is a nonlinear and dynamic (epoch-dependent) perturbation of $\mathbf{x}^q$. The rationale for using (3.12) is that it greatly increases the training diversity by allowing the network to see T new perturbations of each training point, without greatly increasing the total training size since it only requires Q additional gradient computations per epoch. Moreover, (3.12) is in sync with the iterative scheme of RPGD, where the output of the CNN is processed with a gradient descent and is again fed back into itself.

3.4.1 Architecture

Our CNN architecture is the same as in [17], which is a U-net [109] with intrinsic skip connections among its layers and an extrinsic skip connection between the input and the output. The intrinsic skip connections help to eliminate singularities during the training [110]. The extrinsic skip connections make this network a residual net; *i.e.*, $\mathit{CNN} = \mathrm{Id} + \mathit{Unet}$, where Id denotes the identity operator and $\mathit{Unet} : \mathbb{R}^N \to \mathbb{R}^N$ denotes U-net as a function. Therefore, U-net actually provides the projection error (negative perturbation) that should be added to the input to get the projection.

Residual nets have been shown to be effective for image recognition [111] and for solving inverse problems [17]. While the residual-net architecture does not increase the capacity or the approximation power of the CNN, it does help in learning functions that are close to an identity operator, as is the case in our setting.

3.4.2 Sequential Training Strategy

We train the CNN in three stages. In Stage 1, we train it for T_1 epochs with respect to the partial-loss function J_2 in (3.9) which only uses the ensemble $\{\tilde{\mathbf{x}}^{q,2}\}$ generated by (3.11). In Stage 2, we add the ensemble $\{\tilde{\mathbf{x}}^{q,3}\}$ according to (3.12) at every epoch and then train the CNN with respect to the loss function $J_2 + J_3$; we repeat this procedure for T_2 epochs. Finally, in Stage 3, we train the CNN for

T_3 epochs with all three ensembles $\{\tilde{\mathbf{x}}^{q,1}, \tilde{\mathbf{x}}^{q,2}, \tilde{\mathbf{x}}^{q,3}\}$ to minimize the original loss function $J = J_1 + J_2 + J_3$ from (3.9).

We shall see in Section 3.7.2 that this sequential procedure speeds up the training without compromising the performance. The parameters of $\mathscr{U}$ are initialized by a normal distribution with a very low variance. Since *CNN* $= \mathrm{Id} + \mathscr{U}$, this function acts close to an identity operator in the initial epochs and makes it redundant to use $\{\tilde{\mathbf{x}}^{q,1}\}$ for the initial training stages. Therefore, $\{\tilde{\mathbf{x}}^{q,1}\}$ is only added at the last stage when the *CNN* is no longer close to an identity operator. After training with only $\{\tilde{\mathbf{x}}^{q,2}\}$ in Stage 1, $\tilde{\mathbf{x}}^{q,3}$ will be close to $\mathbf{x}^q$ since it is the output of the CNN for the input $\tilde{\mathbf{x}}^{q,2}$. This eases the training for $\{\tilde{\mathbf{x}}^{q,3}\}$ in the second and third stage.

3.5 Experiments

We validate the proposed method on the challenging case of sparse-view CT reconstruction. Conventionally, CT imaging requires many views to obtain good quality reconstruction. We call this scenario full-dose reconstruction. Our main aim in these experiments is to reduce the number of views (or dose) for CT imaging while retaining the quality of full-dose reconstructions. We denote a k-times reduction in views by $\times k$.

The measurement operator $\mathbf{H}$ for our experiments is the Radon transform. It maps an image to the values of its integrals along a known set of lines [2]. In 2D, the measurements are indexed by the angle and offset of each lines and arranged in a 2D sinogram. We implemented $\mathbf{H}$ and $\mathbf{H}^{\mathrm{T}}$ with Matlab's `radon` and `iradon` (normalized to satisfy the adjoint property), respectively. The Matlab code for the RPGD and the sequential-strategy-based training are made publically available[2].

3.5.1 Datasets

We use two datasets for our experiments.

1) Mayo Clinic Dataset. It consists of 500 clinically realistic, (512 × 512) CT images from the lower lungs to the lower abdomen of 10 patients. Those were obtained from the Mayo clinic AAPM Low Dose CT Grand Challenge [112].

2) Rat Brain Dataset. We use a real (1493 px × 720 view × 377 slice) sinogram from a CT scan of a single rat brain. The data acquisition was performed at the

---

[2] `https://github.com/harshit-gupta-epfl/CNN-RPGD`

Paul Scherrer Institute in Villigen, Switzerland at the TOMCAT beam line of the Swiss Light Source. During pre-processing, we split this sinogram slice-by-slice and downsampled it to create a dataset of 377 (729 px × 720 view) sinograms. CT images of size (512×512) were then generated from these full-dose sinograms (using the FBP, see Section 3.5.3). For the qth z-slice, we denote the corresponding image $\mathbf{x}^q_{\mathrm{FD}}$. For experiments based on this dataset, the first 327 and the last 25 slices are used for training and testing, respectively. This left a gap of 25 slices in between the training and testing data.

3.5.2 Experimental Setups

We now describe three experimental setups. We use the first dataset for the first experiment and the second for the last two.

1) Experiment 1. We split the Mayo dataset into 475 images from 9 patients for training and 25 images from the remaining patient for testing. We assume these images to be the ground truth. From the qth image $\mathbf{x}^q$, we generated the sparse-view sinogram $\mathbf{y}^q = \mathbf{H}\mathbf{x}^q$ using several different experimental conditions. Our task is to reconstruct the image from the sinogram.

The sinograms always have 729 offsets per view, but we varied the number of views and the level of measurement noise for different cases. We took 144 views and 45 views, which corresponds to ×5 and ×16 dosage reductions (assuming a full-view sinogram has 720 views). We added Gaussian noise to the sinograms to make the SNR equal to 35, 40, 45, 70, and infinity dB, where we refer to the first three as *high measurement noise* and the last two as *low measurement noise*. The SNR of the sinogram $\mathbf{y} + \mathbf{n}$ is defined as

$$\mathrm{SNR}(\mathbf{y} + \mathbf{n}, \mathbf{y}) = 20 \log_{10} \left(\|\mathbf{y}\|_2 / \|\mathbf{n}\|_2 \right). \tag{3.13}$$

For testing with the low and high measurement noise, we trained the CNNs without noise and at the 40-dB level of noise, respectively (see Section 3.5.4 for details).

To make the experiments more realistic and to reduce the inverse crime, the sinograms were generated by slightly perturbing the angles of the views by a zero-mean additive white Gaussian noise (AWGN) with standard deviation of 0.05 degrees. This creates a deliberate mismatch between the actual measurement process and the forward model.

2) Experiment 2. We used images $\mathbf{x}^q_{\mathrm{FD}}$ from the rat-brain dataset to generate Poisson-noise-corrupted sinograms $\mathbf{y}^q$ with 144 views. Just as in Experiment 1, the task is to reconstruct $\mathbf{x}_{\mathrm{FD}}^q$ back from $\mathbf{y}^q$. Sinograms were generated with 25, 30, and 35 dB SNR with respect to $\mathbf{H}\mathbf{x}_{\mathrm{FD}}^q$. To achieve this, in (A.35) and (A.36), we assume the readout noise to be zero and $\{b_1, \dots, b_m\} = b_0 = 1.66 \times 10^5, 5.24 \times 10^5$, and 1.66×10^6, respectively. More details about this process is given in Appendix A.2.2. The CNNs were trained at only the 30-dB level of noise. Again, our task is to reconstruct the images from the sinograms.

3) Experiment 3. We downsampled the views of the original, (729 × 720) rat-brain sinograms by 5 to obtain sparse-view sinograms of size (729 × 144). For the qth z-slice, we denote the corresponding sparse-view sinograms $\mathbf{y}_{\mathrm{Real}}^q$. Note that, unlike in Experiments 1 and 2, the sinogram was not generated from an image but was obtained experimentally.

3.5.3 Comparison Methods

Given the ground truth $\mathbf{x}$, our figure of merit for the reconstructed $\mathbf{x}^*$ is the regressed SNR given by

$$\mathrm{SNR}(\mathbf{x}^*, \mathbf{x}) = \arg\max_{a,b} \mathrm{SNR}(a\mathbf{x}^* + b, \mathbf{x}), \tag{3.14}$$

where the purpose of a and b is to adjust for contrast and offset. We also evaluate the performance using the structural similarity index (SSIM) [113]. We compare five reconstruction methods.

1) ***FBP****.* FBP is the classical direct inversion of the Radon transform $\mathbf{H}$, here implemented in Matlab by the `iradon` command with the `ram-lak` filter and `linear` interpolation as options.

2) ***Total-Variation*** *Reconstruction.* TV solves

$$\mathbf{x}_{\mathrm{TV}} = \min_{\mathbf{x}} \left(\frac{1}{2} \|\mathbf{H}\mathbf{x} - \mathbf{y}\|_2^2 + \lambda \|\mathbf{x}\|_{\mathrm{TV}} \right) \text{ s.t. } \mathbf{x} \geq 0, \tag{3.15}$$

where

$$\|\mathbf{x}\|_{\mathrm{TV}} = \sum_{i=1}^{N-1} \sum_{j=1}^{N-1} \sqrt{(\mathrm{D}_{\mathrm{h};i,j}(\mathbf{x}))^2 + (\mathrm{D}_{\mathrm{v};i,j}(\mathbf{x}))^2},$$

$\mathrm{D}_{\mathrm{h};i,j}(\mathbf{x}) = [\mathbf{x}]_{i,j+1} - [\mathbf{x}]_{i,j}$, and $\mathrm{D}_{\mathrm{v};i,j}(\mathbf{x}) = [\mathbf{x}]_{i,j+1} - [\mathbf{x}]_{i,j}$. The optimization is carried out via ADMM [13].

*3) **Dictionary Learning (DL)**.* DL [96] solves

$$\mathbf{x}_{\mathrm{DL}} = \arg\min_{\mathbf{x},\boldsymbol{\alpha}} \Bigg(\|\mathbf{H}\mathbf{x}-\mathbf{y}\|^2 + \lambda \sum_{j=1}^{J} \|\mathbf{E}_j\mathbf{x}-\mathbf{D}\boldsymbol{\alpha}_j\|^2 + \lambda\nu_j\|\boldsymbol{\alpha}_j\|_0 \Bigg), \tag{3.16}$$

where $\mathbf{E}_j : \mathbb{R}^{N\times N} \to \mathbb{R}^{L^2}$ extracts and vectorizes the jth patch of size $(L \times L)$ from the image $\mathbf{x}$, $\mathbf{D} \in \mathbb{R}^{L^2\times 256}$ is the dictionary, α_j is the jth column of $\alpha \in \mathbb{R}^{256\times R}$, and $R = (N-L+1)^2$. Note that the patches are extracted with a sliding distance of one pixel.

For a given $\mathbf{y}$, the dictionary $\mathbf{D}$ is learned from the corresponding ground truth using the procedure described in [114]. The objective (3.16) is then solved iteratively by first minimizing it with respect to $\mathbf{x}$ using gradient descent as described in [96] and then with respect to $\boldsymbol{\alpha}$ using orthogonal matching pursuit (OMP) [115]. Since $\mathbf{D}$ is learned from the testing ground truth itself, the performance that we report here is an upper bound to the one that would be achieved by learning it using the training images.

*4) **FBPconv**.* FBPconv [17] is a state-of-the-art deep-learning technique, in which a residual CNN with U-net architecture is trained to directly denoise the FBP . It has been shown to outperform other deep-learning-based direct reconstruction methods for sparse-view CT. In our proposed method, we use a CNN with the same architecture as in FBPconv. As a result, in our framework, FBPconv corresponds to training with only the ensemble in (3.11). In the testing phase, the FBP of the measurements is fed into the trained CNN to output the reconstruction image.

*5) **RPGD**.* RPGD is our proposed method. It is described in Algorithm 1. There the nonlinear operator F is the CNN trained as a projector (as discussed in Section 3.4). For experiments with Poisson noise, we use the slightly modified RPGD described in Appendix A.2.2. For all the experiments, FBP is used for the operator $\mathbf{A}$.

3.5.4 Training and Selection of Parameters

1) Experiment 1. For TV, the regularization parameter λ is selected via a golden-section search over 20 values so as to maximize the SNR of $\mathbf{x}_{\mathrm{TV}}$ with respect to the ground truth. We set the additional penalty parameter inside ADMM (see Equation (2.6) in [13]) equal to λ. The rationale for this heuristic is that it puts the soft-threshold parameter in the same order of magnitude as the image gradients. We set the number of iterations to 100, which was enough to show good empirical convergence.

For DL, the parameters are selected via a parameter sweep, roughly following the approach described in [96, Table 1]. Specifically: The patch size is $L = 8$.

During dictionary learning, the sparsity level is set to 5 and 10. During reconstruction, the sparsity level for OMP is set to 5, 8, 10, 12, 20, and 25, while the tolerance level is taken to be 10, 100, and 1000. This, in effect, is the same as sweeping over ν_j in (3.16). For each of these $2 \times 6 \times 3 = 36$ parameter settings, λ in (3.16) is chosen by a golden-section search over 7 values.

As discussed earlier, the CNNs for both the ×5 and ×16 cases are trained separately for high and low measurement noise.

i) *Training with Noiseless Measurements.* The training of the projector for RPGD follows the sequential procedure described in Section 3.4, with the configurations

- ×5, no noise: $T_1 = 80$, $T_2 = 49$, $T_3 = 5$;
- ×16, no noise: $T_1 = 71$, $T_2 = 41$, $T_3 = 11$.

We use the CNN obtained right after the first stage for FBPconv, since during this stage, only the training ensemble in (3.11) is taken into account. We empirically found that the training error J_2 converged in T_1 epochs of Stage 1, yielding an optimal performance for FBPconv.

ii) *Training with 40-dB Measurement Noise.* This includes replacing the ensemble in (3.11) with $\{\mathbf{A}\mathbf{y}^q\}$ where $\mathbf{y}^q = \mathbf{H}\mathbf{x}^q + \mathbf{n}$, has a 40-dB SNR with respect to $\mathbf{H}\mathbf{x}^q$. With 20% probability, we also perturb the views of the measurements with an AWGN of 0.05 standard deviation so as to enforce robustness to model mismatch. These CNNs are initialized with the ones obtained after the first stage of the noiseless training and are then trained with the configurations

- ×5, 40-dB noise: $T_1 = 35$, $T_2 = 49$, $T_3 = 5$;
- ×16, 40-dB noise: $T_1 = 32$, $T_2 = 41$, $T_3 = 11$.

Similarly to the previous case, the CNNs obtained after the first and the third training stage are used in FBPconv and RPGD, respectively. For clarity, these variants will be referred to as **FBPconv40** and **RPGD40**.

The learning rate is decreased in a geometric progression from 10^{-2} to 10^{-3} in Stage 1 and kept at 10^{-3} for Stages 2 and 3. Recall that the last two stages contain the ensemble with dynamic perturbation (3.12) which changes in every epoch. The lower learning rate, therefore, avoids drastic changes in parameters between the epochs. The batch size is fixed to 2. The other hyper-parameters follow [17]. For stability, gradients above 10^{-2} are clipped and the momentum is set to 0.99. The total training time for the noiseless case is around 21.5 hours on a Titan X GPU (Pascal architecture).

The hyper-parameters for RPGD are chosen as follows: The relaxation parameter α_0 is initialized with 1, the sequence $\{c_k\}$ is set to the constant $C = 0.99$ for RPGD and $C = 0.8$ for RPGD40. For each noise level and views number, the only free parameter γ is swept over 20 values geometrically spaced between 10^{-2} and 10^{-5}. We pick the γ which gives the best average SNR over the 25 test images. Note that, for TV and DL, the value of the optimum λ generally increases as the measurement noise increases; however, no such obvious relation exists for γ. This is mainly because it is the step size of the gradient descent in RPGD and not a regularization parameter. In all experiments, the gradient step is skipped during the first iteration.

On the GPU, one iteration of RPGD takes less than 1 second. The algorithm is stopped when the residual $\|\mathbf{x}_{k+1} - \mathbf{x}_k\|_2$ reaches a value less than 1, which is sufficiently small compared to the dynamic range $[0, 350]$ of the image. It takes around 1-2 minutes to reconstruct an image with RPGD.

2) *Experiment 2.* For this case the CNNs are trained similarly to the CNN for RPGD40 in Experiment 1. Perturbations (3.10)-(3.12) are used with the replacement of $\mathbf{AHx}^q_{\text{FD}}$ in (3.11) by $\mathbf{Ay}^q$, where $\mathbf{y}^q$ had 30 dB Poisson noise. The $\mathbf{x}^q_{\text{FD}}$ and $\mathbf{Ay}^q_{\text{Real}}$ are multiplied with a constant so that their maximum pixel value is 480.

The CNN obtained after the first stage is used as FBPconv.

While testing, we keep $C = 0.4$. Other training hyper-parameters and testing parameters of the RPGD are kept the same as the RPGD40 for ×5 case in Experiment 1.

3) Experiment 3. The CNNs are trained using the perturbations (3.10)-(3.12) with two modifications: (i) $\mathbf{x}^q$ is replaced with $\mathbf{x}^q_{\text{FD}}$ because the actual ground truth

Table 3.2: Reconstruction results for Experiment 1 with low measurement noise (Gaussian). Gray cells indicate that the method was tuned/trained for the corresponding noise level.

Case	Measurement SNR (dB)	Quality Index	Method				
			FBP	TV	DL	FBPconv	RPGD
×16	∞	SNR	12.74	24.21	23.11	26.19	**27.02**
		SSIM	0.178	0.277	0.231	0.323	**0.374**
	70	SNR	12.73	24.20	23.10	26.18	**26.94**
		SSIM	0.178	0.277	0.231	0.324	**0.325**
×5	∞	SNR	24.19	30.80	29.36	32.09	**32.62**
		SSIM	0.434	0.511	0.424	0.480	**0.554**
	70	SNR	24.15	30.74	29.24	32.08	**32.56**
		SSIM	0.432	0.507	0.422	0.483	**0.553**

was unavailable; and (ii) $\mathbf{AHx}^q$ in (3.11) is replaced with $\mathbf{Ay}^q_{\mathrm{Real}}$ because we have now access to the actual sinogram.

All other training hyper-parameters and testing parameters are kept the same as RPGD for the ×5 case in Experiment 1. Similar to Experiment 1, the CNN obtained after the first stage of the sequential training is used as the FBPconv.

3.6 Results and Discussions

3.6.1 Experiment 1

We report in Tables 3.2 and 3.3 the results for low and high measurement noise, respectively. FBPconv and RPGD are used for low noise, while FBPconv40 and RPGD40 are used for high noise. The reconstruction SNRs and SSIMs are averaged over the 25 test images. The gray cells indicate that the method was optimized for that level of noise. As discussed earlier, adjusting λ for TV and DL indirectly implies tuning for the measurement noise; therefore, all of the cells in these columns are gray. This is different for the learning methods, where tuning for the measurement noise requires retraining.

1) Low Measurement Noise. In the low-noise cases (Table 3.2), the proposed

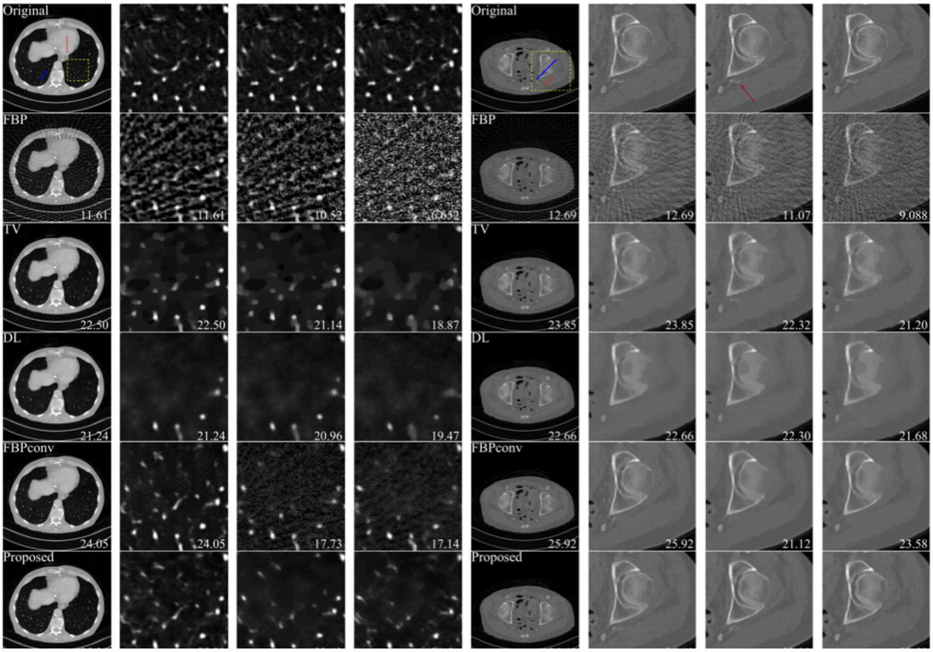

∞-dB zoom(∞-dB) 45-dB 35-dB ∞-dB zoom(∞-dB) 45-dB 40-dB

Figure 3.2: Comparison of reconstructions using different methods for the ×16 case in Experiment 1. First column: reconstruction from noiseless measurements of a lung image. Second column: zoomed version of the area marked by the box in the original in the first column. Third and fourth columns: zoomed version for the case of 45 and 35 dB, respectively. Fifth to eighth columns: corresponding results for an abdomen image. Seventh and eighth column correspond to 45 and 40 dB, respectively.

Table 3.3: Reconstruction results for Experiment 1 with high measurement noise (Gaussian). Gray cells indicate that the method was tuned/trained for the corresponding noise level.

Case	Measurement SNR (dB)	Quality Index	Method FBP	TV	DL	FBPconv40	RPGD40
×16	40+5	SNR	11.08	22.59	22.74	20.87	**24.16**
		SSIM	0.127	0.238	0.222	0.161	**0.262**
	40	SNR	9.09	21.40	22.13	23.26	**23.73**
		SSIM	0.096	0.210	0.209	0.205	**0.252**
	40-5	SNR	6.51	20.01	20.93	16.20	**22.59**
		SSIM	0.066	0.179	0.187	0.128	**0.221**
×5	40+5	SNR	18.85	27.18	**27.82**	22.56	27.17
		SSIM	0.241	0.367	0.364	0.201	**0.384**
	40	SNR	14.96	25.46	26.26	**28.24**	27.61
		SSIM	0.167	0.314	0.315	0.324	**0.361**
	40-5	SNR	10.76	23.44	22.24	18.90	**24.58**
		SSIM	0.110	0.261	0.263	0.193	**0.300**

RPGD method outperforms all the others for both ×5 and ×16 reductions in terms of SNR and SSIM indices. FBP performs the worst but is able to retain enough information to be utilized by FBPConv and RPGD. Due to the convexity of the iterative scheme, TV is able to perform well but tends to smooth textures and edges. DL performs worse than TV for ×16 case but is equivalent to it for ×5 case. On one hand, FBPConv outperforms both TV and DL. but it is surpassed by RPGD. This is mainly due to the feedback mechanism in RPGD which lets RPGD use the information in the given measurements to increase the quality of the reconstruction. In fact, for the ×16, no noise, case, the SNRs of the sinogram of the reconstructed images for TV, FBPconv, and RPGD are around 47 dB, 57 dB, and 62 dB, respectively. This means that reconstruction using RPGD has both better image quality and more reliability since it is consistent with the given noiseless measurement.

2) High Measurement Noise. In the noisier cases (Table 3.3), RPGD40 yields a better SNR than other methods in the low-view cases (×16) and is more consistent in performance than the others in the high-view (×5) cases. In terms of the SSIM index, it outperforms all of them. The performance of DL and TV are robust to the noise level with DL performing better than others in terms of SNR for the 45-dB, ×5, case. FBPconv40 substantially outperforms DL and TV in the two scenarios with 40-dB noise measurement, over which it was actually trained. For this noise level and ×5 case, it even performs slightly better than RPGD40 but only in terms of SNR. However, as the level of noise deviates from 40 dB, the performance of FBPconv40 degrades significantly. Surprisingly, its performances in the 45-dB cases are much worse than those in the corresponding 40-dB cases. In fact, its SSIM index for the 45-dB, ×5, case is even worse than FBP. This implies that FBPConv40 is highly sensitive to the difference between the training and testing conditions. By contrast, RPGD40 is more robust to this difference due to its iterative correction. In the ×16 case with 45-dB and 35-dB noise level, it outperforms FBPconv40 by around 3.5 dB and 6 dB, respectively.

3) Case Study The reconstructions of lung and abdomen images for the case of ×16 downsampling and noiseless measurements are illustrated in Figure 3.2 (first and fifth columns). FBP is dominated by line artifacts, while TV and DL satisfactorily removes those but blurs the fine structures. FBPConv and RPGD are able to reconstruct these details. The zoomed version (second and sixth columns) suggests that RPGD is able to reconstruct the fine details better than the other methods. This observation remains the same when the measurement quality degrades. The remaining columns, contain the reconstructions for different noise levels. For the abdomen image it is noticeable that only TV is able to retain the small bone structure marked by an arrow in the zoomed version of the lung image (seventh column). Possible reason for this could be that the structure similar to this were rare in the training set. Increasing the training data size with suitable images could be a solution.

Figure 3.3 contains the profiles of high- and low-contrast regions of the reconstructions for the two images. These regions are marked by line segments inside the original image in the first column of Figure 3.2. The FBP profile is highly noisy and the TV and DL profiles overly smooth the details. FBPconv40 is able to accommodate the sudden transitions in the high-contrast case. RPGD40 is slightly better in this regard. For the low-contrast case, RPGD40 is able to follow the structures of the original (GT) profile better than the others. A similar analysis holds for the ×5 case (see Figure A.3 in the Appendix).

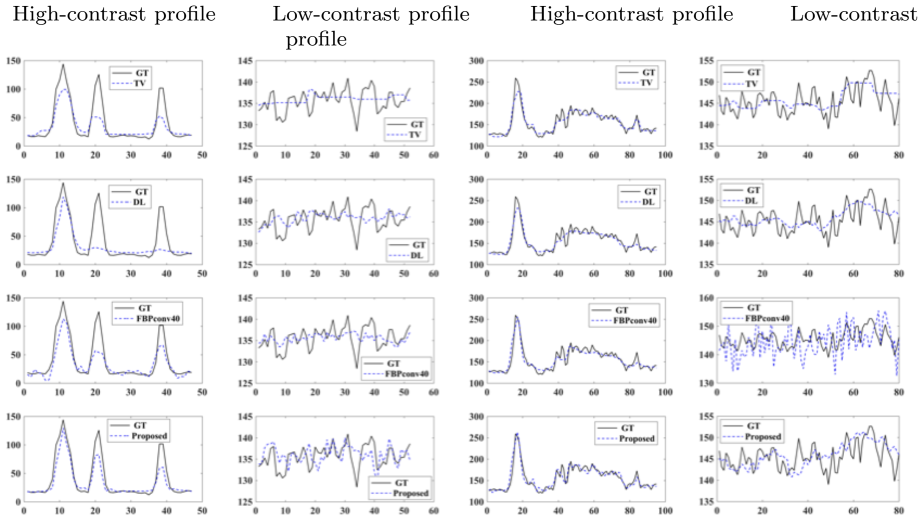

High-contrast profile

Low-contrast profile profile

High-contrast profile

Low-contrast

Figure 3.3: Profile of the high- and low-contrast regions marked in the first and fifth columns of Figure 3.2 by solid and dashed line segments, respectively. First and second columns: ×16, 45-dB noise case for the lung image. Third and fourth columns: ×16, 40-dB noise case for the abdomen image.

3.6.2 Experiment 2

We show in Table 3.4 the regressed SNR and SSIM indices averaged over the 25 reconstructed slices. RPGD outperforms both FBP and FBPconv in terms of SNR and SSIM. Similar to the Experiment 1, its performance is also more robust with respect to noise mismatch. Fig. A.4 in the Appendix compares the reconstructions for a given test slice.

Table 3.4: Reconstruction results for Experiment 2 with Poisson noise and ×5 views reduction. Grey cell indicate that the method was trained for the corresponding noise level.

Measurement SNR (dB)	Quality Index	Method		
		FBP	FBPconv	RPGD
30-5	SNR	4.61	7.45	**8.21**
	SSIM	0.112	0.134	**0.154**
30	SNR	5.96	9.18	**9.22**
	SSIM	0.200	0.174	**0.246**
30+5	SNR	7.75	9.50	**9.75**
	SSIM	0.305	0.132	**0.332**

3.6.3 Experiment 3

In Figure 3.4, we show the reconstruction result for one slice for $\gamma = 10^{-5}$. Since the ground truth is unavailable, we show the reconstructions without a quantitative comparison. It can be seen that the proposed method is able to reconstruct images with reasonable perceptual quality.

3.7 Behavior of Algorithms

We now explore the behavior of the proposed method in more details, including its empirical convergence and the effect of sequential training.

3.7.1 Convergence of RPGD

In Figure 3.5, we show the behavior of RPGD with respect to the iteration number k for Experiment 1. The evolution of the SNR of images $\mathbf{x}_k$ and their measurements $\mathbf{H}\mathbf{x}_k$ computed with respect to the ground truth image and the ground-truth measurement are shown in Figures 3.5 (a) and (b), respectively. We give α_k with respect to the iteration k in Figure 3.5 (c). The results are averaged over 25 test images for ×16, no noise, case and $C = 0.99$. RPGD outperforms all the other methods in the context of both image quality and measurement consistency.

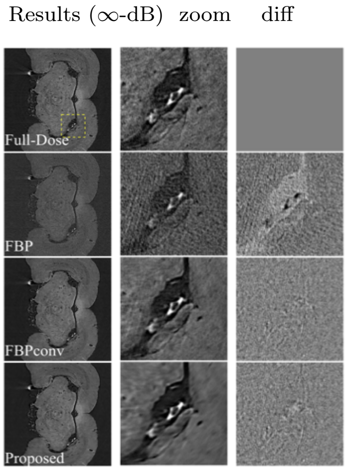

Figure 3.4: Reconstruction results for a test slice in Experiment 3. Full-dose image is obtained by taking FBP of the full-view sinogram. The rest of the reconstructions are obtained from the sparse-view (×5) sinogram. The last column shows the difference between the reconstruction and the full-dose image.

Due to the high value of the step size ($\gamma = 2 \times 10^{-3}$) and the large difference ($\mathbf{H}\mathbf{x}_k - \mathbf{y}$), the initial few iterations have large gradients and result in the instability of the algorithm. The reason is that the CNN is fed with ($\mathbf{x}_k - \gamma \mathbf{H}^{\mathrm{T}}(\mathbf{H}\mathbf{x}_k - \mathbf{y})$), which is drastically different from the perturbations on which it was trained. In this situation, α_k decreases steeply and stabilizes the algorithm. At convergence, $\alpha_k \neq 0$; therefore, according to Theorem 3.3.1, $\mathbf{x}_{100}$ is the fixed point of (3.7) where $F = \mathit{CNN}$.

3.7.2 Advantages of Sequential Training

Here, we experimentally verify the advantages of the sequential-training strategy discussed in Section 3.5. Using the setup of Experiment 1, we compare the training time and performance of the CNNs trained with and without this strategy for the ×16 downsampling and no noise case. For the gold standard (systematic training of CNN), we train a CNN as a projector with the 3 types of perturbation in every epoch. We use 135 epochs for training which is roughly equal to $\{T_1 + T_2 + T_3\}$ used during training for the corresponding sequential-training-based CNN. This number was sufficient for the convergence of the training error. The reconstruction performance of RPGD using this gold standard CNN is 26.86 dB, compared to 27.02 dB for RPGD using the sequentially trained CNN. The total training times are 48 and 22 hours, respectively. This demonstrates that the sequential strategy reduces the training time (in this case more than 50%), while preserving (or even slightly increasing) the reconstruction performance.

3.8 Summary

In this chapter, we present a new image reconstruction method that replaces the projector in a projected gradient descent (PGD) with a convolutional neural network (CNN). Recently, CNNs trained as image-to-image regressors have successfully been used to solve inverse problems in imaging. However, unlike existing iterative image reconstruction algorithms, these CNN-based approaches usually lack a feedback mechanism to enforce that the reconstructed image is consistent with the measurements. We propose a relaxed version of PGD wherein gradient descent enforces measurement consistency, while a CNN recursively projects the solution closer to the space of desired reconstruction images. We show that this algorithm is

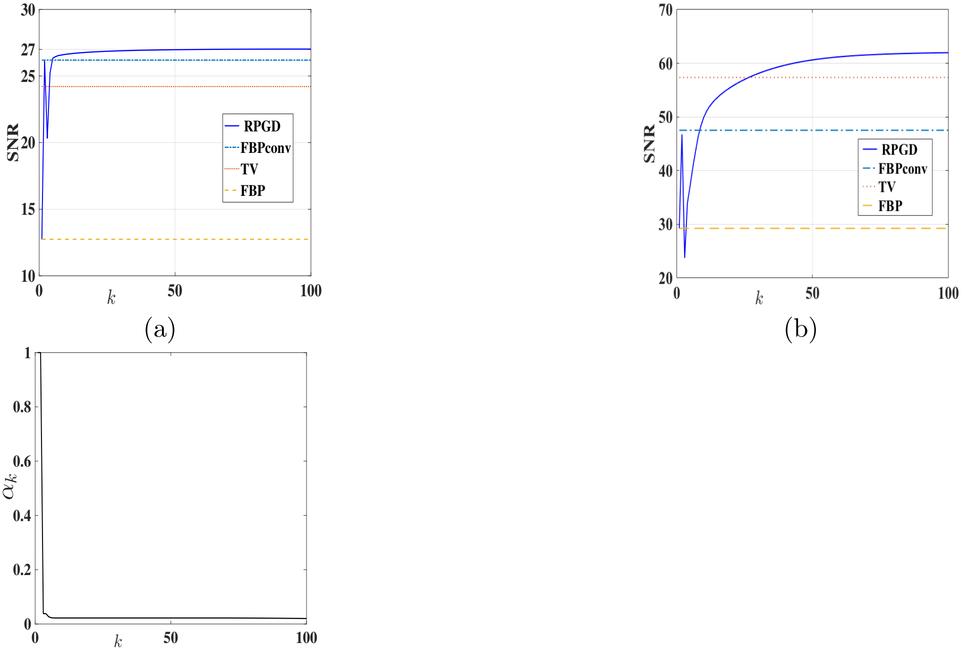

(a)

(b)

(c)

Figure 3.5: Convergence with iteration k of RPGD for the Experiment 1, ×16, no-noise case when $C = 0.99$. Results are averaged over 25 test images. (a) SNRs of $\mathbf{x}_k$ with respect to the ground-truth image. (b) SNRs of $\mathbf{H}\mathbf{x}_k$ with respect to the ground-truth sinogram. (c) Evolution of the relaxation parameters α_k. In (a) and (b), the FBP, FBPconv, and TV results are independent of the RPGD iteration k but have been shown for the sake of comparison.

guaranteed to converge and, under certain conditions, converges to a local minimum of a non-convex inverse problem. Finally, we propose a simple scheme to train the CNN to act like a projector. Our experiments on sparse-view computed-tomography reconstruction show an improvement over total variation-based regularization, dictionary learning, and a state-of-the-art deep learning-based direct reconstruction technique.

Part III

Third Generation

Chapter 4

Time-Dependent Deep Image Prior

[1]Although widely popular, classical and supervised deep-learning-based methods discussed this far have limitations. The quality of the reconstruction of the former is limited and the latter requires training data. Recently, a new unsupervised deep-learning method has emerged which are based on deep-image-prior. These methods requires no additional training nor training data, while their reconstruction quality surpasses the classical methods and is comparable to that of the supervised deep-learning approaches. In this chapter, we extend the deep image prior methods for time dependent inverse problems. Although the chapter targets the application of dynamic MRI, the method is applicable to other modalities as well.

4.1 Overview

There are currently three main approaches to accelerate the magnetic resonance imaging (MRI) of a static image. All three methods rely on a partial sampling of the k-space to reduce the acquisition time. The resulting partial loss of data must then be compensated to maintain the quality of the image. Once compensation is

---

[1] The content of this chapter is based on our work [116].

achieved, the accelerated methods capture accurate motions of fast moving organs such as the heart.

i. In parallel MRI (pMRI), the simultaneous use of several hardware coils results in spatial redundancy that enables algorithms to reconstruct clean images [117,118].

ii. In compressed sensing (CS) MRI, the data are assumed to be sparse in certain transform domains [119,120]. This ultimately leads to regularized computational methods that compensate for the partial sampling. Their success suggests that, in particular, a Fourier-based forward model matches well the assumption of sparsity.

iii. In the context of trainable deep artificial neural networks, learning approaches have already achieved fast and accurate reconstructions of partially sampled MRI data [15,121]. Similarly to CS MRI, dynamic accelerated reconstructions have also been proposed in the literature [122–124], possibly in combination with pMRI in the learning loop [24]. These approaches depend on training datasets [17,20,22,84,125,126].

In the context of dynamic MRI, the approach that consists in the acquisition of a sequence of frames is suboptimal. Instead, it is more efficient to take advantage of the time dependencies between frames to gain additional improvements in terms of temporal resolution [127,128]. For instance, [129,130] design non-overlapping sampling masks at each frame to restore a dynamic volume—a method that demands far fewer k-space samples than would a sequence of static MRI. Indeed, the CS theory explains the possibility of perfect reconstruction despite a sub-Nyquist sampling [131]. The capture of temporal redundancies has also been handled through low-dimensional manifolds [132,133]. In the specific case of cardiac applications, the overall motion of the heart is expected to be approximately cyclic. This periodicity can be exploited to bring additional gains in terms of temporal resolution, but the length and phase of the cycles must be determined first. This is usually achieved either through electrocardiograms (ECG) or self-gating [134,135]. Under the restrictive assumption of ideal periodicity, these methods allow one to prefix the cardiac phases and to reorder temporally the acquired frames, effectively producing a stroboscope-inspired analysis of the motion. Motion irregularities are not captured by those methods.

4.1.1 Contribution

In this chapter, we propose an unsupervised learning framework in which a neural network reconstructs fast dynamic MRI from golden-angle radial lines in k-space, also called spokes. To reconstruct a single image, we feed one realization of low-dimensional latent variables to a artificial neural network. A nonuniform fast Fourier transform[2] (NuFFT) is then applied to the output of the neural network and simulates the MRI measurement process [136]. Inspired by deep image priors [27], we fit the simulated measurements to the real measurements. The fit is controlled by adjusting the weights of the neural network until the Euclidean loss between the simulated and real measurements is minimized; this fitting process is referred to as the learning stage.

In the context of dynamic MRI, we extend the fitting process in such a way that the weights of the network are learned by taking simultaneously into consideration the joint collection of all acquisitions, which yields time-independent weights. Time dependence is recovered by controlling the latent variables. Given some temporal interval, we synthesize two independent random realizations of low-dimensional latent variables and associate one to the initial and one to the final bound of the interval. Timestamped contributions to learning are obtained by taking as ground-truth the real measurements acquired over a quasi-instantaneous observation period (say, five spokes), while we let the activation of the neural network be the intermediate realization of the latent variables obtained by linear interpolation of the two latent endpoints. This approach allows us to impose and exploit temporal dependencies in the latent space.

In short, the action of our neural network is to map the manifold of latent variables onto a manifold of dynamic images. Importantly, our approach is purely unsupervised; moreover, priors are imposed only indirectly, arising from the mere structure of a convolutional network. We demonstrate the performance of our neural network by comparing its reconstructions to those obtained from CS algorithms [129, 130].

4.1.2 Related Works

Deep image priors have been introduced in [27]. They capitalize on the structure of convolutional artificial neural networks to adjust priors to the data, thus avoiding

---

[2] https://github.com/marchdf/python-nufft

the limitations and pitfalls of hand-crafted priors. They have been deployed in [137] to build an unsupervised learning scheme for accelerated MRI; but, contrarily to ours, the task addressed therein is static. Other researchers have used deep image priors to reconstruct positron emission tomography images, albeit again in a non-dynamic fashion [138].

4.2 Methods

Let $\mathcal{R}$ be the Radon transform of the complex-valued continuously defined image $x : \mathbb{R}^2 \to \mathbb{C}$, so that

$$\mathcal{R}\{x\}(r, \vartheta) = \int_{\mathbb{R}^2} x(\boldsymbol{\xi})\, \delta(\mathbf{u}_\vartheta^{\mathsf{T}}\, \boldsymbol{\xi} - r)\, \mathrm{d}\xi, \tag{4.1}$$

where r and ϑ are the spatial and angular Radon arguments, respectively, and where $\mathbf{u}_\vartheta$ is a unit vector in the ϑ direction. Moreover, let $\mathcal{F}$ denote the one-dimensional continuous Fourier transform that follows the convention

$$\mathcal{F}\{x\}(\omega) = \int_{\mathbb{R}} x(r)\, \mathrm{e}^{-\mathrm{j}\,\omega\, r}\, \mathrm{d}r \tag{4.2}$$

at spatial pulsation ω, otherwise known as a coordinate in k-space. Then, our conceptual model of non-Cartesian MRI is the concatenated linear transform

$$\mathcal{H}_\vartheta\{x\}(\omega) = \mathcal{F}\{\mathcal{R}\{x\}(\cdot, \vartheta)\}(\omega) \tag{4.3}$$

which maps a two-dimensional static image onto its measurements at continuously defined direction ϑ and pulsation ω. Mathematically, $\mathcal{H}_\vartheta\{x\}(\omega)$ is invertible because so are $\mathcal{F}$ and $\mathcal{R}$. Consequently, provided we know $\mathcal{H}_\vartheta\{x\}(\omega)$ for every direction ϑ and every pulsation ω, we can in principle recover the value of $x(\boldsymbol{\xi})$ at every spatial argument $\boldsymbol{\xi}$.

4.2.1 Static Discretization

Unfortunately, $\mathcal{H}_\vartheta\{x\}(\omega)$ can be known in practice only at finitely many discrete directions and pulsations. This discretization, which amounts to modeling MRI by a NuFFT, makes the discrete transform an unfaithful surrogate of our conceptual

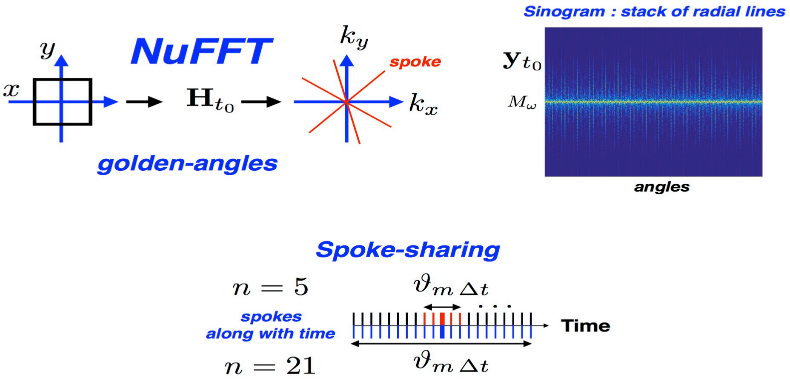

Figure 4.1: Nonuniform fast Fourier transform with golden-angle scheme and spoke sharing.

model, in particular because of aliasing concerns, and also because the discrete version is no more invertible.

Formally, let $\mathbf{x} \in \mathbb{C}^N$ be a vectorized version of the samples of x seen as an image of finite size $(N_1 \times N_2)$, with $N = N_1 \, N_2$. Likewise, let $\mathbf{y} \in \mathbb{C}^M$ be a vectorized version of the samples of the sinogram $\mathcal{H}_\vartheta\{x\}(\omega)$, with measurements of finite size $(M_\vartheta \times M_\omega)$ taken over M_ϑ orientations and M_ω elements of the k-space, with $M = M_\vartheta \, M_\omega$. Then, by linearity of the transformation we write that

$$\mathbf{y} = \mathbf{G} \, \mathbf{x}, \tag{4.4}$$

where $\mathbf{G}$ is an M-rows by N-columns matrix that combines discrete Fourier and discrete Radon transforms.

4.2.2 Spoke Sharing

In dynamic MRI, it is acknowledged that the image $\mathbf{x}$ changes continuously through time. We assume however that the measurements of a spoke (a radial line in k-space, as suggested in Figure 4.1) are instantaneous and indexed at discrete times $t_0 \in \mathbb{Z}\,\Delta t$, taken regularly at temporal interval Δt. The spoke orientations follow the golden-angle strategy

$$\vartheta_t = \vartheta_0 + \omega_0\, t, \tag{4.5}$$

where ϑ_t gives the orientation of a spoke at continuous time $t \in \mathbb{R}$, with ω_0 its angular velocity. The golden-angle specificity is the irrationality condition $(\omega_0\, \Delta t/\pi) \notin \mathbb{Q}$, which is approximated by setting $(\omega_0\, \Delta t) \approx 111.25°$ [129]. We finally denote an image frame at time t as x_t and its spatially discretized version of length N as $\mathbf{x}_t$. Its associated sinogram is s_t. As it turns out, however, it is only natural to set $M_\vartheta = 1$ for the discretization of s_t because of our assumption of ideal temporal sampling. Then, $\mathbf{s}_t$ is a direct representation of a spoke and has length $M = M_\omega$. Its dependence on $\mathbf{x}_t$ is encoded in the time-dependent M_ω-rows by N-columns system matrix $\mathbf{G}_t$, with $\mathbf{s}_t = \mathbf{G}_t\, \mathbf{x}_t$.

In accelerated dynamic MRI, one acquires $\mathbf{s}_{t_0}$ for $t_0 \in \mathbb{Z}\,\Delta t$ and wants to reconstruct $\mathbf{x}_{t_0}$ or, possibly, $\mathbf{x}_t$ for $t \in \mathbb{R}$. Clearly, however, a single orientation does not provide sufficient data for the the recovery of the two-dimensional $\mathbf{x}_{t_0}$. To overcome this issue, we assume next that the changes are slow over some small n, so that $\mathbf{x}_t \approx \mathbf{x}_{t_0}$ for all t in the half-open interval $\mathbb{T}_{t_0} = [t_0 - n\,\Delta t/2, t_0 + n\,\Delta t/2)$. In practice, the odd $n \in 2\,\mathbb{N} + 1$ corresponds to the number of radial lines used for reconstruction and is related to the temporal resolution of dynamic imaging. We then collect n neighboring spokes[3] and concatenate them in the vector $\mathbf{y}_{t_0} = (\mathbf{s}_{m\,\Delta t})_{m \in \mathbb{Z} \cap (\mathbb{T}_{t_0}/\Delta t)}$ of length $(n\, M_\omega)$. Through this mechanism, there are spokes $\mathbf{s}_{t_0}$ that are shared between, say, $\mathbf{y}_{t_0}$ and $\mathbf{y}_{t_0+\Delta t}$.

The dependence of $\mathbf{y}_{t_0}$ on $\mathbf{x}_{t_0}$ is encoded now in the time-dependent $(n\, M_\omega)$-rows by N-columns matrix $\mathbf{H}_{t_0} = (\mathbf{G}_{m\,\Delta t})_{m \in \mathbb{Z} \cap (\mathbb{T}_{t_0}/\Delta t)}$, so that (4.4) becomes time-dependent by writing that

$$\mathbf{y}_{t_0} = \mathbf{H}_{t_0}\, \mathbf{x}_{t_0}. \tag{4.6}$$

Because of the irrationality condition of the golden-angle approach, no direction will ever be measured twice and $\mathbf{H}_{t_0}$ acquires effective time dependence.

---

[3] In the practical case of finite-time acquisitions, it may happen that no $\mathbf{s}_{t_0}$ is available for some $t_0 \in \mathbb{T}_0$. In this case, which happens near the beginning and the end of the temporal range, we use one-sided nearest spokes.

4.2.3 Regularization

Even with $n > 1$, it is observed that $(n\, M_\omega) \ll N$, which makes severely ill-posed the recovery of $\mathbf{x}_{t_0}$ given $\mathbf{y}_{t_0}$. To truly resolve this issue, practitioners often choose to regularize the problem over some extended temporal range. From a notational perspective, K vectors $\mathbf{y}_{t_0}$ are concatenated over a large duration $(K\, \Delta t)$ to build $\mathbf{Y} = (\mathbf{y}_{k\,\Delta t})_{k \in [0 \ldots K-1]}$. Likewise, we write that $\mathbf{X} = (\mathbf{x}_{k\,\Delta t})_{k \in [0 \ldots K-1]}$ and $\mathbf{H} = [\mathbf{H}_{k\,\Delta t}]_{k \in [0 \ldots K-1]}$. The length of $\mathbf{Y}$ and $\mathbf{X}$ are $(K\, n\, M_\omega)$ and $(K\, N)$, respectively. The size of $\mathbf{H}$ ensues.

In the context of CS dynamic imaging, the traditional regularization of the forward model (4.6) is established as a search for the solution

$$\mathbf{X}^* = \arg\min_{\mathbf{X}} \left(\|\mathbf{H}\,\mathbf{X} - \mathbf{Y}\|_2^2 + \lambda\, \|\mathbf{D}\,\mathbf{X}\|_p \right), \tag{4.7}$$

where $\mathbf{D}$ is a sparsifying transform along the temporal domain. Typically, this transform is a finite-difference operator, used as surrogate of the conceptual first-order derivative $\mathcal{D}$. The corresponding regularization term encourages temporal dependency and counterbalances the ill-posedness of (4.6).

By contrast, there exists no explicit regularizer in the context of dynamic MRI reconstruction by traditional neural networks [122, 139]. In return, image priors are data-driven, imposed by the supervised learning of ground-truth pairs $(\mathbf{Y}, \mathbf{X})$. Letting $f_\phi : \mathbb{C}^M \to \mathbb{C}^N$ represent the function that the network implements, where ϕ gives network parameters such as weights and biases, learning will return the solution

$$\phi^* = \arg\min_{\phi} \mathbb{E}_{\{(\mathbf{Y},\mathbf{X})\}} \left\| f_\phi(\mathbf{H}^{\mathsf{H}}\, \mathbf{Y}) - \mathbf{X} \right\|_2^2. \tag{4.8}$$

There, $\mathbf{H}^{\mathsf{H}}$ is the Hermitian transpose of $\mathbf{H}$. With some abuse of notation, the learning process is represented by the expectation operator $\mathbb{E}$. The set $\{(\mathbf{Y}, \mathbf{X})\}$ provides the learning data. After learning has converged to some ϕ^*, reconstructed images are obtained as $f_{\phi^*}(\mathbf{H}^{\mathsf{H}}\, \mathbf{Y})$.

4.2.4 Deep Image Prior with Interpolated Latent Variables

Supervised learning cannot be performed in the absence of ground truth; unsupervised learning methods must be used instead. To that effect, neural networks associated with deep image priors have been proposed in [140], while a cost function that is appropriate to unsupervised learning has been developed in [27]. In

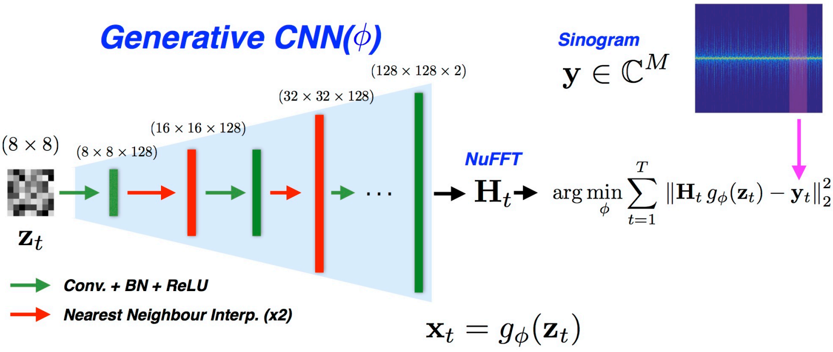

Figure 4.2: Flowchart of our framework. Conv.: convolutional layers, BN: batch normalization. The details are described in Table 4.1.

this chapter, we propose an extension whereby we address the needs of dynamic MRI by performing temporal interpolation of the latent variables. The sketch of our approach is provided in Figure 4.2.

Inpainting and Denoising

As introduction, we address first the static applications of inpainting and denoising. There, the purpose of deep-image-prior algorithms is to reconstruct a clean signal $\mathbf{x}^* = f_{\theta^*}(\mathbf{z})$ given the perturbed measurement $\mathbf{x}$. The neural network with optimal parameter

$$\theta^* = \arg\min_{\theta} \|A(f_\theta(\mathbf{z})) - \mathbf{x}\|_2^2 \tag{4.9}$$

minimizes a data-fidelity term characterized by the forward model A, which is assumed to be known. The Z-dimensional latent variable $\mathbf{z} \in \mathbb{R}^Z$ is maintained at a fixed value during the whole optimization process. Often, the minimizer θ^* is calculated using some stochastic gradient-descent method with random initial parameter [141]. In summary, deep-image-prior methods achieve strong priors from the structure of a generator network and capture advanced image statistics in a purely unsupervised fashion.

Accelerated Dynamic MRI

Our main contribution in this chapter is to use interpolated latent variables as inputs of a neural network. We start by obtaining two realizations of random discrete images; these realizations are kept for the entire duration of the procedure. The first realization $\mathbf{z}_0$ takes the role of the latent variable associated to the beginning of the dynamic MRI sequence. The second realization is denoted by $\mathbf{z}_{(K-1)\,\Delta t}$ and is associated to the end of the sequence. Intermediate realizations are built like shown in Figure 4.3 and are also used as latent variables.

Choosing the dimension Z of the latent variables to be small, we conjecture that the linearly interpolated latent variables span a low-dimensional, smooth manifold. In other words, we interpret our proposal as a way to impose data-driven temporal dependency in the latent space. A convolutional neural network will then transfer this low-dimensional manifold to some corresponding low-dimensional manifold in the MRI space, the central tenet of this work being that the frames from dynamic imaging do span a low-dimensional manifold, too.

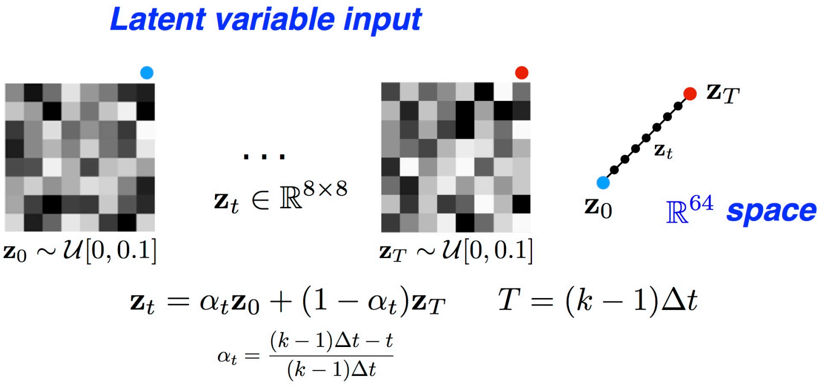

Figure 4.3: Interpolated latent variables.

For a single coil, our deep prior minimizes an Euclidean loss and results in the solution

$$\phi^* = \arg\min_{\phi} \sum_{k=0}^{K-1} \|\mathbf{H}_{k\,\Delta t}\, g_\phi(\mathbf{z}_{k\,\Delta t}) - \mathbf{y}_{k\,\Delta t}\|_2^2\,, \tag{4.10}$$

where $g_\phi : \mathbb{C}^Z \to \mathbb{C}^N$ represents the deep neural network of parameter ϕ. For C coils in pMRI, we establish the solution

$$\phi^* = \arg\min_{\phi} \sum_{c=1}^{C} \sum_{k=0}^{K-1} \|\mathbf{H}_{k\,\Delta t}\, \left(\mathbf{C}_c \odot g_\phi(\mathbf{z}_{k\,\Delta t})\right) - \mathbf{y}_{c,k\,\Delta t}\|_2^2\,, \tag{4.11}$$

where $\mathbf{C}_c$ gives the sensitivity map of the cth coil, $\odot$ is a pixel-wise multiplication operator in the spatial domain which relates true magnetization image to coil sensitivities, and $\mathbf{y}_{c,t}$ concatenates n instantaneous acquisitions of spokes for the cth coil. Once an optimal ϕ^* has been found in either (4.10) or (4.11), we can produce the final estimate $\mathbf{x}_t^*$ for all values of t, including for $t \notin \mathbb{Z}\,\Delta t$ if desired, as

$$\mathbf{x}_t^* = g_{\phi^*}(\mathbf{z}_t). \tag{4.12}$$

4.2.5 Architectures, Datasets, and Training

Architectures

We design our neural network as shown in Table 4.1. The neural network consists of convolutional layers, batch normalization layers, ReLU, and nearest-neighbor interpolations. We apply zero-padding before convolution to let the size of the output mirror that of the input. At the last layer, ReLU is not used. The input of the network is a small-size random variable generated from the uniform distribution $\mathcal{U} \sim (0, 0.1)$, as explained in Section 4.2.4. The output has two channels because MRI images take complex values.

Datasets

All experimental datasets are breath-hold. We use golden-angle radial sparse parallel (GRASP) MRI as a common baseline [129]. Spoke-sharing is not applied for

Table 4.1: Architecture of our convolutional network. Conv.: convolution; BN: batch normalization; NN interp.: nearest-neighbor interpolation.

Operation Layer	Number of Filters	Size of Each Filter (XYC)	Strides (XY)	Zero Padding (XY)	Size of Output Image (XYC)
Input					8×8
Conv+BN+ReLU	128	$3 \times 3 \times 1$	1×1	1×1	$8 \times 8 \times 128$
Conv+BN+ReLU	128	$3 \times 3 \times 128$	1×1	1×1	$8 \times 8 \times 128$
NN interp.			2×2		$16 \times 16 \times 128$
2×(Conv+BN+ReLU)	128	$3 \times 3 \times 128$	1×1	1×1	$16 \times 16 \times 128$
NN interp.			2×2		$32 \times 32 \times 128$
2×(Conv+BN+ReLU)	128	$3 \times 3 \times 128$	1×1	1×1	$32 \times 32 \times 128$
NN interp.			2×2		$64 \times 64 \times 128$
2×(Conv+BN+ReLU)	128	$3 \times 3 \times 128$	1×1	1×1	$64 \times 64 \times 128$
NN interp.			2×2		$128 \times 128 \times 128$
2×(Conv+BN+ReLU)	128	$3 \times 3 \times 128$	1×1	1×1	$128 \times 128 \times 128$
Conv.	2	$3 \times 3 \times 128$	1×1	1×1	$128 \times 128 \times 2$

Algorithm 2 Time-Dependent Deep-Image-Prior

```
Input: Number of iterations $n_{\text{iter}}$, Batch size $B$, Measurement sequence $\{\mathbf{y}_k\}_{k=0}^{K-1}$, Untrained neural network $g_\phi$, Latent variable sequence $\{\mathbf{z}_k\}_{k=0}^{K-1}$.
Output: Reconstructed image sequence $\{\mathbf{x}_k^*\}_{k=0}^{K-1}\}$
for $n_{\text{iter}}$ do
    • Sample a batch $\{b_0, \dots, b_{B-1}\}$ of size $B$ from $\{0, \dots, K-1\}$.
    • Compute batch Loss
          $L_B(\phi) = \frac{1}{B}\sum_{k=0}^{B-1} \|\mathbf{y}_{b_k} - \mathbf{H}_{b_k} g_\phi(\mathbf{z}_{b_k})\|^2.$
    • Update $\phi$ with gradient $\nabla_\theta L_B(\phi)$.
end for
return Reconstruct image series $\{\mathbf{x}_k^* = g_{\phi^*}(\mathbf{z}_k)\}_{k=0}^{K-1}$.
```

GRASP. We assume two-fold upsampling of measurements for every dataset. Therefore, the size of the reconstructed fields of view is half that of the first dimension of the measurements.

- **Retrospective Simulation** A cardiac cine data set was acquired using a 3T whole-body MRI scanner (Siemens; Tim Trio) equipped with a 32-element cardiac coil array. The acquisition sequence was bSSFP and prospective cardiac gating was used. The imaging parameters were as follows: FOV=$300 \times 300\,\text{mm}^2$, acquisition matrix size=128×128, TE/TR=1.37/2.7 ms, receiver bandwidth=1184 Hz/pixel, and flip angle=40°. The number of cardiac phases was 23 and the temporal resolution was 43.2 ms. We then used NuFFT to implement the forward model in a golden-angle context, resulting in fully sampled Cartesian trajectories. Then, sinograms were obtained as shown in Figure 4.1. The number of spokes per phase is 13. The dimension of sinograms is $(K \times M_\omega \times C) = 23 \cdot 13 \times 256 \times 32$. In this simulation, the cardiac motion is discrete; thus, no spoke-sharing strategy is applied.

- **Golden-angle Reconstruction of Fetal Cardiac Motion** Fetal cardiac MRI data were acquired on a 1.5 T clinical MR scanner (MAGNETOM Aera, Siemens AG, Healthcare Sector, Erlangen, Germany) with an 18-channel body array coil and a 32-channel spine coil for signal reception. Images were acquired with an untriggered continuous 2D bSSFP sequence that was modified to acquire radial readouts with a golden-angle trajectory [135]. The acquisition parameters were: FOV = 260 × 260 mm^2, acquisition matrix size = 256 × 256 pixels, slice thickness = 4.0 mm, TE/TR = 1.99/4.1 ms, RF excitation angle = 70deg, radial readouts = 1600, acquisition time = 6.7 s and bandwidth = 1028 Hz/pixel.

Reconstruction Experiments

We use an Intel i7-7820X (3.60GHz) CPU and an Nvidia Titan X (Pascal) GPU. Pytorch 1.0.0 on Python 3.6 is used to implement our method. All experiments are performed in single-batch mode. The input size is (8×8). The cost function used to train our neural network is (4.11). The learning rate is 10^{-3}, with [141] as optimizer.

- **Retrospective Simulation** The number of iterations is 10,000. The scheduling for the learning rate is 0.5 multiplier per 2,000 iterations. We set $\mathcal{U} \sim (0, 0.1)$.

- **Golden-angle Reconstruction of Fetal Cardiac Motion** The number of iterations is 20,000. No scheduling for the learning rate is applied. We used 14 cycles of latent variables. We set $\mathcal{U} \sim (0, 10)$ to obtain the results shown in Figure 4.6.

We use the regressed signal-to-noise ratio (RSNR) as a quantitative metric in our retrospective simulations. With $\mathbf{x}$ the oracle and $\mathbf{x}^*$ the reconstructed image, RSNR is given by

$$\text{RSNR} = \max_{a,b \in \mathbb{R}} 20 \log \frac{\|\mathbf{x}\|_2}{\|\mathbf{x} - a\,\mathbf{x}^* + b\|_2}, \tag{4.13}$$

where a higher RSNR corresponds to a better reconstruction.

For real datasets, a CS approach with self-gated signals takes the role of a state-of-the-art baseline [134, 135].

4.3 Results

4.3.1 Retrospective Simulation

In this experiment, the acquisition process is simulated, which allows us to build the ground truth from a fully sampled k-space. We use $n = 13$ spokes for the reconstruction and present the results in Figure 4.4. There, the bandpass method (BP) corresponds to a zero-filled DFT while GRASP is the baseline against which we compare the performance of our method. We see in Figure 4.4 (A) that GRASP leads to blurring artifacts, while the residual map discloses the occurrence of errors around the wall of the heart. By contrast, our proposed method gives better results. This is confirmed in Figure 4.4 (B), where the cardiac motions are captured better by our model than by GRASP. The systolic phase of our reconstruction is well described and very close to the ground-truth, whereas the systolic phase captured by GRASP is too flat.

4.3.2 Golden-Angle Reconstruction of Fetal Cardiac Motion

In this experiment, the acquisition process is real, so that no ground truth is available. We use $n = 5$ spokes for most reconstructions ($n = 15$ for GRASP) and present the results in Figure 4.5. The synthesis of the sequence of latent variables is explained in Section 4.4.1. The overlapped method (OV) corresponds to frames generated from all spokes, while the reordered method (RD) is able to reconstruct multiple cardiac phases by reordering k-t sparse SENSE by self-gated signals [134, 135]. Self-gating, driven by correlation coefficients between approximately reconstructed images, is necessary because it is impractical to capture the electrocardiogram signals of a fetal heart. This ultimately prevents the reordering of golden angles in accordance with their phase.

In the absence of ground truth, we shall take OV to be the reference image for navigation purposes. However, because OV considers all spokes simultaneously, it is a static image that is of high quality only in regions that are not moving. We see in Figure 4.5 (A) that BP and GRASP are noisier than RD and our method. Comparing now RD to our approach, we see that ours produces better-resolved features, particularly for the hyper-intense dot-like structures.

In Figure 4.5 (B), it becomes apparent that BP fails altogether to capture the fetal cardiac beats, while GRASP is less noisy but still mostly fails to reconstruct

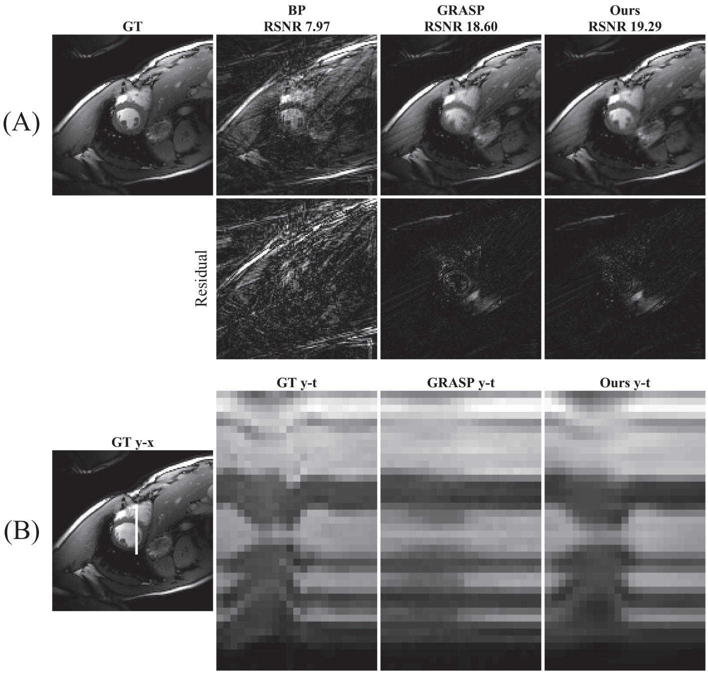

Figure 4.4: Retrospective reconstructions of cine dynamic MRI, with $n = 13$. Top rows (A), from left to right: ground truth; field of view of reconstructed frames from BP, GRASP, and ours. Bottom row (B), from left to right: ground-truth with a white line indicating the (y-t) location of cross sections; cross sections from BP, GRASP, and ours.

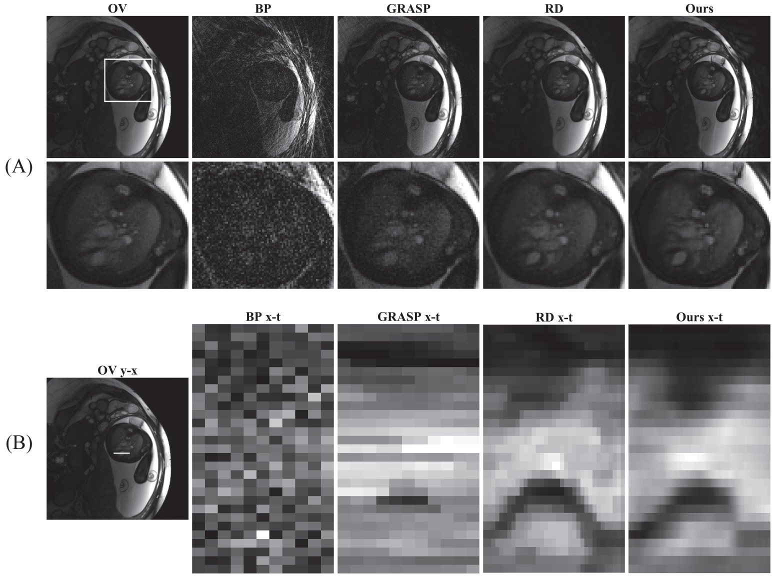

Figure 4.5: Dynamic reconstructions of a fetal heart for one beating cycle. Top rows (A), from left to right: field of view from OV, BP, GRASP, RD, and ours. Bottom row (B), from left to right: OV with a white line indicating the (y-t) location of cross sections; cross sections from BP, GRASP, RD, and ours.

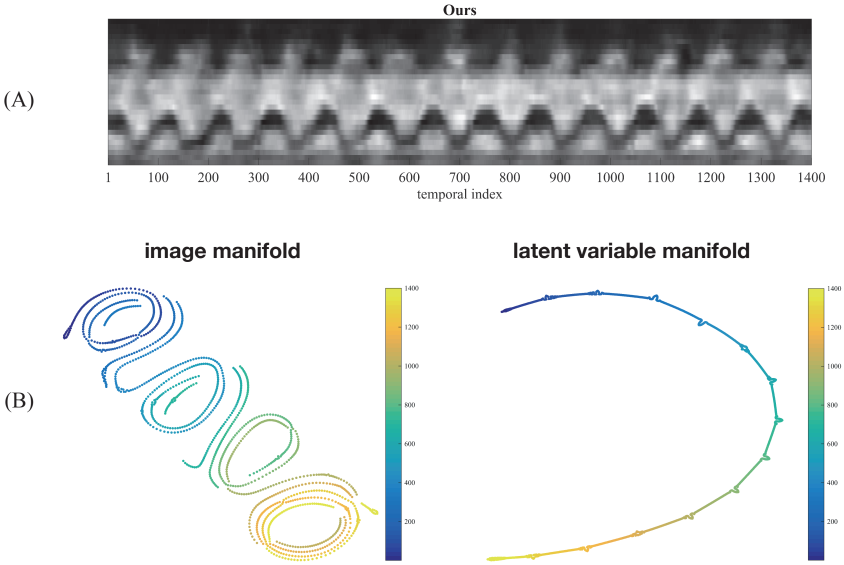

Figure 4.6: Top (A): Series of (y-t) cross sections of our reconstruction from the region of interest in Figure 4.5 (B). Bottom (B): t-SNE embedding from image frames (left) and latent variables (right). The temporal index is color-coded.

motions. RD fares better; unfortunately, its reordering process may lead it to superpose in the same frame spokes that may belong indeed to different cardiac cycles. By contrast, our method[4] reconstructs each frame with data from just a few neighboring spokes, thus avoiding the mingling of different cycles. Consequently, we expect our reconstructed systolic phase to capture well the true motion of the heart. The cross section from our method is similar to that of RD but the motion is smoother in our case, which is the expected behavior of a beating heart.

We provide in Figure 4.6 (A) our complete reconstructed sequence of cardiac cycles, in the form of a (y-t) cross section. The quasi-periodicity of the cardiac motion is clearly visible along the temporal axis, while motion variations can be discerned from cycle to cycle. In Figure 4.6 (B), we also explore over the image and latent domains the structure of the manifold of the cardiac motion. The visualization proceeds through a (t-SNE) embedding [142]. We observe that the continuous trajectory for the image manifold is well aligned with the temporal index, while the input latent variables lie on a smooth manifold in the latent space. The latent variables that correspond to Figure 4.6 (B) have been cut in fourteen chunks; the reason for this is explained in Section 4.4.1. The importance of smoothness in the latent variables is described in Section 4.4.2.

4.4 Discussion

4.4.1 Latent Encoding for Acyclic Data

Letting K in Section 4.2.3 be such that truly all data in Figure 4.6 (A) were taken jointly, and interpolating the latent variables between the only two endpoint realizations $\mathbf{z}_0$ and $\mathbf{z}_{(K-1)\,\Delta t}$, we observed that the reconstruction $\mathbf{x}_t^*$ of the fetal cardiac motion took a constant, time-independent value. We surmise that this failure is due to the overly strong presence of non-periodic components in the data. To cope with them, we adapted our scheme slightly and proceeded by piecewise interpolation, the pieces being made of temporal chunks in the latent space. More precisely, we generated fourteen realizations $\{\mathbf{z}^{(\tau)}\}_{\tau\in[0\ldots13]}$ of the latent variables, equi-spaced in time; then, instead of building $\mathbf{z}_t$ as a linear combination of $\mathbf{z}_0 = \mathbf{z}^{(0)}$

---

[4] For display purposes, we show only one cycle of our cross section. In fact, our reconstructed data have as many frames (K = 1,400) than there are spokes, in reason of the spoke-sharing mechanism of Section 4.2.2.

and $\mathbf{z}_{(K-1)\,\Delta t} = \mathbf{z}^{(13)}$, we built $\mathbf{z}_t$ as a linear combination of $\mathbf{z}^{(\tau)}$ and $\mathbf{z}^{(\tau+1)}$, with an appropriate τ that depends on t. Note that, while the latent variables evolve now chunk-wise, the network is still time-independent and trained over all data jointly. The chunk boundaries are made visible in Figure 4.6 (B).

4.4.2 Smoothness in the Manifold of Latent Variables

To explore the importance of smoothness in the manifold of latent variables, we have trained separately two neural networks, each with a different configuration for the latent variables. On one hand, we have considered independent realizations of random latent variables at each frame, which resulted in an uncorrelated, non-smooth latent spatio-temporal manifold. On the other hand, we have followed the interpolation procedure of Figure 4.3. After examination of the outcome shown in Figure 4.7, we conclude that the non-smooth manifold fails to reconstruct dynamic images, while our proposed approach succeeds.

4.4.3 Size of Latent Variables

Until now, we have fixed the size of the latent variables as (8×8), which corresponds to $Z = 64$. In this section instead, we let Z vary and report the resulting RSNR in Table 4.2, where we observe an acceptable quality of reconstruction for latent variables of size ranging between (2×2) and (16×16). For larger sizes, the reconstruction are corrupted by significant artifacts. We expected the smallest, (1×1) size to reflect the fact that time is a one-dimensional variable. Yet, the learning process failed to converge to a desirable solution in this case, the corresponding sequence of reconstructed images taking the constant, time-independent value shown in the last row of Figure 4.8. In conclusion, this experiment suggests that the (8×8) size provides a good tradeoff between the dimension of the manifold spanned by the latent variables and the convergence issues inherent with any optimization procedure.

4.4.4 Variations on Latent Variables

In this section, we explore various scenarii beyond linear interpolation, which we summarize in Figure 4.8.

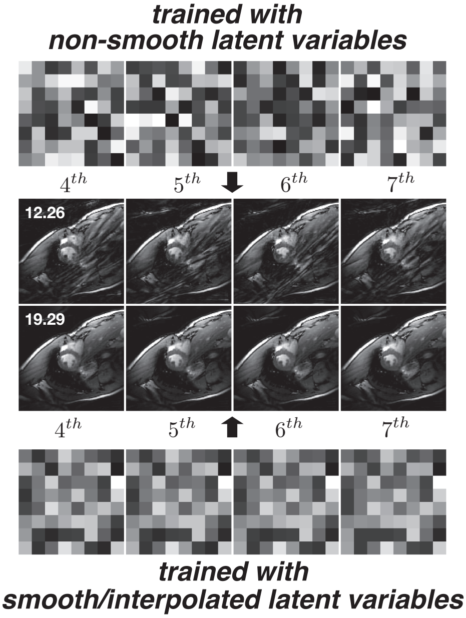

Figure 4.7: Importance of smoothness in a latent manifold.

Table 4.2: Regressed SNR in terms of the size of the latent variables.

Latent size	1×1	2×2	4×4	8×8	16×16	32×32	64×64	128×128
RSNR	16.9	19.1	18.9	19.1	18.9	17.5	13.7	13.6

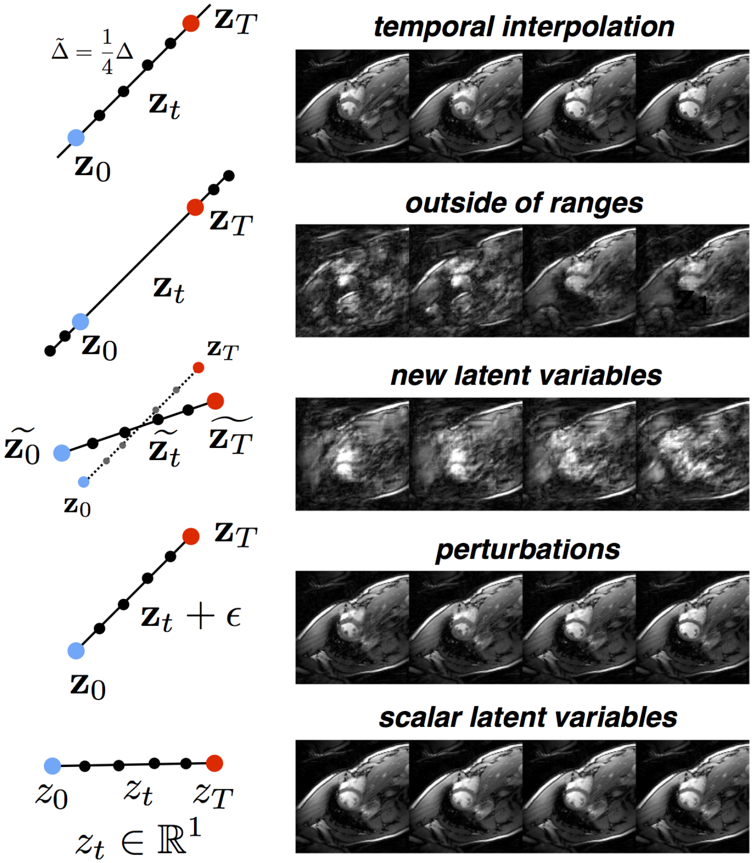

Figure 4.8: Latent-variable scenarii. Here, $T = (k-1)\Delta t$.

i. **Temporal Interpolation** In the first scenario, we consider $t \notin \mathbb{Z}\,\Delta t$, which corresponds to a fine interpolation of the intermediate latent variables. This gives us access to temporally interpolated images at the output of our neural network.

ii. **Temporal Extrapolation** In the second scenario, we extrapolate the latent variables outside of the temporal range that was available for learning. This gives distorted images.

iii. **New Latent Variables** In the third scenario, we make use of random realizations of latent variables that differ from those used while learning. This gives severely perturbed images.

iv. **Perturbations** In the fourth scenario, we perturb the latent variables by adding uniform noise whose energy amounts to 10 %. This setting outputs clean images.

v. **Scalar Latent Variables** In the fifth and last scenario, we deploy scalar latent variables. This results in a time-independent, non-moving sequence of images.

4.4.5 Memory Savings

In a CS context, the gradient updates of the iterative optimization process would need one to allocate enough memory to hold a whole target reconstruction volume. For example, the reconstruction of 5,000 frames with spatial size (256×256) would need one to handle data of size $(256 \times 256 \times 5000)$, which demands for over a gigabyte of memory.

By contrast, our approach is much less memory-hungry. It optimizes the neural network using batches, which requires the simultaneous handling of only those frames that correspond to the batch size. In short, the fact that our proposed approach handles few 2D images whereas CS handles a 2D+t extended sequence leads to substantial savings, particularly for golden-angle dynamic MRI with many frames. In our approach, we only save a nerual network; for example, its memory demands for the spatial size (256×256) are about half-a-dozen megabytes. This cost is negligible compared to that of the CS approach.

4.4.6 Benefits of Our Approach

Several competing methods aim at synthesizing a single cardiac cycle out of spokes collected over several cycles at various cardiac phases. There, a synchronized ECG could in principle allow one to associate a phase to each spoke; however, the deployment of this ancillary ECG would make the practical apparatus more involved, which is often undesirable. Furthermore, there are applications such as fetal cardiology where no ECG can be deployed at all. In traditional ECG-free approaches instead, one deploys self-gating methods for the purpose of phase assignment. They proceed on the basis of either human inspection or heuristic decisions, which makes them arduous, non-reproducible, and prone to errors. (Sometimes, the assignment is no more advanced than a simple manual sorting.) One specific additional difficulty that self-gating methods must deal with originates with the necessary Cartesian-to-polar conversion inherent in radial sampling trajectories, which ultimately results in streaking artifacts that tend to confound phase assignments, particularly those based on visual assessments in the spatial domain.

By contrast, in our proposed approach we relax the hypothesis of ideal periodicity. As presented in Section 4.4.1, this allows us to take into consideration cycle-to-cycle variations, thus providing access to clinical insights that are not available with traditional accelerated methods. Moreover, our ECG-free approach has the major benefit that no phase assignment is needed, thus providing access to a continuously defined time variable. Being fully automated, our approach is also reproducible. Another advantage is that the streaking artifacts associated to radial trajectories play no role since the reference data used for training in (4.10) or (4.11) are the spokes themselves, as opposed to reconstructed images. Finally, when compared to CS, our approach achieves better reconstruction, with a gain of 0.69 dB, as seen in Figure 4.4. It also leads to a simpler optimization task with fewer hyper-parameters. For instance, k-t SENSE requires three interdependent hyper-parameters whose optimal value is found only after some substantial grid-search effort, while the two hyper-parameters of our approach are easier to interpret since they trivially consist of just an initial learning rate, along with a number of iterations.

4.4.7 Limitations and Future Work

One major limitation of our proposed approach is its computational complexity. At times, no fewer than 10,000 iterations are required before convergence is observed. This imposes a large computational burden since every iteration involves both forward and adjoint operations. From a technological perspective, we implemented the forward model in Python 3.6, with Pytorch (v1.0.0) as main library. Unfortunately, NuFFT is currently optimized neither for Python nor for GPU usage, which lead to a marked slowdown of our implementation. For instance, our method spent a whole day letting one GPU (TITAN X) process the cardiac dataset presented in Section 4.2.5, whereas GRASP terminated within ten minutes. This bottleneck will be improved by a better integration of the NuFFT library.

In our future work, we want to explore the dynamics of the latent variables. In this work indeed, we simply took advantage of linear interpolation to build trivial intermediate states of the latent variables; however, refined approaches may better capture the temporal correlation between frames.

4.5 Summary

In this chapter, we develop a novel unsupervised deep-learning-based algorithm to solve the inverse problem found in dynamic magnetic resonance imaging (MRI). Our method needs neither prior training nor additional data; in particular, it does not require either electrocardiogram or spokes-reordering in the context of cardiac images. It generalizes to sequences of images the recently introduced deep-image-prior approach. The essence of the proposed algorithm is to proceed in two steps to fit k-space synthetic measurements to sparsely acquired dynamic MRI data. In the first step, we deploy a convolutional neural network (CNN) driven by a sequence of low-dimensional latent variables to generate a dynamic series of MRI images. In the second step, we submit the generated images to a nonuniform fast Fourier transform that represents the forward model of the MRI system. By manipulating the weights of the CNN, we fit our synthetic measurements to the acquired MRI data. The corresponding images from the CNN then provide the output of our system; their evolution through time is driven by controlling the sequence of latent variables whose interpolation gives access to the sub-frame—or even continuous—temporal control of reconstructed dynamic images. We perform experiments on simulated and real cardiac images of a fetus acquired through 5-spoke-based golden-angle measurements. Our results show improvement over the current state-of-the-art. To the best of our knowledge, this work is the first approach of unsupervised learning in accelerated dynamic MRI.

Chapter 5

CryoGAN: Cryo-EM Reconstruction using GAN Framework

[1]In the previous chapter we discussed deep image prior, which, in principle is applicable to any inverse problem where the forward model is exactly known. In this chapter, we propose a new unsupervised deep-learning-based algorithm for Cryo-EM, an imaging problem where only a parameteric form of the forward model is known. Our method, as we shall see, approaches this problem from a distributional perspective.

5.1 Overview

Single-particle cryo-electron microscopy (Cryo-EM) is a powerful method to determine the atomic structure of macro-molecules by imaging them with electron rays at cryogenic temperatures [143–145]. Its popularity has rocketed in recent years, culminating in 2017 with the Nobel Prizes of Jacques Dubochet, Richard Henderson,

---

[1] This chapter uses content from our work [28].

and Joachim Frank. In Cryo-EM one obtains many 2D noisy tomographic projections from separate instance of the same but randomly oriented 3D biomolecule[2]. There exists a multitude of software packages to produce high-resolution 3D structure(s) from these 2D measurements [146–153]. These sophisticated algorithms, which include projection-matching approaches and maximum-likelihood optimization frameworks, enable the determination of structures with unprecedented atomic resolution.

Yet reconstruction procedures in single-particle Cryo-EM still face complex obstacles. The task involves a high-dimensional, nonconvex optimization problem with numerous local minima. Hence, the outcome of the global process is predicated on the quality of the initial reconstruction [154, 155]. Moreover, one still often relies on the input of an expert user for appropriate processing decisions and parameter tuning [156]. Even for more automated methods, the risk of outputting incorrect and misleading 3D reconstructions is ever-present. A key reason behind such complexity is that the imaged particles have unknown poses. To handle this, most software packages rely on a marginalized maximum-likelihood (ML) formulation [157] that is solved through an expectation-maximization algorithm [151, 153]. The latter involves calculations over the discretized space of poses for each projection, a computationally demanding procedure.

To bypass these limitations, we introduce CryoGAN, an unsupervised reconstruction algorithm for single-particle Cryo-EM that exploits the remarkable ability of generative adversarial networks (GANs) to capture data distributions [36]. Similar to GANs, CryoGAN is driven by the competitive training of two entities: one that tries to capture the distribution of real data, and another that discriminates between generated samples and samples from the real dataset (Figure 5.1). In a classical GAN, the two entities are each a convolutional neural network (CNN). They are known as the generator and the discriminator and are trained simultaneously using backpropagation. The important twist with CryoGAN is that we replace the generator network by a Cryo-EM physics simulator. By doing so, CryoGAN learns the 3D density map whose simulated projections are the most consistent with a given dataset of 2D measurements in a distributional sense.

The CryoGAN architecture represents a complete change of paradigm for single-particle Cryo-EM reconstruction. No estimation of the poses is attempted during

---

[2] In this work we consider the homogenuous case where the biomolecule exhibit only a single conformation. The heterogeneuous (multiple conformation) case is the topic of the next chapter.

the learning procedure; rather, the reconstruction is obtained through distributional matching performed in a likelihood-free manner. Hence, CryoGAN sidesteps many of the computational drawbacks associated with likelihood-based methods.

In practice, CryoGAN requires no prior knowledge of the 3D structure; its learning process is purely unsupervised. The user needs only to feed the particle images and estimates of the contrast transfer function (CTF) parameters. No initial volume is needed: the algorithm starts with a volume initialized with zeros. The CryoGAN framework is backed up by a comprehensive mathematical framework that provides guarantees on the recovery of the correct structure under a given set of assumptions that are often met in practice, at least to some degree of approximation.

We first assessed the performance and stability of CryoGAN on a synthetic β-galactosidase dataset, where we generated noisy projections *in silico*. The results demonstrate that our unsupervised reconstruction paradigm permits the recovery of a 8.6 Å structure (Figure 5.3). We then deployed CryoGAN on a real experimental β-galactosidase dataset (EMPIAR-10061) [158], reaching a resolution of 12.1 Å in 150 minutes in far more challenging conditions (Figure 5.5). These preliminary results provide a strong indication of the suitability of the CryoGAN framework for the reconstruction of real structures. On the implementation side, we expect to be able to improve the resolution of the reconstructions by taking advantage of the many recent technical developments and advances in the area of GANs. In the meantime, the preliminary results obtained with CryoGAN are encouraging and demonstrate the potential of adversarial-learning schemes in image reconstruction. The proposed paradigm opens many new perspectives in single-particle Cryo-EM and paves the way for more applications beyond the present one.

5.2 Image-Formation Model in Single-Particle Cryo-EM

We detail here the Cryo-EM image formation model used in our implementation of the CryoGAN physics simulator.

We use the standard model for the single-particle Cryo-EM imaging procedure [159]

$$\mathbf{y} = \mathbf{H}_{\boldsymbol{\varphi}}\mathbf{x} + \mathbf{n}, \tag{5.1}$$

where $\mathbf{y} \in \mathbb{R}^M$ is a (vectorized) 2D projection; $\mathbf{x} \in \mathbb{R}^V$ is the (vectorized) 3D

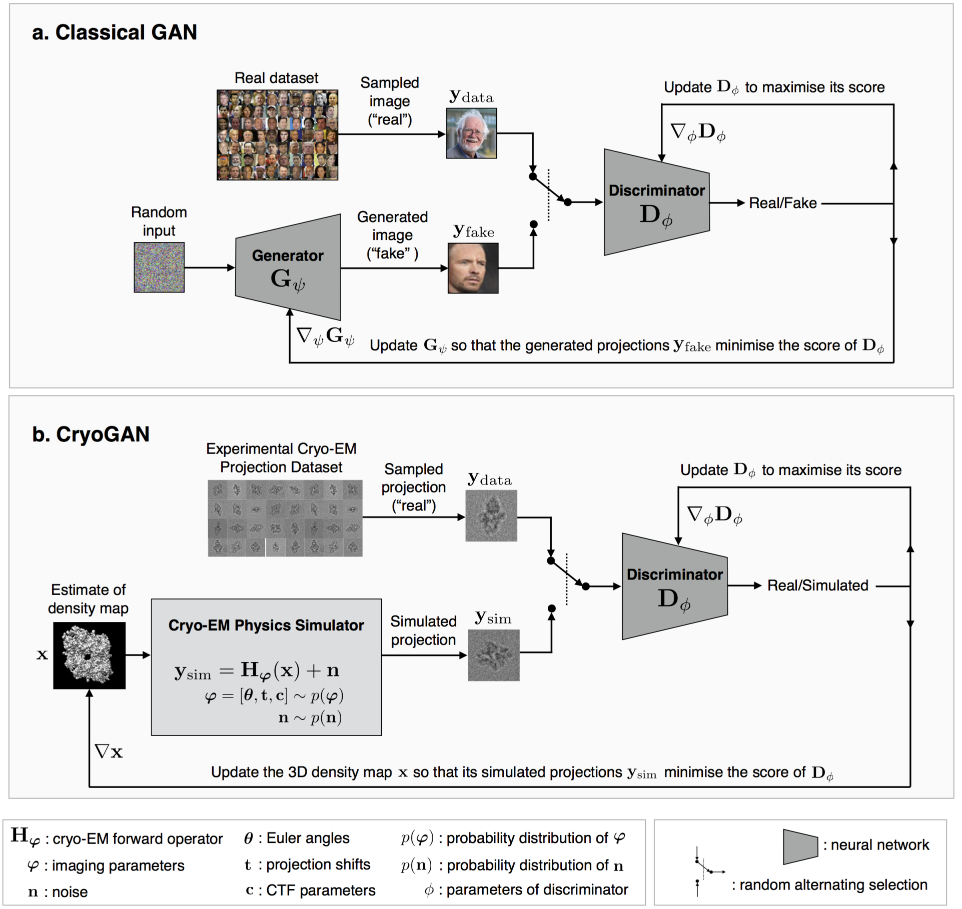

Figure 5.1: A schematic comparison between **(a)** a classical GAN architecture and **(b)** the CryoGAN architecture. Both frameworks rely on a deep adversarial learning scheme to capture the distribution of the real data. CryoGAN exploits this ability to look for the volume whose simulated projections have a distribution that matches the real data distribution. This is achieved by adding a "Cryo-EM physics simulator" that produces measurements following a mathematical model of the Cryo-EM imaging procedure. Importantly, CryoGAN does not rely on a first low-resolution volume estimate, but is initialized with a zero-valued volume. Note that, for both architectures, the updates involve backpropagating through the neural networks; those actions are not indicated here for the sake of clarity.

density; $\mathbf{H}_{\boldsymbol{\varphi}} \in \mathbb{R}^{M \times V}$ is the forward operator with parameters $\boldsymbol{\varphi} \sim p_{\boldsymbol{\varphi}}$; and $\mathbf{n} \in \mathbb{R}^M$ is an additive noise following a distribution $p_{\mathbf{n}}$. In Cryo-EM one obtains $10^4 - 10^7$ measurements of the biomolecule. Each of these measurement is obtained with an unknown $\boldsymbol{\varphi}$. The imaging parameters $\boldsymbol{\varphi}$ comprise the projection (Euler) angles $\boldsymbol{\theta} = (\theta_1, \theta_2,\ \theta_3)$, the projection shifts $\mathbf{t} = (t_1, t_2)$, and the CTF parameters $\mathbf{c} = (d_1, d_2, \alpha_{\mathrm{ast}})$, where d_1 is the defocus-major, d_2 is the defocus-minor, and α_{ast} is the angle of astigmatism.

The forward operator $\mathbf{H}_{\boldsymbol{\varphi}}$ is given by

$$\mathbf{H}_{\boldsymbol{\varphi}} = \mathbf{C}_{\mathbf{c}} \mathbf{S}_{\mathbf{t}} \mathbf{P}_{\boldsymbol{\theta}}, \tag{5.2}$$

where $\mathbf{P}_{\boldsymbol{\theta}} : \mathbb{R}^V \to \mathbb{R}^M$ is a projection operator (mathematically speaking, the X-ray transform [160]), $\mathbf{S}_{\mathbf{t}} : \mathbb{R}^M \to \mathbb{R}^M$ is a shift operator, and $\mathbf{C}_{\mathbf{c}} : \mathbb{R}^M \to \mathbb{R}^M$ is a convolution operator. We discuss in more detail the continuous-domain physics behind the image formation model $\mathbf{H}_{\boldsymbol{\varphi}}$.

5.2.1 Image Formation in Continuous-Domain

Our continuous-domain forward model is based on [159] (2.1-2.10), [161], and [162], which model the relationship between the 3D density map and a 2D measurement as linear.

The continuous-domain measurement can be given in the following form

$$y = H_{\boldsymbol{\varphi}}\{f\} + n \tag{5.3}$$

where $y : \mathbb{R}^2 \to \mathbb{R}$ is the intensity measured on the image plane, $n : \mathbb{R}^2 \to \mathbb{R}$ is the noise, $f : \mathbb{R}^3 \to \mathbb{R}$ is the density map we aim to recover, and $H_{\boldsymbol{\varphi}}$ is the measurement operator dependent on the imaging parameters $\boldsymbol{\varphi}$. This model can be further decomposed as

$$y = C_{\mathbf{c}} S_{\mathbf{t}} P_{\boldsymbol{\theta}}\{f\} + n \tag{5.4}$$

where $\boldsymbol{\varphi} = (\boldsymbol{\theta}, \mathbf{t}, \mathbf{c})$. We now discuss the three involved operators in more details.

Projection Operator

The projection operation is given by [160]

$$P_{\boldsymbol{\theta}}\{f\}(x_1, x_2) = \int_{-\infty}^{\infty} \mathbf{R}_{\theta}\{f\}(x_1, x_2, x_3)\, \mathrm{d}x_3 \tag{5.5}$$

where $\mathbf{R}_{\boldsymbol{\theta}}$ is the rotation matrix associated with $\boldsymbol{\theta}$ and $R_{\boldsymbol{\theta}}\{f\}(\mathbf{x}) = f(R_{\boldsymbol{\theta}}^{-1}\mathbf{x})$.

Shift Operator

The projection measurements are picked from the micrographs and can thus be off-centered. This is modelled via the shift operator which, for any $y_s : \mathbb{R}^2 \to \mathbb{R}$, yields

$$S_{\mathbf{t}}\{y_s\}(x_1, x_2) = y(x_1 - t_1, x_2 - t_2) \tag{5.6}$$

where $\mathbf{t} = (t_1, t_2)$.

Convolution by CTF

The effect of the operator of $C_{\mathbf{c}}$ on any $y_c : \mathbb{R}^2 \to \mathbb{R}$ is given in Fourier domain as

$$\mathcal{F}\{C_{\mathbf{c}}\{y_c\}\}(\boldsymbol{\omega}) = \hat{C}_{\mathbf{c}}(\boldsymbol{\omega}).\hat{y}_c(\boldsymbol{\omega}) \tag{5.7}$$

where $\mathcal{F}$ is the Fourier transform and $\hat{y}_c = \mathcal{F}\{y_c\}$. Its Fourier transform $\hat{C}_{\mathbf{c}}$ (*i.e.*, the CTF) is given by

$$\hat{C}_{\mathbf{c}}(\boldsymbol{\omega}) = \hat{C}^p_{\mathbf{c}}(\boldsymbol{\omega})\hat{E}(\boldsymbol{\omega})\hat{A}(\boldsymbol{\omega}). \tag{5.8}$$

There, $\hat{C}^p_{\mathbf{c}} : \mathbb{R}^2 \to \mathbb{R}$ is the phase-contrast transfer function that takes the form

$$\hat{C}^p_{\mathbf{c}}(\boldsymbol{\omega}) = -\sqrt{1 - A^2}\sin(\gamma_{\mathbf{c}}(\boldsymbol{\omega})) - A^2\cos(\gamma_{\mathbf{c}}(\boldsymbol{\omega})), \tag{5.9}$$

with

$$\gamma_{\mathbf{c}}(\boldsymbol{\omega}) = \pi\lambda\left(d_{\mathbf{c}}(\alpha)\|\boldsymbol{\omega}\|^2 - \frac{1}{4}\lambda^3 c_{\mathrm{s}}\|\boldsymbol{\omega}\|^4\right), \tag{5.10}$$

where λ is the electron wavelength, c_s is third-order spherical aberration constant, α is the phase of the vector $\boldsymbol{\omega}$, and $d_{\mathbf{c}}(\alpha)$ is the defocus arising at the phase α. This defocus is given as

$$d_{\mathbf{c}}(\alpha) = d_1\cos^2(\alpha - \alpha_{\mathrm{ast}}) + d_2\sin^2(\alpha - \alpha_{\mathrm{ast}}), \tag{5.11}$$

where d_1 and d_2 are the horizontal and vertical defocus, respectively, and α_{ast} is the reference angle that defines the azimuthal direction of axial astigmatism. The objective aperture function $\hat{A} : \mathbb{R}^2 \to \mathbb{R}$ is given by

$$\hat{A}(\boldsymbol{\omega}) = \begin{cases} 1, & \|\boldsymbol{\omega}\| \leq \omega_{\mathrm{cutoff}} \\ 0, & \|\boldsymbol{\omega}\| > \omega_{\mathrm{cutoff}}, \end{cases} \tag{5.12}$$

where $\omega_{\mathrm{cutoff}} = \frac{2\pi d_{\mathrm{ap}}}{f_0 \lambda}$ is the cutoff frequency, f_0 is the focal length of the objective lens, and d_{ap} corresponds to the diameter of the aperture. The spatial and chromatic envelope function $\hat{E} : \mathbb{R}^2 \to \mathbb{R}$ is given by

$$\hat{E}(\boldsymbol{\omega}) = \exp\left(-B(\|\boldsymbol{\omega}\|^2)\right), \tag{5.13}$$

where $B : \mathbb{R}^2 \to \mathbb{R}$ is a function influenced by chromatic aberration and spatial incoherence.

Discretization

The discretization of $H_{\boldsymbol{\varphi}}$ results in $\mathbf{H}_{\boldsymbol{\varphi}}$. This discretized measurement operator is itself decomposed of the discretized projection, shift, and convolution operation which are denoted by $\mathbf{P}_{\boldsymbol{\theta}}$, $\mathbf{S}_{\mathbf{t}}$, and $\mathbf{C}_{\mathbf{c}}$, respectively. The input to the operator $\mathbf{H}_{\boldsymbol{\varphi}}$ is a discretized version of the continuous-domain 3D volume. This discretization of the 3D volume is done using a suitable basis function [163].

5.3 Mathematical Framework of CryoGAN

The goal of single-particle Cryo-EM reconstruction is to estimate a 3D density map $\mathbf{x}_{\mathrm{rec}}$ whose projections are consistent with the observed projections (data) of the true density map $\mathbf{x}_{\mathrm{true}}$.[3]

We write the conditional probability density function of a measurement $\mathbf{y}$ given a volume $\mathbf{x}$ by marginalizing over the imaging parameters $\boldsymbol{\varphi}$

$$p(\mathbf{y}|\mathbf{x}) = \int_{\boldsymbol{\varphi} \in \Phi} p_{\mathbf{n}}(\mathbf{y} - \mathbf{H}_{\boldsymbol{\varphi}}\mathbf{x}) p_{\boldsymbol{\varphi}}(\boldsymbol{\varphi}) \mathrm{d}\boldsymbol{\varphi}, \tag{5.14}$$

---

[3] In this forward model, we assume conformational homogeneity of the underlying structure. The extension of CryoGAN to multiple conformations is discussed in the next chapter.

where Φ is the set of all the possible imaging parameters. We denote a noiseless projection as $\mathbf{y}_{\text{noiseless}} = \mathbf{H}_{\boldsymbol{\varphi}}\mathbf{x}$. In our formulation, the projections in the real dataset are samples of a distribution p_{data}; hence, $p(\mathbf{y}|\mathbf{x}_{\text{true}}) = p_{\text{data}}(\mathbf{y})$ assuming that the forward model is correct.

We demonstrate in Theorem 5.5.5 in Section 5.5 that two 3D volumes have identical conditional distributions if and only if they are identical, up to rotation and reflection. Hence, Theorem 5.5.5 implies that, for the reconstruction $\mathbf{x}_{\text{rec}}$ to satisfy $\mathbf{x}_{\text{rec}} = \mathbf{x}_{\text{true}}$, it must also satisfy $p(\mathbf{y}|\mathbf{x}_{\text{rec}}) = p(\mathbf{y}|\mathbf{x}_{\text{true}})$. Thus, we can formulate the reconstruction task as the minimization problem

$$\mathbf{x}_{\text{rec}} = \arg\min_{\mathbf{x}} D\big(p(\mathbf{y}|\mathbf{x}), p(\mathbf{y}|\mathbf{x}_{\text{true}})\big), \tag{5.15}$$

where D is some distance between two distributions. In essence, (5.15) states that the appropriate reconstruction is the 3D density map whose projection distribution is the most similar to the real dataset in a distributional sense. For the sake of conciseness, we shall henceforth use the notation $p(\mathbf{y}|\mathbf{x}) = p_{\mathbf{x}}(\mathbf{y})$.

As distance in (5.15), we use the Wasserstein distance defined as

$$D(p_1, p_2) = \inf_{\gamma \in \Pi(p_1, p_2)} \mathbb{E}_{(\mathbf{y}_1, \mathbf{y}_2) \sim \gamma}[\|\mathbf{y}_1 - \mathbf{y}_2\|], \tag{5.16}$$

where $\Pi(p_1, p_2)$ is the set of all the joint distributions $\gamma(\mathbf{y}_1, \mathbf{y}_2)$ whose marginals are p_1 and p_2, respectively. Our choice is driven by works demonstrating that the Wasserstein distance is more amenable to minimization than other popular distances (*e.g.*, total-variation or Kullback-Leibler divergence) for this kind of application [164]. Using (5.16), the minimization problem (5.15) expands as

$$\mathbf{x}_{\text{rec}} = \operatorname*{argmin}_{\mathbf{x}} \inf_{\gamma \in \Pi(p_{\mathbf{x}}, p_{\text{data}})} \mathbb{E}_{(\mathbf{y}_1, \mathbf{y}_2) \sim \gamma}[\|\mathbf{y}_1 - \mathbf{y}_2\|]. \tag{5.17}$$

By using the formalism of [164–166], this minimization problem can also be stated in its dual form

$$\mathbf{x}_{\text{rec}} = \operatorname*{argmin}_{\mathbf{x}} \max_{f : \|f\|_L < 1} \Big(\mathbb{E}_{\mathbf{y} \sim p_{\text{data}}}[f(\mathbf{y})] - \mathbb{E}_{\mathbf{y} \sim p_{\mathbf{x}}}[f(\mathbf{y})]\Big), \tag{5.18}$$

where $\|f\|_L$ denotes the Lipschitz constant of the function $f : \mathbb{R}^M \to \mathbb{R}$.

5.3.1 Connection with Wasserstein GANs

Equation ((5.18)) falls under the framework of the generative adversarial networks (GANs) [36] called Wasserstein GANs (WGANs) [164]. In the standard WGAN representation, the function f is parameterized by a neural network $\mathbf{D}_\phi$ with parameters ϕ that is called the discriminator. The task of this discriminator is to learn to differentiate between real samples (typically coming from an experimental dataset) and fake samples. The latter are produced by another neural network, called the generator, which aims at producing samples that are realistic enough to fool the discriminator. This adversarial-learning scheme progressively drives the WGAN to capture the distribution of the experimental data.

In CryoGAN, we exploit this capability of adversarial schemes to learn the volume $\mathbf{x}$ whose simulated projections follow the captured real-data distribution. To do so, we rely on a Cryo-EM physics simulator, whose role is to produce projections of a volume estimate $\mathbf{x}$ using (5.1). These simulated projections then follow a distribution $\mathbf{y} \sim p_{\mathbf{x}}$. Hence, (5.18) translates into

$$\mathbf{x}_{\mathrm{rec}} = \operatorname*{argmin}_{\mathbf{x}} \max_{\mathbf{D}_\phi : \|\mathbf{D}_\phi\|_L \leq 1} \Big(\mathbb{E}_{\mathbf{y}\sim p_{\mathrm{data}}}[\mathbf{D}_\phi(\mathbf{y})] - \mathbb{E}_{\mathbf{y}\sim p_{\mathbf{x}}}[\mathbf{D}_\phi(\mathbf{y})] \Big). \tag{5.19}$$

As proposed in [167], the Lipschitz constraint $\|\mathbf{D}_\phi\|_L \leq 1$ can be enforced by penalizing the norm of the gradient of $\mathbf{D}_\phi$ with respect to its input. This gives the final formulation of our reconstruction problem as

$$\mathbf{x}_{\mathrm{rec}} = \operatorname*{argmin}_{\mathbf{x}} \max_{\mathbf{D}_\phi} \Big(\mathbb{E}_{\mathbf{y}\sim p_{\mathrm{data}}}[\mathbf{D}_\phi(\mathbf{y})] - \mathbb{E}_{\mathbf{y}\sim p_{\mathbf{x}}}[\mathbf{D}_\phi(\mathbf{y})] - \lambda \cdot \mathbb{E}_{\mathbf{y}\sim p_{\mathrm{int}}}[(\|\nabla_{\mathbf{y}}\mathbf{D}_\phi(\mathbf{y})\| - 1)^2] \Big). \tag{5.20}$$

Here, p_{int} denotes the uniform distribution along the straight line between points sampled from p_{data} and $p_{\mathbf{x}}$, while $\lambda \in \mathbb{R}_+$ is an appropriate penalty coefficient (see [167], Section 4).

5.4 The CryoGAN Algorithm

Equation (5.20) is a min-max optimization problem. By replacing the expected values with their empirical counterparts (sums) [167], we reformulate it as the minimization of

$$L_S(\mathbf{x},\mathbf{D}_\phi)=\sum_{n\in S}\mathbf{D}_\phi(\mathbf{y}^n_{\mathrm{data}})-\sum_{n\in S}\mathbf{D}_\phi(\mathbf{y}^n_{\mathrm{sim}})-\lambda\sum_{n\in S}(\|\nabla_{\mathbf{y}}\mathbf{D}_\phi(\mathbf{y}^n_{\mathrm{int}})\|-1)^2), \quad (5.21)$$

where S consists of either the full dataset $S_{\mathrm{full}}=\{1,\dots,N_{\mathrm{tot}}\}$ or a batch $B\subseteq S_{\mathrm{full}}$; $\mathbf{y}^n_{\mathrm{data}}$ is a projection sampled from the acquired experimental dataset; $\mathbf{y}^n_{\mathrm{sim}}\sim p_{\mathbf{x}}$ is a projection of the current estimate $\mathbf{x}$ generated by the Cryo-EM physics simulator; and $\mathbf{y}^n_{\mathrm{int}}=\alpha_n\cdot\mathbf{y}^n_{\mathrm{data}}+(1-\alpha_n)\cdot\mathbf{y}_{\mathrm{sim}}$, where α_n is sampled from a uniform distribution between 0 and 1.

In practice, we minimize (5.21) through stochastic gradient descents (SGD) using batches. We alternatively update the discriminator $\mathbf{D}_\phi$ (for n_{discr} iterations) using an Adam optimizer [141] with gradient

$$\nabla_\phi L_B(\mathbf{x},\mathbf{D}_\phi)=\nabla_\phi\left(\sum_{n=1}^{N}\mathbf{D}_\phi(\mathbf{y}^n_{\mathrm{batch}})-\sum_{n=1}^{N}\mathbf{D}_\phi(\mathbf{y}^n_{\mathrm{sim}})-\lambda\sum_{n=1}^{N}(\|\nabla_{\mathbf{y}}\mathbf{D}_\phi(\mathbf{y}^n_{\mathrm{int}})\|-1)^2\right), \quad (5.22)$$

and the volume $\mathbf{x}$ (for 1 iteration) using the batch gradient,

$$\nabla_{\mathbf{x}}L_B(\mathbf{x},\mathbf{D}_\phi)=\nabla_{\mathbf{x}}\left(-\sum_{n=1}^{N}\mathbf{D}_\phi(\mathbf{y}^n_{\mathrm{sim}})\right). \quad (5.23)$$

The pseudocode and a schematic view of the CryoGAN algorithm are given in Algorithm 3 and Figure 5.1b, respectively. We provide further details of the CryoGAN physics simulator and discriminator network in the next two sections.

5.4.1 The Cryo-EM Physics Simulator

The goal of the physics simulator is to sample $\mathbf{y}_{\mathrm{sim}}\sim p_{\mathbf{x}}(\mathbf{y})$; this is done in three steps. First, we sample the imaging parameters $\boldsymbol{\varphi}$ from the distribution $p_{\boldsymbol{\varphi}}$: $\boldsymbol{\varphi}\sim p_{\boldsymbol{\varphi}}$. Second, we generate noiseless CTF-modulated and shifted projections from the current volume estimate with $\mathbf{H}_{\boldsymbol{\varphi}}(\mathbf{x})$. Third, we sample the noise model to simulate noisy projections $\mathbf{y}=\mathbf{H}_{\boldsymbol{\varphi}}(\mathbf{x})+\mathbf{n}$, where $\mathbf{n}\sim p_{\mathbf{n}}$. We detail these steps in the following paragraphs, and a pseudocode of this Cryo-EM physics simulator is given in Algorithm 4.

Algorithm 3 Pseudocode for CryoGAN

Parameters: number of training iterations, n_{train}; number of iterations of the discriminator per training iteration, n_{discr}; size of the batches used for SGD, N; penalty parameter, λ

```
for $n_{\text{train}}$ do
    for $n_{\text{discr}}$ do
        sample real projections: $\{\mathbf{y}^1_{\text{batch}}, \dots, \mathbf{y}^N_{\text{batch}}\} = \{\mathbf{y}^n_{\text{data}}\}_{n \in B}$
        sample projections simulated from current $\mathbf{x}$: $\{\mathbf{y}^1_{\text{sim}}, \dots, \mathbf{y}^N_{\text{sim}}\} \sim p_{\mathbf{x}}$ (see Algorithm 2)
        sample $\{\alpha_1, \dots, \alpha_n\} \sim U[0,1]$
        for all $n \in \{1, \dots, N\}$, compute $\mathbf{y}^n_{\text{int}} = \alpha_n \cdot \mathbf{y}^n_{\text{batch}} + (1 - \alpha_n) \cdot \mathbf{y}^n_{\text{sim}}$
        update the parameters $\phi$ of the discriminator $\mathbf{D}_\phi$ using (5.22)
    sample $\{\mathbf{y}^1_{\text{sim}}, \dots, \mathbf{y_{sim}}^N\} \sim p_{\mathbf{x}}$
    update the volume $\mathbf{x}$ using (5.23)
```

Recall that the set of imaging parameters is given by $\boldsymbol{\varphi} = (\theta_1, \theta_2, \theta_3, t_1, t_2, d_1, d_2, \alpha_{\text{ast}})$. We first sample the Euler angles $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3)$ from a distribution $p_{\boldsymbol{\theta}}$ decided *a priori* based on the acquired dataset. Similarly, the projection shifts $\mathbf{t} = (t_1, t_2)$ are sampled from the prior distribution $p_{\mathbf{t}}$. The CTF parameters $\mathbf{c} = (d_1, d_2, \alpha_{\text{ast}})$ are sampled from the prior distribution $p_{\mathbf{c}}$. In practice, we exploit the fact that the CTF parameters can often be efficiently estimated for all micrographs. We then uniformly sample from the whole set of extracted CTF parameters.

We generate noiseless projections $\mathbf{y}_{\text{noiseless}}$ by applying $\mathbf{H}_{\boldsymbol{\varphi}}$ to the current volume estimate $\mathbf{x}$. The projection operator $\mathbf{P}_{\boldsymbol{\theta}}$ in (5.2) is implemented using the ASTRA toolbox [168].

The precise modeling of the noise is a particularly challenging feat in single-particle Cryo-EM. To produce noise realizations that are as realistic as possible, we extract random background patches directly from the micrographs themselves, at locations where particles do not appear. For consistency, the noise patch added to a given noiseless projection is taken from the same micrograph that was used to estimate the CTF parameters previously applied to that specific projection. Additional details for this implementation are given in Section A.3.

Algorithm 4 Pseudocode for Cryo-EM Physics Simulator

```
Parameters: current volume estimate, $\mathbf{x}$
sample the Euler angles $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3) \sim p_{\boldsymbol{\theta}}$.
sample the 2D shifts $\mathbf{t} = (t_1, t_2) \sim p_{\mathbf{t}}$.
sample the CTF parameters $\mathbf{c} = (d_1, d_2, \alpha_{\mathrm{ast}}) \sim p_{\mathbf{c}}$.
generate a synthetic noiseless projection based on (5.2), with $\mathbf{y}_{\mathrm{noiseless}} = \mathbf{H}_{\boldsymbol{\varphi}}\mathbf{x}$.
sample the noise $\mathbf{n} \sim p_{\mathbf{n}}$. Add to the projection as $\mathbf{y}_{\mathrm{sim}} = \mathbf{y}_{\mathrm{noiseless}} + \mathbf{n}$.
```

5.4.2 The CryoGAN Discriminator Network

The role of the discriminator is to differentiate between projections from the experimental dataset and projections generated by the Cryo-EM physics simulator. The gradients from the discriminator (see (5.23) in Algorithm 3) carry information on the difference between the real and simulated projections at a given run-time. Those gradients are used by the simulator to update itself, thus improving on the realism of the simulated projections.

The discriminator network takes an image as input and outputs a scalar value. Its architecture is illustrated in Figure 5.2. It is composed of 8 layers: 6 convolutional blocks followed by 2 fully connected (FC) layers. Each convolutional block is made up of a convolutional layer followed by a max-pooling and a leaky ReLU (with negative slope of 0.1). The number of channels in each convolutional layer is 96, 192, 384, 768, 1536, and 3072, respectively. The filters in these layers are of size 3, and the padding size is 1. The max-pooling layer uses a kernel of size 2 with a stride of 2. This leads to a downsampling by a factor of 2. The output of the final convolutional block is then reshaped, fed into the FC layer with 10 neurons, and finally processed by a leaky ReLU. The resulting activations are fed to the last FC layer to output a scalar.

5.4.3 Overall Scheme

In summary, CryoGAN is like a standard GAN, except that the generator network is replaced by a Cryo-EM physics simulator (Figure 5.1b). This simulator implements a mathematical model of the imaging procedure to produce a simulated measurement based on 1) the current volume estimate and 2) a random projection orientation. This image-formation model considers that the Cryo-EM

2D measurement is the projection of the volume at that orientation, modulated by microscopy-related effects and corrupted by substantial additive noise.

The Cryo-EM physics simulator is paired with a discriminator network whose architecture is similar to those used in standard GANs. The role of this discriminator in CryoGAN is to encourage the simulator to learn the 3D volume $\mathbf{x}_{\text{rec}}$ whose simulated dataset distribution $p(\mathbf{y}|\mathbf{x}_{\text{rec}})$ matches that of the real dataset, $p_{\text{data}}(\mathbf{y})$. Mathematically, this equates to

$$\mathbf{x}_{\text{rec}} = \arg\min_{\mathbf{x}} D\big(p(\mathbf{y}|\mathbf{x}), p_{\text{data}}(\mathbf{y})\big), \tag{5.24}$$

where D is an appropriate measure of distance between distributions, which the discriminator allows us to compute; in our implementation, D is the Wasserstein distance [165]. This formulation is based on a sound mathematical framework that provides guarantees on the recovery of the volume. In Theorem 5.5.5 (Section 5.5), we show that, under certain constraints, the reconstructed volume is the same as the true volume up to rotation and reflection.

To perform the minimization (5.24), the CryoGAN algorithm alternates between updates of the discriminator and of the volume with stochastic gradient descent (SGD). The code for our implementation of CryoGAN is written in Python using the PyTorch [169] package.

5.5 Theoretical Guarantee of Recovery

The CryoGAN paradigm is supported by Theorem 5.5.5, which states that perfect recovery is possible from continuous-domain measurements. In the continuous domain, we have $y = H_{\boldsymbol{\varphi}} f + n$ where $y : \mathbb{R}^2 \to \mathbb{R}$ is the 2D measurement obtained from the 3D volume f and n is the noise. Here $H_{\boldsymbol{\varphi}} = C_{\mathbf{c}} S_{\mathbf{t}} P_{\boldsymbol{\theta}}$ is the continuous-domain forward operator, $P_{\boldsymbol{\theta}}$ is the projection operator, $S_{\mathbf{t}}$ is the shift operator, and $C_{\mathbf{c}}$ is the operator for convolution with the CTF (see *Image Formation*). We assume that $\boldsymbol{\theta} \sim p_{\boldsymbol{\theta}}$, $\mathbf{c} \sim p_{\mathbf{c}}$, $\mathbf{t} \sim p_{\mathbf{t}}$, and $\mathbb{P}_n$ be the probability measure associated with n. We assume that

i. the characteristic functional $\hat{\mathbb{P}}_n$ of the noise probability measure $\mathbb{P}_n$ is non-zero everywhere in its domain and n is pointwise defined everywhere in $\mathbb{R}^2$;

ii. the support of $p_{\mathbf{c}}$ is such that, for any $\mathbf{c}_1, \mathbf{c}_2 \in \text{Support}\{p_{\mathbf{c}}\}$ and $\mathbf{c}_1 \neq \mathbf{c}_2$, the fourier transform $\mathcal{F}\{C_{\mathbf{c}_1} + C_{\mathbf{c}_2}\}$ is non-zero everywhere;

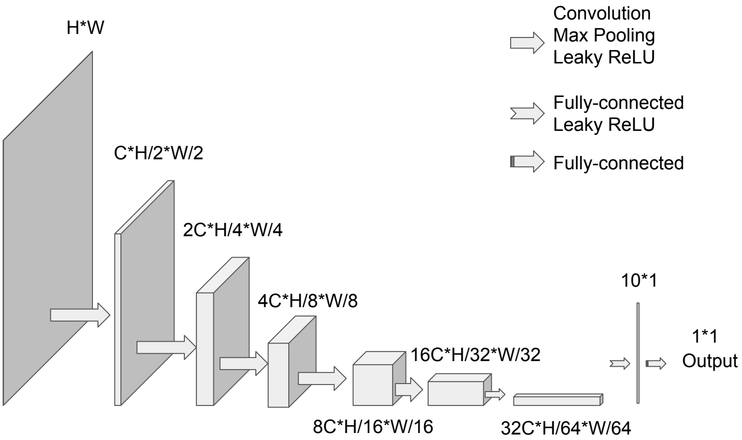

Figure 5.2: Architecture of the discriminator. The parameter for the channel size is $C = 96$ in every experiment. The input image with size $H \times W$ is successively processed and downsampled to output a scalar.

iii. the volume f is nonnegative everywhere and has a bounded support; and

iv. the probability distributions $p_{\boldsymbol{\theta}}$, $p_{\mathbf{c}}$, and $p_{\mathbf{t}}$ are bounded.

Before proving Theorem 5.5.5, we first comment on the assumptions. Assumption 1) is true for many noise distributions, including a white Gaussian noise filtered with any kernel of arbitrarily non-zero compact support. Assumption 2) is generally true as well. In fact, it is used to justify the application of Wiener filter to the clustered projections in classical Cryo-EM reconstruction pipelines. Assumption 3) is true since the volume represents the density map, which is nonnegative. Also, the biological structures considered in Cryo-EM have finite sizes.

In the following two sections, we build up to Theorem 5.5.5. First by analyzing the problem in the absence of the CTF and noise, and then by adding the CTF.

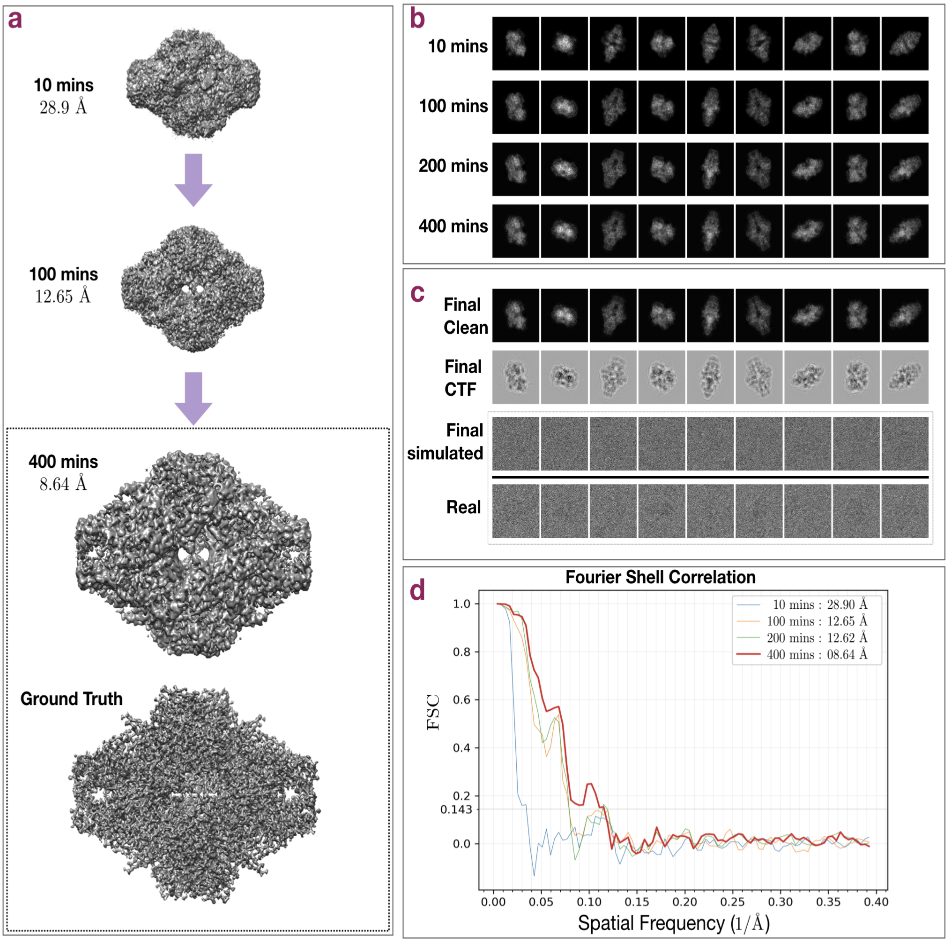

Figure 5.3: CryoGAN is applied on a synthetic projection dataset generated from a 2.5Å β-galactosidase volume. We refer to these synthetic projections as "real," in contrast to the projections coming from CryoGAN, which we term "simulated." **(a)** The volume is initialized with zeros and is progressively updated to produce projections whose distribution matches that of the real projections. **(b)** Evolution during training of some clean projections (*i.e.*, before CTF and noise) generated by the Cryo-EM physics simulator. **(c)** *Row 1*: Clean, simulated projections (before CTF and noise) generated at the final stage of training. *Row 2*: CTF-modulated simulated projections (before noise) generated at the final stage of training. *Row 4*: Real projections, for comparison. **(d)** FSC curves between the two reconstructed half-maps at different points during training.

Notations and Preliminaries

Let $SO(3)$ be the space of the special orthonormal matrices and $\mathcal{D}$ be the Borel $\sigma-$algebra induced using the standard Riemannian metric on $SO(3)$. Then, $(SO(3), \mathcal{D})$ describes the measurable space of orthonormal matrices. Let $\Delta_N^W = \{\mathbf{x} \in \mathbb{R}^N : \|\mathbf{x}\|_2 \leq W\}$ for some $W \in \mathbb{R}^+$. By $(\mathcal{L}_2, \mathcal{B})$, we denote the measurable space of all the square-integrable functions supported in Δ_2^W with Borel $\sigma-$algebra $\mathcal{B}$ induced by the L_2-norm. We denote by $\mathbb{F}$ the set of all the functions supported in Δ_3^W, which are nonnegative and essentially bounded.

For any $f \in \mathbb{F}$ and $\mathbf{A} \in SO(3)$, we denote $y = \mathcal{P}_{\mathbf{A}}\{f\} = \int_{-\infty}^{\infty} \mathbf{A}f(x_1, x_2, x_3)\, \mathrm{d}x_3$ where $\mathbf{A}f(\mathbf{x}) = f(\mathbf{A}^{-1}\mathbf{x})$. Let $p_{\mathbf{A}}$ be a probability density on the space $(SO(3), \mathcal{D})$. Note that there is a bijective mapping from $\boldsymbol{\theta}$ in Theorem 1 and $\mathbf{A}$. In fact, $\mathbf{A}$ represents the rotation matrix associated with the projection angle $\boldsymbol{\theta}$.

We denote by Ψ the normalized Haar measure on $(SO(3), \mathcal{D})$ and by $\Psi_{\mathbf{A}}$ the measure associated with $p_{\mathbf{A}}$ such that $\Psi_{\mathbf{A}}[\cdot] = \int_{(\mathbf{a}\in\cdot)} p_{\mathbf{A}}(\mathbf{a})\Psi[\,\mathrm{d}\mathbf{a}]$.

For a given $f \in \mathbb{F}$, the density $p_{\mathbf{A}}$ induces a probability measure $\mathbb{P}_{\text{proj}}(\cdot|f)$ on the space $(\mathcal{L}_2, \mathcal{B})$ through the mapping $\mathcal{P}_{\mathbf{A}}\{f\}$ such that

$$\mathbb{P}_{\text{proj}}(\cdot|f) = \Psi_{\mathbf{A}}[\{\mathbf{A} \in SO(3) : \mathcal{P}_{\mathbf{A}}\{f\} \in \cdot\}]. \tag{5.25}$$

When $p_{\mathbf{A}}$ is uniform on $SO(3)$, one has that

$$\mathbb{P}_{\text{proj}}(\cdot|f) = \mathbb{P}_{\text{proj}}(\cdot|\mathbf{R}f),\ \forall f \in \mathbb{F} \text{ and } \mathbf{R} \in O(3), \tag{5.26}$$

where O_3 is the space of all orthogonal matrices such that $\det \mathbf{A} \in \{-1, 1\}$. The invariance in (5.26) is true since

$$\begin{aligned}
\mathbb{P}_{\text{proj}}(\cdot|f) &= \Psi[\{\mathbf{A} \in SO(3) : \mathcal{P}_{\mathbf{A}}\{f\} = \cdot\}] \\
&= \Psi[\{\mathbf{A} \in SO(3) : \mathcal{P}_{\mathbf{R}^{-1}\mathbf{A}}\{\mathbf{R}f\} = \cdot\}] \\
&= \Psi[\{\mathbf{R}\boldsymbol{A'} \in SO(3) : \mathcal{P}_{\mathbf{A}'}\{\mathbf{R}f\} = \cdot\}] \\
&= \Psi[\{\mathbf{A}' \in SO(3) : \mathcal{P}_{\mathbf{A}'}\{\mathbf{R}f\} = \cdot\}]
\end{aligned} \tag{5.27}$$

where $\mathbf{A}' = \mathbf{R}^{-1}\mathbf{A}$ and the last equality follows from the right invariance of Haar measure. We define $G\{\mathbb{F}\} = \{\gamma_A : A \in O_3\}$ such that

$$(\gamma_A f)(\cdot) = f(\mathbf{A}^{-1}\cdot), \forall \mathbf{A} \in O(3), f \in \mathbb{F}. \tag{5.28}$$

We define the shape $[f]$ as an orbit of f under the influence of G such that $[f] = \{\gamma_{\mathbf{A}} f : \gamma_A \in G\}$. When $p_{\mathbf{A}}$ is uniform, the shape $[f]$ is composed of all the rotations and reflections of f.

5.5.1 Recovery in the Absence of CTF and Noise

In the absence of CTF and shifts the recoverability of $f : \mathbb{R}^3 \to \mathbb{R}$ from its 2D projections obtained at unknown random poses is guaranteed by [170, Theorem 3.1]. We first go through the notations described in [170] before we state the required foundational result. We then extend [170, Theorem 3.1] to the case when the CTF and shifts are present.

We can now restate [170, Theorem 3.1]. We discuss here the sketch of the proof given in [170].

Theorem 5.5.1 ([170, Theorem 3.1]). *Let $p_{\mathbf{A}}$ be any bounded distribution on $SO(3)$ and let the assumptions of Theorem 5.5.5 be true; then, $\forall f, g \in \mathbb{F}$,*

$$[f] \neq [g] \Rightarrow \mathbb{P}_{\text{proj}}(\cdot|f) \perp \mathbb{P}_{\text{proj}}(\cdot|g). \tag{5.29}$$

Sketch of the Proof. Without loss of generality, we provide the sketch of the proof for the case when $p_{\mathbf{A}}$ is uniform. For the case when $p_{\mathbf{A}}$ is nonuniform the argument remains the same provided that $\Psi_{\mathbf{A}}$ associated with the non-uniform distribution $p_{\mathbf{A}}$ is absolutely continuous with respect to Ψ ($\Psi_{\mathbf{A}} \ll \Psi$). This has been stated in [170]. Since we assume $p_{\mathbf{A}}$ to be bounded, this condition is satisfied. The only difference here with respect to the uniform distribution is that the orbit of f and g are more restricted than O(3).

The proof first uses in [171, Proposition 7.8] which we restate here as Proposition 5.5.2.

Proposition 5.5.2 ([171, Proposition 7.8]). *Let $f \in \mathbb{F}$ and let $S_{\mathbf{A}}$ be an uncountably infinite subset of $SO(3)$, then f is determined by the collection $\{\mathcal{P}_A\{f\}\}_{\mathbf{A} \in S_{\mathbf{A}}}$ ordered with respect to $\mathbf{A} \in S_{\mathbf{A}}$.*

Note that this proposition assumes that the angle of the projections are known. Although in our case the angles are unknown, we shall see that this proposition will be useful.

We now want to determine how different $\mathbb{P}_{\mathrm{proj}}(.|f)$ and $\mathbb{P}_{\mathrm{proj}}(.|g)$ are for any given f and g. For this, we use the equality

$$\mathrm{TV}(\mathbb{P}_1, \mathbb{P}_2) = 2 \inf_{\gamma \in \Pi(\mathbb{P}_1, \mathbb{P}_2)} \mathbb{E}_{(y_1, y_2) \sim \gamma}[\mathbb{1}_{y_1 \neq y_2}], \tag{5.30}$$

where TV is the total variation distance and $\Pi(\mathbb{P}_1, \mathbb{P}_2)$ is the set of all the joint distributions $\gamma(y_1, y_2)$ whose marginals are $\mathbb{P}_1$ and $\mathbb{P}_2$ [165]. In fact, $\mathbb{E}[\mathbb{1}_{y_1 \neq y_2}]$ is equal to the probability of the event $y_1 \neq y_2$. In our context, this translates into

$$\mathrm{TV}(\mathbb{P}_{\mathrm{proj}}(.|f), \mathbb{P}_{\mathrm{proj}}(.|g)) = 2 \inf_{\gamma \in \Pi(\mathbb{P}_{\mathrm{proj}}(.|f), \mathbb{P}_{\mathrm{proj}}(.|g))} \mathrm{Prob}(y_1 \neq y_2), \text{ where } (y_1, y_2) \sim \gamma. \tag{5.31}$$

The optimum is achieved at the extremas which are sparse joint distributions and are such that the variable y_2 is a function of y_1. For any arbitrary joint distribution (or coupling) of this form, the proof then assigns a measurable function $h : SO(3) \rightarrow SO(3)$ such that $(y_1, y_2) = (\mathcal{P}_{\mathbf{A}}\{f\}, \mathcal{P}_{h(\mathbf{A})}\{g\})$ for $\mathbf{A} \sim p_{\mathbf{A}}$.
We can then write that

$$\Psi[\{\mathbf{A} \in SO(3) : \mathcal{P}_{h(\mathbf{A})}\{g\} \in \cdot\}] = \mathbb{P}_{\mathrm{proj}}(\cdot|g). \tag{5.32}$$

The task now is to estimate $\mathrm{Prob}(y_1 \neq y_2)$, where $(y_1, y_2) = (\mathcal{P}_{\mathbf{A}}\{f\}, \mathcal{P}_{h(\mathbf{A})}\{g\})$ for $\mathbf{A} \sim p_{\mathbf{A}}$.
(Continuous h). When h is continuous, Proposition 5.5.2 implies that, if $[f] \neq [g]$, then

$$\Psi[\{\mathbf{A} \in SO(3) : \|\mathcal{P}_{\mathbf{A}}\{f\} - \mathcal{P}_{h(\mathbf{A})}\{g\}\|_2 > 0\}] = 1. \tag{5.33}$$

(General h). When the function h is discontinuous, the proof uses Lusin's theorem to approximate h by a continuous function. Lusin's theorem states that, for any $\delta > 0$, there exists an h_δ such that $h(\mathbf{A}) = h_\delta(\mathbf{A}), \forall \mathbf{A} \in \mathcal{H}_\delta$ and $\Psi[SO(3)|\mathcal{H}_\delta] < \delta$. This then leads to

$$\begin{aligned}\Psi[\{\mathbf{A} \in SO(3) : \|\mathcal{P}_{\mathbf{A}}\{f\} - \mathcal{P}_{h(\mathbf{A})}\{g\}\|_2 > 0\}] &\geq \Psi(H_\delta) \\ &\geq 1 - \delta. \end{aligned} \tag{5.34}$$

Since δ is arbitrarily small, the event $\{\mathcal{P}_{\mathbf{A}}\{f\} \neq \mathcal{P}_{h(\mathbf{A})}\{g\}\}$ has probability 1.

In conclusion, for any arbitrary coupling, the event $\{\mathcal{P}_{\mathbf{A}}\{f\} \neq \mathcal{P}_{h(\mathbf{A})}\{g\}\}$ has probability 1 if $[f] \neq [g]$. This implies that, when $[f]$ and $[g]$ are not the same, the total-variation distance between $\mathbb{P}_{\text{proj}}(\cdot|f)$ and $\mathbb{P}_{\text{proj}}(\cdot|g)$ is 2. This ensures that the two probability measures are mutually singular meaning that the intersection of their support has zero measure. This concludes the proof.

5.5.2 Recovery in the Presence of CTF and Absence of Noise

We now extend the previous result to the case when the CTF is present. For the sake of simplicity we do not take into account shifts in the forward model. However, it is trivial to generalize the results to them since, shifts don't change the information content in the projections but only their location.

We assume that $\mathbf{c} \sim p_{\mathbf{c}}$ such that the support of $p_{\mathbf{c}}$ is in some bounded region $\mathcal{C} \subset \mathbb{R}^3$. We denote $\Psi_{\mathbf{c}}[\cdot]$ as the measure associated with $p_{\mathbf{c}}$ on the space $\mathcal{C}$.

We denote by $(SO(3) \times \mathcal{C})$ the product space of $SO(3)$ and $\mathcal{C}$, while we denote by $\Psi_{\mathbf{A},\mathbf{c}}$ the measure on this product space. We then define

$$\mathbb{P}_{\text{proj,CTF}}(\cdot|f) = \Psi_{\mathbf{A},\mathbf{c}}[\{(\mathbf{A},\mathbf{c}) \in (SO(3) \times \mathcal{C}) : C_{\mathbf{c}} * \mathcal{P}_{\mathbf{A}}\{f\} \in \cdot\}], \tag{5.35}$$

where $C_{\mathbf{c}}$ is the space-domain CTF given in (5.4).

Theorem 5.5.3. *Let $p_{\mathbf{A}}$ be a bounded probability distribution on $SO(3)$, $p_{\mathbf{c}}$ be a distribution of the CTF with parameters $\mathbf{c} \in \mathcal{C}$, and let the assumptions of Theorem 5.5.5 be true; then, $\forall f, g \in \mathbb{F}$,*

$$[f] \neq [g] \Rightarrow \mathbb{P}_{\text{proj,CTF}}(\cdot|f) \perp \mathbb{P}_{\text{proj,CTF}}(\cdot|g). \tag{5.36}$$

Proof. Similarly to the previous proof, we show that the TV distance between $\mathbb{P}_{\text{proj,CTF}}(\cdot|f)$ and $\mathbb{P}_{\text{proj,CTF}}(\cdot|g)$ is 2 when $[f]$ and $[g]$ are distinct. For simplification, we assume that $p_{\mathbf{A}}$ is uniform. (When this is not the case the proof essentially remains the same.) We need to show that $\text{Prob}(y_1 \neq y_2) = 1$, where $(y_1, y_2) \sim \gamma$ for any arbitrary coupling γ of $\mathbb{P}_{\text{proj,CTF}}(\cdot|f)$ and $\mathbb{P}_{\text{proj,CTF}}(\cdot|g)$. For an arbitrary coupling γ such that $\text{Prob}(y_1 \neq y_2)$ is minimum, we again assign $h : (SO(3) \times \mathcal{C}) \rightarrow (SO(3) \times \mathcal{C})$ such that

$$(y_1, y_2) = (C_{\mathbf{c}} * \mathcal{P}_{\mathbf{A}}\{f\}, C_{h_1(\mathbf{A},\mathbf{c})} * \mathcal{P}_{h_0(\mathbf{A},\mathbf{c})}\{g\}), \tag{5.37}$$

where $\mathbf{A} \sim p_{\mathbf{A}}, \mathbf{c} \sim p_{\mathbf{c}}$ and where $h_0 : (SO(3)\times\mathcal{C}) \to SO(3)$ and $h_1 : (SO(3)\times\mathcal{C}) \to \mathcal{C}$ are such that $h(\mathbf{A},\mathbf{c}) = (h_0(\mathbf{A},\mathbf{c}), h_1(\mathbf{A},\mathbf{c}))$. This implies that

$$\mathbb{P}_{\text{proj,CTF}}(\cdot|g) = \Psi_{\mathbf{A},\mathbf{c}}[\{(\mathbf{A},\mathbf{c}) \in (SO(3)\times\mathcal{C}) : C_{h_1(\mathbf{A},\mathbf{c})} * \mathcal{P}_{h_0(\mathbf{A},\mathbf{c})}\{g\} \in \cdot\}]. \quad (5.38)$$

We now show that, for any h, the event $\{y_1 \neq y_2\}$ has probability 1.
(Continuous h). We first assume that h is continuous and use the same kind of technique as in the proof of [170, Theorem 3.1].

Since $SO(3)$ is transitive, we can write that

$$h(\mathbf{A},\mathbf{c}) = (\mathbf{A}\Gamma_{\mathbf{A},\mathbf{c}}, h_1(\mathbf{A},\mathbf{c})). \quad (5.39)$$

As h is continuous, so is $\Gamma_{\mathbf{A},\mathbf{c}}$. Let $\{\mathcal{A}_n^m \times \mathcal{C}_n^m\}_{m=1}^n$ be a collection of n disjoint sets which creates the partition of $(SO(3)\times\mathcal{C})$. These partitions are such that for any m, there exists a k_m such that $\{\mathcal{A}_{n+1}^m \times \mathcal{C}_{n+1}^m\} \subset \{\mathcal{A}_n^{k_m} \times \mathcal{C}_n^{k_m}\}$. This means that, as n increases, the partitions become finer. We now define

$$h_n(\mathbf{A},\mathbf{c}) = (\mathbf{A}\Gamma_n^m, h_{n,1}^m(\mathbf{A},\mathbf{c}))\ \forall\ (\mathbf{A},\mathbf{c}) \in \{\mathcal{A}_n^m \times \mathcal{C}_n^m\}, \quad (5.40)$$

such that

$$\Gamma_n^m = \underset{\Gamma\in\{\Gamma_{\mathbf{A},\mathbf{c}}:(\mathbf{A},\mathbf{c})\in\{\bar{\mathcal{A}}_n^m\times\bar{\mathcal{C}}_n^m\}\}}{\arg\min}\ \min_{(\mathbf{A},\mathbf{c})\in\{\bar{\mathcal{A}}_n^m\times\bar{\mathcal{C}}_n^m\}} \|\mathcal{P}_{\mathbf{A}}\{f\} - \mathcal{P}_{\mathbf{A}\Gamma}\{g\}\|, \quad (5.41)$$

where $\bar{\mathcal{A}}_n^m$ and $\bar{\mathcal{C}}_n^m$ are the closures of $\mathcal{A}_n^m$ and $\mathcal{C}_n^m$, respectively. The sequence h_n converge to h as $n \to \infty$. We denote

$$\mathcal{K} = \{(\mathbf{A},\mathbf{c}) \in (SO(3)\times\mathcal{C}) : \|C_{\mathbf{c}} * \mathcal{P}_{\mathbf{A}}\{f\} - C_{h_1(\mathbf{A},\mathbf{c})} * \mathcal{P}_{\mathbf{A}\Gamma_{\mathbf{A}}}\{g\}\| > 0\}, \quad (5.42)$$

$$\mathcal{K}_n = \{(\mathbf{A},\mathbf{c}) \in (\mathcal{A}_n^m \times \mathcal{C}_n^m) : \|C_{\mathbf{c}} * \mathcal{P}_{\mathbf{A}}\{f\} - C_{h_1(\mathbf{A},\mathbf{c})} * \mathcal{P}_{\mathbf{A}}\{\Gamma_n^m g\}\| > 0\}. \quad (5.43)$$

Similarly to [170, Theorem 3.1], we can then show that

$$\Psi_{\mathbf{A},\mathbf{d}}[\mathcal{K}] = \lim_{n\to\infty} \sum_{m=1}^{m=n} \Psi_{\mathbf{A},\mathbf{d}}[\mathcal{K}_m]. \quad (5.44)$$

We invoke Proposition 5.5.4, which gives that $\Psi_{\mathbf{A},\mathbf{c}}[\mathcal{K}_n] = \Psi_{\mathbf{A},\mathbf{c}}[(\mathcal{A}_n^m \times \mathcal{C}_n^m)]$. Therefore, $\Psi_{\mathbf{A},\mathbf{d}}[\mathcal{K}] = \Psi_{\mathbf{A},\mathbf{c}}[(SO(3)\times\mathcal{C})] = 1$. This means that, when h is continuous, the event $\{y_1 \neq y_2\}$ has probability 1 if $[f] \neq [g]$.

(General h). When h is discontinuous, we can invoke Lusin's theorem to claim the same, similarly to Theorem 5.5.1. This means that, for any h, if $[f] \neq [g]$, then the probability of the event $\{y_1 \neq y_2\}$ is 1. Therefore, the TV distance between $\mathbb{P}_{\text{proj,CTF}}(\cdot|f)$ and $\mathbb{P}_{\text{proj,CTF}}(\cdot|g)$ is 2, yielding that $\mathbb{P}_{\text{proj,CTF}}(\cdot|f) \perp \mathbb{P}_{\text{proj,CTF}}(\cdot|g)$. This concludes the proof.

Proposition 5.5.4. *Let $f, g \in \mathbb{F}$, $\mathcal{A}' \subseteq SO(3)$, $\mathcal{C}' \subseteq \mathcal{C}$, $\Gamma \in SO(3)$, and*

$$\mathcal{K}' = \{(\mathbf{A}, \mathbf{c}) \in (\mathcal{A}' \times \mathcal{C}') : \|C_{\mathbf{c}} * \mathcal{P}_{\mathbf{A}}\{f\} - C_{h_1(\mathbf{A},\mathbf{c})} * \mathcal{P}_{\mathbf{A}}\{\Gamma g\}\| > 0\}. \tag{5.45}$$

Let the assumptions from Theorem 5.5.5 be true. Then, if $[f] \neq [g]$, it holds that

$$\Psi_{\mathbf{A},\mathbf{c}}[\mathcal{K}'] = \Psi_{\mathbf{A},\mathbf{c}}[(\mathcal{A}' \times \mathcal{C}')]. \tag{5.46}$$

Proof. We show that $\Psi_{\mathbf{A},\mathbf{c}}[\mathcal{K}'^c] = 0$, where $(\mathcal{K}'^c \cup \mathcal{K}') = (\mathcal{A}' \times \mathcal{C}')$. We define the set $S_{\mathbf{A}} = \{\mathbf{c} \in \mathcal{C}' : \|C_{\mathbf{c}} * \mathcal{P}_{\mathbf{A}}\{f\} - C_{h_1(\mathbf{A},\mathbf{c})} * \mathcal{P}_{\mathbf{A}}\{\Gamma g\}\| = 0\}$. We define $S_{\mathcal{A}''} = \cup_{\mathbf{A} \in \mathcal{A}''} S_{\mathbf{A}}$ for any $\mathcal{A}'' \subseteq \mathcal{A}'$. We define

$$\mathcal{A}'_1 = \{\mathbf{A} \in \mathcal{A}' : S_{\mathbf{A}} \text{ is an uncountable set}\}, \tag{5.47}$$

$$\mathcal{A}'_2 = \{\mathbf{A} \in \mathcal{A}' : S_{\mathbf{A}} \text{ is a countable non-empty set}\}. \tag{5.48}$$

Note that $\mathcal{K}'^c = \cup_{k=1}^{2} \cup_{\mathbf{A} \in \mathcal{A}'_k} (\mathbf{A} \times S_{\mathbf{A}})$. Then,

$$\Psi_{\mathbf{A},\mathbf{c}}[\{\mathcal{K}'^c] = \sum_{k=1}^{2} \Psi_{\mathbf{A},\mathbf{c}}[\cup_{\mathbf{A} \in \mathcal{A}'_k} (\mathbf{A} \times S_{\mathbf{A}})\}] \tag{5.49}$$

We now look at the two cases.

- *(When $S_{\mathbf{A}}$ is uncountable).* For this case, we show that $\Psi[\mathcal{A}'_1] = 0$. The main argument is that if this is not true, then it contradicts $[f] \neq [g]$.

  For the sake of conciseness, we denote $\mathcal{P}_{\mathbf{A}}\{f\}$ by I_f and $\mathcal{P}_{\mathbf{A}}\{\Gamma g\}$ by I_g. For any $\mathbf{A} \in \mathcal{A}'_1$, it holds that

  $$C_{\mathbf{c}} * I_f = C_{h_1(\mathbf{A},\mathbf{c})} * I_g, \ \forall \mathbf{c} \in S_{\mathbf{A}}, \tag{5.50}$$

  $$\widehat{C}_{\mathbf{c}} \cdot \hat{I}_f = \widehat{C}_{h_1(\mathbf{A},\mathbf{c})} \cdot \hat{I}_g, \ \forall \mathbf{c} \in S_{\mathbf{A}}, \tag{5.51}$$

  where $\widehat{C}, \hat{I}_f$, and $\hat{I}_g$ are the Fourier transforms of $\mathbf{C}, I_f$, and I_g, respectively.

We define $ze(\hat{I}) = \{\boldsymbol{\omega} \in \mathbb{R}^2 : \widehat{I}(\boldsymbol{\omega}) = 0\}$, $\omega_\alpha = \{[(r\cos\alpha, r\sin\alpha)] : r > 0\}$, and $ze_\alpha(\hat{I}) = ze(\hat{I}) \cap \omega_\alpha$. From (5.51), we can write that

$$ze(\widehat{C}_{\mathbf{c}}) \cup ze(\hat{I}_f) = ze(\widehat{C}_{h_1(\mathbf{A},\mathbf{c})}) \cup ze(\hat{I}_g), \ \forall \mathbf{c} \in S_{\mathbf{A}}. \tag{5.52}$$

Two remarks are in order. Firstly, by assumption 2 of Theorem 5.5.5, $ze(\widehat{C}_{\mathbf{c}_1}) \cap ze(\widehat{C}_{\mathbf{c}_2}) = \emptyset$ for $\mathbf{c}_1 \neq \mathbf{c}_2$. (Remember that $ze_\alpha(\widehat{C}_{\mathbf{c}})$ for any $\alpha \in [0, \pi]$ is nonempty (see "Image Formation Theory).) Secondly, by assumption 3 of Theorem 5.5.5, the supports of f and g are compact and nontrivial, so are the supports of I_f and I_g. This means that their Fourier transforms $\hat{I}_f$ and $\hat{I}_g$ are analytic functions, which implies that there are infinitely many α such that the cardinality of the sets $ze_\alpha(\hat{I}_f)$ and $ze_\alpha(\hat{I}_f)$ is countable. We call the set of such α as $\mathcal{S}_\alpha$. Now, we have that

$$\begin{aligned}
&ze_\alpha(\widehat{C}_{\mathbf{c}}) \cap (ze_\alpha(\widehat{C}_{\mathbf{c}}) \cup ze_\alpha(\hat{I}_f)) = ze_\alpha(\widehat{C}_{\mathbf{c}}) \cap (ze_\alpha(\widehat{C}_{h_1(\mathbf{A},\mathbf{c})}) \cup ze_\alpha(\hat{I}_g)),\\
&ze_\alpha(\widehat{C}_{\mathbf{c}}) \cup (ze_\alpha(\widehat{C}_{\mathbf{c}}) \cap ze_\alpha(\hat{I}_f)) = (ze_\alpha(\widehat{C}_{\mathbf{c}}) \cap ze_\alpha(\widehat{C}_{h_1(\mathbf{A},\mathbf{c})})) \cup (ze_\alpha(\widehat{C}_{\mathbf{c}}) \cap ze_\alpha(\hat{I}_g)),\\
&ze_\alpha(\widehat{C}_{\mathbf{c}}) \cup (ze_\alpha(\widehat{C}_{\mathbf{c}}) \cap ze_\alpha(\hat{I}_f)) = ze_\alpha(\widehat{C}_{\mathbf{c}}) \cap ze_\alpha(\hat{I}_g) \qquad (5.53)
\end{aligned}$$

for all $\mathbf{c} \in S_{\mathbf{A}}$ and $\alpha \in [0, \pi]$.

We can now write that

$$\cup_{\mathbf{c} \in S_{\mathbf{A}}} ze_\alpha(\widehat{C}_{\mathbf{c}}) \cup (ze_\alpha(\widehat{C}_{\mathbf{c}}) \cap ze_\alpha(\hat{I}_f)) = \cup_{\mathbf{c} \in S_{\mathbf{A}}} ze_\alpha(\widehat{C}_{\mathbf{c}}) \cap ze_\alpha(\hat{I}_g). \tag{5.54}$$

for any $\alpha \in S_\alpha$. The set on the left hand side of (5.54) has an uncountably infinite cardinality since there are uncountably many $\mathbf{c} \in \mathcal{S}_{\mathbf{A}}$ and for each $\mathbf{c}$ there are distinct $ze_\alpha(\widehat{C}_{\mathbf{c}})$. In return, the set in the right hand side of (5.54) is countable for a given $\alpha \in S_\alpha$. Therefore, for any $\alpha \in S_\alpha$, the two sets have different cardinality, which raises a contradiction. The only possible scenario in which (5.52) is true is when $h_1(\mathbf{A}, \mathbf{c}) = \mathbf{c}$. Using (5.51), we infer that $\mathcal{P}_{\mathbf{A}}\{f\} = \mathcal{P}_{\mathbf{A}}\{\Gamma g\}$. Therefore, for any $\mathbf{A} \in \mathcal{A}_1'$, $\mathcal{P}_{\mathbf{A}}\{f\} = \mathcal{P}_{\mathbf{A}}\{\Gamma g\}$. However, $\Psi[\mathcal{A}_1'] = 0$ since, if this is not true, then $[f] = [g]$ by Proposition 5.5.2.

Now note that

$$\begin{aligned}
\Psi_{\mathbf{A},\mathbf{c}}[\cup_{\mathbf{A} \in \mathcal{A}_1'}(\mathbf{A} \times S_{\mathbf{A}})] &\leq \underbrace{\Psi[\mathcal{A}_1']}_{0} \underbrace{\Psi_{\mathbf{c}}[\cup_{\mathbf{A} \in \mathcal{A}_1'} \mathcal{S}_{\mathbf{A}}]}_{\text{finite}}\\
&= 0. \qquad (5.55)
\end{aligned}$$

- *(When $S_{\mathbf{A}}$ is countable and nonempty).* Since $S_{\mathbf{A}}$ is a countable set in this case, its elements have a bijection with natural numbers. We denote this bijection by $b : \mathbb{Z} \times \mathcal{A}'_2 \to S_{\mathbf{A}}$. We denote by $q(z) = \cup_{\mathbf{A} \in \mathcal{A}'_2} (\mathbf{A}, b_{\mathbf{A}}(z))$, $\forall z \in \mathbb{Z}$. Note that $q(z)$ is a graph of the function $b(z, \cdot)$. Since it is a graph, $\Psi_{\mathbf{A},\mathbf{c}}[q(z)] = 0$.

  We also have that $\Psi_{\mathbf{A},\mathbf{c}}[\cup_{\mathbf{A} \in \mathcal{A}'_2} (\mathbf{A} \times S_{\mathbf{A}})] = \Psi_{\mathbf{A},\mathbf{c}}[\sum_{z \in \mathbb{Z}} q(z)]$. The latter vanishes since it is the measure of a countable addition of sets of measure zero. Hence, $\Psi_{\mathbf{A},\mathbf{c}}[\cup_{\mathbf{A} \in \mathcal{A}'_2} (\mathbf{A} \times S_{\mathbf{A}})] = 0$.

This gives that $\Psi_{\mathbf{A},\mathbf{c}}[\mathcal{K}'^c] = \sum_{k=1}^{2} \Psi_{\mathbf{A},\mathbf{c}}[\cup_{\mathbf{A} \in \mathcal{A}'_k} (\mathbf{A} \times S_{\mathbf{A}})] = 0$, which concludes the proof.

5.5.3 Recovery in the presence of CTF and Noise

Theorem 5.5.5. *Let $y = H_{\boldsymbol{\varphi}} f + n$ as given in (5.4) with $\boldsymbol{\varphi} = (\boldsymbol{\theta}, \mathbf{t}, \mathbf{c})$, where $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3)$ are the projection angles, $\mathbf{t} = (t_1, t_2)$ are the shifts, and $\mathbf{c} = (d_1, d_2, \alpha_{\mathrm{ast}})$ are the CTF parameters (defocus-major, defocus-minor, and angle of astigmatism, respectively), f is the continuous-domain 3D volume, and y, n are continuous-domain 2D images.*

Then, it holds that

$$\mathbb{P}(\cdot|f_1) = \mathbb{P}(\cdot|f_2) \Leftrightarrow f_1 = G(f_2), \tag{5.56}$$

where G is some member of the set of rotation-reflection operations.

Proof of Theorem 5.5.5. We denote the probability measure of $y_{\mathrm{noiseless}} = H_{\boldsymbol{\varphi}} f$ with $\mathbb{P}_{\mathrm{noiseless}}(\cdot|f)$. We shall prove the following in sequence:

i. $\mathbb{P}(\cdot|f_1) = \mathbb{P}(\cdot|f_2) \Leftrightarrow \mathbb{P}_{\mathrm{noiseless}}(\cdot|f_1) = \mathbb{P}_{\mathrm{noiseless}}(\cdot|f_2)$,

ii. $\mathbb{P}_{\mathrm{noiseless}}(\cdot|f_1) = \mathbb{P}_{\mathrm{noiseless}}(\cdot|\mathbf{f}_2) \Leftrightarrow f_2 = G(f_1)$.

For the first part we progress by noting that $y = y_{\mathrm{noiseless}} + n$. Recall that the characteristic function of the probability measure associated to the sum of two random variables is the product of their characteristic functions. Mathematically,

$$\hat{\mathbb{P}}(\cdot|f) = \hat{\mathbb{P}}_{\mathrm{noiseless}}(\cdot|f)\, \hat{\mathbb{P}}_n. \tag{5.57}$$

By Assumption (1, Theorem 5.5.5), we can now write that

$$\hat{\mathbb{P}}_{\text{noiseless}}(\cdot|f) = \frac{\hat{\mathbb{P}}(\cdot|f)}{\hat{\mathbb{P}}_n}. \tag{5.58}$$

From (5.58), it is easy to see that $\mathbb{P}(\cdot|f_1) = \mathbb{P}(\cdot|f_2) \Leftrightarrow \mathbb{P}_{\text{noiseless}}(\cdot|f_1) = \mathbb{P}_{\text{noiseless}}(\cdot|f_2)$. This concludes the first part.

For the second part, we now invoke the result from Theorem 5.5.3. It states that if $f_2 \neq G(f_1)$ for any G in the set of rotation and reflection operation, then the corresponding $\mathbb{P}_{\text{noiseless}}(\cdot|f_1)$ and $\mathbb{P}_{\text{noiseless}}(\cdot|f_2)$ are mutually singular. This means that the support of their intersection has zero measure. Since, we have $\mathbb{P}_{\text{noiseless}}(\cdot|f_1) = \mathbb{P}_{\text{noiseless}}(\cdot|f_2)$ this means they are not mutually singular. This implies that $f_2 = G(f_1)$ for some G in the set of rotation and reflection operation. This concludes the proof.

Discrete-domain Extension. In practice, the Cryo-EM measurements are acquired on a detector grid and are therefore discrete. CryoGAN reconstructs a voxel-domain $\mathbf{x}_{\text{rec}}$ such that $p(\mathbf{y}|\mathbf{x}_{\text{rec}}) = p(\mathbf{y}|\mathbf{x}_{\text{GT}})$, meaning that the corresponding probability measures are equal as well. Theorem 5.5.5 then holds approximately, up to some error that results from discretization of the measurements, forward model, and 3D density map. We leave a more thorough analysis of this error for future work.

5.6 Related Works

Related works fall into two main categories: current Cryo-EM reconstruction methods and deep learning techniques that may apply to the Cryo-EM pipeline; we now discuss each of these.

Cryo-EM Reconstruction

The main challenge in Cryo-EM reconstruction is that every particle has an unknown pose in its micrograph—if the poses were known, maximum-likelihood (ML) or maximum *a posteriori* (MAP) estimation of the volume could be performed by solving a standard linear inverse problem, where robustness would result from the large number of measurements which would counteract the low SNR of each measurement. One approach is to attempt to estimate the unknown poses iteratively. Pose estimation can be achieved with a variety of strategies, including the popular projection-matching approach [172, 173]. Whatever the method used, pose estimation is challenging because the SNR of individual projection images is extremely low. It also requires the estimation of additional parameters and the projection of the current reconstructed volume at a large number of poses and at every iteration of the reconstruction pipeline; ultimately, this is very computationally demanding.

Another approach is to formulate the reconstruction as a ML (or MAP) estimation problem in which the unknown poses are marginalized away [151, 153, 174]. This is attractive in that no extra parameter need to be estimated. The problem can then be solved using the expectation-maximization algorithm (e.g., [151, 153]), where marginalization over poses during the so-called E-step is computationally expensive. Alternatively, the problem can be minimized using stochastic gradient descent (e.g., during the *ab initio* phase of [153]); here, the challenge is that the involved gradients require computations over all poses. For a more in-depth discussion, see [156, 157, 175]. For additional mathematical details on the relationship between likelihood-based methods and CryoGAN, see *CryoGAN vs. Likelihood-based Methods*.

Likelihood-free methods for Cryo-EM reconstruction are relatively few. An early approach is [176], which proposes to reconstruct an *ab initio* structure such that the first few moments of the distribution of its theoretical Cryo-EM measurements match the ones of the particles. However, the method assumes that the poses of the particles have a uniform distribution. This moment-matching technique has been recently extended in [177] to reconstruct an *ab initio* structure in the case of nonuniform pose distributions.

By contrast, our CryoGAN framework proposes to match the distribution of the theoretical Cryo-EM measurements to that of the real projections, by which we mean all the moments and not just the first few. Moreover, our method works for any pose distribution of the particles provided it is known beforehand. Alternatively, one could rely on a parametric model of the pose distribution and use the backpropagation mechanism of neural networks to learn its parameters during the CryoGAN run; this technique is explored in [177].

Deep Learning for Cryo-EM

Deep learning has already had a profound impact in a wide range of image-reconstruction applications [178–180]; however, their current utilization in Cryo-EM is mostly restricted to preprocessing steps such as micrograph denoising [181] or particle picking [182–186]. A recent work uses neural networks to model continuous generative factors of structural heterogeneity [187]. However, the algorithm necessitates a pose-estimation procedure that relies on a merely conventional approach. Another recent work [188] uses a variational autoencoder trained using a discriminator-based objective to find a low-dimensional latent representation of the particles. These representations are then used to estimate the poses.

Deep learning is now extensively used to solve inverse problems in imaging [17, 21, 26, 179]. However, most methods are based on supervised learning and thus rely on training data. A GAN-based scheme that recovers the underlying distribution of the data from its noisy partial observations through a forward model was recently proposed in [189]. Finally, the reconstruction of a 3D structure (implicitly or explicitly) from its 2D viewpoints (and not projections) is an important problem in computer vision [190]. Many recent deep-learning algorithms have been used in this regard [191, 192]. While this problem is ostensibly similar to Cryo-EM reconstruction, the measurement model for these problems is much less complicated than it is for Cryo-EM and bares no relation to this modality.

5.6.1 CryoGAN vs. Likelihood-based Methods

As discussed in [157], most reconstruction approaches in SPA currently rely a on maximum-likelihood (ML) formulation [174] written as

$$\mathbf{x}_{\mathrm{rec}} = \arg\max_{\mathbf{x}} \sum_{n=1}^{N} \log p(\mathbf{y}_{\mathrm{data}}^{n}|\mathbf{x}). \tag{5.59}$$

This formulation can be solved by using an expectation-maximization (ML-EM) algorithm or gradient descent. The former is preferred for iterative refinement [151], while a stochastic version of the latter is used for generating initial volume estimates in [153]. These techniques all involve the likelihood-estimation for each projection given the current volume estimate, which requires marginalizing the joint distribution $p(\mathbf{y}, \boldsymbol{\theta}|\mathbf{x})$ over the space of poses. In ML-EM, this is side-stepped in the expectation step (E-step) by computing the conditional distribution on the poses for each projection given the volume estimate. This conditional distribution is then used to update the volume in the marginalization step (M-step). Hence, all ML techniques require, either implicitly or explicitly, computations over a large number of poses for each projection.

For another perspective, ML techniques can be viewed as distribution-matching approaches. Specifically, (5.59) minimizes an empirical estimate of the Kullback-Leibler (KL) divergence between the real distribution $p_{\text{data}}(\mathbf{y})$ and the simulated distribution $p(\mathbf{y}|\mathbf{x})$, such that

$$\mathbf{x}_{\text{rec}} = \arg\min_{\mathbf{x}} \text{KL}(\text{p}_{\text{data}}(\mathbf{y})||\text{p}(\mathbf{y}|\mathbf{x})) \tag{5.60}$$

$$= \arg\min_{\mathbf{x}} \mathbb{E}_{\mathbf{y}\sim p_{\text{data}}} \left[\log \frac{p_{\text{data}}(\mathbf{y})}{p(\mathbf{y}|\mathbf{x})}\right] \tag{5.61}$$

$$= \arg\min_{\mathbf{x}} \mathbb{E}_{\mathbf{y}\sim p_{\text{data}}} \left[-\log p(\mathbf{y}|\mathbf{x})\right] \tag{5.62}$$

$$\approx \arg\max_{\mathbf{x}} \sum_{n=1}^{N} \log p(\mathbf{y}^n_{\text{data}}|\mathbf{x}). \tag{5.63}$$

Hence, ML methods aim at finding the reconstruction whose simulated projection distribution matches that of the real data. In practice, this specific goal can be achieved by minimizing any suitable distance between these two distributions. By changing the distance, one can avoid the challenging likelihood computations that are inherent to the current ML methods, while preserving the theoretical guarantees that come with distribution matching (*e.g.*, Theorem 5.5.5).

This is precisely the philosophy behind CryoGAN, which relies on the Wasserstein distance to formulate the distribution-matching task as

$$\mathbf{x}_{\text{rec}} = \operatorname{argmin}_{\mathbf{x}} \max_{\mathbf{D}_\phi:\|\mathbf{D}_\phi\|_L\leq 1} \Big(\mathbb{E}_{\mathbf{y}\sim p_{\text{data}}}[\mathbf{D}_\phi(\mathbf{y})] - \mathbb{E}_{\mathbf{y}\sim p_{\mathbf{x}}}[\mathbf{D}_\phi(\mathbf{y})]\Big). \tag{5.64}$$

In this formulation, the two distributions indirectly interact through a common function: the discriminator network D_ϕ. As a result, likelihood estimation is avoided; only a reliable sampler for each of the two distributions [193] is required. The samples from the real data distribution are readily available in the form of the acquired projection dataset, while the ones from the simulated distribution are generated by the Cryo-EM physics simulator. Hence, CryoGAN is a likelihood-free technique with theoretical properties that are at least as good—if not better—than the ML ones. In fact, the Wasserstein distance is often easier to minimize than the KL divergence (*e.g.*, due to the smoothness of the former) [164].

5.7 Results

5.7.1 Results on Synthetic Data

We first assessed the viability and performance of CryoGAN on a synthetic dataset that consists of 41,000 β-galactosidase projections, designed to mimic the EMPIAR-10061 data [158] in terms of noise level and CTF parameters. To create this dataset, we generated a 2.5 Å-resolution density map from the PDB entry (*5a1a*) of the protein and applied the forward model described in *Online Methods* to obtain projections modulated by CTF effects and corrupted by noise. We then randomly divided this dataset in two and applied the CryoGAN algorithm separately on both halves to generate half-maps. In the context of this experiment, we refer to these synthetic projections as "real," in contrast to the projections coming from CryoGAN, which we term "simulated." The details of the experimental setup are given in Appendix A.3.

We ran the CryoGAN algorithm for 400 minutes on an NVIDIA V100 GPU and obtained a reconstruction with a resolution of 8.64 Å (Figure 5.3a). Starting from a zero-valued volume, CryoGAN progressively updates the 3D structure so that its simulated projections (Figure 5.3b) reach a distribution that matches that of the real projections. These gradual updates are at the core of the deep adversarial learning scheme of CryoGAN. At each iteration of the algorithm, the gradients from the discriminator carry information about the current difference between the real projections and the simulated projections. These gradients are used by the Cryo-EM physics simulator to update the volume so as to improve the fidelity of the simulated projections. Hence, at the end of its run, the volume learned by CryoGAN has simulated projections (Figure 5.3.c, Rows 1-3) that are similar to the real projections (Figure 5.3.c, Row 4) in a *distributional* sense. The evolution of the Fourier-shell correlation (FSC) between the reconstructed half-maps (Figure 5.3.d) testifies to the progressive increase in resolution that derives from this adversarial learning scheme.

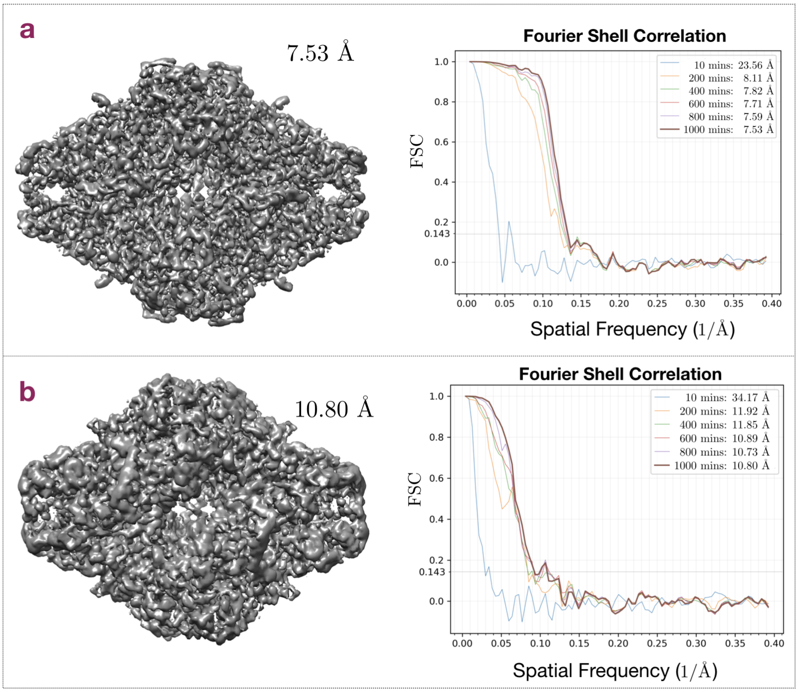

Figure 5.4: Additional CryoGAN reconstructions for synthetic datasets with different imaging conditions. **(a)** Reconstruction for a low noise case (-5.2 dB SNR). The corresponding evolution of the FSC curves with time is shown on the right. **(b)** Reconstruction for a realistic noise level (-20 dB SNR) and with translations (3% of the image size) in the data. The corresponding evolution of the FSC curves with time is shown on the right.

5.7.2 Results on Additional Synthetic Data

The main synthetic experiment given in Figure 5.3 considers imaging conditions with a realistic noise level (-20 dB SNR). We performed additional experiments on synthetic β-galactosidase datasets to understand the effect of different imaging conditions on the quality of CryoGAN reconstruction. More precisely, we considered 1) the case with a lower noise level (-5.2 dB SNR) and 2) the case with a realistic noise level (-20 dB SNR) and translations (3% of the image size) in the projections. All the other conditions are identical to the main synthetic experiment. More details on these two additional experiments are provided in Appendix A.3.

The reconstructions obtained in the two cases reach a resolution of 7.53 Å and 10.8 Å, respectively, as shown in Figure 5.4. As expected, the decrease of the noise level from -20 dB to -5.2 dB improves the reconstruction quality (from 8.64 Å to 7.53 Å). For the second case, the challenging presence of translations results in a slightly lower resolution. We believe that using a discriminator architecture that is invariant to shifts in the input image would improve the result in this case. Indeed, the discriminator would then be blind to such shifts in the real and simulated data, which would yield a reconstruction with a quality similar to the no-shift case.

5.7.3 Results on Experimental Data (EMPIAR-10061)

We then deployed CryoGAN on 41,123 β-galactosidase projections (obtained from EMPIAR-10061 [158]) to assess its capacity to reconstruct real, experimental data. Here as well, we randomly divided the dataset in two and applied CryoGAN separately on both halves. The details of this experimental setup are given in Appendix A.3.

We ran CryoGAN for 150 minutes to obtain a 12.1 Å-resolution reconstruction using an NVIDIA V100 GPU. The results are displayed in Figure 5.5. The flexible architecture of CryoGAN permits the straightforward injection of prior knowledge on this specific imaging procedure into the reconstruction pipeline (*e.g.*, the assumption of uniform pose distribution). Using this prior knowledge and its adversarial learning scheme, CryoGAN converges toward the reconstruction that best explains the statistics of the dataset (Figure 5.5a). As with the synthetic experiments, this is achieved by exploiting the gradients of the discriminator to update the simulator and the current volume estimate, so that, at later iterations, the simulated projections (Figure 5.5b) follow a distribution that better approaches that of the real dataset. Higher-resolution details are thus progressively introduced in the estimated volume throughout the run, as illustrated by the evolution of the FSC curves between the reconstructed half-maps (Figure 5.5d). This resulted in a 12.08 Å β-galactosidase structure whose simulated projections closely resemble the real ones (Figure 5.5c).

5.8 Summary

In this chapter, we present CryoGAN, a new paradigm for single-particle Cryo-EM reconstruction based on unsupervised deep adversarial learning. The major challenge in single-particle Cryo-EM is that the imaged particles have unknown poses. Current reconstruction techniques are based on a marginalized maximum-likelihood formulation that requires calculations over the set of all possible poses for each projection image, a computationally demanding procedure. CryoGAN sidesteps this problem by using a generative adversarial network (GAN) to learn the 3D structure that has simulated projections that most closely match the real data in a distributional sense. The architecture of CryoGAN resembles that of standard GAN, with the twist that the generator network is replaced by a model of the Cryo-EM image acquisition process. CryoGAN is an unsupervised algorithm that only demands projection images and an estimate of the contrast transfer function parameters. No initial volume estimate or prior training is needed. Moreover, CryoGAN requires minimal user interaction and can provide reconstructions in a matter of hours on a high-end GPU. In addition, we provide sound mathematical guarantees on the recovery of the correct structure. CryoGAN currently achieves a 8.6 Å resolution on a realistic synthetic dataset. Preliminary results on real β-galactosidase data demonstrate CryoGAN's ability to exploit data statistics under standard experimental imaging conditions. We believe that this paradigm opens the door to a family of novel likelihood-free algorithms for Cryo-EM reconstruction.

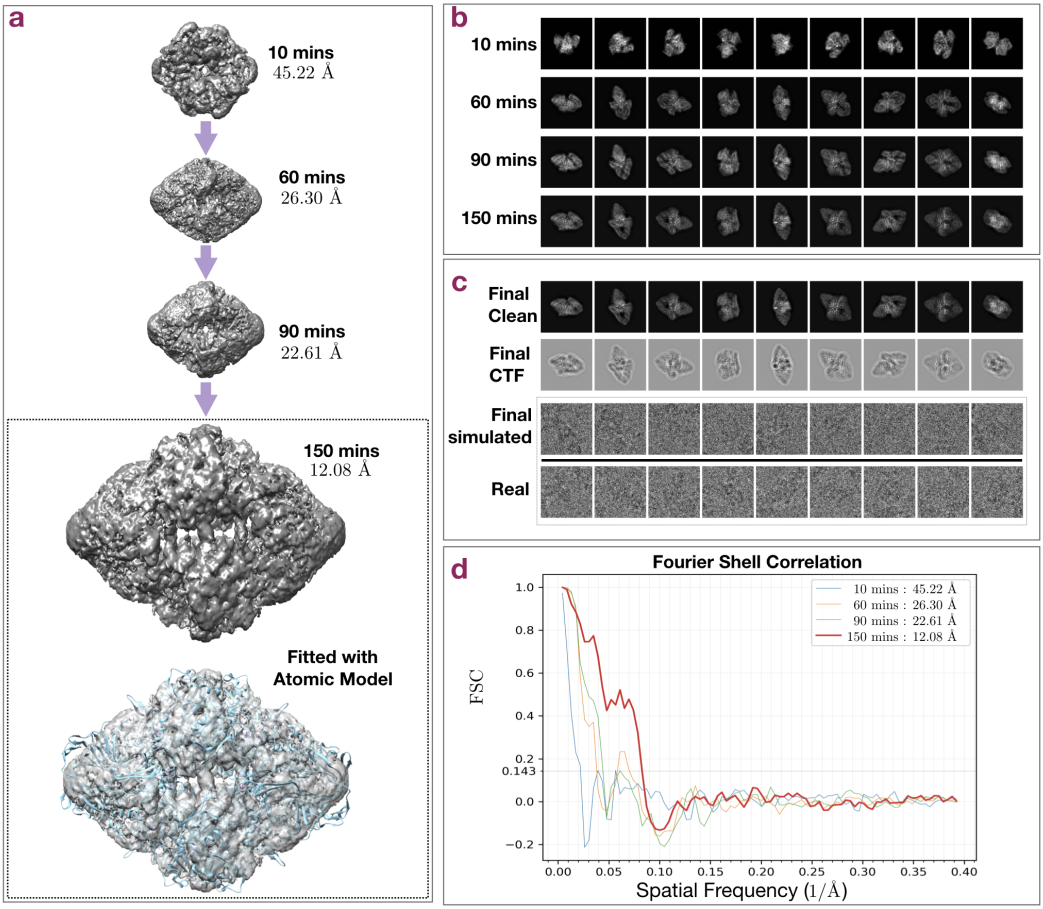

Figure 5.5: Evolution of CryoGAN while reconstructing the experimental β-galactosidase dataset (EMPIAR-10061) from [158]. **(a)** The volume is initialized with zeros and is progressively updated to produce projections whose distribution matches that of the experimental dataset. **(b)** Evolution during the training of the clean projections (*i.e.*, before CTF and noise) generated by the Cryo-EM physics simulator. **(c)** *Row 1*: Clean, simulated projections generated at the final stage of training. *Row 2*: CTF-modulated, simulated projections (before noise) generated at the final stage of training. *Row 3*: Simulated projections (with CTF and noise) generated at the final stage of training. *Row 4*: Real projections, for comparison. **(d)** FSC curves between the two reconstructed half-maps at different points during training.

Chapter 6

Reconstructing Continuous Conformations in CryoEM using GANs

[1]In the previous chapters, we have only discussed inverse problems in which at least the signal or the forward model is deterministic. In this chapter, we deal with the challenging case of heterogeneous Cryo-EM. In this case the 3D structure exhibits an unknown conformation variability while being imaged by the stochastic forward model. In order to solve this problem, we devise a third-generational algorithm which is based on the extension of the distributional perspective proposed in CryoGAN.

6.1 Overview

The determination of the structure of nonrigid macromolecules is an important aspect of structural biology and is fundamental in our understanding of biological mechanisms and in drug discovery [156]. Among other popular techniques such as X-ray crystallography and nuclear magnetic resonance spectroscopy, Cryo-EM

---

[1] This chapter uses content from our work [194].

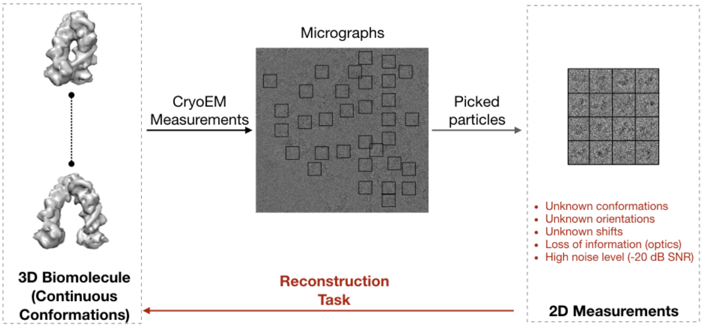

Figure 6.1: Reconstruction task of Cryo-EM. Many samples of a biomolecule (which may exhibit continuously varying conformations) are frozen in vitreous ice. These are then imaged/projected using an electron beam to get 2D micrographs. The 2D images containing projection of a single sample are then picked out (black-box). The task then is to reconstruct the conformations of the biomolecule from these measurements.

has emerged as a unique method to determine molecular structures at unprecedented high resolutions. It is widely applicable to proteins that are difficult to crystallize or have large structures. Cryo-EM produces a large number (from 10^4 to 10^7) of tomographic projections of the molecules dispersed in a frozen solution. The reconstruction of 3D molecular structures from these data involves three main challenges: possible structural heterogeneity of the molecule, random locations and orientations of the molecules in the ice, and an extremely poor signal-to-noise ratio (SNR), which can be as low as to -20 dB (Figure 6.1). In fact, the reconstruction of continuously varying conformations of a nonrigid molecule is still an open problem in the field [195, 196]. A solution would considerably enhance our understanding of the functions and behaviors of many biomolecules.

Most current methods [151, 153] find the 3D structure by maximizing the likelihood of the data. They employ an iterative optimization scheme that alternatively estimates the distribution of poses (or orientations) and reconstructs a 3D structure

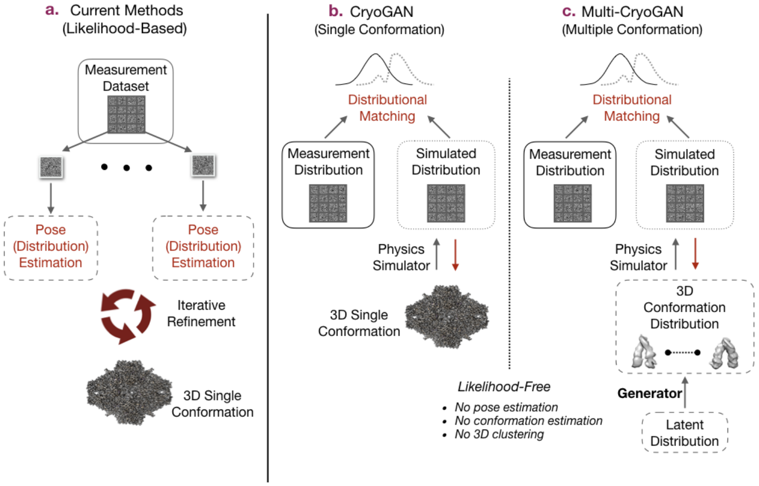

Figure 6.2: Schematic overview of the reconstruction methods in Cryo-EM. (a) Current methodels; (b) CryoGAN; (c) proposed method (Multi-CryoGAN).

until a criterion is satisfied (Figure 6.2(a)). To address the structural variability of protein complexes, these methods typically use discrete clustering approaches. However, the pose estimation and clustering steps are computationally heavy and include heuristics. This makes these methods inefficient when the molecule has continuous variations or a large set of discrete conformations.

Recently, two deep-learning-based reconstruction methods that require no prior-training nor additional training data have been introduced. On one hand, Cryo-DRGN [197] uses a variational auto-encoder (VAE) to model continuous structural variability, avoiding the heuristic clustering step. It is a likelihood-based method that requires pose estimation using an external routine like a branch-and bound-method [153]. This additional processing step can complicate the reconstruction procedure and limit the flexibility of the model. On the other hand, Gupta *et al.* [28] have recently proposed CryoGAN. It addresses the problem under a generative adversarial framework [36]. CryoGAN learns to reconstruct a 3D structure whose randomly projected 2D Cryo-EM images match the acquired data in a distributional sense (Figure 6.2(b)). Due to this likelihood-free characteristic, CryoGAN does not require any additional processing step such as pose estimation, while it can be directly deployed on the Cryo-EM measurements. This largely helps simplify the reconstruction procedure. However, its application is limited to the reconstruction of a single conformation.

In this work, we combine the advantages of CryoDRGN and CryoGAN. We propose an unsupervised deep-learning-based method, called Multi-CryoGAN. It can reconstruct continuously varying conformations of a molecule in a truly standalone and likelihood-free manner. Using a convolutional neural network (CNN), it directly learns a mapping from a latent space to the 3D conformation distribution. Unlike current methods, it requires no calculation of pose or conformation estimation for each projection, while it has the capacity to reconstruct low-dimensional but complicated conformation manifolds [198].

Using synthetic Cryo-EM data as our benchmark, we show that our method can reconstruct the conformation manifold for both continuous and discrete conformation distributions. In the discrete case, it also reconstructs the corresponding probabilities. To the best of our knowledge, this is the first standalone method that can construct whole manifold of the biomolecule conformations.

6.2 Related Work

Traditional Cryo-EM Image Reconstruction. A detailed survey of the classical methods is provided in [157,199]. Most of them fall into the maximum-likelihood (ML) framework and rely on either expectation-maximization (ML-EM) [151] or gradient descent (the first stage of [153]). In the context of heterogeneous conformation reconstruction, a conjugate-gradient descent is used to estimate the volume covariance matrix [200]. The eigendecomposition of this matrix contains information about the conformation distribution which is then input to the ML framework. In [201], a conformation manifold is generated for each group of projections with similar poses. This data-clustering approach assumes orientation rather than structural heterogeneity to be the dominant cause for variations among the projection images, a strong constraint. In addition, the reconstruction of 3D movies from multiple 2D manifolds can be computationally expensive. In another method, Moscovich *et al.* [202] compute the graph Laplacian eigenvectors of the conformations using covariance estimation. In [203], the problem of heterogeneous reconstructions is reformulated as the search for a homogeneous high-dimensional structure that represents all the states, called a hypermolecule, which is characterized by a basis of hypercomponents. This allows for reconstruction of high-dimensional conformation manifolds but requires assumptions on the variations of the conformations as a prior in their Bayesian formulation.

One of the main drawbacks of these methods is that they require marginalization over the space of poses for each projection image, which is computationally demanding and potentially inaccurate. In addition, because they rely on 3D clustering to deal with structural variations of protein complexes, these methods become inefficient for a large set of discrete conformations and struggle to recover a continuum of conformations.

Deep Learning for Cryo-EM Reconstructions. In addition to CryoDRGN and CryoGAN that have already been discussed in the introduction, there is a third described in [188]. It uses a VAE and a framework based on a generative adversarial network (GAN) to learn the latent distribution of the acquired data. This representation is then used to estimate the orientation and other important parameters for each projection image.

Deep Learning to Recover a 3D Object from 2D Projections. The implicit or explicit recovery of 3D shapes from 2D views is an important problem in computer vision. Many deep-learning algorithms have been proposed for this [190–192]. Taking inspiration from compressed sensing, Bora *et al.* [189] have recently introduced a GAN framework that can recover an original distribution from the measurements through a forward model. While these approaches would in principle be applicable, they consider a forward model that is too simple for Cryo-EM, where a contrast transfer function (CTF) must be taken into account and where the noise is orders of magnitude stronger (e.g. with a typical SNR of -10 to -20 *dB*.

6.3 Background and Preliminaries

6.3.1 Image-Formation Model

The aim of Cryo-EM is to reconstruct the 3D molecular structure from the measurements $\{\mathbf{y}^1_{\text{data}}, \ldots, \mathbf{y}^Q_{\text{data}}\}$, where Q is typically between 10^4 to 10^7. Each measurement $\mathbf{y}^q \in \mathbb{R}^{N\times N}$ is given by

$$\mathbf{y}^q = \underbrace{\mathbf{C}_{\mathbf{c}^q} * \mathbf{S}_{\mathbf{t}^q}\mathbf{P}_{\boldsymbol{\theta}^q}}_{\mathbf{H}_{\boldsymbol{\varphi}^q}}\{\mathbf{x}^q\} + \mathbf{n}^q, \tag{6.1}$$

where

- $\mathbf{x}^q \in \mathbb{R}^{N\times N\times N}$ is a separate instance of the 3D molecular structure;
- $\mathbf{n}^q \in \mathbb{R}^{N\times N}$ is the noise;
- $\mathbf{H}_{\boldsymbol{\varphi}^q}$ is the measurement operator which depends on the imaging parameters $\boldsymbol{\varphi}^q = (\boldsymbol{\theta}^q, \mathbf{t}^q, \mathbf{c}^q) \in \mathbb{R}^8$ and involves three operations.
  - The term $\mathbf{P}_{\boldsymbol{\theta}^q}\{\mathbf{x}^q\}$ is the tomographic projection of $\mathbf{x}^q$ rotated by $\boldsymbol{\theta}^q = (\theta_1^q, \theta_2^q, \theta_3^q)$.
  - The operator $\mathbf{S}_{\mathbf{t}^q}$ shifts the projected image by $\mathbf{t}^q = (t_1^q, t_2^q)$. This shift arises from off-centered particle picking.
  - The Fourier transform of the resulting image is then modulated by the CTF $\hat{\mathbf{C}}_{\mathbf{c}^q}$ with defocus parameters $\mathbf{c}^q = (d_1^q, d_2^q, \alpha_{\text{ast}}^q)$ and thereafter subjected to inverse Fourier Transform.

For more details, please see Section on Image Formation in the previous chapter. The challenge of Cryo-EM is that, for each measurement $\mathbf{y}^q$, the structure $\mathbf{x}^q$ and the imaging parameters $(\boldsymbol{\theta}^q, \mathbf{t}^q)$ are unknown, the CTF is a band pass filter with multiple radial zero frequencies that incur irretrievable loss of information, and the energy of the noise is multiple times ($\sim$10 to 100 times) that of the signal which corresponds to SNRs of -10 to -20 dB. In the homogeneous case (single conformation), all $\mathbf{x}^q$ are identical. But in the heterogeneous case (multiple conformations), each $\mathbf{x}^q$ represents a different conformation of the same biomolecule.

Stochastic Modeling. We denote the probability distribution over the conformation landscape by $p_{\text{conf}}(\mathbf{x})$ from which a conformation $\mathbf{x}^q$ is assumed to be sampled from. We assume that the imaging parameters and the noise are sampled from known distributions $p_{\boldsymbol{\varphi}} = p_{\boldsymbol{\theta}}p_{\mathbf{t}}p_{\mathbf{c}}$ and $p_{\mathbf{n}}$, respectively. For a given conformation distribution $p_{\text{conf}}(\mathbf{x})$, this stochastic forward model induces a distribution over the measurements which we denote by $p(\mathbf{y})$. We denote by $p_{\text{conf}}^{\text{data}}(\mathbf{x})$ the true conformation distribution from which the data distribution $p_{\text{data}}(\mathbf{y})$ is acquired such that

$\{\mathbf{y}^1_{\text{data}}, \ldots, \mathbf{y}^Q_{\text{data}}\} \sim p_{\text{data}}(\mathbf{y})$. The distribution $p^{\text{data}}_{\text{conf}}(\mathbf{x})$ is unknown and needs to be recovered.

The classical methods are likelihood-based and rely on the estimation of imaging parameters $(\boldsymbol{\theta}^q, \mathbf{t}^q)$ (or a distribution over them) and of the conformation class for each measurement image $\mathbf{y}^q$. This information is then utilized to reconstruct the multiple discrete conformations. Our method, in contrast, is built upon the insight that, to recover $p^{\text{data}}_{\text{conf}}(\mathbf{x})$ it is sufficient to find a $p^{\text{gen}}_{\text{conf}}(\mathbf{x})$ whose corresponding measurement distribution $p_{\text{gen}}(\mathbf{y})$ is equal to $p_{\text{data}}(\mathbf{y})$ (see Theorem 6.5.1). This does away with pose (or distributions over the poses) estimation and conformation clustering for each measurement.

6.3.2 CryoGAN

Our scheme is extension of the CryoGAN [28] method, which is applicable only for the homogeneous case $p^{\text{data}}_{\text{conf}}(\mathbf{x}) = \delta(\mathbf{x} - \mathbf{x}_{\text{data}})$, where $\mathbf{x}_{\text{data}}$ is the true 3D structure. CryoGAN tackles the challenge by casting the reconstruction problem as a distribution-matching problem (Figure 6.2(b)). More specifically, it learns to reconstruct the 3D volume $\mathbf{x}^*$ whose simulated projection set (measurement distribution) is most similar to the real projection data in a distributional sense, such that

$$\mathbf{x}^* = \arg\min_{\mathbf{x}} \text{WD}(\text{p}_{\text{data}}(\mathbf{y})||\text{p}_{\text{gen}}(\mathbf{y};\mathbf{x})). \tag{6.2}$$

Here, $p_{\text{gen}}(\mathbf{y};\mathbf{x})$ is the distribution generated from the Cryo-EM physics simulator and WD refers to the Wasserstein distance [165]. This goal is achieved by solving the min-max optimization problem:

$$\mathbf{x}^* = \arg\min_{\mathbf{x}} \underbrace{\max_{D_\phi:\|D_\phi\|_L\leq 1} (\mathbb{E}_{\mathbf{y}\sim p_{\text{data}}(\mathbf{y})}[D_\phi(\mathbf{y})] - \mathbb{E}_{\mathbf{y}\sim p_{\text{gen}}(\mathbf{y};\mathbf{x})}[D_\phi(\mathbf{y})])}_{\text{WD}(\text{p}_{\text{data}}(\mathbf{y})||\text{p}_{\text{gen}}(\mathbf{y};\mathbf{x}))}, \tag{6.3}$$

where D_ϕ is a neural network with parameters ϕ that is constrained to have Lipschitz constant $\|D_\phi\|_L \leq 1$ [167] (Figure 6.3(b)). Here, D_ϕ learns to differentiate between the real projection $\mathbf{y}$ and the simulated projection $\mathbf{H}_\varphi\{\mathbf{x}\}$ and scores the realness of given samples. As the discriminative power of D_ϕ becomes stronger (`maximization`) and the underlying volume estimate $\mathbf{x}$ is updated accordingly (`minimization`), $p_{\text{data}}(\mathbf{y})$ and $p_{\text{gen}}(\mathbf{y};\mathbf{x})$ become indistinguishable so that the algorithm recovers $\mathbf{x}^* = \mathbf{x}_{\text{data}}$.

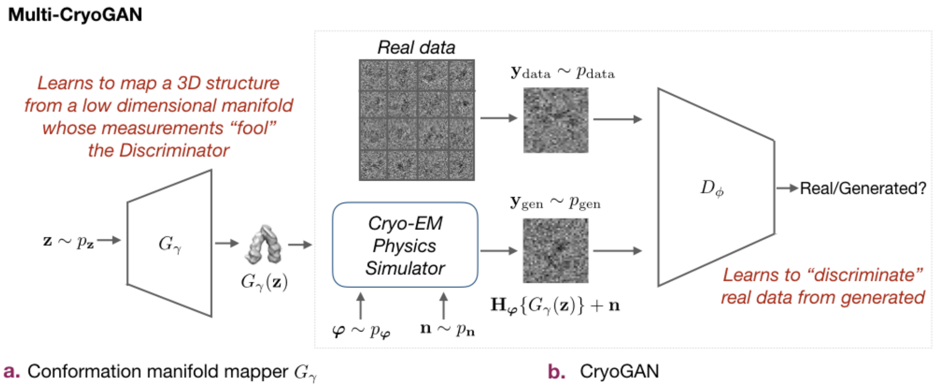

Figure 6.3: Schematic illustration of Multi-CryoGAN and its components. (a) Conformation manifold mapper; (b) CryoGAN.

6.4 Method

6.4.1 Parameterization of the Conformation Manifold

CryoGAN successfully reconstructs the volumetric structure of a protein by finding a single volume $\mathbf{x}$ that explains the entire set of projections, which is adequate when all the imaged particles are identical (homogeneous case). However, in reality, many biomolecules have nonrigid structures, which carry vital information.

To address this, we introduce a manifold-learning module G_γ that uses a CNN with learnable weights γ (Figure 6.3 (a)). Sampling from $p_{\text{conf}}(\mathbf{x})$ is then equivalent to getting $G_\gamma(\mathbf{z})$, where $\mathbf{z}$ is sampled from a prior distribution. Therefore, $\mathbf{y} \sim p_{\text{gen}}(\mathbf{y})$ is obtained by evaluating $\mathbf{H}_{\boldsymbol{\varphi}}\{G_\gamma(\mathbf{z})\} + \mathbf{n}$, where $(\mathbf{n}, \mathbf{z}, \boldsymbol{\varphi})$ are sampled from their distributions (see Algorithm 5 and Figure 6.3). To explicitly show this dependency to G_γ, we hereafter denote the generated distribution of projection data by $p_{\text{gen}}(\mathbf{y}; G_\gamma)$.

Algorithm 5 Samples from the generated distribution $p_{\text{gen}}(\mathbf{y}; G_\gamma)$.

```
Input: Latent distribution $p_{\mathbf{z}}$; angle distribution $p_{\boldsymbol{\theta}}$; translation distribution $p_{\mathbf{t}}$; CTF parameters distribution $p_{\mathbf{c}}$; noise distribution $p_{\mathbf{n}}$
Output: Simulated projection $\mathbf{y}_{\text{gen}}$
1. Sample $\mathbf{z} \sim p_{\mathbf{z}}$.
2. Feed $\mathbf{z}$ into generator network to get $\mathbf{x} = G_\gamma(\mathbf{z})$.
3. Sample the imaging parameters $\boldsymbol{\varphi} = [\boldsymbol{\theta}, \mathbf{t}, \mathbf{c}]$.
    • Sample the Euler angles $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3) \sim p_{\boldsymbol{\theta}}$.
    • Sample the 2D shifts $\mathbf{t} = (t_1, t_2) \sim p_{\mathbf{t}}$.
    • Sample the CTF parameters $\mathbf{c} = (d_1, d_2, \alpha_{\text{ast}}) \sim p_{\mathbf{c}}$.
4. Sample the noise $\mathbf{n} \sim p_{\mathbf{n}}$.
5. Generate $\mathbf{y}_{\text{gen}} = \mathbf{H}_{\boldsymbol{\varphi}}\mathbf{x} + \mathbf{n}$ based on (6.6).
return $\mathbf{y}_{\text{gen}}$
```

6.4.2 Optimization Scheme

We now find G_{γ^*} such that the distance between $p_{\text{data}}(\mathbf{y})$ and $p_{\text{gen}}(\mathbf{y}; G_\gamma)$ is minimized, which results in the min-max optimization problem [164]

$$\begin{aligned} G_{\gamma^*} &= \arg\min_{G_\gamma} \mathrm{WD}(\mathrm{p}_{\mathrm{data}}(\mathbf{y})||\mathrm{p}_{\mathrm{gen}}(\mathbf{y}; \mathrm{G}_\gamma)) && (6.4) \\ &= \arg\min_{G_\gamma} \underbrace{\max_{D_\phi : \|D_\phi\|_L \le 1} \left(\mathbb{E}_{\mathbf{y}\sim p_{\text{data}}}[D_\phi(\mathbf{y})] - \mathbb{E}_{\mathbf{y}\sim p_{\text{gen}}(\mathbf{y}; G_\gamma)}[D_\phi(\mathbf{y})]\right)}_{\mathrm{WD}(\mathrm{p}_{\mathrm{data}}(\mathbf{y})||\mathrm{p}_{\mathrm{gen}}(\mathbf{y}; \mathrm{G}_\gamma))}. && (6.5) \end{aligned}$$

As will be discussed in Theorem 6.5.1, the global minimizer G_{γ^*} of (6.5) indeed captures the true conformation landscape $p_{\text{conf}}(\mathbf{x})$, which is achieved when D_ϕ is no longer able to differentiate the samples from $p_{\text{data}}(\mathbf{y})$ and $p_{\text{gen}}(\mathbf{y}; G_{\gamma^*})$.

It is crucial to note the difference between Multi-CryoGAN and conventional generative adversarial frameworks [36]. In the latter, G directly outputs the samples from $p_{\text{gen}}(\mathbf{y})$, whereas ours outputs the samples $\mathbf{x}$ from the conformation distribution $p_{\text{conf}}(\mathbf{x})$ whose stochastic projections are the samples of $p_{\text{gen}}(\mathbf{y})$. The conventional schemes only helps one to generate samples which are similar to the real data but does not recover the underlying conformation landscape. Our proposed scheme includes the physics of Cryo-EM, which ties $p_{\text{gen}}(\mathbf{y})$ with the conformation landscape $p_{\text{conf}}(\mathbf{x})$ and is thus able to recover it (See Theorem 6.5.1 in Section 6.5).

6.5 Theoretical Guarantee of Recovery

Our experiments illustrate that enforcing a match between the distribution of the simulated measurements and that of the data is sufficient to reconstruct the true conformations. We now prove this mathematically. For the homogeneous case, the proof is already discussed in [28, Theorem 1] which we now extend to the heterogeneous case. We switch to a continuous-domain formulation of the Cryo-EM problem while noting that the result is transferable to the discrete-domain as well, albeit with some discretization error.

Notations and Preliminaries. We denote by $\mathcal{L}_2(\mathbb{R}^3)$ the space f 3D structures $f:\mathbb{R}^3 \to \mathbb{R}$ with finite energy $\|f\|_{L_2} < \infty$. The imaging parameters $\boldsymbol{\varphi}$ are assumed to lie in $\mathcal{B} \subset \mathbb{R}^8$. We denote by $\mathcal{L}_2(\mathbb{R}^2)$ the space of 2D measurements with finite energy. Each individual continuous-domain Cryo-EM measurement $y \in \mathcal{L}_2(\mathbb{R}^2)$ is given by

$$y = H_{\boldsymbol{\varphi}}\{f\} + n, \tag{6.6}$$

where $f \in \mathcal{L}_2(\mathbb{R}^3)$ is some conformation of the biomolecule sampled from the probability measure $\tilde{\mathbb{P}}_{\text{conf}}$ on $\mathcal{L}_2(\mathbb{R}^3)$, the imaging parameters $\boldsymbol{\varphi}$ are sampled from $p_{\boldsymbol{\varphi}}$, and n is sampled from the noise probability measure $\mathbb{P}_n$ on $\mathcal{L}_2(\mathbb{R}^2)$.

We define $[f] := \{r_{\mathbf{A}}\{f\} : r_{\mathbf{A}} \in O\}$ as the set of all the rotated-reflected versiond of f. There, O is the set of all the rotated-reflected versions over $\mathcal{L}_2(\mathbb{R}^3)$. We define the space $\sum \mathcal{L}_2(\mathbb{R}^3) = \mathcal{L}_2(\mathbb{R}^3)/O$ as the quotient space of the shapes. For any $\tilde{\mathbb{P}}_{\text{conf}}$ defined over $\mathcal{L}_2(\mathbb{R}^3)$, an equivalent $\mathbb{P}_{\text{conf}}$ exists over $\sum \mathcal{L}_2(\mathbb{R}^3)$. Since we are interested only in the shape of conformations of the biomolecule, we will only focus on recovering $\mathbb{P}_{\text{conf}}$. We denote by Ψ the probability measure on $\mathcal{B} \in \mathbb{R}^8$. The measure Ψ is associated to the density function $p_{\boldsymbol{\varphi}}$. Both of these induce a probability measure $\mathbb{P}_{\text{clean}}$ on the space $\mathcal{L}_2^2 = \{f : \mathbb{R}^3 \to \mathbb{R}^2 \text{ s.t. } \|f\|_{L_2} < \infty\}$ through the forward operator. This is given by $\mathbb{P}_{\text{clean}}[A] = (\mathbb{P}_{\text{conf}} \times \Psi)[([f], \boldsymbol{\varphi}) \in (\sum \mathcal{L}_2 \times \mathcal{B}) : H_{\boldsymbol{\varphi}} f \in A]$ for any Borel measure set $A \in \mathcal{L}_2(\mathbb{R}^2)$. We denote $\mathbb{P}_{\text{meas}}$ as the probability measure of the noisy measurements.

Theorem 6.5.1. *Let* $\mathbb{P}_{\mathrm{conf}}^{\mathrm{data}}$ *and* $\mathbb{P}_{\mathrm{conf}}^{\mathrm{gen}}$ *be the true and the reconstructed conformation probability measures on the quotient space of 3D structures* $\sum \mathcal{L}_2(\mathbb{R}^3)$*, respectively. We assume that they are atomic and that they are supported only on nonnegative-valued shapes. Let* $\mathbb{P}_{\mathrm{meas}}^{\mathrm{data}}$ *and* $\mathbb{P}_{\mathrm{meas}}^{\mathrm{gen}}$ *be the probability measures of the noisy Cryo-EM measurements obtained from* $\mathbb{P}_{\mathrm{conf}}^{\mathrm{data}}$ *and* $\mathbb{P}_{\mathrm{conf}}^{\mathrm{gen}}$*, respectively.*

Make the following physical assumptions:

i. the noise probability measure $\mathbb{P}_n$ *is such that its characteristic functional vanishes nowhere in its domain and that its sample* n *is pointwise-defined everywhere;*

ii. the distributions $p_{\boldsymbol{\theta}}$*,* $p_{\mathbf{t}}$*, and* $p_{\mathbf{c}}$ *are bounded;*

iii. for any two $\mathbf{c}_1, \mathbf{c}_2 \sim p_{\mathbf{c}}, \mathbf{c}_1 \neq \mathbf{c}_2$*, the CTFs* $\hat{\mathbf{C}}_{\mathbf{c}_1}$ *and* $\hat{\mathbf{C}}_{\mathbf{c}_2}$ *share no common zero frequencies.*

Then, it holds that

$$\mathbb{P}_{\mathrm{meas}}^{\mathrm{data}} = \mathbb{P}_{\mathrm{meas}}^{\mathrm{gen}} \Rightarrow \mathbb{P}_{\mathrm{conf}}^{\mathrm{data}} = \mathbb{P}_{\mathrm{conf}}^{\mathrm{gen}}. \tag{6.7}$$

Proof. We first prove that $\mathbb{P}_{\mathrm{meas}}^{\mathrm{data}} = \mathbb{P}_{\mathrm{meas}}^{\mathrm{gen}} \Rightarrow \mathbb{P}_{\mathrm{clean}}^{\mathrm{data}} = \mathbb{P}_{\mathrm{clean}}^{\mathrm{gen}}$. Note that, due to the independence of clean measurements and noise, we have that

$$\begin{aligned} \hat{\mathbb{P}}_{\mathrm{meas}}^{\mathrm{data}} &= \hat{\mathbb{P}}_{\mathrm{clean}}^{\mathrm{data}}\, \hat{\mathbb{P}}_n \\ \hat{\mathbb{P}}_{\mathrm{meas}}^{\mathrm{gen}} &= \hat{\mathbb{P}}_{\mathrm{clean}}^{\mathrm{gen}}\, \hat{\mathbb{P}}_n. \end{aligned} \tag{6.8}$$

From the assumption that $\hat{\mathbb{P}}_n$ is nonzero everywhere, we deduce that $\hat{\mathbb{P}}_{\mathrm{clean}}^{\mathrm{data}} = \hat{\mathbb{P}}_{\mathrm{clean}}^{\mathrm{gen}}$. This proves the first step.

To prove the next step, we invoke Theorem 4 in [28] which states that any two probability measures $\mathbb{P}_{\mathrm{clean}}^{1}$ and $\mathbb{P}_{\mathrm{clean}}^{2}$ that correspond to Dirac probability measures $\mathbb{P}_{\mathrm{conf}}^{1}$ and $\mathbb{P}_{\mathrm{conf}}^{2}$ on $\sum \mathcal{L}_2(\mathbb{R}^3)$, respectively, are mutually singular (zero measure of the common support) if and only if the latter are distinct. We denote the relation of mutual singularity by $\perp$.

Since $\mathbb{P}_{\mathrm{conf}}^{\mathrm{data}}$ is an atomic measure (countable weighted sum of distinct Dirac measures), the corresponding $\mathbb{P}_{\mathrm{clean}}^{\mathrm{data}}$ is composed of a countable sum of mutually singular measures. The same is true for $\mathbb{P}_{\mathrm{clean}}^{\mathrm{gen}}$ since it is equal to $\mathbb{P}_{\mathrm{clean}}^{\mathrm{data}}$.

We proceed by contradiction. We denote by $\mathrm{Supp}\{\mathbb{P}\}$ the support of the measure $\mathbb{P}$. Assume that $\mathrm{Supp}\{\mathbb{P}_{\mathrm{conf}}^{\mathrm{data}}\} \neq \mathrm{Supp}\{\mathbb{P}_{\mathrm{conf}}^{\mathrm{gen}}\}$. Let us define $\mathcal{S}_1 = \mathrm{Supp}\{\mathbb{P}_{\mathrm{conf}}^{\mathrm{gen}}\} \cap$

$\mathrm{Supp}\{\mathbb{P}^{\mathrm{data}}_{\mathrm{conf}}\}^C$. For any $[f] \in \mathcal{S}_1$, we denote by $\mathbb{P}^{f}_{\mathrm{clean}}$ its noiseless probability measure. Since $f \in \mathcal{S}_1$, it is distinct from any constituent Dirac measure in $\mathbb{P}^{\mathrm{data}}_{\mathrm{conf}}$. Therefore, by using [28, Theorem 4], $\mathbb{P}^{f}_{\mathrm{clean}}$ is mutually singular to each of the constituent mutually singular measures of $\mathbb{P}^{\mathrm{data}}_{\mathrm{clean}}$, implying that $\mathbb{P}^{f}_{\mathrm{clean}} \perp \mathbb{P}^{\mathrm{data}}_{\mathrm{clean}}$.

From $\mathrm{Supp}\{\mathbb{P}^{\mathrm{f}}_{\mathrm{clean}}\} \subset \mathrm{Supp}\{\mathbb{P}^{\mathrm{gen}}_{\mathrm{clean}}\}$, it follows that $\mathbb{P}^{\mathrm{data}}_{\mathrm{clean}} \neq \mathbb{P}^{\mathrm{gen}}_{\mathrm{clean}}$, which raises a contradiction. Therefore, the set $\mathcal{S}_1$ is empty. The same can be proved for the set $\mathcal{S}_2 = \mathrm{Supp}\{\mathbb{P}^{\mathrm{gen}}_{\mathrm{conf}}\}^C \cap \mathrm{Supp}\{\mathbb{P}^{\mathrm{data}}_{\mathrm{conf}}\}$. Therefore, $\mathrm{Supp}\{\mathbb{P}^{\mathrm{gen}}_{\mathrm{conf}}\} = \mathrm{Supp}\{\mathbb{P}^{\mathrm{data}}_{\mathrm{conf}}\}$, which means that the location of their constituent Dirac measures are the same. To maintain $\mathbb{P}^{\mathrm{data}}_{\mathrm{clean}} = \mathbb{P}^{\mathrm{gen}}_{\mathrm{clean}}$, the weight of their constituent Dirac measures have to be the same, too. This concludes the proof. □ □

In essence, Theorem 6.5.1 claims that a reconstructed manifold of conformations recovers the true conformations if its measurements match the acquired data in a distributional sense. Though the result assumes the true conformation landscape to be discrete (atomic measure), it holds for an infinite number of discrete conformations which could be arbitrarily close/similar to each other and is thus relevant to continuously varying conformations. We leave the proof of the latter case to future works.

6.6 Experiments and Results

We evaluate the performance of the proposed algorithm on synthetic datasets obtained from a protein with multiple conformations. We synthesize two datasets: one continuum of configurations and one where the particles can only take a discrete number of states. During reconstruction, no assumption is made of their continuous or discrete nature, which suggests that our method is capable of learning different conformation distribution behaviors.

Dataset. For each dataset, we generate 100,000 simulated projections from the *in vivo* conformation variation of the heat-shock protein *Hsp90*. The Coulomb density maps of each conformation are created by the code provided in [204] with slight modifications. The conformation variation of this protein is represented by the bond-angle β, which describes the work cycle of the molecule, where the two subunits continuously vary between fully closed ($\beta = 0°$, protein database entry *2cg9*) and fully opened ($\beta = 20°$). We sample $\beta \sim \mathrm{Uniform}\,(0°, 20°)$ for the continuous case and $\beta \sim 20° * \mathrm{Bernoulli}\,(0.75)$ for the discrete case. Here, $\mathrm{Uniform}\,(a, b)$ is the uniform distribution between a and b, and Bernoulli (p) denotes the Bernoulli distribution with parameter p. A conformation is generated with $(32 \times 32 \times 32)$ voxels, where the size of each voxel is 5 Å. A 2D projection with random orientation of this conformation is obtained ((32×32) image, Figure 6.4b). The orientation is sampled from a uniform distribution over $SO(3)$. Then, the CTF is applied to this projection image with a defocus uniformly sampled between $[1.0\ \mu\text{m}, 2.0\ \mu\text{m}]$, assuming that the horizontal and vertical defocus values are the same and there is no astigmatism. Translations/shifts are disabled in these experiments. Finally, Gaussian noise was added to the CTF-modulated images, resulting in an SNR of approximately -10 dB.

Implementation Details. The reconstruction of the conformation is done by solving (6.5) using Algorithm 6. For both continuous and discrete conformations, we use the same distribution $p_{\mathbf{z}} \sim \text{Uniform}\,(\mathbf{z}_0, \mathbf{z}_1)$, where $\mathbf{z}_0, \mathbf{z}_1 \in \mathbb{R}^{32\times32\times32}$ are randomly chosen from $\text{Uniform}(0, 0.025)$ and fixed throughout the process. Thus, we do not impose any prior knowledge whether the landscape is continuous or discrete. As we shall see later, this latent distribution is sufficiently rich to represent the variation of interest in the synthetic datasets. The architecture of D, G, and training details are provided in the Appendix A.4.

Optimization Details The models are trained end-to-end on the synthetic datasets with the usual WGAN loss (gradient-penalty regularizer $\lambda = 0.6$) on a TITAN X GPU. In all experiments, G_γ, D_ϕ and the noise parameters are optimized using three separate Adam optimizers with learning rate of 10^{-3} and gradient norm clipping value of $1, 10^3, 1$, respectively. Between each generator step, there are $n_{disc} = 5$ discriminator steps. The batch size is kept at 16 samples and the algorithm is run for 30 epochs which was sufficient for the convergence.

Metric. We deploy two metrics based on Fourier-Shell Correlation (FSC). The FSC between two structures $\mathbf{x}_1$ and $\mathbf{x}_2$ is given by

$$\text{FSC}(\omega, \mathbf{x}_1, \mathbf{x}_2) = \frac{\langle \mathbf{V}^{\omega}_{\hat{\mathbf{x}}_1}, \mathbf{V}^{\omega}_{\hat{\mathbf{x}}_2} \rangle}{\|\mathbf{V}^{\omega}_{\hat{\mathbf{x}}_1}\| \|\mathbf{V}^{\omega}_{\hat{\mathbf{x}}_2}\|} \tag{6.9}$$

where $\mathbf{V}^{\omega}_{\hat{\mathbf{x}}}$ is the vectorization of the shell of $\hat{\mathbf{x}}$ at radius ω and $\hat{\mathbf{x}}$ is the 3D Fourier transform of $\mathbf{x}$. As first metric, we use the FSC between a reconstructed conformation and the corresponding ground truth conformation. This metric encapsulates the structural quality of an individual reconstructed conformation.

To evaluate the landscape of the conformations, we propose a second metric that we call the matrix $\mathbf{M} \in \mathbb{R}^{L\times L}$ of FSC cross conformations (FSCCC). Its entries are given by

$$\mathbf{M}[m,n] = \mathrm{AreaFSC}(\mathbf{x}_m, \mathbf{x}_n) = \int_{0<\omega\leq\omega_c} \mathrm{FSC}(\omega, \mathbf{x}_m, \mathbf{x}_n)\, \mathrm{d}\omega \tag{6.10}$$

where $\mathbf{x}_m$ and $\mathbf{x}_n$ are samples in the reconstructed conformation manifold and ω_c is the normalized Nyquist frequency. We determine it for the reconstructed landscape by setting $\mathbf{x}_m = G_{\gamma^*}((1-\alpha_m) * \mathbf{z}_0 + \alpha_m * \mathbf{z}_1)$ where $\alpha_m = (m/\mathrm{L})$ for $m \in \{0, \ldots, L\}$. The matrix $\mathbf{M}$ encapsulates how similar $\mathbf{x}_m$ is compared to other structures $\mathbf{x}_n$ across the manifold ($\mathbf{M}[m,n]$ is proportional to the similarity between $\mathbf{x}_m$ and $\mathbf{x}_n$), hence allowing for a visualization of the manifold.

For the continuous conformation, it is useful to compare the FSCCC of our reconstructions with that of the ground truth. To that end, we also evaluate $\mathbf{M}[m,n]$ when $\mathbf{x}_m$ corresponds to the bond-angle $\beta = 20°(m/\mathrm{L})$, where $m \in \{0, \ldots, L\}$. In our experiments, we used $L = 20$ for all FSCCC calculations.

6.6.1 Continuous Conformations

We give in Figure 6.4 a qualitative comparison between the ground truth conformation variation, as the angle β goes from 0° to 20°, and the reconstructions $G(\gamma^*)(\mathbf{z})$, where $\mathbf{z} = (1-\alpha)\mathbf{z}_0 + \mathbf{z}_1$ and α goes from 0 to 1. Our method successfully reconstructs a manifold that exhibits smooth continuous conformation variation (Figure 6.4(a)), where the input parameter α has direct control over the bond-angle for the reconstruction. This shows that not only the true conformation landscape has been captured, but its factor of variation has been meaningfully encoded by the latent variables $\mathbf{z}$. The similarity between simulated projections and the ground truth data in Figure 6.4(b) suggests that the algorithm has achieved $p_{\mathrm{data}} = p_{\mathrm{gen}}$. Moreover, their underlying distributions of noiseless projections are also similar, in accordance to the property discussed in Theorem 6.5.1.

We also evaluate the structural quality of reconstruction for certain representative individual conformations. In Figure 6.4(c), the extreme conformations for the ground truth $\beta = 0°$ and $\beta = 20°$ and the reconstructions $\alpha = 0$ and $\alpha = 1$ are shown. Their FSC plot reach the value 0.5 after the normalized frequency of 0.25, so that at least half of the Nyquist resolution is achieved (Figure 6.4(d)). All these results are further confirmed by the very similar FSCCC matrix for ground truth and reconstruction in Figure 6.5(a). This implies that the reconstruction manifold successfully approximates the continuous ground truth.

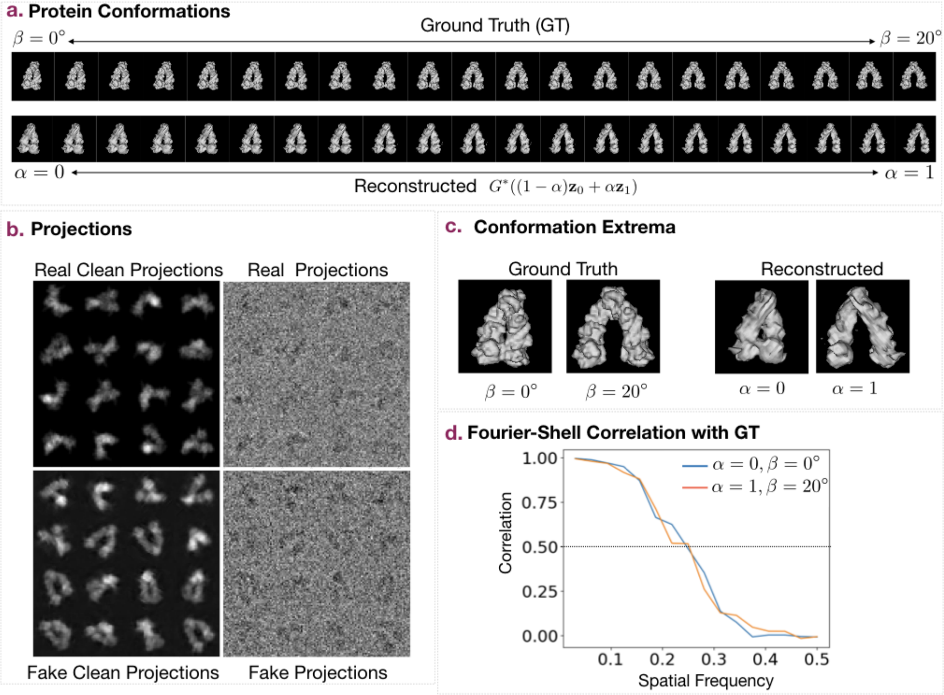

Figure 6.4: Continuous conformations experiment. (a) Comparison between the ground truth conformation manifold and the reconstructed conformation manifold $G^*(\mathbf{z})$ where $\mathbf{z} = (1-\alpha)\mathbf{z}_0 + \alpha\mathbf{z}_1$, $\alpha \in [0,1]$. (b, left-column) Clean projections of random samples from the ground truth and reconstructed manifold; (b, right-column) their CTF-modulated and noise-corrupted projection. These are the real and generated samples that are fed to D_ϕ. (c) Ground truth with angles 0° and 20°, and the reconstruction corresponding to the endpoints in the latent space. (d) The FSC between them.

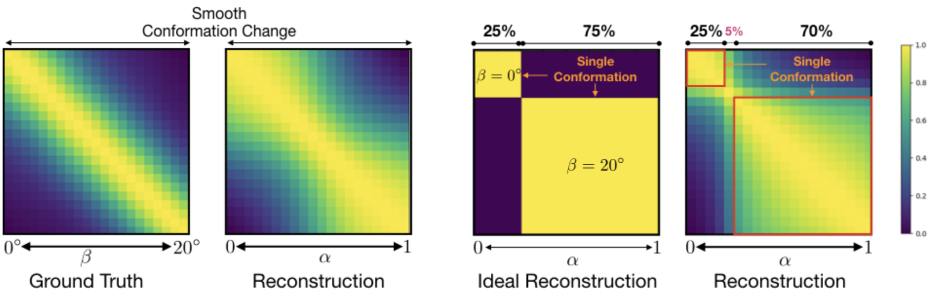

a. **Continuous Conformations (FSCCC Matrix)**

b. **Discrete Conformations (FSCCC Matrix)**

Figure 6.5: The FSC cross-conformation (FSCCC) matrix in (6.10). (a) Continuous conformation with (left) ground truth and (right) reconstruction. It shows that the reconstructed conformations smoothly vary (without forming clusters) similar to the ground truth case. (b) Discrete conformation case with (left) ideal reconstruction and (right) obtained reconstruction. The ideal reconstruction describes the case where 25% and 75% of latent space would have mapped to the two distinct conformations without any transitions. The obtained reconstruction case can be seen to be very similar to the ideal case with 25% and 70% being the latent space occupied by the two conformations.

6.6.2 Discrete Conformations

We present in Figure 6.6 the reconstruction results for the discrete case, where our proposed method successfully recovers not only the conformations but also their probabilities. About 70% of the reconstructed landscape matches the configuration for $\beta = 20°$, while 25% matches $\beta = 0°$. The remaining 5% of the landscape corresponds to a relatively abrupt transition between them. This suggests that our model distribution $p_{\text{conf}}(\mathbf{x})$ closely follows the ground truth Bernoulli distribution. This is further supported in Figure Figure 6.5(b), where the reconstructed FSCCC matrix greatly resembles the ideal reconstruction case. Ideally, one would expect the two conformations to occupy 25% and 75% of the latent space without having

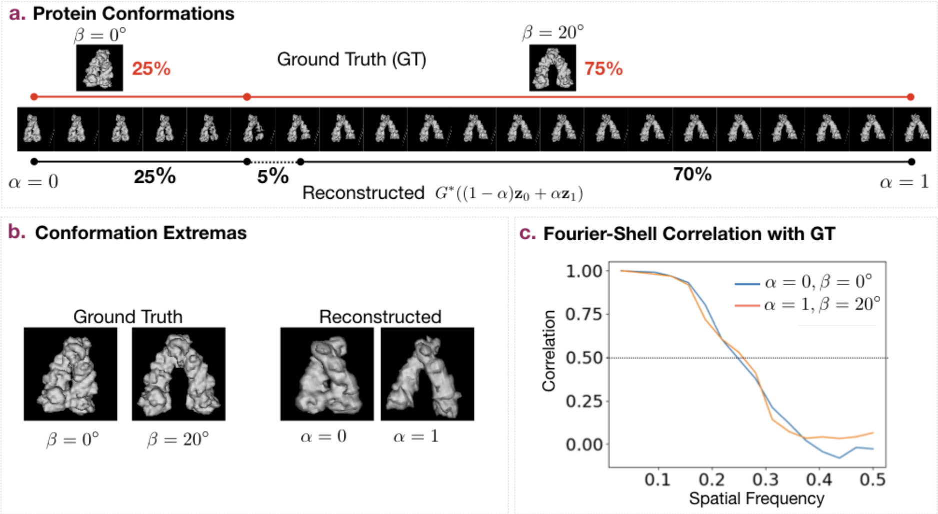

Figure 6.6: Discrete conformations experiment. (a) Comparison between the ground truth (GT) taking only two conformations with probability 0.25 and 0.75 and the reconstructed conformations $G_{\gamma^*}(\mathbf{z})$, where $\mathbf{z} = (1-\alpha)\mathbf{z}_0 + \alpha\mathbf{z}_1$, $\alpha \in [0, 1]$. (b) The GT with bond angles 0° and 20° and their reconstruction. (c) FSC of the structures in (b).

any transition conformations. The structural quality of these two recovered configurations with respect to the corresponding ground truth are given in Figure 6.6(b). Their FSC show that at least half of the Nyquist resolution is achieved (Figure 6.6(c)).

The FSCCC reconstruction matrix (Figure 6.5(b)) validates the fact that the reconstructed structures cluster into two main conformations. We use it to determine the probabilities of these cluster/conformations. We determine the probability of a conformation using its first and last row (similarity of the conformations with respect to the extreme conformations). We consider that a conformation $\mathbf{x}_n$ belongs to the first cluster/conformation if $\mathbf{M}[0, n] > 0.5$ and $\mathbf{M}[20, n] < 0.5$. If the case is reversed ($\mathbf{M}[0, n] < 0.5$ and $\mathbf{M}[20, n] > 0.5$), then it belongs to the second cluster/conformation. Otherwise, it is considered as a transitioning conformation. This yields that the first 25% and the last 70% structures cluster together to form the $\beta = 0^\circ$ and $\beta = 20^\circ$ conformations, respectively, and the middle 5% are the transitioning conformations.

6.7 Summary

In this chapter, we propose a deep-learning-based reconstruction method for cryo-electron microscopy (Cryo-EM) that can model multiple conformations of a nonrigid biomolecule in a standalone manner. Cryo-EM produces many noisy projections from separate instances of the same but randomly oriented biomolecule. Current methods rely on pose and conformation estimation which are inefficient for the reconstruction of continuous conformations that carries valuable information. We introduce Multi-CryoGAN, which sidesteps the additional processing by casting the volume reconstruction into the distribution matching problem. By introducing a manifold mapping module, Multi-CryoGAN can learn continuous structural heterogeneity *without pose estimation* nor *clustering*. We also give a theoretical guarantee of recovery of the true conformations. Our method can successfully reconstruct 3D protein complexes on synthetic 2D Cryo-EM datasets for both continuous and discrete structural variability scenarios. To the best of our knowledge, Multi-CryoGAN is the first model that can reconstruct continuous conformations of a biomolecule from Cryo-EM images in a fully unsupervised and standalone manner.

Algorithm 6 Reconstruction of multiple conformations using Multi-CryoGAN.

```
Input: Dataset $\{\mathbf{y}^1_{\text{data}}, \dots, \mathbf{y}^Q_{\text{data}}\}$; training parameters: number of steps $k$ to apply to the discriminator and penalty parameter $\lambda$
Output: A mapping $G$ from the latent space to the 3D conformation space.
for $n_{\text{train}}$ training iterations do
    for $k$ steps do
        • sample from real data: $\{\mathbf{y}^1_{\text{data}}, \dots, \mathbf{y}^B_{\text{data}}\}$.
        • sample from generated data: $\{\mathbf{y}^1_{\text{gen}}, \dots, \mathbf{y}^B_{\text{gen}}\} \sim p_{\text{gen}}(\mathbf{y}; G_\gamma)$ (see Algorithm 1).
        • sample from $\{\kappa_1, \dots, \kappa_B\} \sim \text{U}[0, 1]$.
        • compute $\mathbf{y}^b_{\text{int}} = \kappa_b \cdot \mathbf{y}^b_{\text{batch}} + (1 - \kappa_b) \cdot \mathbf{y}^b_{\text{gen}}$ for all $b \in \{1, \dots, B\}$.
        • update the discriminator $D_\phi$ by gradient ascent on the loss (6.5) complemented with the gradient penalty term from [167].
    end for
    • sample generated data: $\{\mathbf{y}^1_{\text{gen}}, \dots, \mathbf{y}^B_{\text{gen}}\} \sim p_{\text{gen}}(\mathbf{y}; G_\gamma)$ (see Algorithm 1).
    • update the Generator $G_\gamma$ by gradient descent on the loss (6.5).
end for
return $\mathbf{G}_{\gamma^*}$
```

Chapter 7

Conclusion and Outlook

Reconstruction methods for linear problems range from classical to unsupervised-deep-learning methods. This thesis brings contribution across these methods for inverse problems ranging from 1D sampling to heterogeneous Cryo-EM.

Continuous-domain solutions of inverse problems

In chapter 2, we develop the continuous-domain extension of classical methods and devise numerical approaches to solve them. We have shown that the formulation of continuous-domain linear inverse problems with Tikhonov- and total-variation-based regularizations leads to spline solutions. The nature of these splines is dependent on the Green's function of the regularization operator $(\mathrm{L}^*\mathrm{L})$ and L for Tikhonov and total variation, respectively. The former is better to reconstruct smooth signals; the latter is an attractive choice to reconstruct signals with sparse innovations. Representer theorems for the two cases come handy in the numerical reconstruction of the solution. They allow us to reformulate the infinite-dimensional optimization as a finite-dimensional parameter search.

Future Work

We expect that similar results exist in higher dimensions since the theory can be generalized. However, the computations can also be expected to be challenging for signals defined over $\mathbb{R}^d$ with $d > 1$, for example, when considering images rather than signals. Our scheme to solve the gTV regularized 1D inverse problems has been made more efficient in [205]. This formulation can be useful for cases when a function needs to be optimized with a constraint on its sparsity. For example, in [206] this numerical scheme has been used to learn activations of a neural network in order to increase its capacity while maintaining its stability.

Deep-learning-based iterative methods

We propose a deep-learning-based iterative framework in chapter 3 to solve imaging problems such as CT in order to bring more robustness and quality to the reconstruction. The purpose of this chapter is to develop a simple yet effective iterative scheme (RPGD) where one step of enforcing measurement consistency is followed by a CNN that tries to project the solution onto the set of desired reconstruction images. The whole scheme is ensured to be convergent. We also introduced a novel method to train a CNN that acts like a projector using a reasonably small dataset. For sparse-view CT reconstruction, our method outperforms the previous techniques for both noiseless and noisy measurements.

Future Work

The proposed framework is generic and can be used to solve a variety of inverse problems including superresolution, deconvolution, accelerated MRI, *etc.* This can bring more robustness and reliability to the direct deep-learning-based techniques. Infact, this method has contributed fundamentally in the emergence of the family of such deep-learning-based iterative algorithms [21–25].

Time-dependent deep image prior

In chapter 4, we develop an unsupervised deep learning method for time-dependent inverse problem. To this end, we introduce time-dependent deep image prior which is the first deep learning method to reconstruct dynamic MRI without the requirement of additional data or training. This framework fits well the learning of spatio-temporal manifolds that are smooth temporally; it is purely unsupervised. It is also particularly appropriate in the context of inverse problems where no ground-truth is available. In practice, it results in significant memory savings when compared to compressed-sensing (CS) approaches. Our study shows that our proposed method has the potential to reconstruct dynamic magnetic resonance images (MRI) in the absence of an electrocardiogram signal.

Future Work

The method results in state-of-the-art reconstruction and could be utilized for other similar temporal inverse problems. Recently, the research community has started attempts to theoretically understand deep image prior [207–209]. We believe that this continued effort could bring perspectives that could even further improve the proposed method.

The current bottleneck of our method is the slow forward model, NuFFT. Currently, the NuFFT package[1] is optimized neither for Python nor for GPU usage, which leads to a major reason of slowdown in our implementation. With an efficient implementation, we can speed-up the algorithm multiple times.

In future, we could explore more architectures to further improve the performance. Because the network architecture explored in this work is never exhaustive, there can be some architectural variations that would bring more improvement than the currently reported results. For example, we can explore advanced architectures or good initialization techniques of the network parameters. Similar to [210], which have proposed a progressive learning strategy, we can also progressively increase the spatial resolution of the network and the forward model to achieve better reconstruction. This would not only bring faster convergence but could also result in improved performance and stability. Similar improvements could be achieved by applying the strategy in temporal dimension. This could be achieved, for example, by initializing with a high number of spoke-sharing and decreasing it over the course of optimization. This could be done in parallel with a similar aggregating strategy being followed in the latent variable domain.

In addition, instead of finding an entire image, we can let the network focus on

---

[1] https://github.com/marchdf/python-nufft

finding the residual of the static gold standard image. This facilitates to find the dynamics of the dataset, which can bring better reconstruction [211]. This could result in faster convergence to the solution. The convergence could even further be accelerated by employing temporal and spatial multi-resolution approaches for this routine.

Currently, our framework aims to optimize the network on the measurement space from the beginning. Instead, as a warm start, in future we can first optimize the network on the image space to fit the gold standard image. Because this initial routine would not require the repetitive use of the forward model, it could potentially even speed-up the reconstruction. Currently, the network is initialized randomly which outputs random image series. Instead, it could be initialized to output a more coherent image series. For example, the network could be first directly optimized to output precomputed backprojection images $\mathbf{H}_k^T \mathbf{y}_k$. This could result in faster convergence to the solution. The convergence could even further be accelerated by employing temporal and spatial multi-resolution approaches for this routine. Moreover, since this initial routine wouldn't require the repetitive use of the forward model, it could potentially even speed-up the reconstruction.

CryoGAN for Cryo-EM reconstruction

In chapter 5, we propose CryoGAN which reconstructs the 3D structure of a biomolecule from its multiple noisy tomographic projections obtained from unknown random orientations. The work utilizes the GAN framework to learn a biomolecule whose measurements resemble the acquired measurements in a distributional sense. This method sidesteps the need to involve calculations involving poses. We believe that with better network architecture and multi-resolution approaches this paradigm could produce excellent reconstructions. Moreover, we believe that this work could lead to a completely novel family of algorithms dedicated to improve or even remove certain stages of Cryo-EM pipeline, for example, particle picking.

Future Work

Our implementation of CryoGAN is bound to further improve. Beyond simple engineering tweaks (*e.g.*, tuning the number of layers in the discriminator, testing different optimization strategies, or using Fourier methods to accelerate projection), we expect that several interesting developmental steps lie ahead.

A promising direction of research is the use of a coarse-to-fine strategy to reconstruct the volume progressively as the resolution improves. The motivation is that increased robustness during the low-resolution regime tends to have a positive impact on the convergence of the higher-resolution steps. Several GAN architectures rely on such approaches, such as the progressive GANs [210] and the style-GANs [212]. The benefits of multi-scale reconstruction could be considerable for CryoGAN, given the extremely difficult imaging conditions that prevail in single-particle Cryo-EM and that make the convergence of optimization algorithms to good solutions particularly challenging. The core idea here would be to have the discriminator learn to differentiate between real and simulated distributions at a low resolution first, and then at successively higher ones. The impact on CryoGAN could be as important as the one it had on GANs, which progressed in just a few years from generating blurry facial images [36] to simulated images indistinguishable from real facial images [210, 212]. More generally, the upcoming tools and extensions in GAN architectures could bring significant gain in resolution to the CryoGAN implementation.

The performance of the Cryo-EM physics simulator should also improve hand-in-hand with our ability to precisely model the physics behind single-particle Cryo-EM with computationally tractable entities. At the moment, CryoGAN assumes that the noise is additive in its image-formation model. One could alternatively consider a Poisson-noise-based forward model [162,213]. This would however require backpropagating through a Poisson distribution, a nontrivial operation.

Another interesting extension of the simulator would be to directly simulate the patches of nonaligned micrographs/frames (rather than the individual projections) and match their distribution to that of the raw dataset. Doing so would allow CryoGAN to bypass additional preprocessing tasks, in particular particle picking.

Similar to likelihood-based methods, the CryoGAN algorithm requires the specification of the distribution of poses. In the case of CryoGAN, one could also parameterize this distribution and learn its parameters during the reconstruction procedure, along the lines of [177]. The same approach could be used to calibrate the distribution of the translations of the projections.

On the theoretical side, we currently have mathematical guarantees on the recovery of volumes for which the assumed distribution of poses (be it uniform or not) matches the distribution of the real data. We have prior mathematical indications that this can also be achieved when there is a certain mismatch between the assumed distribution of poses and the actual one, given that an appropriate GAN loss is used.

Like all reconstruction algorithms, CryoGAN can fail if the dataset contains too many corrupted particle images, *e.g.*, those with broken structures or strong optical aberrations. Several solutions could be deployed to handle excessive outliers in the data distribution. One approach would be to include a step that automatically spots and discards corrupted data so that the discriminator never sees them. Recent deep-learning-based approaches able to track outliers in data could prove useful in this regard [188].

While the spatial resolution of the CryoGAN reconstructions from real data is not yet competitive with the state-of-the-art, the algorithm is already able to steadily perform the hardest part of the job, which is to obtain a reasonable structure by using nothing save the particle dataset and CTF estimations. We believe that the aforementioned developments will help to bring the CryoGAN algorithm to the stage where it becomes a relevant contributor for high-resolution reconstruction in single-particle Cryo-EM. We have laid out a roadmap for future improvements that should get us to this stage, and may eventually help us reconstruct dynamic structures.

Multi-CryoGAN for reconstructing continuous conformations in Cryo-EM

We propose an extension of CryoGAN called Multi-CryoGAN in chapter 6, which is applicable to the much sought after problem of heterogeneous Cryo-EM. This method reconstructs continuous conformations of a biomolecule without conformation estimation. This is the first method that can perform this task in a standalone manner and has no limit on the complexity of conformation landscape it can reconstruct. By matching the simulated Cryo-EM projections with the acquired data distribution, this method naturally learns to generate a set of 3D conformations in a likelihood-free way. This allows us to reconstruct both continuous and discrete conformations without any prior assumption on the conformation landscape, data preprocessing steps, nor external algorithms such as pose estimation. Our experiments shows that Multi-CryoGAN successfully recovers the molecular conformation manifold, including the underlying distribution.

Future Work

The current experiments have been conducted on synthetic data with simple dynamics. In future, the method could be deployed on real data. We believe that, with a better incorporation of state-of-the art GAN architectures [210,212] and the future work intended for Cryo-GAN, this method could provide excellent results on challenging datasets. Moreover, we also hope that this work will bring the interest of computer vision researchers to the problem of Cryo-EM. In conclusion, with suitable extension this method could, for the first time, let biologists reconstruct the true dynamics of biomolecules in a user friendly and robust manner.

Appendix A

Appendices

A.1 Chapter 2

A.1.1 Proof of Theorem 2.4.1

Abstract Representer Theorem

The result presented in this section is preparatory to Theorem 2.4.1. It is classical for Hilbert spaces [47, Theorem 16.1]. We give its proof for the sake of completeness.

Theorem A.1.1. *Let $\mathcal{X}$ be a Hilbert space equipped with the inner product $\langle\cdot,\cdot\rangle_{\mathcal{X}}$ and let $h_1,\dots,h_M \in \mathcal{X}'$ be linear and continuous functionals. Let $\mathcal{C} \in \mathbb{R}^M$ be a feasible convex compact set, meaning that there exists at least a function $f \in \mathcal{X}$ such that $\mathrm{H}\{f\} \in \mathcal{C}$. Then, the minimizer*

$$f^* = \arg\min_{f\in\mathcal{X}} \|f\|_{\mathcal{X}}^2 \ \textit{s.t.}\ \mathrm{H}\{f\} \in \mathcal{C} \tag{A.1}$$

exists, is unique, and can be written as

$$f^* = \sum_{m=1}^{M} a_m h_m^{\#} \tag{A.2}$$

for some $\{a_m\}_{m=1}^M \in \mathbb{R}$, where $h_m^{\#} = \Pi h_m$ and $\Pi : \mathcal{X}' \to \mathcal{X}$ is the Riesz map of $\mathcal{X}$.

Proof. The feasibility of the set $\mathcal{C}$ implies that the set $\mathcal{C}_{\mathcal{X}} = \mathrm{H}^{-1}(\mathcal{C}) = \{f \in \mathcal{X} : \mathrm{H}\{f\} \in \mathcal{C}\} \in \mathcal{X}$, is nonempty. Since H is linear and bounded and since $\mathcal{C}$ is convex and compact, its preimage $\mathcal{C}_{\mathcal{X}}$ is also convex and closed. By Hilbert's projection theorem [214], the solution f^* exists and is unique as it is the projection of the null function onto $\mathcal{C}_{\mathcal{X}}$. Let the measurement of this unique point f^* be $\mathbf{z}_0 = \mathrm{H}\{f^*\}$. The Riesz representation theorem states that $\langle h_m, f\rangle = \langle h_m^{\#}, f\rangle_{\mathcal{X}}$ for every $f \in \mathcal{X}$, where $h_m^{\#} \in \mathcal{X}$ is the unique Riesz conjugate of the functional h_m. We then uniquely decompose f^* as $f^* = f^{*\perp} + \sum_{m=1}^{M} a_m h_m^{\#}$, where $f^{*\perp}$ is orthogonal to the span of the $h_m^{\#}$ with respect to the inner product on $\mathcal{X}$ *i.e.*, $\mathrm{H}\{f^{*\perp}\} = \mathbf{0}$. The orthogonality also implies that

$$\|f^*\|_{\mathcal{X}}^2 = \|f^{*\perp}\|_{\mathcal{X}}^2 + \left\|\sum_{m=1}^{M} a_m h_m^{\#}\right\|_{\mathcal{X}}^2 . \tag{A.3}$$

This means that the minimum norm is reached when $f^{*\perp} = 0$ while keeping $\mathrm{H}\{f^*\} = \mathbf{z}_0$, implying the form (A.2) of the solution.

□

Proof of Theorem 2.4.1

The proof of Theorem 2.4.1 has two steps. We first show that if Assumption 2 holds, then there is at least one solution and, moreover, if Assumption 2' holds, then the solution is unique. After this, we use Theorem A.1.1 to deduce the form of the solution.

Existence of the Solution. We use the classical result on Hilbert spaces which states that a proper, coercive, lsc, and convex objective functional over a Hilbert space has a nonempty and convex set of minimizers [215].

Properness: By Assumption 2, $E(\mathbf{z}, \cdot)$ is proper. The regularization $\|\mathrm{L}f\|_{L_2}^2$ is proper by the definition of $\mathcal{X}_2$. This means that $\mathrm{J}_2(\cdot|\mathbf{z})$ is proper in $\mathcal{X}_2$.

Lower semi-continuity: $E(\mathbf{z}, \cdot)$ is lsc in $\mathbb{R}^M$, and $\mathrm{H} : \mathcal{X}_2 \to \mathbb{R}^M$ is continuous. Therefore, $E(\mathbf{z}, \mathrm{H}\{\cdot\})$ is lsc over $\mathcal{X}_2$. Similarly, by composition $f \mapsto \|\mathrm{L}f\|_{L_2}$ is continuous, hence lsc over $\mathcal{X}_2$. Since $\mathrm{J}_2(\cdot|\mathbf{z})$ is the sum of two lsc functionals, it is lsc as well.

Convexity: $E(\mathbf{z}, \cdot)$ and $\|\cdot\|_{L_2}^2$ are convex, and H and L are linear. Therefore, $\mathrm{J}_2(\cdot|\mathbf{z}) = E(\mathbf{z}, \mathrm{H}\{\cdot\}) + \|\mathrm{L}\cdot\|_{L_2}^2$ is convex too.

Coercivity: The measurement operator H is continuous and linear from $\mathcal{X}_2$ to

$\mathbb{R}^M$; hence, there exists a constant A such that

$$\|\mathrm{H}\{f\}\|_2 \leq A\|f\|_{\mathcal{X}_2} \tag{A.4}$$

for every $f \in \mathcal{X}_2$. Likewise, the condition $\mathrm{H}\{p\} = \mathrm{H}\{q\} \Rightarrow p = q$ for $p, q \in \mathcal{N}_{\mathrm{L}}$ implies the existence of $B > 0$ such that [46, Proposition 8]

$$\|\mathrm{H}\{p\}\|_2 \geq B\|p\|_{\mathcal{N}_{\mathrm{L}}} \tag{A.5}$$

for every $p \in \mathcal{N}_{\mathrm{L}}$. As presented in the supplementary material (see [46] for more details), the search space $\mathcal{X}_2$ is a Hilbert space for the Hilbertian norm

$$\|f\|_{\mathcal{X}_2} = \|\mathrm{L}f\|_{L_2} + \|\mathrm{P}f\|_{\mathcal{N}_{\mathrm{L}}} \tag{A.6}$$

with P being the projector on $\mathcal{N}_{\mathrm{L}}$ introduced in (A.23). We set $p = \mathrm{P}f$ and $g = f - p$. Then, $g \in \mathcal{X}_2$ satisfies $\mathrm{L}g = \mathrm{L}f$ and $\mathrm{P}g = 0$, and hence

$$\|g\|_{\mathcal{X}_2} = \|\mathrm{L}g\|_{L_2} + \|\mathrm{P}g\|_{\mathcal{N}_{\mathrm{L}}} = \|\mathrm{L}f\|_{L_2}. \tag{A.7}$$

Now consider a sequence of (generalized) functions $f_m \in \mathcal{X}_2, m \in \mathbb{N}$ such that $\|f_m\|_{\mathcal{X}_2} \to \infty$. We set $p_m = \mathrm{P}f_m$ and $g_m = f_m - p_m$. Assume by contradiction that $\mathrm{J}_2(f_m|\mathbf{z})$ is bounded. Then, $\|\mathrm{L}f_m\|_{L_2}$ and $\|\mathrm{H}\{f_m\}\|_2$ are bounded (for the latter, we use that $E(\mathbf{z}, \cdot)$ is coercive). However, we have

$$\begin{aligned} \|\mathrm{H}\{f_m\}\|_2 &\geq \|\mathrm{H}\{p_m\}\|_2 - \|\mathrm{H}\{g_m\}\|_2 && \text{(A.8)} \\ &\geq B\|p_m\|_{\mathcal{N}_{\mathrm{L}}} - A\|g_m\|_{\mathcal{X}_2} && \text{(A.9)} \\ &= B\|f_m\|_{\mathcal{X}_2} - (A+B)\|\mathrm{L}f_m\|_{L_2} && \text{(A.10)} \end{aligned}$$

where we used respectively the triangular inequality in (A.8), the inequalities (A.4) and (A.5) in (A.9), and the relations $\|p_m\|_{\mathcal{N}_{\mathrm{L}}} = \|f_m\|_{\mathcal{X}_2} - \|\mathrm{L}f_m\|_{L_2}$ and $\|g_m\|_{\mathcal{X}_2} = \|\mathrm{L}f_m\|_{L_2}$ in (A.10). Since $\|\mathrm{L}f_m\|_{L_2}$ is bounded and $\|f_m\|_{\mathcal{X}_2} \to \infty$, we deduce that $\|\mathrm{H}\{f_m\}\|_2 \to \infty$, which is known to be false. Finally, we obtain a contradiction, proving the coercivity.

Since, $\mathrm{J}_2(\cdot|\mathbf{z})$ is proper, lsc, convex, and coercive on $\mathcal{X}_2$, therefore, it has at least one minimizer.

Uniqueness of the Solution. We now prove that if $E(\mathbf{z}, \cdot)$ satisfies Assumption 2' then the solution is unique. We first show this for the case when Assumption

2'.i) is satisfied. We already know that the solution set is nonempty. It is then clear that the uniqueness is achieved if $\mathrm{J}_2(\cdot|\mathbf{z})$ is strictly convex. We now prove the convex functional $\mathrm{J}_2(\cdot|\mathbf{z})$ is actually strictly convex.

For $\beta \in (0,1)$, $f_A, f_B \in \mathcal{X}_2$, we denote $f_{AB} = \beta f_A + (1-\beta) f_B$. Then, the equality case $\mathrm{J}_2(f_{AB}|\mathbf{z}) = \beta \mathrm{J}_2(f_A|\mathbf{z}) + (1-\beta)\mathrm{J}_2(f_B|\mathbf{z})$ implies that $E(\mathbf{z}, f_{AB}) = \beta E(\mathbf{z}, f_A) + (1-\beta) E(\mathbf{z}, f_B)$ and $\|\mathrm{L} f_{AB}\|_{L_2} = \beta \|\mathrm{L} f_B\|_{L_2} + (1-\beta)\|\mathrm{L} f_B\|_{L_2}$, since the two parts of the functional are themselves convex. The strict convexity of $E(\boldsymbol{z}, \cdot)$ and the norm $\|\cdot\|_2$ then implies that

$$\mathrm{L} f_A = \mathrm{L} f_B \text{ and } \mathrm{H}\{f_A\} = \mathrm{H}\{f_B\} \tag{A.11}$$

and, therefore, $(f_A - f_B) \in \mathcal{N}_{\mathrm{L}} \cap \mathcal{N}_{\mathrm{H}}$. Since $\mathcal{N}_{\mathrm{L}} \cap \mathcal{N}_{\mathrm{H}} = \mathbf{0}$ by Assumption 1, therefore, $f_A = f_B$. This demonstrates that $\mathrm{J}_2(\cdot|\mathbf{z})$ is strictly convex.

For Assumption 2'.ii), that is when $E(\mathbf{z}, \cdot) = I(\mathbf{z}, \cdot)$, the solution set can be written as

$$\mathcal{V}_2 = \arg \min_{f \in \mathrm{H}^{-1}\{\mathbf{z}\}} \|\mathrm{L} f\|_{L_2}^2. \tag{A.12}$$

where the set $\mathrm{H}^{-1}\{\mathbf{z}\} = \{f \in \mathcal{X}_2 : \mathrm{H}\{f\} = \mathbf{z}\}$ is nonempty since we assumed $I(\mathbf{z}, \cdot)$ to be proper in the range of H.

According to [52, Theorem 1.1 and 1.2] given that the range of $\mathrm{L} : \mathcal{X}_2 \to L_2$ is closed in L_2, $\mathcal{V}_2$ in (A.12) is singleton . As discussed in the supplementary material, given any $w \in L_2$, we can always find an $f \in \mathcal{X}_2$ such that $\mathrm{L} f = w$. This means that the range of L is the whole L_2 which is clearly closed in L_2.

Form of the Minimizer. We first take the case when E satisfies Assumption 2'. Let f_2^* be the unique solution and $\mathbf{z}_0 = \mathrm{H}\{f_2^*\}$. One decomposes again $\mathcal{X}_2$ as the direct sum $\mathcal{X}_2 = \mathcal{Q} \oplus \mathcal{N}_{\mathrm{L}}$, where

$$\mathcal{Q} = \{f \in \mathcal{X}_2 : \ \langle f, p \rangle_{\mathcal{X}_2} = 0, \forall p \in \mathcal{N}_{\mathrm{L}}\}$$

is the Hilbert space with norm $\|\mathrm{L} \cdot\|_{L_2}$. In particular, we have that $f_2^* = q^* + p^*$ with $q^* \in \mathcal{Q}$ and $p^* \in \mathcal{N}_{\mathrm{L}}$.

Consider the optimization problem

$$\min_{g \in \mathcal{Q}} \|\mathrm{L} g\|_{L_2}^2 \text{ s.t. } \mathrm{H}\{g\} = (\mathbf{z}_0 - \mathrm{H}\{p^*\}). \tag{A.13}$$

According to Theorem A.1.1, this problem admits a unique minimizer g^* such that $\Pi^{-1}g^* \in \mathcal{Q}' \cap \mathrm{Span}\{h_m\}_{m=1}^M$ where $\Pi^{-1} : \mathcal{X} \to \mathcal{X}'$ is the inverse of the Riesz map $\Pi : \mathcal{X}' \to \mathcal{X}$ and $Q' = \Pi^{-1}\mathcal{Q}$. The set $\mathcal{Q}' \cap \mathrm{Span}\{h_m\}_{m=1}^M$ is represented by $\sum_{m=1}^M a_m h_m$, where $\sum_m a_m \langle h_m, p\rangle = 0$ for every $p \in \mathcal{N}_\mathrm{L}$.

However, by definition, the function q^* also satisfies $\mathrm{H}\{q^*\} = (\mathbf{z}_0 - \mathrm{H}\{p^*\})$. Moreover, $\|\mathrm{L}q^*\|_{L_2}^2 \leq \|\mathrm{L}g^*\|_{L_2}^2$; otherwise, the function $\tilde{f} = g^* + p^* \in \mathcal{X}_2$ would satisfy $J_2(\tilde{f}|\mathbf{z}) < J_2(f_2^*|\mathbf{z})$, which is impossible. However, since (A.13) has a unique solution, we have $q^* = g^*$.

This proves that $f_2^* = \Pi\left\{\sum_{m=1}^M a_m h_m\right\} + p^*$. For any $q' \in \mathcal{Q}'$ the Riesz map $\Pi q' = q' * \rho_{\mathrm{L}*\mathrm{L}} + p_{q'}$ for some $p_{q'} \in \mathcal{N}_\mathrm{L}$ [47, 52]. Here $\rho_{\mathrm{L}^*\mathrm{L}}$ is the Green's function of the operator $(\mathrm{L}^*\mathrm{L})$ (see Definition 2.2.1). Therefore,

$$f_2^* = p_0 + \rho_{\mathrm{L}^*\mathrm{L}} * \left\{\sum_{m=1}^M a_m h_m\right\} \tag{A.14}$$

where $p_0 = (p_{q'} + p^*) \in \mathcal{N}_\mathrm{L}$ and where $\sum_m a_m \langle h_m, p\rangle = 0$ for every $p \in \mathcal{N}_\mathrm{L}$.

The component $\rho_{\mathrm{L}^*\mathrm{L}} * \left\{\sum_{m=1}^M a_m h_m\right\}$ in (A.14) can be written as, $\sum_{m=1}^M a_m \varphi_m$ provided that $\varphi_m = \rho_{\mathrm{L}^*\mathrm{L}} * h_m = \mathrm{F}^{-1}\left\{\frac{\widehat{h}_m}{|\widehat{\mathrm{L}}|^2}\right\}$ is well-defined. To show that this is the case, we decompose $h_m = \mathrm{Proj}_{\mathcal{Q}'}\{h_m\} + \mathrm{Proj}_{\mathcal{N}_\mathrm{L}'}\{h_m\}$ where $\mathrm{Proj}_{\mathcal{Q}'}$ and $\mathrm{Proj}_{\mathcal{N}_\mathrm{L}'}$ are the projection operators on $\mathcal{Q}'$ and $\mathcal{N}_\mathrm{L}'$, respectively. Since, $\mathrm{Proj}_{\mathcal{Q}'}\{h_m\} \in \mathcal{Q}'$, as discussed earlier, $\rho_{\mathrm{L}*\mathrm{L}} * \mathrm{Proj}_{\mathcal{Q}'}\{h_m\}$ is well-defined.

Now, one can always select a basis $\{p_n\}_{n=1}^{N_0}$ such that $\mathcal{N}_\mathrm{L}' = \mathrm{Span}\{\phi_n\}_{n=1}^{N_0}$ with $\phi_n = \delta(\cdot - x_n)$ and $\langle \phi_m, p_n\rangle = \delta[m-n]$. The other component $\mathrm{Proj}_{\mathcal{N}_\mathrm{L}'}\{h_m\} = \sum_{n=1}^{N_0} c_n \phi_n$ where $c_n = \langle h_m, p_n\rangle$. Therefore, $\rho_{\mathrm{L}*\mathrm{L}} * \mathrm{Proj}_{\mathcal{N}_\mathrm{L}'}\{h_m\}$ is a linear combination of shifted Green's functions, which proves that $\varphi_m = \mathrm{F}^{-1}\left\{\frac{\widehat{h}_m}{|\widehat{\mathrm{L}}|^2}\right\}$ is well defined.

For general case, when Assumption 2 is satisfied, we see that any solution $f_2^* \in \mathcal{V}_2$ also minimizes the following

$$\min_{f \in \mathrm{H}^{-1}\{\mathrm{H}\{f_2^*\}\}} \|\mathrm{L}f\|_{L_2}. \tag{A.15}$$

As discussed earlier, the minimizer of (A.15) is unique so that it is clearly f_2^*. We now use the same reasoning as for the cases of Assumption 2' to show that f_2^* takes the form (2.16). This concludes the proof.

Note that, even in the absence of convexity of $E(\mathbf{z}, \cdot)$, results on the form of the solution can still be obtained.

A.1.2 Proof of Theorem 2.4.2

Similarly to the L_2 case, the proof has two steps. We first show that the set of minimizers is nonempty. We then connect the optimization problem to the one studied in [46, Theorem 2] to deduce the form of the extreme points. The functional to minimize is $\mathrm{J}_1(f|\mathbf{z}) = E(\mathbf{z}, \mathrm{H}\{f\}) + \lambda \|\mathrm{L}f\|_{\mathcal{M}}$, defined over the Banach space $\mathcal{X}_1$.

Existence of Solutions. We first show that $\mathcal{V} = \arg\min_{f \in \mathcal{X}_1} \mathrm{J}_1(f|\mathbf{z})$ is nonempty, convex, and weak*-compact.

We rely on the generalized Weierstrass theorem presented in [215]: Any proper, lower semi-continuous (lsc) functional over a compact topological vector space reaches its minimum, from which we deduce the following result. We recall that the dual space $\mathcal{B}'$ of a Banach space $\mathcal{B}$ can be endowed with the weak*-topology, and that one can define a norm $\|f\|_{\mathcal{B}'} = \sup_{\|x\|_{\mathcal{B}}} \langle f, x \rangle$ for which $\mathcal{B}'$ is a Banach space.

Proposition A.1.2. *Let $\mathcal{B}$ be a Banach space. Then, a functional $\mathrm{J} : \mathcal{B}' \to \mathbb{R}^+ \cup \{\infty\}$ which is proper, convex, coercive, and weak*-lsc is lower bounded and reaches its infimum. Moreover, the set $\mathcal{V} = \arg\min \mathrm{J}$ is convex and weak*-compact.*

Proof. Let $\alpha > \inf \mathrm{J}$. The coercivity implies that there exists $r > 0$ such that $\mathrm{J}(f) \geq \alpha$ as soon as $\|f\|_{\mathcal{B}'} > r$. The infimum of J can only be reached on $B_r = \{f \in \mathcal{B}', \|f\|_{\mathcal{B}'} \leq r\}$, hence we restrict our analysis to it. The Banach-Alaoglu theorem implies that B_r is weak*-compact. As a consequence, the functional J is proper and lsc on the compact space B_r endowed with the weak*-topology. According to the generalized Weierstrass theorem [215, Theorem 7.3.1], J reaches its infimum on B_r, hence on $\mathcal{X}'$.

Let $\mathcal{V} = \arg\min \mathrm{J}$ and $\alpha_0 = \min \mathrm{J}$. The convexity of J directly implies the one of $\mathcal{V}$. The set $\mathcal{V}$ is included in the ball B_{α_0} which is weak*-compact. Therefore, it suffices to show that $\mathcal{V}$ is weak*-closed to deduce that it is weak*-compact. Moreover, the weak*-lower semi-continuity is equivalent to the weak*-closedness of the level sets $\{f \in \mathcal{B}' : \mathrm{J}(f) \leq \alpha\}$ are weak*-closed. Applying this to $\alpha = \alpha_0$, we deduce that $\mathcal{V} = \{f \in \mathcal{B}' : \mathrm{J}(f) \leq \alpha_0\}$ is weak*-closed, as expected. ☐

We apply Proposition A.1.2 to $\mathcal{B}' = \mathcal{X}_1$, which is the dual of the Banach space

$\mathcal{B} = C_{\mathrm{L}}(\mathbb{R})$ introduced in [46] and recapped in the supplementary material. One has to show that the functional $\mathrm{J} = \mathrm{J}_1(\cdot|\mathbf{z})$ is coercive and weak*-lsc over $\mathcal{X}_1$. The coercivity is deduced exactly in the same way as for Theorem 2.4.1. The weak*-lower semi-continuity is deduced as follows. First, H is weak*-continuous by assumption and $E(\mathbf{z}, \cdot)$ is lsc; hence, the composition $f \mapsto E(\mathbf{z}, \mathrm{H}\{f\})$ is weak*-lsc. Similarly, the norm $\|\cdot\|_{\mathcal{M}}$ is weak*-lsc on $\mathcal{M}$ and $\mathrm{L} : \mathcal{X}_1 \to \mathcal{M}$ is continuous, hence $f \mapsto \|\mathrm{L}f\|_{\mathcal{M}}$ is weak*-continuous, and therefore weak*-lsc over $\mathcal{X}_1$. Finally, $\mathrm{J}_1(\cdot|\mathbf{z})$ is weak*-lsc over $\mathcal{X}_1$ as it is a sum of two weak*-lsc functionals.

Form of the Extreme Points. Let f_e be an extreme point of the set $\mathcal{V}_1$ and $\mathbf{z}_e = \mathrm{H}f_e$. Then f_e is also a member of the solution set

$$\mathcal{V}_e = \arg \min_{f=\mathrm{H}^{-1}\{\mathbf{z}_e\}} \|\mathrm{L}f\|_{\mathcal{M}}. \tag{A.16}$$

Since $\mathbf{z}_e$ is convex and compact, and the set $\mathrm{H}^{-1}\{\mathbf{z}_e\}$ is nonempty, we can apply Theorem 2 of [46] to deduce that $\mathcal{V}_e$ is convex and weak*-compact, together with the general form (2.19) of the extreme-points of $\mathcal{V}_e$.

Since $\mathcal{V}_e \subseteq \mathcal{V}_1$, and $f_e \in \mathcal{V}_e$ it can be easily shown that f_e is also an extreme point of $\mathcal{V}_e$. This proves that the extreme points of $\mathcal{V}_1$ admit the form (2.19).

Measurement of the solution set. We now show that in the case of Assumption 2' the measurement of the solution set is unique. We first prove this for the case of Assumption 2'.i). Let J_1^* be the minimum value attained by the solutions. Let f_A^* and f_B^* be two solutions. Let e_A, e_B be their corresponding E functional value and let r_A, r_B be their corresponding regularization functional value. Since the cost function is convex, any convex combination $f_{AB} = \beta f_A^* + (1-\beta) f_B^*$ is also a solution for $\beta \in [0, 1]$ with functional value J_1^*. Let us assume that $\mathrm{H}\{f_A^*\} \neq \mathrm{H}\{f_B^*\}$. Since $E(\mathbf{z}, \cdot)$ is strictly convex and $\mathrm{R}_1(f) = \|\mathrm{L}f\|_{\mathcal{M}}$ is convex, we get that

$$\begin{aligned} rcl J_1^* &= E(\mathbf{z}, \mathrm{H}\{\beta f_A^* + (1-\beta) f_B^*\}) + \lambda \mathrm{R}_1(\beta f_A^* + (1-\beta) f_B^*) \\ &< \underbrace{\beta e_A + (1-\beta) e_B + \lambda \beta r_A + \lambda (1-\beta) r_B}_{J_1^*}. \end{aligned}$$

This is a contradiction. Therefore, $\mathrm{H}\{f_A^*\} = \mathrm{H}\{f_B^*\} = \mathrm{H}\{f_{AB}\}$.

In the case of Assumption 2'.ii), $E(\mathbf{z}, \cdot)$ is an indicator function. It is therefore obvious that all the solutions have the same measurement $\mathbf{z}$.

A.1.3 Proof of Theorem 2.6.2

We first state two propositions that are needed for the proof. Their proofs are given after this section.

Proposition A.1.3 (Adapted from [45, Theorem 5])**.** *Let* $\mathbf{z} \in \mathbb{R}^M$ *and* $\mathbf{H} \in \mathbb{R}^{M \times N}$, *where* $M < N$*. Then, the solution set* α_λ *of*

$$\mathbf{a}^* = \arg \min_{\mathbf{a} \in \mathbb{R}^N} \left(\|\mathbf{z} - \mathbf{H}\mathbf{a}\|_2^2 + \lambda \|\mathbf{a}\|_1 \right) \tag{A.17}$$

is a compact convex set and $\|\mathbf{a}\|_0 \leq M$, $\forall \mathbf{a} \in \alpha_{E,\lambda}$, *where* $\alpha_{E,\lambda}$ *is the set of the extreme points of* α_λ*.*

Proposition A.1.4. *Let the convex compact set* α_λ *be the solution set of Problem (2.48) and let* $\alpha_{E,\lambda}$ *be the set of its extreme points. Let the operator* $\mathrm{T} : \alpha_\lambda \to \mathbb{R}^N$ *be such that* $\mathrm{T}\mathbf{a} = \mathbf{u}$ *with* $u_m = |a_m|, m \in [1, \ldots, N]$*. Then, the operator is linear and invertible over the domain* α_λ *and the range* $\mathrm{T}\alpha_\lambda$ *is convex compact such that the image of any extreme point* $\mathbf{a}_E \in \alpha_{E,\lambda}$ *is also an extreme point of the set* $\mathrm{T}\alpha_\lambda$*.*

The linear program corresponding to (2.50) is

$$\begin{aligned} (\mathbf{a}^*, \mathbf{u}^*) = \min_{\mathbf{a},\mathbf{u}} \sum_{n=1}^{N} u_n, \text{ subject to } \mathbf{u} + \mathbf{a} &\geq \mathbf{0}, \\ \mathbf{u} - \mathbf{a} &\geq \mathbf{0}, \\ \mathbf{P}\mathbf{a} &= \mathbf{z}. \end{aligned} \tag{A.18}$$

By putting $\mathbf{u} + \mathbf{a} = \mathbf{s}_1$ and $(\mathbf{u} - \mathbf{a}) = \mathbf{s}_2$, the standard form of this linear program is

$$\begin{aligned} (\mathbf{s}_1^*, \mathbf{s}_2^*) = \min_{\mathbf{s}_1, \mathbf{s}_2} \left(\sum_{n=1}^{N} s_{1n} + s_{2n} \right), \text{ s.t. } \mathbf{s}_1 &\geq \mathbf{0}, \\ \mathbf{s}_2 &\geq \mathbf{0}, \\ \mathbf{P}\mathbf{s}_1 - \mathbf{P}\mathbf{s}_2 &\leq \mathbf{z} \\ -\mathbf{P}\mathbf{s}_1 + \mathbf{P}\mathbf{s}_2 &\leq -\mathbf{z}. \end{aligned} \tag{A.19}$$

Any solution $\mathbf{a}^*$ of (A.18) is equal to $(\mathbf{s}_1^* - \mathbf{s}_2^*)$ for some solution pair (A.19). We denote the concatenation of any two independent points $\mathbf{s}_1^r, \mathbf{s}_2^r \in \mathbb{R}^N$ by the variable $\mathbf{s}^r = (\mathbf{s}_1^r, \mathbf{s}_2^r) \in \mathbb{R}^{2N}$. Then, the concatenation of the feasible pairs $\mathbf{s}^f = \left(\mathbf{s}_1^f, \mathbf{s}_2^f\right)$

that satisfies the constraints of the linear program (A.19) forms a polytope in $\mathbb{R}^{2N}$. Given that (A.19) is solvable, it is known that at least one of the extreme points of this polytope is also a solution. The simplex algorithm is devised such that its solution $\mathbf{s}^*_{\mathrm{SLP}} = \left(\mathbf{s}^*_{1,\mathrm{SLP}}, \mathbf{s}^*_{2,\mathrm{SLP}}\right)$ is an extreme point of this polytope [68]. Our remaining task is to prove that $\mathbf{a}^*_{\mathrm{SLP}} = \left(\mathbf{s}^*_{1,\mathrm{SLP}} - \mathbf{s}^*_{2,\mathrm{SLP}}\right)$ is an extreme point of the set α_λ, the solution set of the problem (2.48).

Proposition A.1.3 claims that the solution set α_λ of the LASSO problem is a convex set with extreme points $\alpha_{E,\lambda} \in \mathbb{R}^N$. As α_λ is convex and compact, the concatenated set $\zeta = \{\boldsymbol{w} \in \mathbb{R}^{2N} : \boldsymbol{w} = (\mathbf{a}^*, \mathbf{u}^*), \mathbf{a}^* \in \alpha_\lambda\}$ is convex and compact by Proposition A.1.4. The transformation $(\mathbf{a}^*, \mathbf{u}^*) = (\mathbf{s}^*_1 - \mathbf{s}^*_2, \mathbf{s}^*_1 + \mathbf{s}^*_2)$ is linear and invertible. This means that the solution set of (A.19) is convex and compact, too. The simplex solution corresponds to one of the extreme points of this convex compact set.
Since the map $(\mathbf{a}^*, \mathbf{u}^*) = (\mathbf{s}^*_1 - \mathbf{s}^*_2, \mathbf{s}^*_1 + \mathbf{s}^*_2)$ is linear and invertible, it also implies that an extreme point of the solution set of (A.19) corresponds to an extreme point of ζ. Proposition A.1.4 then claims that this extreme point of ζ corresponds to an extreme point $\mathbf{a}_{\mathrm{SLP}} \in \alpha_{\lambda,E}$.

A.1.4 Proof of Proposition A.1.3

Using Lemma 2.6.1, it is clear that α_λ is also a solution set of

$$\alpha_\lambda = \arg\min \|\mathbf{a}\|_1 \quad \text{s.t.} \quad \mathbf{Ha} = \mathbf{z}_{0,\lambda} \tag{A.20}$$

for some $\mathbf{z}_{0,\lambda}$. The solution of the problem akin to (A.20) has been discussed in [45] and is proven to be convex and compact such that the extreme points $\alpha_{E,\lambda}$ of the convex set α_λ satisfy $\|\mathbf{a}\|_0 \leq M$ for any $\mathbf{a} \in \alpha_{E,\lambda}$.

A.1.5 Proof of Proposition A.1.4

Proof. We use the Karush-Kuhn-Tucker (KKT) conditions for the lasso problem derived in [70]. For a given $\mathbf{a} \in \alpha_\lambda$, these conditions state that their exists a

$\gamma \in \mathbb{R}^N$, such that

$$\mathbf{H}^T(\mathbf{z} - \mathbf{H}\mathbf{a}) = \lambda\gamma, \text{ and} \tag{A.21}$$

$$\gamma_m \in \begin{cases} \operatorname{sign}(a_m), & \text{if } a_m \neq 0 \\ [-1,1], & \text{if } a_m = 0, \end{cases} \tag{A.22}$$

for any $m \in [1, \dots, N]$. The γ is unique since $\mathbf{Ha}$ is unique for all $\mathbf{a} \in \alpha_\lambda$. Condition (A.22) implies that $|a_m| = \gamma_m . a_m$ for any $m \in [1, \dots, N]$ and $\mathbf{a} \in \alpha_\lambda$.

Therefore, for any $\mathbf{a} \in \alpha_\lambda$, $\mathrm{T}\mathbf{a} = \mathbf{R}\mathbf{a}$, where $\mathbf{R} \in \mathbb{R}^{N\times N}$ is a diagonal matrix with entries $\mathbf{R}_{mm} = \gamma_m$. Thus, the operation of T is linear in the domain α_λ. Also, $\mathbf{a} = \mathbf{RRa}$ for $\mathbf{a} \in \alpha$ implying that the operator T is invertible.
This ensures that the image of the convex compact set $\mathrm{T}\alpha_\lambda$ is also convex compact and the image of any extreme point $\mathbf{a}_E \in \alpha_{E,\lambda}$ is also an extreme point of the set $\mathrm{T}\alpha_\lambda$. Similarly, it can be proved that the concatenated set $\zeta = \{\boldsymbol{w} \in \mathbb{R}^{2N} : \boldsymbol{w} = (\mathbf{a}, \mathrm{T}\mathbf{a}), \mathbf{a} \in \alpha_\lambda\}$ is the image of a linear and invertible concatenation operation on α. Thus, it is convex and compact, and the image of any extreme point through the inverse operation of the concatenation $w_E \in \zeta_{E,\lambda}$ is also an extreme point of α_λ. □

A.1.6 Structure of the Search Spaces

Decomposition of $\mathcal{X}_1$ and $\mathcal{X}_2$. The set $\mathcal{X}_1$ is the search space, or native space, for the gTV case. It is defined and studied in [46, Section 6], from which we recap the main results. Note that the same construction is at work for $\mathcal{X}_2$, which is then a Hilbert space.

Let $\boldsymbol{p} = (p_1, \dots, p_{N_0})$ be a basis of the finite-dimensional null space $\mathcal{N}_{\mathrm{L}}$ of L. If $\boldsymbol{\phi} = (\phi_1, \dots, \phi_{N_0})$ and $\boldsymbol{p} = (p_1, \dots, p_{N_0})$ form a biorthonormal system such that $\langle \phi_{n_1}, p_{n_2} \rangle = \delta[n_1 - n_2]$, and if ϕ_n is in $\mathcal{S}(\mathbb{R})$, then

$$\mathrm{P}f = \sum_{n=1}^{N_0} \langle f, \phi_n \rangle p_n \tag{A.23}$$

is a well-defined projector from $\mathcal{X}_1$ to $\mathcal{N}_{\mathrm{L}}$. The finite-dimensional null space of L is a Banach (and even a Hilbert) space for the norm

$$\|p\|_{\mathcal{N}_{\mathrm{L}}} = \left(\sum_{n=1}^{N_0} \langle p, \phi_n \rangle^2 \right)^{1/2}. \tag{A.24}$$

Moreover, $f \in \mathcal{X}_1$ is uniquely determined by $w = \mathrm{L}f \in \mathcal{M}(\mathbb{R})$ and $p = \mathrm{P}f \in \mathcal{N}_{\mathrm{L}}$. More precisely, there exists a right-inverse operator L_{ϕ}^{-1} of L such that [46, Theorem 4]

$$f = \mathrm{L}_{\phi}^{-1} w + p. \tag{A.25}$$

In other words, $\mathcal{X}_1$ is isomorphic to the direct sum $\mathcal{M}(\mathbb{R}) \oplus \mathcal{N}_{\mathrm{L}}$, from which we deduce that it is a Banach space for the norm [46, Theorem 5]

$$\|f\|_{\mathcal{X}_1} = \|\mathrm{L}f\|_{\mathcal{M}} + \|\mathrm{P}f\|_{\mathcal{N}_{\mathrm{L}}} = \|w\|_{\mathcal{M}} + \|p\|_{\mathcal{N}_{\mathrm{L}}}. \tag{A.26}$$

The same construction is at work for $\mathcal{X}_2$ by replacing $\mathcal{M}(\mathbb{R})$ with $L_2(\mathbb{R})$. Then, $\mathcal{X}_2$ is a Hilbert space for the Hilbertian norm

$$\|f\|_{\mathcal{X}_2} = \|\mathrm{L}f\|_{L_2} + \|\mathrm{P}f\|_{\mathcal{N}_{\mathrm{L}}}. \tag{A.27}$$

Predual of $\mathcal{X}_1$. The space $\mathcal{M}(\mathbb{R})$ is the topological dual of the space $C_0(\mathbb{R})$ of continuous and vanishing functions. The space $\mathcal{X}_1$ inherits this property: It is the topological dual of $C_{\mathrm{L}}(\mathbb{R})$, defined as the image of $C_0(\mathbb{R})$ by the adjoint L^* of L according to [46, Theorem 6].

We can therefore define a weak*-topology on $\mathcal{X}_1$: It is the topology for which $f_n \to 0$ if $\langle f_n, \varphi \rangle \to 0$ for every $\varphi \in C_{\mathrm{L}}(\mathbb{R})$. The weak*-topology is crucial to ensure the existence of solutions of (2.18); see [46] for more details.

Weak-continuity of h_m.* The weak*-continuity of h_m is equivalent to its inclusion in the predual space $C_{\mathrm{L}}(\mathbb{R})$ [216, Theorem IV.20, p. 114].

A.2 Chapter 3

A.2.1 Proof of Theorem 3.3.1

(i) Set $\mathbf{r}_k = (\mathbf{x}_{k+1} - \mathbf{x}_k)$. On one hand, it is clear that

$$\mathbf{r}_k = (1 - \alpha_k)\mathbf{x}_k + \alpha_k \mathbf{z}_k - \mathbf{x}_k = \alpha_k \left(\mathbf{z}_k - \mathbf{x}_k\right). \tag{A.28}$$

On the other hand, from the construction of $\{\alpha_k\}$,

$$\begin{aligned} & \alpha_k \left\|\mathbf{z}_k - \mathbf{x}_k\right\|_2 \le c_k \alpha_{k-1} \left\|\mathbf{z}_{k-1} - \mathbf{x}_{k-1}\right\|_2 \\ \Leftrightarrow \quad & \left\|\mathbf{r}_k\right\|_2 \le c_k \left\|\mathbf{r}_{k-1}\right\|_2. \end{aligned} \tag{A.29}$$

Iterating (A.29) gives

$$\|\mathbf{r}_k\|_2 \leq \|\mathbf{r}_0\|_2 \prod_{i=1}^{k} c_i, \quad \forall k \geq 1. \tag{A.30}$$

We now show that $\{\mathbf{x}_k\}$ is a Cauchy sequence. Since $\{c_k\}$ is asymptotically upper-bounded by $C < 1$, there exists K such that $c_k \leq C, \forall k > K$. Let m, n be two integers such that $m > n > K$. By using (A.30) and the triangle inequality,

$$\begin{aligned} \|\mathbf{x}_m - \mathbf{x}_n\|_2 &\leq \sum_{k=n}^{m-1} \|\mathbf{r}_k\|_2 \leq \|\mathbf{r}_0\|_2 \prod_{i=1}^{K} c_i \sum_{k=n-K}^{m-1-K} C^k \\ &\leq \left(\|\mathbf{r}_0\|_2 \prod_{i=1}^{K} c_i \right) \frac{C^{n-K} - C^{m-K}}{1 - C}. \end{aligned} \tag{A.31}$$

The last inequality proves that $\|\mathbf{x}_m - \mathbf{x}_n\|_2 \to 0$ as $m \to \infty, n \to \infty$, or $\{\mathbf{x}_k\}$ is a Cauchy sequence in the complete metric space $\mathbb{R}^N$. As a consequence, $\{\mathbf{x}_k\}$ must converge to some point $\mathbf{x}^* \in \mathbb{R}^N$.

(ii) Assume from now on that $\{\alpha_k\}$ is lower-bounded by $\varepsilon > 0$. By definition, $\{\alpha_k\}$ is also non-increasing and, thus, convergent to $\alpha^* > 0$. Next, we rewrite the update of $\mathbf{x}_k$ in Algorithm 1 as

$$\mathbf{x}_{k+1} = (1 - \alpha_k)\mathbf{x}_k + \alpha_k G_\gamma(\mathbf{x}_k), \tag{A.32}$$

where G_γ is defined by (3.7). Taking the limit of both sides of (A.32) leads to

$$\mathbf{x}^* = (1 - \alpha^*)\mathbf{x}^* + \alpha^* \lim_{k \to \infty} G_\gamma(\mathbf{x}_k). \tag{A.33}$$

Moreover, since the nonlinear operator F is continuous, G_γ is also continuous. Hence,

$$\lim_{k \to \infty} G_\gamma(\mathbf{x}_k) = G_\gamma\left(\lim_{k \to \infty} \mathbf{x}_k \right) = G_\gamma(\mathbf{x}^*). \tag{A.34}$$

By plugging (A.34) into (A.33), we get that $\mathbf{x}^* = G_\gamma(\mathbf{x}^*)$, which means that $\mathbf{x}^*$ is a fixed point of the operator G_γ.

(iii) Now that $F = P_{\mathcal{S}}$ satisfies (3.4), we invoke Proposition 3.2.1 to infer that $\mathbf{x}^*$ is a local minimizer of (3.1), thus completing the proof.

A.2.2 RPGD for Poisson Noise in CT

In the case where the CT measurements are corrupted by Poisson noise, the data-fidelity term in (3.1) should be replaced by weighted least squares [96, 217, 218]. For the sake of completeness, we show a sketch of the derivation. Let $\mathbf{x}$ represent the distribution of linear attenuation coefficient of an object and $[\mathbf{H}\mathbf{x}]_m$ represents their line integral. The mth CT measurement, y_m, is a Poisson random variable with parameters

$$p_m \sim \text{Poisson}\left(b_m \mathrm{e}^{-[\mathbf{H}\mathbf{x}]_m} + r_m\right) \tag{A.35}$$

$$y_m = -\log\left(\frac{p_m}{b_m}\right) \tag{A.36}$$

where b_m is the blank scan factor and r_m is the readout noise. Since logarithm is bijective, the negative log-likelihood of $\mathbf{y}$ given $\mathbf{x}$ is equal to the one of $\mathbf{p}$ given $\mathbf{x}$. After removing the constants, we use this negative log-likelihood as the data-fidelity term

$$E(\mathbf{H}\mathbf{x}, \mathbf{y}) = \sum_{m=1}^{M} \left(\hat{p}_m - p_m \log \hat{p}_m\right), \tag{A.37}$$

where $\hat{p}_m = b_m \mathrm{e}^{-[\mathbf{H}\mathbf{x}]_m} + r_m$ is the expected value of p_m. We then perform a quadratic approximation of E with respect to $\mathbf{H}\mathbf{x}$ around the point $(-\ln(\frac{\hat{p}_m - r_m}{b_m}))$ using a Taylor expansion. After ignoring the higher-order terms, this yields

$$E(\mathbf{H}\mathbf{x}, \mathbf{y}) = \sum_{m=1}^{M} \frac{w_m}{2} \left(\mathbf{H}\mathbf{x} - \log\left(\frac{b_m}{p_m - r_m}\right)\right)^2, \tag{A.38}$$

where $w_m = \frac{(p_m - r_m)^2}{p_m}$.

In the case when the readout noise r_m is insignificant, (A.38) can be written as

$$\begin{aligned} E(\mathbf{H}\mathbf{x}, \mathbf{y}) &= \sum_{m=1}^{M} \frac{w_m}{2} \left([\boldsymbol{H}\boldsymbol{x}]_m - y_m\right)^2 & \text{(A.39)} \\ &= \frac{1}{2} \|\mathbf{W}^{\frac{1}{2}} \mathbf{H}\mathbf{x} - \mathbf{W}^{\frac{1}{2}} \mathbf{y}\|^2 & \text{(A.40)} \\ &= \frac{1}{2} \|\mathbf{H}'\mathbf{x} - \mathbf{y}'\|^2, & \text{(A.41)} \end{aligned}$$

where $\mathbf{W} \in \mathbb{R}^{M\times M}$ is a diagonal matrix with $[\mathbf{diag}(\mathbf{W})]_m = w_m = p_m$, $\mathbf{H}' = \mathbf{W}^{\frac{1}{2}}\mathbf{H}$, and $\mathbf{y}' = \mathbf{W}^{\frac{1}{2}}\mathbf{y}$.
Imposing the data manifold prior, we get the equivalent of Problem (3.1) as

$$\min_{\mathbf{x}\in\mathcal{S}} \frac{1}{2}\|\mathbf{H}'\mathbf{x} - \mathbf{y}'\|^2. \tag{A.42}$$

Note that all the results discussed in Section 3.2 and 3.3 apply to Problem (A.42). As a consequence, we use Algorithm 1 to solve the problem with the following small change in the gradient step:

$$\mathbf{z}_k = F(\mathbf{x}_k - \gamma\mathbf{H}'^T\mathbf{H}'\mathbf{x}_k + \gamma\mathbf{H}'^T\mathbf{y}'). \tag{A.43}$$

A.2.3 Proof of Proposition 3.2.1

Suppose that (3.4) is fulfilled and let $\mathbf{x}^* \in \mathcal{S}$ be a fixed point of G_γ. We show that $\mathbf{x}^*$ is also a local minimizer of (3.1). Indeed, setting $\mathbf{x} = \mathbf{x}^* - \gamma\mathbf{H}^{\mathrm{T}}\mathbf{H}\mathbf{x}^* + \gamma\mathbf{H}^{\mathrm{T}}\mathbf{y}$ leads to $P_{\mathcal{S}}\mathbf{x} = \mathbf{x}^*$. Then, there exists $\varepsilon > 0$ such that, for all $\mathbf{z} \in \mathcal{S} \cap \mathcal{B}_\varepsilon(\mathbf{x}^*)$,

$$\begin{aligned} 0 &\geq \langle \mathbf{z} - P_{\mathcal{S}}\mathbf{x} \,,\, \mathbf{x} - P_{\mathcal{S}}\mathbf{x} \rangle \\ &= \gamma \left\langle \mathbf{z} - \mathbf{x}^* \,,\, \mathbf{H}^{\mathrm{T}}\mathbf{y} - \mathbf{H}^{\mathrm{T}}\mathbf{H}\mathbf{x}^* \right\rangle \\ &= \frac{\gamma}{2}\left(\|\mathbf{H}\mathbf{x}^* - \mathbf{y}\|_2^2 - \|\mathbf{H}\mathbf{z} - \mathbf{y}\|_2^2 + \|\mathbf{H}\mathbf{x}^* - \mathbf{H}\mathbf{z}\|_2^2 \right). \end{aligned}$$

Since $\gamma > 0$, the last inequality implies that

$$\|\mathbf{H}\mathbf{x}^* - \mathbf{y}\|_2^2 \leq \|\mathbf{H}\mathbf{z} - \mathbf{y}\|_2^2\,, \quad \forall \mathbf{z} \in \mathcal{S} \cap \mathcal{B}_\varepsilon(\mathbf{x}^*),$$

which means that $\mathbf{x}^*$ is a local minimizer of (3.1).

Assume now that $P_{\mathcal{S}}$ satisfies (3.5). By just removing the ε-ball in the above argument, one easily verifies that

$$\|\mathbf{H}\mathbf{x}^* - \mathbf{y}\|_2^2 \leq \|\mathbf{H}\mathbf{z} - \mathbf{y}\|_2^2\,, \quad \forall \mathbf{z} \in \mathcal{S},$$

which means that $\mathbf{x}^*$ is a solution of (3.1).

A.2.4 Proof of Proposition 3.2.2

We prove by contradiction. Assuming that $\mathcal{S}$ is non-convex, there must exist $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{S}$ and $\alpha \in (0,1)$ such that $\mathbf{x} = \alpha \mathbf{x}_1 + (1-\alpha)\mathbf{x}_2 \notin \mathcal{S}$. Since $P_{\mathcal{S}}\mathbf{x} \in \mathcal{S}$, it must be that

$$\begin{aligned} 0 < \|\mathbf{x} - P_{\mathcal{S}}\mathbf{x}\|_2^2 &= \langle \mathbf{x} - P_{\mathcal{S}}\mathbf{x} \,, \mathbf{x} - P_{\mathcal{S}}\mathbf{x} \rangle \\ &= \alpha \langle \mathbf{x}_1 - P_{\mathcal{S}}\mathbf{x} \,, \mathbf{x} - P_{\mathcal{S}}\mathbf{x} \rangle \\ &\quad + (1-\alpha) \langle \mathbf{x}_2 - P_{\mathcal{S}}\mathbf{x} \,, \mathbf{x} - P_{\mathcal{S}}\mathbf{x} \rangle . \end{aligned}$$

Thus, there exists $i \in \{0,1\}$ such that

$$\langle \mathbf{x}_i - P_{\mathcal{S}}\mathbf{x} \,, \mathbf{x} - P_{\mathcal{S}}\mathbf{x} \rangle > 0,$$

which violates (3.5). So, $\mathcal{S}$ is convex.

A.2.5 Proof of Proposition 3.2.3

Suppose that $\mathcal{S} = \bigcup_{i=1}^{n} \mathcal{C}_i$, where $\mathcal{C}_i$ is a closed convex set for all $i = 1, \ldots, n$. The statement is trivial when $n = 1$; assume now that $n \geq 2$. Let $\mathbf{x} \in \mathbb{R}^N$ and $\hat{\mathbf{x}}$ be the orthogonal projection of $\mathbf{x}$ onto $\mathcal{S}$. Consider two cases.

Case 1: $\hat{\mathbf{x}} \in \bigcap_{i=1}^{n} \mathcal{C}_i$.
It is then clear that

$$\|\hat{\mathbf{x}} - \mathbf{x}\|_2 \leq \|\mathbf{z} - \mathbf{x}\|_2 \,, \forall \mathbf{z} \in \mathcal{C}_i, \forall i.$$

This means that $\hat{\mathbf{x}}$ is the orthogonal projection of $\mathbf{x}$ onto each $\mathcal{C}_i$. Consequently,

$$\langle \mathbf{z} - \hat{\mathbf{x}} \,, \mathbf{x} - \hat{\mathbf{x}} \rangle \leq 0, \forall \mathbf{z} \in \mathcal{C}_i, \forall i \leq n,$$

which implies that (3.4) holds true for all $\varepsilon > 0$.

Case 2: $\hat{\mathbf{x}} \notin \bigcap_{i=1}^{n} \mathcal{C}_i$.
Without loss of generality, there exists $m < n$ such that

$$\hat{\mathbf{x}} \in \bigcap_{i=1}^{m} \mathcal{C}_i, \quad \hat{\mathbf{x}} \notin \bigcup_{i=m+1}^{n} \mathcal{C}_i. \tag{A.44}$$

Let d be the distance from $\hat{\mathbf{x}}$ to the set $\mathcal{T} = \bigcup_{i=m+1}^{n} \mathcal{C}_i$. Since each $\mathcal{C}_i$ is closed, $\mathcal{T}$ must be closed too and, so, $d > 0$. We now choose $0 < \varepsilon < d$. Then, $\mathcal{B}_\varepsilon(\hat{\mathbf{x}}) \cap \mathcal{T} = \emptyset$. Therefore,

$$\mathcal{S} \cap \mathcal{B}_\varepsilon(\hat{\mathbf{x}}) = \bigcup_{i=1}^{m} \left(\mathcal{C}_i \cap \mathcal{B}_\varepsilon(\hat{\mathbf{x}})\right) = \bigcup_{i=1}^{m} \tilde{\mathcal{C}}_i, \tag{A.45}$$

where $\tilde{\mathcal{C}}_i = \mathcal{C}_i \cap \mathcal{B}_\varepsilon(\hat{\mathbf{x}})$ is clearly a convex set, for all $i \leq m$. It is straightforward that $\hat{\mathbf{x}}$ is the orthogonal projection of $\mathbf{x}$ onto the set $\bigcup_{i=1}^{m} \tilde{\mathcal{C}}_i$ and that $\mathbf{x} \in \bigcap_{i=1}^{m} \tilde{\mathcal{C}}_i$. We are back to Case 1 and, therefore,

$$\langle \mathbf{z} - \hat{\mathbf{x}}\, ,\, \mathbf{x} - \hat{\mathbf{x}} \rangle \leq 0, \forall \mathbf{z} \in \tilde{\mathcal{C}}_i, \forall i \leq m. \tag{A.46}$$

From (A.45) and (A.46), (3.4) is fulfilled for the chosen ε.

A.2.6 Proof of Theorem 3.2.4

Let $\{\lambda_i\}$ denote the set of eigenvalues of $\mathbf{H}^{\mathrm{T}}\mathbf{H}$. We first have that, for all $\mathbf{x} \in \mathbb{R}^N$,

$$\left\| \mathbf{x} - \gamma \mathbf{H}^{\mathrm{T}}\mathbf{H}\mathbf{x} \right\|_2 \leq \left\| \mathbf{I} - \gamma \mathbf{H}^{\mathrm{T}}\mathbf{H} \right\|_2 \|\mathbf{x}\|_2 \,, \tag{A.47}$$

where the spectral norm of $\mathbf{I} - \gamma \mathbf{H}^{\mathrm{T}}\mathbf{H}$ is given by

$$\left\| \mathbf{I} - \gamma \mathbf{H}^{\mathrm{T}}\mathbf{H} \right\|_2 = \max_i \{|1 - \gamma\lambda_i|\}. \tag{A.48}$$

On the other hand, choosing $\gamma = 2/(\lambda_{\max} + \lambda_{\min})$ yields

$$\begin{aligned} & \frac{2\lambda_{\min}}{\lambda_{\max} + \lambda_{\min}} \leq \gamma\lambda_i \leq \frac{2\lambda_{\max}}{\lambda_{\max} + \lambda_{\min}}, \quad \forall i \\ \Leftrightarrow \quad & |1 - \gamma\lambda_i| \leq \frac{\lambda_{\max} - \lambda_{\min}}{\lambda_{\max} + \lambda_{\min}}, \quad \forall i. \end{aligned} \tag{A.49}$$

By combining (A.47), (A.48), and (A.49),

$$\left\| \mathbf{x} - \gamma \mathbf{H}^{\mathrm{T}}\mathbf{H}\mathbf{x} \right\|_2 \leq \frac{\lambda_{\max} - \lambda_{\min}}{\lambda_{\max} + \lambda_{\min}} \|\mathbf{x}\|_2 \,, \quad \forall \mathbf{x}. \tag{A.50}$$

Combining (A.50) with the Lipschitz continuity of $P_{\mathcal{S}}$ gives

$$\begin{aligned} \|G_\gamma(\mathbf{x}) - G_\gamma(\mathbf{z})\|_2 &\leq L \left\| (\mathbf{x}-\mathbf{z}) - \gamma \mathbf{H}^{\mathrm{T}}\mathbf{H}(\mathbf{x}-\mathbf{z}) \right\|_2 \\ &\leq L \, \frac{\lambda_{\max} - \lambda_{\min}}{\lambda_{\max} + \lambda_{\min}} \|\mathbf{x}-\mathbf{z}\|_2 , \quad \forall \mathbf{x}, \forall \mathbf{z}. \end{aligned} \tag{A.51}$$

Since $L < (\lambda_{\max} + \lambda_{\min})/(\lambda_{\max} - \lambda_{\min})$, (A.51) implies that G_γ is a contractive mapping. By the Banach-Picard fixed point theorem [219, Thm. 1.48], $\{\mathbf{x}_k\}$ defined by $\mathbf{x}_{k+1} = G_\gamma(\mathbf{x}_k)$ converges to the unique fixed point $\mathbf{x}^*$ of G_γ, for every initialization $\mathbf{x}_0$. Finally, since $P_{\mathcal{S}}$ satisfies (3.4), by Proposition 3.2.1, $\mathbf{x}^*$ is also a local minimizer of (3.1).

A.2.7 Proof of Theorem 3.2.5

Again, let $\{\lambda_i\}$ be the set of eigenvalues of $\mathbf{H}^{\mathrm{T}}\mathbf{H}$. With $\gamma < 2/\lambda_{\max}$, one readily verifies that, $\forall \mathbf{x} \in \mathbb{R}^N$,

$$\left\| \mathbf{x} - \gamma \mathbf{H}^{\mathrm{T}}\mathbf{H}\mathbf{x} \right\|_2 \leq \max_i \{|1-\gamma\lambda_i|\} \cdot \|\mathbf{x}\|_2 \leq \|\mathbf{x}\|_2 .$$

Combining this with the non-expansiveness of $P_{\mathcal{S}}$ leads to

$$\begin{aligned} \|G_\gamma(\mathbf{x}) - G_\gamma(\mathbf{z})\|_2 &\leq \left\| (\mathbf{x}-\mathbf{z}) - \gamma \mathbf{H}^{\mathrm{T}}\mathbf{H}(\mathbf{x}-\mathbf{z}) \right\|_2 \\ &\leq \|\mathbf{x}-\mathbf{z}\|_2 , \quad \forall \mathbf{x}, \mathbf{z} \in \mathbb{R}^N . \end{aligned}$$

Now that G_γ is a non-expansive operator with a nonempty fixed-point set, we invoke the Krasnosel'skiĭ-Mann theorem [219, Thm. 5.14] to deduce that the iteration (3.6) must converge to a fixed point of G_γ which is, by Proposition 3.2.1, also a local minimizer of (3.1).

A.2.8 Experiments

1) Experiment 1

Figure A.1 show the difference images between the reconstructions and the ground truth for lung and abdomen images for ×16 case.

Fig. A.3 compares the reconstructions for the ×5 case when the noise levels are ∞-dB (first and second columns), 45 dB (third columns) and 35 dB (fourth column). It is visible that FBPconv40 results in a noisy image and TV is again blurred. DL and RPGD40 reconstruction for 45-dB case have similar reconstruction performance and outperform others. For the other cases, RPGD40 is the best performer. Fig. A.2 shows the profile for the reconstructions for ×5, 45-dB noise case.

2) Experiment 2

Fig. A.4 shows the reconstructions from measurements corrupted with Poisson noise discussed in Section 3.6.2.

Results(∞-dB) diff(∞-dB) diff(45-dB) diff(35-dB) Results(∞-dB) diff(∞-dB) diff(45-dB) diff(40-dB)

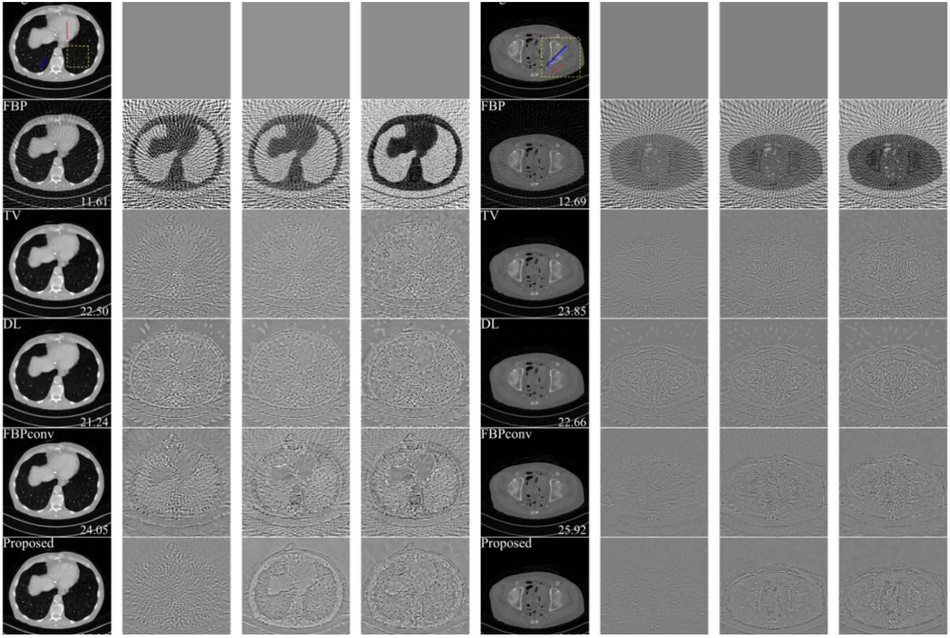

Figure A.1: Comparison of reconstructions using different methods for the ×16 case in Experiment 1 for lung and abdomen images. The first and fifth column shows the reconstruction from noiseless measurements for lung and abdomen images, respectively. Second-fourth columns: show the difference between the reconstructions and the lung image for noiseless, 45-dB noise, and 35-dB noise cases, respectively. Sixth-eighth columns: show the difference between the reconstructions and the abdomen image for noiseless, 45-dB noise, and 40-dB noise cases, respectively.

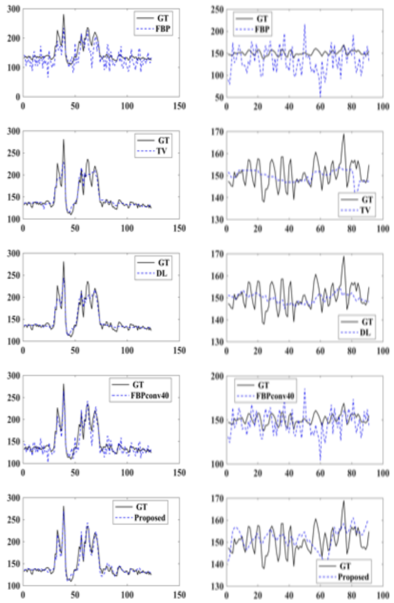

High-contrast profile Low-contrast profile

Figure A.2: Profile of the high and low contrast regions marked by solid and dashed line segments, respectively, inside the original image in the first column of Fig. A.3. This case corresponds to ×5, 35-dB noise case.

∞-dB zoom(∞-dB) 45-dB 35-dB diff(∞-dB) 45-dB
35-dB

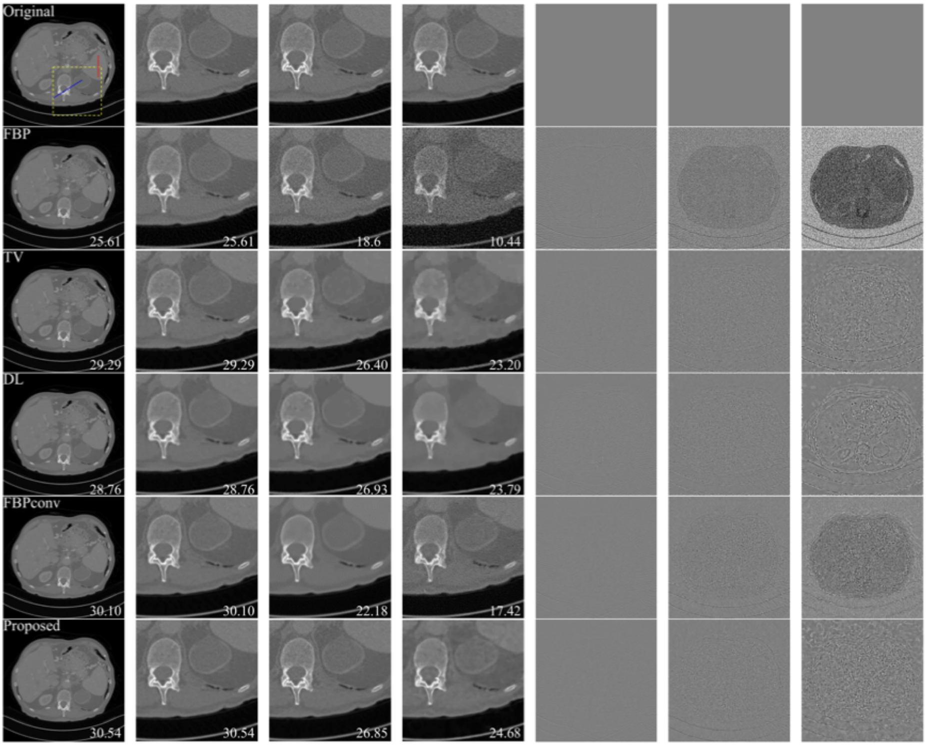

Figure A.3: Comparison of reconstructions using different methods for the ×5 case. The first column shows the reconstruction from noiseless measurements. Second column shows the zoomed version of a the box area given in the original in the first column. Third and fourth column show the zoomed version for the case of 45 and 35 dB, respectively. The next three columns show the difference between the reconstruction and the ground truth for noiseless, 45-dB noise, and 35-dB noise case, respectively.

Results (25-dB) zoom (25-dB) zoom (30-dB) zoom (35-dB) diff (25-dB) diff (30-dB) diff(35-dB)

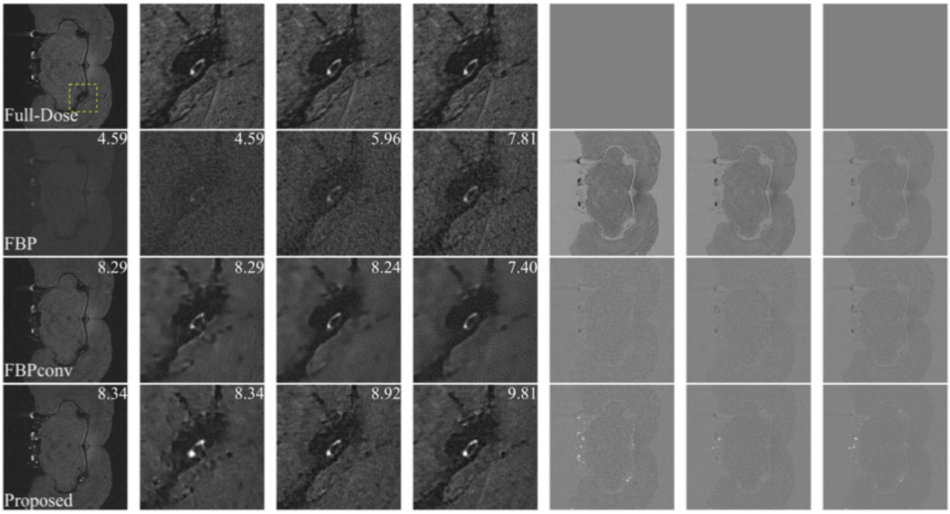

Figure A.4: Reconstruction results for Experiment 2. The first column shows the result when the measurement noise is 25 dB. Second column shows the zoomed version of the box area given in the full-dose image in the first column. Third and fourth columns show the zoomed version for the case of 30 and 35 dB respectively. The next three columns show the difference between the reconstruction and the ground truth for 25dB, 30dB and 35 dB respectively.

A.3 Chapter 5

A.3.1 Synthetic Data (Figure 5.3)

Dataset. We construct a synthetic Cryo-EM dataset that mimics the experimental β-galactosidase dataset (EMPIAR-10061) from [158]. We generate 41,000 synthetic β-galactosidase projections of size 192 × 192 using our Cryo-EM image-formation model (see *Online Methods*). For the ground-truth volume, we generate a 2.5 Å density map from PDB-*5a1a* atomic model using Chimera [220]. This gives a volume of size $(302 \times 233 \times 163)$ with voxel size 0.637 Å. The volume is then padded, averaged, and downsampled to a size $(180 \times 180 \times 180)$ with voxel size 1.274 Å. This corresponds to a Nyquist resolution of 2.548 Å for the reconstructed volume.

The projection poses are sampled from a uniform distribution over $SO(3)$, where $SO(3)$ is the group of 3D rotations around the origin of $\mathbb{R}^3$.

In order to apply random CTFs and noise, we randomly pick a micrograph in the EMPIAR- 10061 dataset. We extract its CTF paramters using CTFFIND4 [221] and apply them to a clean projection. The parameter B of the envelope function of the CTF (see (5.13)) is chosen such that it decays to 0.2 at the Nyquist frequency. We then randomly select a background patch from the same micrograph to simulate noise. The noise is downsampled to size 192 × 192, and normalized to zero-mean, scaled and added to the projection. The scaling is such that the ratio of the energy of the signal to the energy of the noise (SNR) is kept at 0.01, which is equivalent to -20 dB. The dataset is then randomly divided into two halves. The algorithm is applied separately on both halves to generate the half-maps. The FSC between the two half maps is then reported using FOCUS [222, 223].

Generator Settings. We reconstruct a volume of size $(184 \times 184 \times 184)$ voxels for each half dataset. The pixel size is 1.274 Å. The volumes are initialized with zeros. The D2 symmetry of β-galactosidase is enforced during reconstruction.

We use our image-formation model to simulate realistic projections from the current volume estimate at every CryoGAN iteration. The distribution of the imaging parameters is identical to the one used to generate the dataset. To add the noise on the CTF-modulated projections, we also keep the same approach as the one used to generate the dataset. However, we assume that the final SNR of each projection is unknown, leading us to learn the scaling parameter that controls the ratio between the projections and the noise patches.

We apply a binary spherical mask of size $(184 \times 184 \times 184)$ on the learned volume. To handle the sharp transitions at the mask borders, we restrict the voxel values to lie above a certain value. This value changes as a function of position and iteration number: it increases linearly with the distance from the center of the volume to the border of the mask, from V_{min} to 0, and the value V_{min} changes as a function of iteration number, starting at 0 and decreasing to -2% of the current maximum voxel value. This promotes nonnegativity during the initial phases of reconstruction and increases the stability of the algorithm.

Discriminator Architecture. The architecture of the discriminator network is detailed in *Online Methods*. The discriminator is initialized identically for the two half datasets. All projections (*i.e.*, the real projections and the ones generated by the simulator) are normalized to zero-mean and a unit standard-deviation before being fed to the discriminator.

General Settings. The adversarial learning scheme is implemented in Pytorch [169]. For the optimization, we use [141] ($\beta_1 = 0.5$, $\beta_2 = 0.9$, $\epsilon = 10^{-8}$) with a learning rate of 10^{-3} and a batch size of 8. The algorithm is run for 40 epochs and the learning rate decreases by 1% at every epoch. The parameter for the gradient-penalty term (see (5.20)) is kept to $\lambda = 0.001$. The discriminator is trained 4 times for every training of the generator ($n_{\mathrm{discr}} = 4$ in Algorithm 3).

For the back-propagations, the norm of the gradients for the discriminator are clipped to a maximal value of 10^8. For the generator, the gradients for each pixel are clipped to a maximal value of 10^4. The clipping values increase linearly from zero to those maxima in the first two epochs. Doing so improves on the stability of the adversarial learning scheme in the start, in particular, on that of the discriminator. All parameters are tuned for a fixed value range that follows from the normalization of all projections.

Computational Resources. The reconstruction is run on a Nvidia V100 GPU with 18GB memory. Each epoch lasts 10 minutes. The algorithm is run for 40 epochs which, in the current implementation, takes 400 minutes.

A.3.2 Additional Synthetic Data (Figure 5.4)

Synthetic dataset. The data is generated similarly to the main synthetic experiment, with the exception of some changes to the imaging conditions. In a first case, the noise level is set to a SNR of -5.2 dB. This corresponds to the energy of the noise being almost four times that of the signal. In the second case, the SNR is kept at -20 dB (the same as for the main experiment), but the projections are also translated. The translation (both horizontal and vertical) for each projection is sampled from a zero-mean symmetric triangular distribution whose total width is 6% of the image size from the centre. This corresponds to a shift of maximum 5 pixels from the centre in each direction.

Reconstruction Settings. The reconstruction settings for both cases are the same than the ones used in the main experiment, except for the few following differences. For the second case, the translations are also imposed, and the translation distribution is kept the same as the one used for generating the dataset. Furthermore, in both cases, the lower bound of the clipping value at the centre reaches -5% of the maximum voxel value of the volume. Finally, the algorithm is run for 100 epochs for both experiments.

A.3.3 Experimental Data (Figure 5.5)

Dataset. The dataset consists of 41,123 β-galactosidase (EMPIAR-10061) projections extracted from 1,539 micrographs [158]. The projections of size (384×384) are downsampled to (192×192), with a pixel size of 1.274 Å. This corresponds to a Nyquist resolution of 2.548 Å for a reconstructed volume of size $(180 \times 180 \times 180)$.

The dataset is randomly divided into two halves. The algorithm is applied separately on both halves to generate the half-maps. The defocus and astigmatism parameters of the CTF are estimated from each micrograph using Relion [151].

Generator Settings. We reconstruct a volume of size $(180 \times 180 \times 180)$ voxels for each half dataset. The pixel size is 1.274 Å. The volumes are initialized with zeros. The D2 symmetry of β-galactosidase is enforced during reconstruction. A uniform distribution is assumed for the poses. The CTF parameters estimated with CTFFIND4 [221] are used in the forward model of the Cryo-EM physics simulator. The parameter B of the envelope function of the CTF (see (5.13)) decays to 0.4 at the Nyquist frequency. The translations are set to zero.

To handle the noise, we randomly extract (prior to the learning procedure) 41,123 patches of size (384×384) from the background of the micrographs at locations where particles do not appear; this is done by identifying patches with the lowest variance. We extract as many noise patches per micrograph as we have particle images. Each noise patch is then downsampled to size (192×192) and normalized. Then, during run-time, the noise patches are sampled from this collection, scaled, and added to the simulated projections. For consistency, the noise patch added to a given simulated projection is taken from the same micrograph as the one that was used to estimate the CTF parameters previously applied to that specific projection. The scaling operation weights the contribution of the noise with respect to the projection signal. This is handled by multiplying the pixel values of the noise patches and the projections by two scalars that are learnt throughout the procedure. These two scalar values are the same for every pair of noise/projection images, so that the same amount of extracted noise is added to every simulated projection.

We apply a binary spherical mask of size ($171 \times 171 \times 171$) on the learned volume. To handle the sharp transitions at the mask borders, we enforce constraints on the masked volume. These are similar to those used in the synthetic experiment, with difference that the lower bound, $V_{\min}$ decreases to -5% of the maximum voxel value.

Discriminator Architecture. The architecture of the discriminator network is detailed in *Online Methods*. The discriminator is initialized identically for the two half datasets. The projection images (real and simulated) are smoothed with a Gaussian kernel before being fed to the discriminator. The standard deviation of the kernel is initially set at 2 and changes in every iteration so that it decreases by a total of 2% in each epoch.

General Settings. The adversarial learning scheme is implemented in Pytorch [169]. For the optimization, we use [141] ($\beta_1 = 0.5$, $\beta_2 = 0.9$, $\epsilon = 10^{-8}$) with a learning rate of 10^{-3} and a batch size of 8. The algorithm is run for 15 epochs and the learning rate decreases by 8% at every epoch. The parameter for the gradient-penalty term is kept to $\lambda = 1$ (see (5.20)). The discriminator is trained 4 times for every training of the generator ($n_{\text{discr}} = 4$ in Algorithm 3).

For the back-propagations, the norm of the gradients for the discriminator are clipped to a maximal value of 10^7. For the generator, the gradients for each pixel are clipped to a maximal value of 10^3. The clipping values increase linearly from zero to those maxima in the first two epochs. Doing so improves on the stability of the adversarial learning scheme in the start, in particular, on that of the discriminator. The gradients that correspond to the learning of the scaling ratios between the noise and projection images are clipped to a value of 10.

Computational Resources. The reconstruction is run on a Nvidia V100 GPU with 18GB memory. Each epoch lasts 10 minutes. The algorithm is run for 15 epochs, which takes 150 minutes.

A.4 Chapter 6

A.4.1 Neural Network Architectures

LAYER ID	LAYER	RESAMPLE	OUTPUT SHAPE
0	Input	-	1 × 32 × 32
1	Conv2d	MaxPool	96 × 16 × 16
2	Conv2d	MaxPool	192 × 8 × 8
3	Conv2d	MaxPool	384 × 4 × 4
4	Conv2d	MaxPool	768 × 2 × 2
5	Flatten	-	3072 × 1 × 1
6	FC	-	50 × 1 × 1
7	FC	-	1 × 1 × 1

Table A.1: 2D Discriminator architecture. LeakReLU(0.1) is used after every MaxPool and in layer 6.

LAYER ID	LAYER	RESAMPLE	NORM	OUTPUT SHAPE (C, D, H, W)
0	Input	-	-	$1 \times 32 \times 32 \times 32$
1	Conv3d	-	BN	$16 \times 32 \times 32 \times 32$
2	Conv3d	MaxPool	BN	$16 \times 16 \times 16 \times 16$
3	Conv3d	-	BN	$32 \times 16 \times 16 \times 16$
4	Conv3d	MaxPool	BN	$32 \times 8 \times 8 \times 8$
5	Conv3d	-	BN	$64 \times 8 \times 8 \times 8$
6	Conv3d	MaxPool	BN	$64 \times 4 \times 4 \times 4$
7	Conv3d	-	BN	$128 \times 4 \times 4 \times 4$
8	Conv3d	MaxPool	BN	$128 \times 2 \times 2 \times 2$
9	Conv3d	-	BN	$256 \times 2 \times 2 \times 2$
10	Conv3d	-	BN	$256 \times 2 \times 2 \times 2$
11	Conv3d	Upsample	BN	$128 \times 4 \times 4 \times 4$
12	Concat(layer 8)	-	-	$256 \times 4 \times 4 \times 4$
13	Conv3d	-	BN	$128 \times 4 \times 4 \times 4$
14	Conv3d	-	BN	$128 \times 4 \times 4 \times 4$
15	Conv3d	Upsample	BN	$64 \times 8 \times 8 \times 8$
16	Concat(layer 6)	-	-	$128 \times 8 \times 8 \times 8$
17	Conv3d	-	BN	$64 \times 8 \times 8 \times 8$
18	Conv3d	-	BN	$64 \times 8 \times 8 \times 8$
19	Conv3d	Upsample	BN	$32 \times 16 \times 16 \times 16$
20	Concat(layer 4)	-	-	$64 \times 16 \times 16 \times 16$
21	Conv3d	-	BN	$32 \times 16 \times 16 \times 16$
22	Conv3d	-	BN	$32 \times 16 \times 16 \times 16$
23	Conv3d	Upsample	BN	$16 \times 32 \times 32 \times 32$
24	Concat(layer 2)	-	-	$32 \times 32 \times 32 \times 32$
25	Conv3d	-	BN	$16 \times 32 \times 32 \times 32$
26	Conv3d	-	BN	$16 \times 32 \times 32 \times 32$
27	Conv3d	-	BN	$1 \times 32 \times 32 \times 32$

Table A.2: 3D Generator architecture. ReLU is used after every BatchNorm (BN). Concatenation is with the values before pooling.

Bibliography

[1] Mario Bertero and Patrizia Boccacci, *Introduction to inverse problems in imaging*, CRC press, 1998.

[2] A. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*, Classics in Applied Mathematics. Society for Industrial and Applied Mathematics, New York, NY, 2001.

[3] Otmar Scherzer, *Handbook of mathematical methods in imaging*, Springer Science & Business Media, 2010.

[4] M Bertero and M Piana, "Inverse problems in biomedical imaging: modeling and methods of solution," in *Complex systems in biomedicine*, pp. 1–33. Springer, 2006.

[5] C. Bouman and K. Sauer, "A generalized Gaussian image model for edge-preserving MAP estimation," *IEEE Trans. Image Process.*, vol. 2, no. 3, pp. 296–310, 1993.

[6] P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud, "Deterministic edge-preserving regularization in computed imaging," *IEEE Trans. Image Process.*, vol. 6, no. 2, pp. 298–311, 1997.

[7] M. Lustig, D. L. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, Dec. 2007.

[8] E. Candès and J. Romberg, "Sparsity and incoherence in compressive sampling," *Inverse Problems*, vol. 23, no. 3, pp. 969–985, Jun. 2007.

[9] S. Ramani and J.A. Fessler, “Parallel MR image reconstruction using augmented Lagrangian methods,” *IEEE Trans. Med. Imag.*, vol. 30, no. 3, pp. 694–706, 2011.

[10] M. A. T. Figueiredo and R. D. Nowak, “An EM algorithm for wavelet-based image restoration,” *IEEE Transactions on Image Processing*, vol. 12, no. 8, pp. 906–916, Aug. 2003.

[11] Ingrid Daubechies, Michel Defrise, and Christine De Mol, “An iterative thresholding algorithm for linear inverse problems with a sparsity constraint,” *Commun. Pure Appl. Math*, vol. 57, no. 11, pp. 1413–1457, 2004.

[12] Amir Beck and Marc Teboulle, “A fast iterative shrinkage-thresholding algorithm for linear inverse problems,” *SIAM J. Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.

[13] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein, “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[14] Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep learning*, MIT Press, 2016.

[15] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, 2012, pp. 1097–1105.

[16] Michael T. McCann, Kyong Hwan Jin, and Michael Unser, “Convolutional neural networks for inverse problems in imaging: A review,” *IEEE Signal Process. Mag.*, vol. 34, no. 6, pp. 85–95, Nov. 2017.

[17] Kyong Hwan Jin, Michael T McCann, Emmanuel Froustey, and Michael Unser, “Deep convolutional neural network for inverse problems in imaging,” *IEEE Trans. Image Process.*, vol. 26, no. 9, pp. 4509–4522, 2017.

[18] Yo Seob Han, Jaejun Yoo, and Jong Chul Ye, “Deep learning with domain adaptation for accelerated projection reconstruction MR,” arXiv:1703.01135 [cs.CV], 2017.

[19] Stephan Antholzer, Markus Haltmeier, and Johannes Schwab, "Deep learning for photoacoustic tomography from sparse data," arXiv:1704.04587 [cs.CV], 2017.

[20] Shanshan Wang, Zhenghang Su, Leslie Ying, Xi Peng, Shun Zhu, Feng Liang, Dagan Feng, and Dong Liang, "Accelerating magnetic resonance imaging via deep learning," in *Proc. IEEE Int. Symp. Biomed. Imaging (ISBI)*, 2016, pp. 514–517.

[21] Jonas Adler and Ozan Öktem, "Solving ill-posed inverse problems using iterative deep neural networks," arXiv:1704.04058 [math.OC], 2017.

[22] Kerem C Tezcan, Christian F Baumgartner, Roger Luechinger, Klaas P Pruessmann, and Ender Konukoglu, "MR image reconstruction using deep density priors," *IEEE Trans. on Med. Imag.*, vol. 38, no. 7, pp. 1633–1642, July 2019.

[23] Yong Chun and Jeffrey A Fessler, "Deep bcd-net using identical encoding-decoding cnn structures for iterative image recovery," in *2018 IEEE 13th Image, Video, and Multidimensional Signal Processing Workshop (IVMSP)*. IEEE, 2018, pp. 1–5.

[24] K. Hammernik, T. Klatzer, E. Kobler, M.P. Recht, D.K. Sodickson, T. Pock, et al., "Learning a variational network for reconstruction of accelerated MRI data," *Magn. Reson. in Med.*, vol. 79, no. 6, pp. 3055–3071, June 2018.

[25] JH Rick Chang, Chun-Liang Li, Barnabas Poczos, BVK Vijaya Kumar, and Aswin C Sankaranarayanan, "One network to solve them all–solving linear inverse problems using deep projection models," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 5888–5897.

[26] Harshit Gupta, Kyong Hwan Jin, Ha Q Nguyen, Michael T McCann, and Michael Unser, "CNN-based projected gradient descent for consistent ct image reconstruction," *IEEE transactions on medical imaging*, vol. 37, no. 6, pp. 1440–1453, 2018.

[27] V. Lempitsky, A. Vedaldi, and D. Ulyanov, "Deep image prior," in *Proc. of the IEEE Comput. Soc. Conf. on Comput. Vision and Pattern Recognition (CVPR)*, Salt Lake City UT, USA, July 18-23 2018, pp. 9446–9454.

[28] Harshit Gupta, Michael T McCann, Laurene Donati, and Michael Unser, “Cryogan: A new reconstruction paradigm for single-particle cryo-em via deep adversarial learning,” *BioRxiv*, 2020.

[29] Harshit Gupta, Julien Fageot, and Michael Unser, “Continuous-domain solutions of linear inverse problems with tikhonov versus generalized tv regularization,” *IEEE Transactions on Signal Processing*, vol. 66, no. 17, pp. 4670–4684, 2018.

[30] X. Pan, E. Y. Sidky, and M. Vannier, “Why do commercial CT scanners still employ traditional, filtered back-projection for image reconstruction?,” *Inverse Probl.*, vol. 25, no. 12, pp. 123009, 2009.

[31] I. Daubechies, M. Defrise, and C. De Mol, “An iterative thresholding algorithm for linear inverse problems with a sparsity constraint,” *Communications on Pure and Applied Mathematics*, vol. 57, no. 11, pp. 1413–1457, 2004.

[32] A. Beck and M. Teboulle, “A fast iterative shrinkage-thresholding algorithm for linear inverse problems,” *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, Jan. 2009.

[33] Michael Elad and Michal Aharon, “Image denoising via sparse and redundant representations over learned dictionaries,” *IEEE Trans. Image Process.*, vol. 15, no. 12, pp. 3736–3745, 2006.

[34] Emmanuel J Candès, Yonina C Eldar, Deanna Needell, and Paige Randall, “Compressed sensing with coherent and redundant dictionaries,” *Applied and Computational Harmonic Analysis*, vol. 31, no. 1, pp. 59–73, 2011.

[35] S. Ravishankar, R. R. Nadakuditi, and J. Fessler, “Efficient sum of outer products dictionary learning (SOUP-DIL) and its application to inverse problems,” *IEEE Trans. Comput. Imaging*, vol. 3, no. 4, pp. 694–709, 2017.

[36] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, “Generative adversarial nets,” in *Advances in neural information processing systems*, 2014, pp. 2672–2680.

[37] Diederik P Kingma and Max Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.

[38] A. N. Tikhonov, “Solution of incorrectly formulated problems and the regularization method,” *Soviet Mathematics*, vol. 4, pp. 1035–1038, 1963.

[39] Mario Bertero and Patrizia Boccacci, *Introduction to Inverse Problems in Imaging*, CRC press, 1998.

[40] D. L. Donoho, “Compressed sensing,” *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[41] M.A.T. Figueiredo, R.D. Nowak, and S.J. Wright, “Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 1, no. 4, pp. 586–597, Dec. 2007.

[42] Arthur E. Hoerl and Robert W. Kennard, “Ridge regression: Biased estimation for nonorthogonal problems,” *Technometrics*, vol. 12, no. 1, pp. 55–67, Feb. 1970.

[43] R. Tibshirani, “Regression shrinkage and selection via the Lasso,” *Journal of the Royal Statistical Society. Series B*, vol. 58, no. 1, pp. 265–288, 1996.

[44] Bradley Efron, Trevor Hastie, and Robert Tibshirani, “Discussion: The Dantzig selector: Statistical estimation when p is much larger than n,” *The Annals of Statistics*, vol. 35, no. 6, pp. 2358–2364, Dec. 2007.

[45] M. Unser, J. Fageot, and H. Gupta, “Representer theorems for sparsity-promoting ℓ_1-regularization,” *IEEE Transactions on Information Theory*, vol. 62, no. 9, pp. 5167–5180, Sep. 2016.

[46] Michael Unser, Julien Fageot, and John Paul Ward, “Splines are universal solutions of linear inverse problems with generalized TV regularization,” *SIAM Review*, vol. 59, no. 4, pp. 769–793, Dec. 2017.

[47] Holger Wendland, *Scattered Data Approximation*, vol. 17, Cambridge University press, 2004.

[48] Bernhard Schölkopf and Alexander J. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*, MIT Press, Cambridge, MA, USA, 2001.

[49] Bernhard Schölkopf, Ralf Herbrich, and Alex J Smola, "A generalized representer theorem," *Lecture Notes in Computer Science*, vol. 2111, pp. 416–426, 2001.

[50] Grace Wahba, *Spline Models for Observational Data*, vol. 59, SIAM, 1990.

[51] Grace Wahba, "Support vector machines, reproducing kernel Hilbert spaces and the randomized GACV," *Advances in Kernel Methods-Support Vector Learning*, vol. 6, pp. 69–87, 1999.

[52] Anatoly Yu Bezhaev and Vladimir Aleksandrovich Vasilenko, *Variational theory of splines*, Springer, 2001.

[53] J. Kybic, T. Blu, and M. Unser, "Generalized sampling: A variational approach—Part I: Theory," *IEEE Transactions on Signal Processing*, vol. 50, no. 8, pp. 1965–1976, Aug. 2002.

[54] J. Kybic, T. Blu, and M. Unser, "Generalized sampling: A variational approach—Part II: Applications," *IEEE Transactions on Signal Processing*, vol. 50, no. 8, pp. 1977–1985, Aug. 2002.

[55] Leonid I. Rudin, Stanley Osher, and Emad Fatemi, "Nonlinear total variation based noise removal algorithms," *Physics D*, vol. 60, no. 1-4, pp. 259–268, Nov. 1992.

[56] Gabriele Steidl, Stephan Didas, and Julia Neumann, "Splines in higher order TV regularization," *International Journal of Computer Vision*, vol. 70, no. 3, pp. 241–255, Dec. 2006.

[57] SD Fisher and JW Jerome, "Spline solutions to L_1 extremal problems in one and several variables," *Journal of Approximation Theory*, vol. 13, no. 1, pp. 73–83, Jan. 1975.

[58] K. Bredies and H.K. Pikkarainen, "Inverse problems in spaces of measures," *ESAIM: Control, Optimisation and Calculus of Variations*, vol. 19, no. 1, pp. 190–218, Jan. 2013.

[59] E.J. Candès and C. Fernandez-Granda, “Super-resolution from noisy data,” *Journal of Fourier Analysis and Applications*, vol. 19, no. 6, pp. 1229–1254, Dec. 2013.

[60] Quentin Denoyelle, Vincent Duval, and Gabriel Peyré, “Support recovery for sparse super-resolution of positive measures,” *Journal of Fourier Analysis and Applications*, vol. 23, no. 5, pp. 1153–1194, Oct. 2017.

[61] Antonin Chambolle, Vincent Duval, Gabriel Peyré, and Clarice Poon, “Geometric properties of solutions to the total variation denoising problem,” *Inverse Problems*, vol. 33, no. 1, pp. 015002, Dec. 2016.

[62] A. Flinth and P. Weiss, “Exact solutions of infinite dimensional total-variation regularized problems,” arXiv:1708.02157 [math.OC], 2017.

[63] Michael Unser, “A representer theorem for deep neural networks,” arXiv:1802.09210 [stat.ML], 2018.

[64] Andrea Braides, *Gamma-convergence for Beginners*, vol. 22, Clarendon Press, 2002.

[65] Vincent Duval and Gabriel Peyré, “Sparse regularization on thin grids I: the Lasso,” *Inverse Problems*, vol. 33, no. 5, pp. 055008, 2017.

[66] Gongguo Tang, Badri Narayan Bhaskar, and Benjamin Recht, “Sparse recovery over continuous dictionaries-just discretize,” in *Asilomar Conference on Signals, Systems and Computers*. IEEE, 2013, pp. 1043–1047.

[67] George B Dantzig, Alex Orden, and Philip Wolfe, “The generalized simplex method for minimizing a linear form under linear inequality restraints,” *Pacific Journal of Mathematics*, vol. 5, no. 2, pp. 183–195, Oct. 1955.

[68] David G Luenberger, *Introduction to Linear and Nonlinear Programming*, vol. 28, Addison-Wesley Reading, MA, 1973.

[69] Vincent Duval and Gabriel Peyré, “Exact support recovery for sparse spikes deconvolution,” *Foundations of Computational Mathematics*, vol. 15, no. 5, pp. 1315–1355, 2015.

[70] Ryan J Tibshirani, "The LASSO problem and uniqueness," *Electronic Journal of Statistics*, vol. 7, pp. 1456–1490, 2013.

[71] Holger Rauhut, Karin Schnass, and Pierre Vandergheynst, "Compressed sensing and redundant dictionaries," *IEEE Transactions on Information Theory*, vol. 54, no. 5, pp. 2210–2219, Apr. 2008.

[72] Simon Foucart and Holger Rauhut, *A Mathematical Introduction to Compressive Sensing*, Springer, 2013.

[73] Antonin Chambolle and Charles Dossal, "On the convergence of the iterates of FISTA," *Journal of Optimization Theory and Applications*, vol. 166, no. 3, pp. 25, 2015.

[74] E. Mammen and S. van de Geer, "Locally adaptive regression splines," *Annals of Statistics*, vol. 25, no. 1, pp. 387–413, 1997.

[75] Emmanuel J Candès and Carlos Fernandez-Granda, "Towards a mathematical theory of super-resolution," *Communications on Pure and Applied Mathematics*, vol. 67, no. 6, pp. 906–956, 2014.

[76] Gongguo Tang, Badri Narayan Bhaskar, Parikshit Shah, and Benjamin Recht, "Compressed sensing off the grid," *IEEE Transactions on Information Theory*, vol. 59, no. 11, pp. 7465–7490, 2013.

[77] Yohann De Castro, Fabrice Gamboa, Didier Henrion, and J-B Lasserre, "Exact solutions to super resolution on semi-algebraic domains in higher dimensions," *IEEE Transactions on Information Theory*, vol. 63, no. 1, pp. 621–630, 2017.

[78] M. Unser and T. Blu, "Generalized smoothing splines and the optimal discretization of the Wiener filter," *IEEE Transactions on Signal Processing*, vol. 53, no. 6, pp. 2146–2159, Jun. 2005.

[79] M. Unser and P. D. Tafti, "Stochastic models for sparse and piecewise-smooth signals," *IEEE Transactions on Signal Processing*, vol. 59, no. 3, pp. 989–1006, Mar. 2011.

[80] M. Unser and P. D. Tafti, *An Introduction to Sparse Stochastic Processes*, Cambridge University Press, 2014.

[81] J. Fageot, V. Uhlmann, and M. Unser, "Gaussian and sparse processes are limits of generalized Poisson processes," arXiv:1702.05003 [math.PR], 2017.

[82] Ali Mousavi and Richard G Baraniuk, "Learning to invert: Signal recovery via deep convolutional networks," arXiv:1701.03891 [stat.ML], 2017.

[83] Karol Gregor and Yann LeCun, "Learning fast approximations of sparse coding," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2010, pp. 399–406.

[84] Yan Yang, Jian Sun, Huibin Li, and Zongben Xu, "Deep ADMM-Net for compressive sensing MRI," in *Adv. Neural Inf. Process. Syst. (NIPS)*, pp. 10–18. 2016.

[85] Patrick Putzky and Max Welling, "Recurrent inference machines for solving inverse problems," arXiv:1706.04008 [cs.NE], 2017.

[86] Jo Schlemper, Jose Caballero, Joseph V Hajnal, Anthony Price, and Daniel Rueckert, "A deep cascade of convolutional neural networks for MR image reconstruction," in *International Conference on Information Processing in Medical Imaging*. Springer, 2017, pp. 647–658.

[87] Singanallur V Venkatakrishnan, Charles A Bouman, and Brendt Wohlberg, "Plug-and-play priors for model based reconstruction," in *Proc. IEEE Glob. Conf. Signal Inform. Process. (GlobalSIP)*, 2013, pp. 945–948.

[88] Stanley H Chan, Xiran Wang, and Omar A Elgendy, "Plug-and-Play ADMM for image restoration: Fixed-point convergence and applications," *IEEE Trans. Comput. Imaging*, vol. 3, no. 1, pp. 84–98, 2017.

[89] Suhas Sreehari, S Venkat Venkatakrishnan, Brendt Wohlberg, Gregery T Buzzard, Lawrence F Drummy, Jeffrey P Simmons, and Charles A Bouman, "Plug-and-play priors for bright field electron tomography and sparse interpolation," *IEEE Tran. Comput. Imaging*, vol. 2, no. 4, pp. 408–423, 2016.

[90] Yaniv Romano, Michael Elad, and Peyman Milanfar, "The little engine that could: Regularization by denoising (red)," *SIAM Journal on Imaging Sciences*, vol. 10, no. 4, pp. 1804–1844, 2017.

[91] JH Chang, Chun-Liang Li, Barnabás Póczos, BVK Kumar, and Aswin C Sankaranarayanan, “One network to solve them all—Solving linear inverse problems using deep projection models,” arXiv:1703.09912 [cs.CV], 2017.

[92] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G Dimakis, “Compressed sensing using generative models,” arXiv:1703.03208 [stat.ML], 2017.

[93] Brendan Kelly, Thomas P Matthews, and Mark A Anastasio, “Deep learning-guided image reconstruction from incomplete data,” arXiv:1709.00584 [cs.CV], 2017.

[94] Jerome Zhengrong Liang, Patrick J. La Riviere, Georges El Fakhri, Stephen J. Glick, and Jeff Siewerdsen, “Guest editorial low-dose CT: What has been done, and what challenges remain?,” *IEEE Trans. Med. Imag.*, vol. 36, no. 12, pp. 2409–2416, 2017.

[95] S. Ramani and J. A. Fessler, “A splitting-based iterative algorithm for accelerated statistical X-ray CT reconstruction,” *IEEE Trans. Med. Imag.*, vol. 31, no. 3, pp. 677–688, 2012.

[96] Qiong Xu, Hengyong Yu, Xuanqin Mou, Lei Zhang, Jiang Hsieh, and Ge Wang, “Low-dose X-ray CT reconstruction via dictionary learning,” *IEEE Trans. Med. Imag.*, vol. 31, no. 9, pp. 1682–1697, 2012.

[97] Shanzhou Niu, Yang Gao, Zhaoying Bian, Jing Huang, Wufan Chen, Gaohang Yu, Zhengrong Liang, and Jianhua Ma, “Sparse-view X-ray CT reconstruction via total generalized variation regularization,” *Phys. Med. Biol.*, vol. 59, no. 12, pp. 2997–3017, 2014.

[98] Lars Gjesteby, Qingsong Yang, Yan Xi, Ye Zhou, Junping Zhang, and Ge Wang, “Deep learning methods to guide CT image reconstruction and reduce metal artifacts,” in *Medical Imaging 2017: Physics of Medical Imaging*, Orlando, Fl, 2017.

[99] Hu Chen, Yi Zhang, Mannudeep K. Kalra, Feng Lin, Yang Chen, Peixi Liao, Jiliu Zhou, and Ge Wang, “Low-dose CT with a residual encoder-decoder convolutional neural network,” *IEEE Trans. Med. Imag.*, vol. 36, no. 12, pp. 2524–2535, 2017.

[100] Eunhee Kang, Junhong Min, and Jong Chul Ye, "A deep convolutional neural network using directional wavelets for low-dose X-ray CT reconstruction," *Med. Phys.*, vol. 44, no. 10, pp. e360–e375, 2017.

[101] Yoseop Han, Jaejoon Yoo, and Jong Chul Ye, "Deep residual learning for compressed sensing CT reconstruction via persistent homology analysis," arXiv:1611.06391 [cs.CV], 2016.

[102] Bertolt Eicke, "Iteration methods for convexly constrained ill-posed problems in Hilbert space," *Numer. Funct. Anal. Optim.*, vol. 13, no. 5-6, pp. 413–429, 1992.

[103] L. Landweber, "An iteration formula for Fredholm integral equations of the first kind," *Amer. J. Math.*, vol. 73, no. 3, pp. 615–624, 1951.

[104] Dimitri P. Bertsekas, *Nonlinear Programming*, Athena Scientific, Cambridge, MA, 2 edition, 1999.

[105] Patrick L. Combettes and V. Wajs, "Signal recovery by proximal forward-backward splitting," *Multiscale Modeling and Simulation*, vol. 4, no. 4, pp. 1168–1200, 2005.

[106] Patrick L. Combettes and Jean-Christophe Pesquet, *Proximal Splitting Methods in Signal Processing*, pp. 185–212, Springer, New York, NY, 2011.

[107] J. Bect, L. Blanc-Feraud, G. Aubert, and A. Chambolle, "A ℓ_1-unified variational framework for image restoration," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2004, pp. 1–13.

[108] A. Aldroubi and R. Tessera, "On the existence of optimal unions of subspaces for data modeling and clustering," *Found. Comput. Math.*, vol. 11, no. 3, pp. 363âĂŞ–379, 2011.

[109] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proc. Med. Image. Comput. Comput. Assist. Interv. (MICCAI)*, 2015, pp. 234–241.

[110] A. Emin Orhan and Xaq Pitkow, "Skip connections eliminate singularities," arXiv:1701.09175 [cs.NE], 2017.

[111] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2016, pp. 770–778.

[112] C. McCollough, "TU-FG-207A-04: Overview of the Low Dose CT Grand Challenge," *Med. Phys.*, vol. 43, no. 6-part-35, pp. 3759–3760, 2016.

[113] Zhou Wang, Eero P Simoncelli, and Alan C Bovik, "Multiscale structural similarity for image quality assessment," in *Proc. IEEE Asilomar Conf. Signals, Syst., Comput.*, Pacific Grove, CA, Nov. 2003, vol. 2, pp. 1398–1402.

[114] Julien Mairal, Francis Bach, Jean Ponce, and Guillermo Sapiro, "Online learning for matrix factorization and sparse coding," *Journal of Machine Learning Research (JMLR)*, vol. 11, pp. 19–60, 2010.

[115] Joel A Tropp and Anna C Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Trans. Inf. Theory*, vol. 53, no. 12, pp. 4655–4666, 2007.

[116] Kyong Hwan Jin, Harshit Gupta, Jerome Yerly, Matthias Stuber, and Michael Unser, "Time-dependent deep image prior for dynamic mri," *arXiv preprint arXiv:1910.01684*, 2019.

[117] M.A. Griswold, P.M. Jakob, R.M. Heidemann, M. Nittka, V. Jellus, J. Wang, et al., "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magn. Reson. in Med.*, vol. 47, no. 6, pp. 1202–1210, June 2002.

[118] K.P. Pruessmann, M. Weiger, M.B. Scheidegger, and P. Boesigner, "SENSE: Sensitivity encoding for fast MRI," *Magn. Reson. in Med.*, vol. 42, no. 5, pp. 952–962, November 1999.

[119] M. Lustig, D. Donoho, and J.M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magn. Reson. in Med.*, vol. 58, no. 6, pp. 1182–1195, December 2007.

[120] J. Fessler, "Model-based image reconstruction for MRI," *IEEE Sig. Process. Mag.*, vol. 27, no. 4, pp. 81–89, July 2010.

[121] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 27 2015.

[122] J. Schlemper, J. Caballero, J.V. Hajnal, A.N. Price, and D. Rueckert, "A deep cascade of convolutional neural networks for dynamic MR image reconstruction," *IEEE Trans. on Med. Imag.*, vol. 37, no. 2, pp. 491–503, February 2018.

[123] A. Hauptmann, S. Arridge, F. Lucka, V. Muthurangu, and J.A. Steeden, "Real-time cardiovascular MR with spatio-temporal artifact suppression using deep learning—Proof of concept in congenital heart disease," *Magn. Reson. in Med.*, vol. 81, no. 2, pp. 1143–1156, February 2019.

[124] S. Biswas, H.K. Aggarwal, and M. Jacob, "Dynamic MRI using model-based deep learning and SToRM priors: MoDL-SToRM," *Magn. Reson. in Med.*, vol. 82, no. 1, pp. 485–494, July 2019.

[125] Y. Han, J. Yoo, H.H. Kim, H.J. Shin, K. Sung, and J.C. Ye, "Deep learning with domain adaptation for accelerated projection-reconstruction MR," *Magn. Reson. in Med.*, vol. 80, no. 3, pp. 1189–1205, September 2018.

[126] M. Mardani, E. Gong, J.Y. Cheng, S.S. Vasanawala, G. Zaharchuk, L. Xing, et al., "Deep generative adversarial neural networks for compressive sensing MRI," *IEEE Trans. on Med. Imag.*, vol. 38, no. 1, pp. 167–179, January 2019.

[127] H. Jung, K. Sung, K.S. Nayak, E.Y. Kim, and J.C. Ye, "k-t FOCUSS: A general compressed sensing framework for high resolution dynamic MRI," *Magn. Reson. in Med.*, vol. 61, no. 1, pp. 103–116, January 2009.

[128] S.G. Lingala, Y. Hu, E. DiBella, and M. Jacob, "Accelerated dynamic MRI exploiting sparsity and low-rank structure: k-t SLR," *IEEE Trans. on Med. Imag.*, vol. 30, no. 5, pp. 1042–1054, 2011.

[129] L. Feng, R. Grimm, K.T. Block, H. Chandarana, S. Kim, J. Xu, et al., "Golden-angle radial sparse parallel MRI: Combination of compressed sensing, parallel imaging, and golden-angle radial sampling for fast and flexible dynamic volumetric MRI," *Magn. Reson. in Med.*, vol. 72, no. 3, pp. 707–717, September 2014.

[130] L. Feng, L. Axel, H. Chandarana, K.T. Block, D.K. Sodickson, and R. Otazo, "XD-GRASP: Golden-angle radial MRI with reconstruction of extra motion-state dimensions using compressed sensing," *Magn. Reson. in Med.*, vol. 75, no. 2, pp. 775–788, February 2016.

[131] E.J. Candes, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *IEEE Trans. on Info. Theory*, vol. 52, no. 2, pp. 489–509, February 2006.

[132] S. Poddar and M. Jacob, “Dynamic MRI using smoothness regularization on manifolds (SToRM),” *IEEE Trans. on Med. Imag.*, vol. 35, no. 4, pp. 1106–1115, April 2015.

[133] U. Nakarmi, Wang Y., J. Lyu, D. Liang, and L. Ying, “A kernel-based low-rank (KLR) model for low-dimensional manifold recovery in highly accelerated dynamic MRI,” *IEEE Trans. on Med. Imag.*, vol. 36, no. 11, pp. 2297–2307, November 2017.

[134] J. Yerly, G. Ginami, G. Nordio, A.J. Coristine, S. Coppo, P. Monney, et al., “Coronary endothelial function assessment using self-gated cardiac cine MRI and k-t sparse SENSE,” *Magn. Reson. in Med.*, vol. 76, no. 5, pp. 1443–1454, November 2016.

[135] J. Chaptinel, J. Yerly, Y. Mivelaz, M. Prsa, L. Alamo, Y. Vial, et al., “Fetal cardiac cine magnetic resonance imaging *in utero*,” *Scientific Reports*, vol. 7, no. 15540, pp. 1–10, November 14 2017.

[136] J.A. Fessler and B.P. Sutton, “Nonuniform fast Fourier transforms using min-max interpolation,” *IEEE Trans. on Sig. Proc.*, vol. 51, no. 2, pp. 560–574, February 2003.

[137] A.P. Yazdanpanah, O. Afacan, and S.K. Warfield, “Non-learning based deep parallel MRI reconstruction (NLDpMRI),” in *Proc. of the SPIE Conf. on Med. Imag.: Imag. Process.*, San Diego CA, USA, February 16-21 2019, International Society for Optics and Photonics, vol. 10949, pp. 1094904–1094910.

[138] K. Gong, C. Catana, J. Qi, and Q. Li, “PET image reconstruction using deep image prior,” *IEEE Trans. on Med. Imag.*, vol. 38, no. 7, pp. 1655–1665, July 2019.

[139] C. Qin, J. Schlemper, J. Caballero, A.N. Price, J.V. Hajnal, and D. Rueckert, “Convolutional recurrent neural networks for dynamic MR image reconstruction,” *IEEE Trans. on Med. Imag.*, vol. 38, no. 1, pp. 280–290, January 2019.

[140] A. Radford, L. Metz, and S. Chintala, “Unsupervised representation learning with deep convolutional generative adversarial networks,” in *Int. Conf. on Learn. Representations (ICLR)*, San Diego CA, USA, May 7-9 2015.

[141] D.P. Kingma and J. Ba, “ADAM: A method for stochastic optimization,” in *Int. Conf. on Learn. Representations (ICLR)*, San Diego CA, USA, May 7-9 2015.

[142] L. van der Maaten and G. Hinton, “Visualizing data using t-SNE,” *J. of Mach. Learn. Res.*, vol. 9, pp. 2579–2605, November 2008.

[143] Jacques Dubochet, Marc Adrian, Jiin-Ju Chang, Jean-Claude Homo, Jean Lepault, Alasdair W McDowall, and Patrick Schultz, “Cryo-electron microscopy of vitrified specimens,” *Quarterly Reviews of Biophysics*, vol. 21, no. 2, pp. 129–228, 1988.

[144] Richard Henderson, J M_ Baldwin, TA Ceska, F Zemlin, E a Beckmann, and Kenneth H Downing, “Model for the structure of bacteriorhodopsin based on high-resolution electron cryo-microscopy,” *Journal of molecular biology*, vol. 213, no. 4, pp. 899–929, 1990.

[145] Joachim Frank, *Three-dimensional electron microscopy of macromolecular assemblies: visualization of biological molecules in their native state*, Oxford University Press, 2006.

[146] Joachim Frank, Brian Shimkin, and Helen Dowse, “SpiderâĂŤa modular software system for electron image processing,” *Ultramicroscopy*, vol. 6, no. 4, pp. 343–357, 1981.

[147] COS Sorzano, Roberto Marabini, Javier Velázquez-Muriel, José Román Bilbao-Castro, Sjors HW Scheres, José M Carazo, and Alberto Pascual-Montano, “Xmipp: A new generation of an open-source image processing package for electron microscopy,” *Journal of Structural Biology*, vol. 148, no. 2, pp. 194–204, 2004.

[148] Guang Tang, Liwei Peng, Philip R Baldwin, Deepinder S Mann, Wen Jiang, Ian Rees, and Steven J Ludtke, “Eman2: An extensible image processing suite for electron microscopy,” *Journal of Structural Biology*, vol. 157, no. 1, pp. 38–46, 2007.

[149] Nikolaus Grigorieff, “Frealign: High-resolution refinement of single particle structures,” *Journal of Structural Biology*, vol. 157, no. 1, pp. 117–125, 2007.

[150] Michael Hohn, Grant Tang, Grant Goodyear, Philip R Baldwin, Zhong Huang, Pawel A Penczek, Chao Yang, Robert M Glaeser, Paul D Adams, and Steven J Ludtke, “Sparx, a new environment for cryo-em image processing,” *Journal of Structural Biology*, vol. 157, no. 1, pp. 47–55, 2007.

[151] Sjors HW Scheres, “Relion: Implementation of a bayesian approach to cryo-em structure determination,” *Journal of Structural Biology*, vol. 180, no. 3, pp. 519–530, 2012.

[152] J.M. de la Rosa-TrevÃŋn, A. Quintana, L. del Cano, A. ZaldÃŋvar, I. Foche, J. GutiÃľrrez, J. GÃşmez-Blanco, J. Burguet-Castell, J. Cuenca-Alba, V. Abrishami, J. Vargas, J. OtÃşn, G. Sharov, J.L. Vilas, J. Navas, P. Conesa, M. Kazemi, R. Marabini, C.O.S. Sorzano, and J.M. Carazo, “Scipion: A software framework toward integration, reproducibility and validation in 3D electron microscopy,” *Journal of Structural Biology*, vol. 195, no. 1, pp. 93 – 99, 2016.

[153] Ali Punjani, John L Rubinstein, David J Fleet, and Marcus A Brubaker, “cryoSPARC: algorithms for rapid unsupervised cryo-EM structure determination,” *Nature Methods*, vol. 14, no. 3, pp. 290–296, Feb. 2017.

[154] Carlos Oscar Sánchez Sorzano, Roberto Marabini, Alberto Pascual-Montano, Sjors HW Scheres, and José María Carazo, “Optimization problems in electron microscopy of single particles,” *Annals of Operations Research*, vol. 148, no. 1, pp. 133–165, 2006.

[155] Richard Henderson, Andrej Sali, Matthew L. Baker, Bridget Carragher, Batsal Devkota, Kenneth H. Downing, Edward H. Egelman, Zukang Feng, Joachim Frank, Nikolaus Grigorieff, Wen Jiang, Steven J. Ludtke, Ohad Medalia, Pawel A. Penczek, Peter B. Rosenthal, Michael G. Rossmann, Michael F. Schmid, Gunnar F. SchrÃűder, Alasdair C. Steven, David L. Stokes, John D. Westbrook, Willy Wriggers, Huanwang Yang, Jasmine Young, Helen M. Berman, Wah Chiu, Gerard J. Kleywegt, and Catherine L. Lawson, “Outcome of the first electron microscopy validation task force meeting,” *Structure*, vol. 20, no. 2, pp. 205–214, 2012.

[156] Tamir Bendory, Alberto Bartesaghi, and Amit Singer, "Single-particle cryo-electron microscopy: Mathematical theory, computational challenges, and opportunities," *IEEE Signal Processing Magazine*, vol. 37, no. 2, pp. 58–76, 2020.

[157] Amit Singer and Fred J Sigworth, "Computational methods for single-particle electron cryomicroscopy," *Annual Review of Biomedical Data Science*, vol. 3, 2020.

[158] Alberto Bartesaghi, Alan Merk, Soojay Banerjee, Doreen Matthies, Xiongwu Wu, Jacqueline LS Milne, and Sriram Subramaniam, "2.2 å resolution cryo-em structure of β-galactosidase in complex with a cell-permeant inhibitor," *Science*, vol. 348, no. 6239, pp. 1147–1151, 2015.

[159] Joachim Frank, *Electron tomography: methods for three-dimensional visualization of structures in the cell*, Springer Science & Business Media, 2008.

[160] F. Natterer, *The mathematics of computerized tomography*, Society for Industrial and Applied Mathematics, jan 2001.

[161] Miloš Vulović, Raimond B.G. Ravelli, Lucas J. van Vliet, Abraham J. Koster, Ivan Lazić, Uwe Lücken, Hans Rullgård, Ozan Öktem, and Bernd Rieger, "Image formation modeling in cryo-electron microscopy," *Journal of Structural Biology*, vol. 183, no. 1, pp. 19–32, July 2013.

[162] Hans Rullgård, L-G Öfverstedt, Sergey Masich, Bertil Daneholt, and Ozan Öktem, "Simulation of transmission electron microscope images of biological specimens," *Journal of microscopy*, vol. 243, no. 3, pp. 234–256, 2011.

[163] M. Unser, "Sampling—50 years after Shannon," *Proceedings IEEE*, vol. 88, no. 4, pp. 569–587, apr 2000.

[164] Martin Arjovsky, Soumith Chintala, and Léon Bottou, "Wasserstein generative adversarial networks," in *International conference on machine learning*, 2017, pp. 214–223.

[165] Cédric Villani, *Optimal transport: old and new*, vol. 338, Springer Science & Business Media, 2008.

[166] Gabriel Peyré, Marco Cuturi, et al., "Computational optimal transport," *Foundations and Trends® in Machine Learning*, vol. 11, no. 5-6, pp. 355–607, 2019.

[167] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville, "Improved training of wasserstein gans," in *Advances in neural information processing systems*, 2017, pp. 5767–5777.

[168] Wim van Aarle, Willem Jan Palenstijn, Jan De Beenhouwer, Thomas Altantzis, Sara Bals, K Joost Batenburg, and Jan Sijbers, "The astra toolbox: A platform for advanced algorithm development in electron tomography," *Ultramicroscopy*, vol. 157, pp. 35–47, 2015.

[169] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al., "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems*, 2019, pp. 8024–8035.

[170] Victor M Panaretos et al., "On random tomography with unobservable projection angles," *The Annals of Statistics*, vol. 37, no. 6A, pp. 3272–3306, 2009.

[171] Sigurdur Helgason, *The radon transform*, vol. 2, Springer, 1980.

[172] Pawel A. Penczek, Robert A. Grassucci, and Joachim Frank, "The ribosome at improved resolution: New techniques for merging and orientation refinement in 3D cryo-electron microscopy of biological particles," *Ultramicroscopy*, vol. 53, no. 3, pp. 251 – 270, 1994.

[173] T.S. Baker and R.H. Cheng, "A model-based approach for determining orientations of biological macromolecules imaged by cryoelectron microscopy," *Journal of Structural Biology*, vol. 116, no. 1, pp. 120–130, 1996.

[174] Fred J Sigworth, "A maximum-likelihood approach to single-particle image refinement," *Journal of structural biology*, vol. 122, no. 3, pp. 328–339, 1998.

[175] Fred J Sigworth, Peter C Doerschuk, Jose-Maria Carazo, and Sjors HW Scheres, "An introduction to maximum-likelihood methods in cryo-em," in *Methods in enzymology*, vol. 482, pp. 263–294. Elsevier, 2010.

[176] Zvi Kam, “The reconstruction of structure from electron micrographs of randomly oriented particles,” in *Electron Microscopy at Molecular Dimensions*, pp. 270–277. Springer, 1980.

[177] Nir Sharon, Joe Kileel, Yuehaw Khoo, Boris Landa, and Amit Singer, “Method of moments for 3-d single particle ab initio modeling with non-uniform distribution of viewing angles,” *Inverse Problems*, 2019.

[178] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen Awm Van Der Laak, Bram Van Ginneken, and Clara I Sánchez, “A survey on deep learning in medical image analysis,” *Medical image analysis*, vol. 42, pp. 60–88, 2017.

[179] Michael T. McCann, Kyong Hwan Jin, and Michael Unser, “Convolutional neural networks for inverse problems in imaging: A review,” *IEEE Signal Processing Magazine*, vol. 34, no. 6, pp. 85–95, Nov. 2017.

[180] George Barbastathis, Aydogan Ozcan, and Guohai Situ, “On the use of deep learning for computational imaging,” *Optica*, vol. 6, no. 8, pp. 921–943, 2019.

[181] Tristan Bepler, Alex J Noble, and Bonnie Berger, “Topaz-denoise: general deep denoising models for cryoem,” *bioRxiv*, p. 838920, 2019.

[182] Feng Wang, Huichao Gong, Gaochao Liu, Meijing Li, Chuangye Yan, Tian Xia, Xueming Li, and Jianyang Zeng, “Deeppicker: a deep learning approach for fully automated particle picking in cryo-em,” *Journal of structural biology*, vol. 195, no. 3, pp. 325–336, 2016.

[183] Yanan Zhu, Qi Ouyang, and Youdong Mao, “A deep convolutional neural network approach to single-particle recognition in cryo-electron microscopy,” *BMC bioinformatics*, vol. 18, no. 1, pp. 348, 2017.

[184] Dimitry Tegunov and Patrick Cramer, “Real-time cryo-em data pre-processing with warp,” *BioRxiv*, p. 338558, 2018.

[185] Thorsten Wagner, Felipe Merino, Markus Stabrin, Toshio Moriya, Claudia Antoni, Amir Apelbaum, Philine Hagel, Oleg Sitsel, Tobias Raisch, Daniel Prumbaum, et al., “Sphire-cryolo is a fast and accurate fully automated

particle picker for cryo-em," *Communications Biology*, vol. 2, no. 1, pp. 218, 2019.

[186] Tristan Bepler, Andrew Morin, Micah Rapp, Julia Brasch, Lawrence Shapiro, Alex J Noble, and Bonnie Berger, "Positive-unlabeled convolutional neural networks for particle picking in cryo-electron micrographs," *Nature methods*, pp. 1–8, 2019.

[187] Ellen D. Zhong, Tristan Bepler, Joseph H. Davis, and Bonnie Berger, "Reconstructing continuous distributions of 3D protein structure from cryo-em images," in *International Conference on Learning Representations*, 2020.

[188] Nina Miolane, Frédéric Poitevin, Yee-Ting Li, and Susan Holmes, "Estimation of orientation and camera parameters from cryo-electron microscopy images with variational autoencoders and generative adversarial networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 970–971.

[189] Ashish Bora, Eric Price, and Alexandros G Dimakis, "AmbientGAN: Generative models from lossy measurements.," *ICLR*, vol. 2, pp. 5, 2018.

[190] Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo Martin-Brualla, Tomas Simon, Jason Saragih, Matthias Nießner, et al., "State of the art on neural rendering," *arXiv preprint arXiv:2004.03805*, 2020.

[191] Shubham Tulsiani, Alexei A Efros, and Jitendra Malik, "Multi-view consistency as supervisory signal for learning shape and pose prediction," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2897–2905.

[192] Matheus Gadelha, Subhransu Maji, and Rui Wang, "3D shape induction from 2D views of multiple objects," in *2017 International Conference on 3D Vision (3DV)*. IEEE, 2017, pp. 402–411.

[193] Shakir Mohamed and Balaji Lakshminarayanan, "Learning in implicit generative models," *arXiv preprint arXiv:1610.03483*, 2016.

[194] Harshit Gupta, Thong H. Phan, Jaejun Yoo, and Michael Unser, “Multi-cryogan: Reconstruction of continuous conformations in cryo-em using generative adversarial networks,” in *Proc. European Conference on Computer Vision Workshops (ECCVW August 23-28)*, 2020.

[195] Joachim Frank and Abbas Ourmazd, “Continuous changes in structure mapped by manifold embedding of single-particle data in cryo-em,” *Methods*, vol. 100, pp. 61–67, 2016.

[196] Abbas Ourmazd, “Cryo-em, xfels and the structure conundrum in structural biology,” *Nature methods*, vol. 16, no. 10, pp. 941–944, 2019.

[197] Ellen D Zhong, Tristan Bepler, Bonnie Berger, and Joseph H Davis, “Cryodrgn: Reconstruction of heterogeneous structures from cryo-electron micrographs using neural networks,” *bioRxiv*, 2020.

[198] Kurt Hornik, Maxwell Stinchcombe, Halbert White, et al., “Multilayer feedforward networks are universal approximators.,” .

[199] Carlos Oscar S Sorzano, A Jiménez, Javier Mota, José Luis Vilas, David Maluenda, M Martínez, E Ramírez-Aportela, T Majtner, J Segura, Ruben Sánchez-García, et al., “Survey of the analysis of continuous conformational variability of biological macromolecules by electron microscopy,” *Acta Crystallographica Section F: Structural Biology Communications*, vol. 75, no. 1, pp. 19–32, 2019.

[200] Joakim AndÃĪn, Eugene Katsevich, and Amit Singer, “Covariance estimation using conjugate gradient for 3d classification in cryo-EM,” pp. 200–204.

[201] Ali Dashti, Peter Schwander, Robert Langlois, Russell Fung, Wen Li, Ahmad Hosseinizadeh, Hstau Y Liao, Jesper Pallesen, Gyanesh Sharma, Vera A Stupina, et al., “Trajectories of the ribosome as a brownian nanomachine,” *Proceedings of the National Academy of Sciences*, vol. 111, no. 49, pp. 17492–17497, 2014.

[202] Amit Moscovich, Amit Halevi, Joakim Andén, and Amit Singer, “Cryo-em reconstruction of continuous heterogeneity by laplacian spectral volumes,” *Inverse Problems*, vol. 36, no. 2, pp. 024003, 2020.

[203] Roy R. Lederman, Joakim AndÃĺn, and Amit Singer, "Hyper-Molecules: on the Representation and Recovery of Dynamical Structures, with Application to Flexible Macro-Molecular Structures in Cryo-EM," *Inverse Problems*, vol. 36, Apr. 2020.

[204] Evan Seitz, Francisco Acosta-Reyes, Peter Schwander, and Joachim Frank, "Simulation of cryo-em ensembles from atomic models of molecules exhibiting continuous conformations," *BioRxiv*, p. 864116, 2019.

[205] Thomas Debarre, Julien Fageot, Harshit Gupta, and Michael Unser, "B-spline-based exact discretization of continuous-domain inverse problems with generalized tv regularization," *IEEE Transactions on Information Theory*, vol. 65, no. 7, pp. 4457–4470, 2019.

[206] Shayan Aziznejad, Harshit Gupta, Joaquim Campos, and Michael Unser, "Deep neural networks with trainable activations and controlled lipschitz constant," *arXiv preprint arXiv:2001.06263*, 2020.

[207] Mohammad Zalbagi Darestani and Reinhard Heckel, "Can un-trained neural networks compete with trained neural networks at image reconstruction?," *arXiv preprint arXiv:2007.02471*, 2020.

[208] Gauri Jagatap and Chinmay Hegde, "Algorithmic guarantees for inverse imaging with untrained network priors," in *Advances in Neural Information Processing Systems*, 2019, pp. 14832–14842.

[209] Reinhard Heckel and Mahdi Soltanolkotabi, "Compressive sensing with untrained neural networks: Gradient descent finds the smoothest approximation," *arXiv preprint arXiv:2005.03991*, 2020.

[210] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen, "Progressive growing of gans for improved quality, stability, and variation," *arXiv:1710.10196*, 2017.

[211] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee, "Accurate image super-resolution using very deep convolutional networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 1646–1654.

[212] Tero Karras, Samuli Laine, and Timo Aila, "A style-based generator architecture for generative adversarial networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 4401–4410.

[213] Miloš Vulović, Raimond BG Ravelli, Lucas J van Vliet, Abraham J Koster, Ivan Lazić, Uwe Lücken, Hans Rullgård, Ozan Öktem, and Bernd Rieger, "Image formation modeling in cryo-electron microscopy," *Journal of structural biology*, vol. 183, no. 1, pp. 19–32, 2013.

[214] Walter Rudin, *Real and Complex Analysis*, Tata McGraw-Hill Education, 1987.

[215] Andrew J Kurdila and Michael Zabarankin, *Convex functional analysis*, Springer Science & Business Media, 2006.

[216] Reed Michael and Barry Simon, *Methods of modern mathematical physics I: Functional analysis*, Academic Press, 1980.

[217] Ken Sauer and Charles Bouman, "A local update strategy for iterative reconstruction from projections," *IEEE Trans. Signal Process.*, vol. 41, no. 2, pp. 534–548, 1993.

[218] Idris A Elbakri and Jeffrey A Fessler, "Statistical image reconstruction for polyenergetic X-ray computed tomography," *IEEE Trans. Med. Imag.*, vol. 21, no. 2, pp. 89–99, 2002.

[219] Heinz H. Bauschke and Patrick L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, Springer, New York, NY, 2011.

[220] Eric F Pettersen, Thomas D Goddard, Conrad C Huang, Gregory S Couch, Daniel M Greenblatt, Elaine C Meng, and Thomas E Ferrin, "Ucsf chimeraâĂŤa visualization system for exploratory research and analysis," *Journal of computational chemistry*, vol. 25, no. 13, pp. 1605–1612, 2004.

[221] Alexis Rohou and Nikolaus Grigorieff, "Ctffind4: Fast and accurate defocus estimation from electron micrographs," *Journal of structural biology*, vol. 192, no. 2, pp. 216–221, 2015.

[222] Nikhil Biyani, Ricardo D Righetto, Robert McLeod, Daniel Caujolle-Bert, Daniel Castano-Diez, Kenneth N Goldie, and Henning Stahlberg, “Focus: The interface between data collection and data processing in cryo-em,” *Journal of structural biology*, vol. 198, no. 2, pp. 124–133, 2017.

[223] Ricardo D Righetto, Nikhil Biyani, Julia Kowal, Mohamed Chami, and Henning Stahlberg, “Retrieving high-resolution information from disordered 2d crystals by single-particle cryo-em,” *Nature communications*, vol. 10, no. 1, pp. 1–10, 2019.

Curriculum Vitæ

HARSHIT GUPTA

harshit.gupta.cor@gmail.com
BM 4.134, EPFL, Lausanne CH-1015, Switzerland
Homepage ◇ Google Scholar

CURRENT RESEARCH FOCUS

My research is aimed at designing mathematically backed deep-learning algorithms for solving inverse problems in imaging. I am interested in the modalities of Cryo Electron Microscopy (Cryo-EM), Computational Tomography (CT), and Magnetic Resonance Imaging (MRI).

EDUCATION

July 2015 - September 2020 **École polytechnique fédérale de Lausanne (EPFL), Switzerland**
Ph.D. in Electrical Engineering
Thesis: "From Classical to Unsupervised-Deep-Learning Methods for Solving Inverse Problems in Imaging".
Advisor: Prof. Michael Unser

July 2011 - May 2015 **Indian Institute of Technology (IIT), Guwahati, India**
B. Tech in Electronics and Communications Engineering

RESEARCH EXPERIENCES

July 2014 - May 2015 **Indian Institute of Technology (IIT), Guwahati, India**
Bachelor Thesis Project
Topic: "Blind Image Quality Assessment"
Advisor: Prof. Kannan Karthik

May 2014 - July 2014 **École polytechnique fédérale de Lausanne (EPFL), Switzerland**
Research Internship
Topic: "Interpolation using Derivatives"
Advisor: Prof. Michael Unser

May 2013 - July 2013 **Indian Institute of Science (IISc), Bangalore, India**
Research Internship
Topic: "Building a MATLAB GUI on Optic Disk Localization using ℓ_1-minimization"
Advisor: Prof. Chandra Sekhar Seelamantula

PUBLICATIONS

Preprints

8. **Gupta H***, McCann M T*, Donati L, Unser M, "CryoGAN: A New Reconstruction Paradigm for Single-particle Cryo-EM Via Deep Adversarial Learning," bioRxiv 2020.03.20.001016, March 2020. **Co-first authors.* [PDF]

7. Jin K H*, **Gupta H***, Yerly J, Stuber M, Unser M, "Time-Dependent Deep Image Prior for Dynamic MRI," IEEE Transactions on Medical Imaging, in Revision. **Co-first authors.* [PDF]

Journals

6. Aziznejad S, **Gupta H**, Campos J, Unser M, "Deep Neural Networks with Trainable Activations and Controlled Lipschitz Constant," IEEE Transactions on Signal Processing, vol. 68, pp. 4688 - 4699, August 2020. [PDF]
5. Yang F, Pham T, **Gupta H**, Unser M, Ma J, "Deep-learning projector for optical diffraction tomography," Optics Express, vol. 28(3), pp. 3905-3921, February 2020. [PDF]
4. Debarre T, Fageot J, **Gupta H**, Unser M, "B-spline-based exact discretization of continuous-domain inverse problems with generalized TV regularization," IEEE Transactions on Information Theory, vol. 65(7), pp.4457-4470, March 2019. [PDF]
3. **Gupta H**, Jin K H, Nguyen H Q, McCann M T, Unser M, "CNN-based projected gradient descent for consistent CT image reconstruction," IEEE Transactions on Medical Imaging, vol. 37(6), pp. 1440-1453, May 2018. [PDF]
2. **Gupta H**, Fageot J, Unser M, "Continuous-domain solutions of linear inverse problems with Tikhonov versus generalized TV regularization," IEEE Transactions on Signal Processing, vol. 66(17), pp. 4670-4684, July 2018. [PDF]
1. Unser M, Fageot J, **Gupta H**, "Representer Theorems for Sparsity-Promoting ℓ_1 Regularization," IEEE Transactions on Information Theory, vol. 62(9), pp. 5167-5180, August 2016. [PDF]

Conference and Workshop Proceedings

4. **Gupta H**, Phan T H, Yoo J, Unser M, "Multi-CryoGAN: Reconstruction of continuous conformations in Cryo-EM using Generative Adversarial Networks," Proc. European Conference on Computer Vision Workshops (ECCVW 2020) (Online, August 23-28), in press. [PDF]
3. Debarre T, Fageot J, **Gupta H**, Unser M, "Solving Continuous-domain Problems Exactly with Multiresolution B-splines," Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2019) (Brighton, UK, May 12-17), pp. 5122-5126. [PDF]
2. **Gupta H**, Schmitter D, Uhlmann V, Unser M, "General surface energy for spinal cord and aorta segmentation," IEEE Proc. International Symposium on Biomedical Imaging (ISBI 2017), (Sydney, Australia, April 18-21), pp. 319-322. [PDF]
1. Uhlmann V, Fageot J, **Gupta H**, Unser M, "Statistical optimality of Hermite splines," Proc. International Conference on Sampling Theory and Applications (SampTA 2015), (Washington, DC, US, May 25-29), pp. 226-230. [PDF]

TEACHING EXPERIENCES

September 2015 - August 2020	**Teaching Assistant at EPFL** Image Processing I - Autumn 2015, 2016, 2017, 2018, 2019 Image Processing II - Spring 2016, 2017, 2018, 2019, 2020
September 2019 - February 2019	**Supervisor for Master Semester Project** Student: Huy Thong, EPFL Topic: “Reconstructing multiple-conformations of particles in Cryo-Electron Microscopy with deep learning”
January 2019 - June 2019	**Co-supervisor for Master Semester Project** Student: Huy Thong, EPFL Topic: “Implementing Deep-learning-based iterative algorithm to solve inverse problem of MRI”
January 2019 - June 2019	**Supervisor for Master Semester Project** Student: Joaquim Campos, EPFL Topic: “Learning Spline-based activations for very deep learning”
September 2018 - February 2019	**Co-supervisor for Master Thesis** Student: Matthieu Broisin, EPFL in collaboration with MIT, USA Topic: “Segmentation of images using a Deep-Learning-based approach”
April 2017 - Septmeber 2017	**Co-supervisor for Master Thesis** Student: Thomas Debarre, ENS Paris Saclay, Cachan, France Topic: “B-spline-based exact discretization of continuous-domain inverse problems with generalized TV regularization”

TECHNICAL STRENGTHS

Programming Languages	Python, Matlab, Java, C, C++
Libraries	PyTorch, MatConvNet
Softwares	ImageJ, Fiji, Chimera

HONOURS

- Selected for II round of Texas Instruments Innovation Challenge: India Analog Design Contest 2014.
- Selected in national Top-30 in Manthan, CAG, 2014, among more than 150 teams.
- Placed among top 0.5% in 2011 IIT-Joint Entrance Exam (to enroll in undergraduate program) given by 500,000 students.
- Placed among National Top 1% in National Standard Examination in Physics, 2010-11, organized by Indian Association of Physics Teachers.
- Secured AIR-171 in National Level Science Talent Search Examination, 2009.
- Secured 3rd position in SBM Inter School Science and Environment Quiz, 2008.
- Awarded the Talent Scholarship Award by Saraswati Siksha Sansthan, 2008.