
Nonideal Sampling and Regularized Interpolation of Noisy Data

Sathish Ramani

Thèse N° 4305 (January 2009)

*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. René-Paul Salathé, *président*
Prof. Michael Unser, *directeur de thèse*
Prof. Erik Meijering, *rapporteur*
Prof. Sylvain Meignen, *rapporteur*
Prof. Stephan Morgenthaler, *rapporteur*

École polytechnique fédérale de Lausanne—2009

Cover design by Annette Unser
Printing and binding by Repro-EPFL
Typeset with L^AT_EX
Copyright © 2009 by Sathish Ramani
Available at <http://bigwww.epfl.ch/>

Abstract

Conventional sampling (Shannon’s sampling formulation and its approximation-theoretic counterparts) and interpolation theories provide effective solutions to the problem of reconstructing a signal from its samples, but they are primarily restricted to the noise-free scenario. The purpose of this thesis is to extend the standard techniques so as to be able to handle noisy data.

First, we consider a realistic setting where a multidimensional signal is pre-filtered prior to sampling, and the samples are corrupted by additive noise. In order to counterbalance the effect of noise, the reconstruction problem is formulated in a variational framework where the solution is obtained by minimizing a continuous-domain Tikhonov-like L_2 -regularization subject to a ℓ_p -based data fidelity constraint. We present theoretical justification for the minimization of this cost functional and show that the global-minimum solution belongs to a shift-invariant space generated by a function that is generally not bandlimited. The optimal reconstruction space is characterized by a condition that links the generating function to the regularization operator and implies the existence of a B-spline-like basis. We also consider stochastic formulations—min-max and minimum mean-squared error (MMSE/Wiener) formulations—of the nonideal sampling problem and show that they yield the same type of estimators and point towards the existence of optimal shift-invariant spaces for certain classes of stochastic processes. In the stochastic context, we also derive an exact formula for the error of approximating a stationary stochastic signal in the presence of discrete additive noise and justify the noise-reducing effect of regularization through illustrations.

Next, we focus on the use of a much wider class of non-quadratic regularization functionals for the problem of interpolation in the presence of noise. Starting from the affine-invariance of the solution, we show that the L_p -norm ($p \neq 2$) is the most suitable type of non-quadratic regularization for our purpose. We give monotonically convergent numerical algorithms to carry out the

minimization of the non-quadratic cost criterion. We also demonstrate experimentally that the proposed regularized interpolation scheme provides superior interpolation performance compared to standard methods in the presence of noise.

Finally, we address the problem of selecting an appropriate value for the regularization parameter which is most crucial for the working of variational methods in general including those discussed in this thesis. We propose a practical scheme that is based on the concept of risk estimation to achieve minimum MSE performance. In this context, we first review a well known result due to Stein (Stein's unbiased risk estimate—SURE) that is applicable for data corrupted by additive Gaussian noise and also derive a new risk estimate for a Poisson-Gaussian mixture model that is appropriate for certain biomedical imaging applications. Next, we introduce a novel and efficient Monte-Carlo technique to compute SURE for arbitrary nonlinear algorithms. We demonstrate experimentally that the proposed Monte-Carlo SURE yields regularization parameter values that are close to the oracle-optimum (minimum MSE) for all methods considered in this work. We also present results that illustrate the applicability of our technique to a wide variety of algorithms in denoising and deconvolution.

Keywords

Nonideal sampling, interpolation, regularization, regularization parameter, shift-invariant spaces, smoothing splines, Tikhonov criterion, total-variation functional, Wiener solution, Matérn class, Stein's unbiased risk estimate (SURE), Monte-Carlo simulation.

Résumé

Les théories traditionnelles d'échantillonnage et d'interpolation (formule d'échantillonnage de Shannon et ses versions en théorie de l'approximation) offrent des solutions efficaces au problème de la reconstruction d'un signal à partir de ses échantillons, mais elles ne sont applicables que dans un scénario où le signal n'est pas corrompu par du bruit. Le but de cette thèse est d'étendre ces techniques standard de sorte à pouvoir traiter des données bruitées.

Tout d'abord, nous considérons un cas réaliste où un signal multidimensionnel a été préfiltré avant d'avoir été échantillonné, et où les échantillons sont corrompus par un bruit additif. Pour contrecarrer l'effet du bruit, le problème de reconstruction est formulé dans un cadre variationnel où la solution s'obtient par la minimisation dans le domaine continu d'un terme de régularisation L_2 inspiré par Tikhonov, sous la contrainte d'un terme de fidélité aux données en norme ℓ_p . Nous proposons une justification théorique à la minimisation de cette fonction-coût et montrons que la solution optimale appartient à un espace invariant par translation engendré par une fonction qui n'est généralement pas à bande limitée. L'espace de reconstruction optimal se caractérise par une condition qui lie la fonction de génération à l'opérateur de régularisation et sous-entend l'existence d'une base similaire à un B-spline. De plus, nous examinons les formulations stochastiques du problème d'échantillonnage non idéal telles que formulations min-max, erreur aux moindres carrés, et formulation de Wiener, et montrons que, pour certaines classes de processus stochastiques, elles mènent au même type d'estimateur et qu'elles indiquent l'existence d'espaces optimaux invariants par translation. Toujours dans un contexte stochastique, nous établissons la formule exacte de l'erreur d'approximation d'un signal stochastique stationnaire en présence d'un bruit additif échantillonné. Nous justifions par le biais d'illustrations l'effet de réduction de bruit qu'apporte la régularisation.

Ensuite, nous examinons une classe beaucoup plus large de fonctionnelles non

quadratiques de régularisation destinées à l'interpolation en présence de bruit. En partant de l'invariance de la solution à une transformation affine, nous montrons que la norme L_p , avec $p \neq 2$, est le type de régularisation non quadratique qui sied le mieux à notre problème. Nous proposons des algorithmes numériques à convergence monotone pour minimiser le critère de coût non quadratique. En outre, nous démontrons par des expériences que la méthode d'interpolation régularisée que nous proposons offre une meilleure performance d'interpolation en présence de bruit que les méthodes standard.

Enfin, nous nous occupons du problème qui consiste à donner la bonne valeur au paramètre de régularisation, qui est un aspect des plus cruciaux du fonctionnement de toute méthode variationnelle, y compris de celles discutées dans cette thèse. Nous proposons une approche pratique fondée sur le concept d'estimation de risque pour atteindre des performances optimales au sens des moindres carrés. Dans ce contexte, nous commençons par l'examen d'un résultat dû à Stein (estimation de risque non biaisée de Stein, ou "Stein's unbiased risk estimate—SURE") qui s'applique à des données corrompues par du bruit additif gaussien ; nous en déduisons un nouvel estimateur de risque pour un modèle de Gauss-Poisson qui est propice à certaines applications biomédicales. Puis nous proposons une nouvelle technique efficace de Monte-Carlo pour calculer SURE dans le cas d'algorithmes arbitraires, y compris non linéaires. Nous démontrons par des expériences que notre méthode "Monte-Carlo SURE" donne au paramètre de régularisation des valeurs qui sont proches des valeurs optimales au sens des moindres carrés qu'un oracle aurait pu proposer, pour chacun des cas considérés dans ce travail. Nous présentons finalement des résultats qui illustrent la pertinence de notre méthode pour un large éventail d'algorithmes de débruitage et de déconvolution.

Mots-clés

Problème d'échantillonnage non idéal, interpolation, régularisation, paramètre de régularisation, des espaces invariant par translation, smoothing splines, critère de Tikhonov, critère total-variation, solution de Wiener, classe de Matérn, estimation de risque non biaisée de Stein (Stein's unbiased risk estimate—SURE), simulation de type Monte-Carlo.

Acknowledgements

I express my heartfelt thanks to my advisor Prof. Michael Unser for hosting me at Biomedical Imaging Group (BIG) and for his guidance throughout the course of this work. I have benefitted a lot from his sense of intuition for seeing through tough problems and his tutorage about the art of presenting scientific material.

I thank the jury members Prof. René-Paul Salathé (STI-IMT-LOA, EPFL), Prof. Erik Meijering (BIGR, Erasmus MC-University Medical Center Rotterdam), Prof. Sylvain Meignen (MGMI, Laboratoire Jean Kuntzmann), and Prof. Stephan Morgenthaler (SB-IMA-STAP, EPFL) for their constructive criticism and positive feedback about the thesis.

I also take this opportunity to thank the Swiss National Science Foundation (SNSF) for providing financial support.

I am grateful to Prof. Thierry Blu (The Chinese University of Hong Kong), Dr. Dimitri Van De Ville, and Dr. Philippe Thévenaz who were always ready to share their knowledge. I am very happy to note that I have worked with them closely; discussions with them were stimulating and have contributed a lot to this thesis. I thank Dr. Philippe Thévenaz also for the assistance he provided in translating the abstract of this thesis.

I express my sincere appreciation for Dr. Daniel Sage who has been very kind to offer help on scientific and administrative grounds including computer-support. My BIG colleagues have also been very friendly and cooperative. I particularly enjoyed the discussions with Cédric Vonesch, Florian Luisier, Pouya Dehghani Tafti, Aurélien Bourquard, and Matthieu Guerquin-Kern and I am looking forward to collaborating with them in the future.

Many thanks are also due to Ms. Manuelle Borruat and Ms. Yvette Bernhard for the help and support they offered me in handling my administrative duties.

I am thankful to all my friends including Karthick Babu, Ramachandra Rao,

Rajesh Langoju, Abhijit Patil, Venkatesh Sivasubramaniam, Sarath Kumar, Chandra Sekhar and Kunal Narayan Chaudhury for making my stay in Swiss more memorable.

I would also like to express my deep sense of gratitude to Bhagawan Sri Sathya Sai Baba for providing cost-free education during my undergraduate and post-graduate studies.

I am indebted to my family members, especially my wife and my parents, for sharing my enthusiasm and for their encouragement in all my endeavors. I would like to dedicate this work to them in recognition of their invaluable support.

Contents

Abstract	i
Résumé	iii
Acknowledgements	v
1 Introduction	1
1.1 Sampling and Interpolation	1
1.2 Interpolation for Biomedical Imaging	2
1.2.1 Digital Correction	3
1.2.2 Optimal Filtering	3
1.3 Contributions of the Thesis	5
1.3.1 Theoretical Investigation and Numerical Implementation .	5
1.3.2 Selection of the Best Regularization Parameter	6
1.4 Thesis Outline	6
2 Noise-Free Sampling and Linear Approximation Theory	9
2.1 Introduction	9
2.1.1 Shannon's Theorem Reinterpreted	10
2.2 Preliminaries	12
2.2.1 Notations	13
2.2.2 Shift-Invariant Spaces and Riesz Bases	15
2.3 Sampling in Shift-Invariant Spaces	20
2.3.1 Minimum-Error Sampling	22
2.3.2 Consistent Sampling	22
2.3.3 Interpolation	24
2.4 Characterizing the Approximation Error	25

2.5	Approximation Order	29
2.5.1	Strang-Fix Conditions	29
2.5.2	Asymptotic Form of the Error for $d = 1$	30
2.6	Quasi-Interpolation	31
2.7	Summary	31
3	Nonideal Sampling in the Presence of Noise	33
3.1	Introduction	33
3.1.1	Contributions	35
3.1.2	Organization of the Chapter	36
3.2	Preliminaries	37
3.3	Nonideal Sampling and Regularization	39
3.3.1	Problem Specification	39
3.3.2	Consistent Sampling in $V(\varphi)$	42
3.3.3	Space of Admissible Solutions	43
3.3.4	Solution to the Variational Problem	45
3.4	Stability of the Reconstruction	50
3.5	Case Illustrations	54
3.5.1	Ideal versus Nonideal Sampling	54
3.5.2	Regularization Operators and Reconstruction Spaces	55
3.6	Stochastic Formulations	60
3.6.1	Connection with the Minimax Estimator	60
3.6.2	Unification with the Wiener Formulation	62
3.7	Matérn Model for the Nonideal Sampling Problem	65
3.7.1	Anisotropic Matérn B-spline	66
3.7.2	Reconstruction using the Matérn Model	67
3.8	An Approximation Error Formula for the Noisy Scenario	72
3.9	Summary	78
4	Non-quadratic Regularization for Interpolation of Noisy Data	79
4.1	Introduction	79
4.1.1	Contributions	80
4.1.2	Organization of the Chapter	81
4.2	Standard Interpolation	81
4.2.1	Standard Interpolation in the Presence of Noise	82
4.3	Regularized Interpolation	84
4.3.1	Problem Formulation	84
4.3.2	Discretization of the Problem	89
4.3.3	Algorithms for Regularized Interpolation	91

4.3.4	Spline-Based Regularized Interpolation	96
4.4	Results	97
4.4.1	Regularized versus Standard Interpolation	99
4.4.2	Regularized Interpolation: Varying Spline Degree	102
4.4.3	Experiments with Poisson Noise	102
4.5	Discussion	104
4.5.1	Quadratic versus Non-Quadratic Regularization	104
4.5.2	Influence of the Likelihood Model	104
4.5.3	Computational Cost	105
4.6	Summary	106
5	Data-driven Methods for Risk Estimation	109
5.1	Introduction	109
5.2	Contributions and Chapter Organization	110
5.3	Denoising Using Stein's Unbiased Risk Estimate	111
5.3.1	Notation and Problem formulation	111
5.3.2	Stein's Unbiased Risk Estimate—SURE	112
5.3.3	Results	115
5.4	Deconvolution of 3-D Fluorescence Micrographs	122
5.4.1	New Unbiased Risk Estimate for a Poisson-Gaussian Mixture Model	123
5.4.2	Deconvolution Algorithm	125
5.4.3	Estimation of CCD parameters	126
5.4.4	Results	127
5.5	Summary	134
6	Monte-Carlo Risk Minimization for Non-linear Algorithms	137
6.1	Problem Statement	137
6.2	Contributions and Chapter Organization	138
6.3	Monte-Carlo Estimation of the Divergence	139
6.3.1	Linear Algorithm with Unstructured System Matrices	139
6.3.2	General Algorithm for Non-Linear Problems	141
6.4	Description of Some Popular Denoising Algorithms	144
6.4.1	Scale-Dependent Wavelet Soft-Thresholding	145
6.4.2	Total-Variation Denoising (TVD)	147
6.5	Range of Validity of the Proposed Monte-Carlo SURE	148
6.6	Results	150
6.6.1	Results with One-Parameter Optimization	150
6.6.2	Results with Multi-Parameter Optimization	160

6.7 Summary	164
7 Summary and Conclusions	167
A Details of Some Results of Chapter 2	171
A.1 Sampling in the \mathbf{DZ}^d lattice	171
A.2 Condition for the Existence of Riesz Bases	172
A.3 Approximation Error Formula in d -dimensions	173
B Proofs of Some Results of Chapter 3	181
B.1 Sobolev Spaces—Proof of Theorem 3.1	181
B.2 Orthogonality Principle—Proof of Theorem 3.2	183
B.3 Inverse Fourier Transform of d -dimensional Matérn B-spline	184
B.4 Approximation Error for the Noisy Scenario	185
C Poisson-Likelihood-Based Regularized Interpolation	191
D Risk Estimation	195
D.1 Proof of Unbiasedness of SURE	195
D.2 Proof of Unbiasedness of the New URE	196
E Differentiability and Monte-Carlo Divergence Estimation	201
E.1 Verification of Taylor Expansion-Based Hypothesis	201
E.2 Divergence Estimation Under a Weaker Hypothesis	204
Bibliography	210
Curriculum Vitæ	227

Chapter 1

Introduction

1.1 Sampling and Interpolation

Digital computing has radically transformed the fields of signal and image processing over the past few decades and has heralded the era of computed imaging in biomedical applications. This is in tune with the various benefits such as flexibility in terms of programming, high-speed computing, consistency of performance, and effective storage and retrieval mechanisms that digital systems enjoy over their analog counterparts. Since digital systems operate on discrete sequences of samples, their applicability to real-world problems inherently relies on the conversion of continuously-defined signals to discrete samples and vice versa: The forward operation is achieved via *sampling*, while the reverse—reconstruction of a continuously-defined signal from discrete samples—is performed by *interpolation*.

In this context, Shannon’s sampling theorem represents one of the mathematical foundations of digital signal processing, information and communication theories. It states that a continuously-defined signal $f(x)$ that contains no frequencies higher than ω_{\max} cycles per second—that is, a bandlimited signal—is completely determined from its values at time instants spaced $T_{\max} = \pi/\omega_{\max}$ seconds apart [1,2]. The critical sampling period T_{\max} is often referred to as the Nyquist interval [3]. The reconstruction of $f(x)$ is in turn achieved by performing sinc-interpolation of its discrete samples [4–6]—this fits in the framework of interpolation theory which dates back much further and was originally developed for addressing various scientific problems in astronomy [6]. Shannon’s formu-

lation is also closely related to the theory of cardinal series and interpolatory functions [4] due to Whittaker [2, 7–10] and others [11–13].

It is important to keep in mind that Shannon’s work symbolizes an idealization: Real-world signals are seldom bandlimited; in such cases it is suggested to apply an ideal anti-aliasing (low-pass) filter prior to sampling to avoid aliasing artifacts. Again, this ideal low-pass filter is unrealizable in practice. Also, sinc-interpolation is primarily a theoretical construct and is rarely used in practice due to the infinite support and slow decay of the sinc-function. Instead, image-processing practitioners use spline interpolation [14, 15], which is computationally more efficient and provides a good approximation to the sinc-based approach. The reconstruction is then performed by approximating the continuously-defined signal with an interpolating function constructed in such a way as to agree perfectly with the signal at the given sample points. In light of these indicators, it may be argued that interpolation and approximation theories form the essential principles of many developments (including the present work) that propose to circumvent the “idealness” of Shannon’s work.

1.2 Interpolation for Biomedical Imaging

Interpolation is used on a routine basis in image processing and biomedical imaging [6, 15–19]. For instance, it is employed in graphical displays for performing geometric transformation (scaling, rotation, skewing, etc) of discretely-defined objects. In the context of medical imaging, geometric transformations are utilized in registration for aligning multi-modal discrete medical images [20–22]. It is also used for geometrical rectification of discrete images acquired using digital systems with wide-angle lenses and panoramic mirrors. It plays a vital role in achieving sampling grid conversions such as spiral-to-cartesian in k -space for magnetic resonance imaging (MRI) [23]. In volumetric imaging, it is applied for resizing three-dimensional (3-D) volume data [24]. Fitting 3-D data on geometric shapes is also best done by performing interpolation [25]. Other applications in biomedical imaging include volume rendering for visualization of scalar fields [26–28], evaluation of image gradients [29, 30], texture mapping where a 2-D image is painted on a 3-D surface [31, 32], and contour extraction and ray-tracing [33].

While interpolation serves as a computational tool linking the discrete to the continuous, conventional interpolation methods rely on the assumption that the given samples are purely that of the signal-of-interest. However, in practice, rarely do we encounter such a situation: Owing to the physics underlying many

acquisition systems, the response of the acquisition system gets imprinted on the signal during the data-acquisition process. Moreover, due to the imperfect nature of the sensors that record the signal, the measured samples are frequently corrupted by noise. As these nonideal considerations are not part of Shannon's sampling formulation, a direct application of conventional interpolation techniques can be detrimental and lead to loss of performance.

1.2.1 Digital Correction

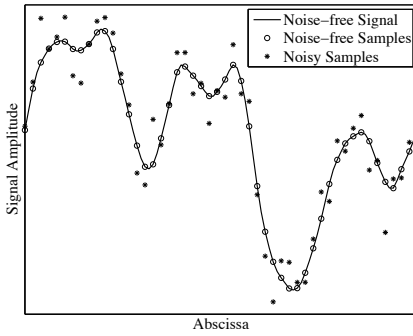
In order to be able to handle nonideal situations, interpolation methods must first compensate for the nonideal nature of the samples. Fortunately, efficient digital processing techniques can be devised to achieve this. For instance, when the samples are noise-free, inverse filtering is typically employed to compensate for the effect of the system response [34]. However, in the presence of noise, this process is usually accompanied by a suitable noise-filtering strategy that permits a stabilized inversion [35].

Probably, the most popular approach for filtering noise in this context is the use of *regularization* [35]: The underlying idea is to enforce certain *smoothness* constraints on the continuously-defined output to counterbalance the fluctuations due to noise. A common practice in spline and regularization [35–39] literature is to use derivatives for characterizing the smoothness (or equivalently, roughness) of a continuously-defined signal. The regularization constraints are then typically prescribed in terms of some norm of a derivative of the output and the problem is solved by minimizing this norm subject to the requirement that the output also fits the given samples closely. Even though the problem is formulated in terms of analog entities, the corresponding solution can be efficiently implemented by means of digital filtering (linear or non-linear) techniques as we shall see in Chapters 3 and 4.

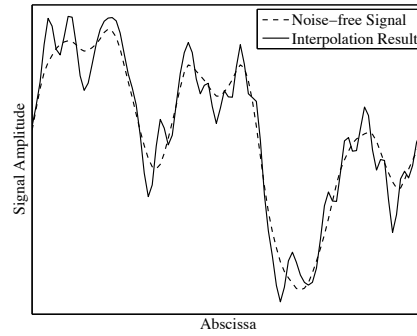
1.2.2 Optimal Filtering

The use of regularization entails a constrained optimization problem where in a *regularization parameter* governs the trade-off between the strength of the applied regularization constraints and the goodness-of-fit of output to the data samples. Depending on the value of this parameter, the output can range anywhere between being high fluctuating or heavily constrained.

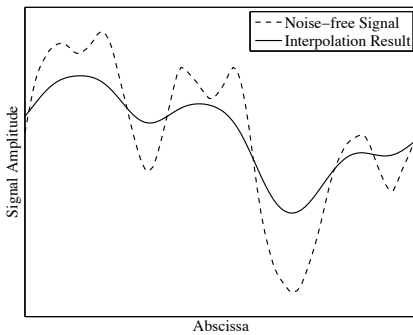
Figure 1.1 illustrates the effect of this parameter on the performance of interpolation in the case of a simple 1-D example. We are given noisy samples marked by “*” in Figure 1.1(a) on which we perform cubic-spline interpolation



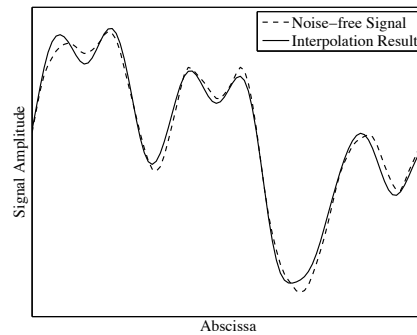
(a) Original signal and noisy samples



(b) No filtering of noise



(c) Heavy filtering of noise



(d) Optimal filtering of noise

Figure 1.1: Illustration of the effect of noise-filtering on interpolation: (a) Arbitrary signal and noisy samples; (b) Exact data fitting results in a highly fluctuating output; (c) Interpolation with heavy filtering results in loss of details (minor peaks have vanished in the output); (d) Interpolation with optimal filtering preserves the primary signal features while being robust to fluctuations due to noise.

with the constraint that the squared-norm of the second-derivative of the spline is minimum. In Figure 1.1(b), the regularization parameter is set to zero—this is equivalent to exact-fitting of the noisy sample points. Correspondingly, the output fluctuates excessively. Figure 1.1(c) illustrates the case where the regularization parameter is set to a very high value. Here, the fluctuations due to noise have significantly reduced, but while doing so, the output fails to capture the finer details of the original signal (such as the minor peaks) as it has been over-constrained by the heavy-handed application of regularization.

Figure 1.1(d) depicts the setting where we have set the regularization parameter “optimally” in the sense that we have achieved significant noise reduction while also simultaneously capturing the subtle attributes of the original signal as much as possible. This example illustrates the crucial role of the regularization parameter in our constrained interpolation problem and also highlights the necessity for finding the right amount of regularization needed for a given noisy scenario.

1.3 Contributions of the Thesis

The main focus of this thesis is on the development of regularization methods for the following two problems that we consider in this work: (1) Nonideal sampling problem, and (2) Robust interpolation of noisy data in multiple dimensions.

1.3.1 Theoretical Investigation and Numerical Implementation

We consider a very general variational¹ criterion (based on quadratic regularization) for the nonideal sampling problem that can be handled analytically. The major results are that the mathematical optimization of the criterion implicitly leads to a shift-invariant² reconstruction which is included in some spline-like space (which need not necessarily be bandlimited) that is matched to the particular sampling setup and/or class of input signals. Moreover, we establish a formal equivalence between the variational formulation and the *minimum mean-squared error* (MMSE) or the Wiener solution of the nonideal sampling problem

¹The term “variational” derives from the field of calculus of variations and is also generally associated with the class of problems involving regularization criteria.

²Here, shift invariance implies that the solutions reconstructed from the given data and that shifted by an arbitrary amount are equivalent up to the arbitrary shift. It must also be noted that shift invariance is an important trait of many sampling and interpolation theories including that of Shannon.

thereby demonstrating that we can perform optimal reconstruction of stationary stochastic signals from their noisy samples.

We also consider the use of non-quadratic regularization criteria which are known to be better suited for signals with sharp features, especially images which contain edges. Here, we address a more specific problem; that is, interpolation in the presence of noise. Starting from shift-, rotation-, and scale-invariant³ requirements on the continuous-domain solution, we show that the L_p -norm⁴ of an appropriate vector derivative is the most suitable choice of regularization. This includes the famous total-variation (TV) regularization (for $p = 1$) which is well-known for its edge-preserving characteristics in image-processing applications. We give efficient numerical algorithms (that converge monotonically) to carry out the optimization of the non-quadratic (for $p > 1$) criterion.

1.3.2 Selection of the Best Regularization Parameter

In both sampling and interpolation problems, we encounter the main issue of the selection of the best regularization parameter. For this, we propose a novel data-driven scheme to compute the *mean-squared error* (MSE) of an arbitrary interpolation algorithm. Then, we demonstrate its applicability for the problems considered in this thesis and show how to select the regularization parameter in an optimal way so as to minimize the MSE of the reconstruction. Our data-driven technique is also applicable to other inverse problems such as denoising and deconvolution where we show how to tune various algorithm parameters to achieve best performance in terms of the MSE.

1.4 Thesis Outline

As this thesis is intended towards the enhancement of conventional sampling and interpolation theories, we will stick to the case of uniformly-spaced samples throughout this work.

In Chapter 2, we will review the primary developments that have occurred in sampling theory since Shannon's work. These advances predominantly view sampling as a projection operation on to a shift-invariant space. The key point

³Rotation and scale invariance are also defined in a fashion similar to shift invariance except that instead of shifts, we now consider rotation by arbitrary angles and dilation by arbitrary factors of the solution, respectively.

⁴We will specify the mathematical details of what we mean by L_p -norm in Chapter 2.

is to relax the perfect reconstruction constraint and aim for a close approximation of the original signal, which allows for the treatment of arbitrary input signals (which need not be bandlimited). We will also review some results in approximation theory that provide tools for the selection of appropriate sampling procedures.

While projection-based approaches are mathematically elegant, they are nevertheless restricted to noise-free sampling conditions. This forms the motivation for us to consider the nonideal sampling problem in the presence of noise which is the topic of Chapter 3. Here, we pose the problem in a variational framework where we minimize a data-fitting term subject to a Tikhonov-like L_2 -regularization (quadratic, continuous-domain regularization) to obtain the continuous-space solution. Apart from the fact that we solve this problem analytically, we also prove the existence of a best reconstruction space and show that it is characterized by a condition that links the generating function of the reconstruction space to the regularization (differential) operator. We justify that this in turn implies the existence of a B-spline-like basis. To make the scheme practical, we specify the generating functions corresponding to the most popular families of regularization operators (derivatives, iterated Laplacian), as well as a new, generalized one that leads to a new brand of Matérn⁵ splines. We demonstrate that our solution is general in the sense that we recover some of the classical results when the sampling is ideal. We also establish equivalences with the minimax and minimum mean square error (MMSE / Wiener) solutions of the nonideal sampling problem.

In Chapter 4, we consider the scenario where the reconstruction space is fixed before hand, but investigate a much wider class of non-quadratic regularization (L_p -norm with $p \neq 2$) for the problem of interpolation in the presence of noise. We provide numerical algorithms based on the *majorize-minimize* (MM) approach for performing the corresponding minimization of the non-quadratic criterion. We also present experimental results that demonstrate the superior performance of our method—which we call *regularized interpolation*—over conventional interpolation methods when applied to noisy images.

In Chapters 5 and 6, we focus on the issue of the choice of the regularization parameter based on the MSE of the reconstruction. Since MSE depends on the original signal (which is unknown *a priori*) it is not directly computable in practice. So, we first review a concept (in Chapter 5) in the statistics literature that goes by the name *Stein's unbiased risk estimate*—SURE—which permits

⁵This term is associated with Bertil Matérn's stochastic model that is well known in kriging theory for interpolation of spatial data [40].

us to unbiasedly estimate the MSE purely from data when the noise is additive Gaussian. We also derive a new risk estimate for a data model that considers a mixture of Poisson and Gaussian noise that is appropriate for bioimaging applications. Then, we show how to compute the risk estimate for linear filtering algorithms (that is, those based on quadratic regularization) with specific application to image deconvolution in fluorescence microscopy.

In Chapter 6, we take up the problem of risk estimation for arbitrary non-linear algorithms. We present a novel and efficient numerical method based on Monte-Carlo simulation for estimating the MSE of arbitrary interpolation / denoising algorithms. We also present experimental results that demonstrate the effectiveness of our Monte-Carlo technique on a variety of denoising and interpolation algorithms. Finally, we draw our conclusions in Chapter 7, where we also give some insight in to possible extensions of this thesis.

Chapter 2

Noise-Free Sampling and Linear Approximation Theory

2.1 Introduction

Shannon's sampling theory provides an elegant method to perfectly reconstruct bandlimited signals from their equidistant samples [1, 2, 4, 5]. Although Shannon introduced his work in an information-theoretic background, sampling has its origin¹ in interpolation theory [6–10] which in itself is a subset of the more broad field of approximation theory [42–47]. In the post-Shannon-era, approximation theory played a key role in the evolution of many extensions of Shannon's result [5]. The basic idea behind the approximation-theoretic approach is to consider sampling and reconstruction as the process of approximating a (not necessarily bandlimited) finite-energy signal in some suitable (shift-invariant) Hilbert-subspace. The minimum-error-approximation (in the sense of squared-norm) is obtained by projecting the input signal orthogonally onto the approximation space [44, 45]. This view point has two main advantages over Shannon's framework: *(i)* It allows for the treatment of non-bandlimited signals, and *(ii)* It circumvents the need for the slowly-decaying sinc-function for reconstruction.

¹A historical overview of sampling theory (including some works prior to Shannon) can be found in [4, 5, 41].

In fact, reconstruction in spline-like spaces provides a computationally efficient approximation to sinc-interpolation [14, 48, 49]. The connection between splines and sampling was formally established by Unser *et al.* in [50] and Aldroubi *et al.* in [51] where they showed that the minimum-error-approximation in spline-like spaces is asymptotic equivalent to Shannon's sampling framework. They also demonstrated that the minimum-error-approximation is equivalent to a three-step processing of the input signal: Optimal prefiltering (which corresponds to filtering by an anti-aliasing filter in the Shannon-framework), sampling, and reconstructing an approximating signal inside the approximation space [50, 51]. In the case where the prefilter is specified *a priori*, a consistent sampling strategy was proposed in [52]; this again has the interpretation of an oblique projection onto the approximation space.

In this chapter, we review some state-of-the-art sampling theories for the noise-free scenario using a unifying projection formalism and extend their scope to arbitrary sampling lattices in d -dimensions. A crucial notion in approximation theory is the error of approximation. For this, we derive an explicit expression for the approximation error in d -dimensions. Our result also supplements standard results due to de Boor *et al.* [44] and Blu *et al.* [53, 54] in approximation theory.

In the sequel, we first present in detail the projection-interpretation of Shannon's sampling theorem (1-D). Next, we discuss signal representation in a general shift-invariant space in d -dimensions and specify a generic approximation operator that characterizes our (generalized) sampling scheme in the noise-free case. Then, we study various sampling strategies (such as minimum-error sampling, consistent sampling, etc) which can all be described mathematically using the same approximation operator, while making appropriate links to related results in the literature. Finally, we conclude the chapter with a discussion of the approximation error.

2.1.1 Shannon's Theorem Reinterpreted

The reconstruction formula accompanying Shannon's sampling theorem (1-D) is [1, 2, 4, 5]

$$f(x) = \sum_{k \in \mathbb{Z}} f(kT) \operatorname{sinc} \left(\frac{x}{T} - k \right), \quad (2.1)$$

where $T > 0$ is the sampling-step, and $f(x)$ is a signal bandlimited to $\omega_{\max} = \frac{\pi}{T}$. The reconstruction is achieved by computing a sum of scaled and shifted versions




Figure 2.1: Shannon’s sampling scheme: The non-bandlimited analog signal $f(x)$ is prefiltered with the ideal low-pass filter $\frac{1}{T}\text{sinc}\left(\frac{x}{T}\right)$ to avoid aliasing. Sampling is modeled as the multiplication of the prefiltered signal with a train of Dirac impulses. The reconstruction is achieved by filtering the resulting samples $\{g[k]\}$ with $\text{sinc}\left(\frac{x}{T}\right)$.

of the sinc-function, $\text{sinc}(x) = \sin(\pi x)/(\pi x)$, weighted by the samples $\{f(kT)\}$.

When the input signal $f(x) \in L_2(\mathbb{R})$ is not bandlimited, then it is first filtered by an ideal low-pass filter with the frequency response $\hat{h}_{\text{ideal}}(\omega T)$ prior to sampling, where²

$$\hat{h}_{\text{ideal}}(\omega) = \text{rect}\left(\frac{\omega}{2\pi}\right).$$

The corresponding impulse response is $\frac{1}{T}h_{\text{ideal}}\left(\frac{x}{T}\right)$ where $h_{\text{ideal}}(x) = \text{sinc}(x)$. The standard Shannon’s sampling model is illustrated in Figure 2.1. The corresponding reconstruction formula is

$$f_{\text{rec}}(x) = \sum_{k \in \mathbb{Z}} \underbrace{\left(f \star \frac{1}{T} h_{\text{ideal}}\left(\frac{\bullet}{T}\right) \right)}_{g[k]}(kT) \text{sinc}\left(\frac{x}{T} - k\right), \quad (2.2)$$

where “ \star ” denotes the continuous-space convolution operation and $f_{\text{rec}}(x)$ is the bandlimited approximation to $f(x)$. Equation (2.2) forms the key link to the modern projection-based framework for the sampling problem [5, 44, 52–54].

²The rect-function is defined as $\text{rect}(x) = \begin{cases} 1, & |x| \leq \frac{1}{2}, \\ 0, & \text{otherwise.} \end{cases}$

Here, Shannon's sampling theorem can be posed as the following optimization problem

$$f_{\text{rec}} = \arg \min_{s \in V_T(\text{sinc})} \|f - s\|_{L_2}^2, \quad (2.3)$$

where $V_T(\text{sinc})$ is the space of bandlimited functions defined as

$$V_T(\text{sinc}) = \left\{ s(x) = \sum_{k \in \mathbb{Z}} c[k] \text{sinc} \left(\frac{x}{T} - k \right) : \{c[k]\} \in \ell_2(\mathbb{Z}^d) \right\}, \quad (2.4)$$

and $f_{\text{rec}}(x)$ represents the minimum L_2 -error approximation of $f(x)$ in $V_T(\text{sinc})$ whenever $f(x)$ is not bandlimited. The minimization in (2.3) is a standard problem in functional analysis [55] whose solution can be found to be

$$\begin{aligned} f_{\text{rec}}(x) &= P_{V_T(\text{sinc})}\{f\}(x) \\ &= \sum_{k \in \mathbb{Z}} \underbrace{\left\langle f, \frac{1}{T} h_{\text{ideal}} \left(\frac{\bullet}{T} - k \right) \right\rangle_{L_2}}_{c[k] \text{ in (2.4)}} \text{sinc} \left(\frac{x}{T} - k \right), \end{aligned} \quad (2.5)$$

where $P_{V_T(\text{sinc})} : L_2(\mathbb{R}^d) \rightarrow V_T(\text{sinc})$ is the orthogonal projection operator which satisfies the property $P_{V_T(\text{sinc})}\{P_{V_T(\text{sinc})}\{f\}\}(x) = P_{V_T(\text{sinc})}\{f\}(x)$. Since h_{ideal} is symmetric, we have that

$$\left\langle f, \frac{1}{T} h_{\text{ideal}} \left(\frac{\bullet}{T} - k \right) \right\rangle_{L_2} = \left(f \star \frac{1}{T} h_{\text{ideal}} \left(\frac{\bullet}{T} \right) \right) (kT),$$

and so the projection formulation (*cf.* Equation (2.5)) is equivalent to Shannon's model (*cf.* Equation (2.2)).

2.2 Preliminaries

Modern approaches to the sampling problem rely on the projection formalism but provide an extra degree of freedom in that they allow for reconstruction in a general (shift-invariant) space that can be specified by the user (not necessarily $V_T(\text{sinc})$). The important point is to relinquish the perfect reconstruction constraint and approach the problem from an approximation-theoretic perspective which offers many advantages (see Section 2.1 at the beginning of this chapter) over the classical framework of Shannon. Our aim, in the rest of the chapter, is

to present an overview of some well-known results in this context. In what follows, we present our notation and then discuss signal representation in a general shift-invariant space in d -dimensions.

2.2.1 Notations

Henceforth, we deal with scalar signals in d -dimensions. Continuous-space signals are denoted with parentheses:

$$f(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}, \quad \mathbf{x} = (x_1, x_2, \dots, x_d)^T \in \mathbb{R}^d,$$

and discrete-space signals (or sequences) with brackets:

$$c[\mathbf{k}] : \mathbb{Z}^d \rightarrow \mathbb{R}, \quad \mathbf{k} = (k_1, k_2, \dots, k_d)^T \in \mathbb{Z}^d,$$

where $(\bullet)^T$ represents the transpose. We write flipped functions or sequences with a bar:

$$\bar{f}(\mathbf{x}) = f(-\mathbf{x}), \quad \bar{c}[\mathbf{k}] = c[-\mathbf{k}],$$

and denote by $\{c[\mathbf{k}]\}$ the sequence c in its entirety (that is, $\forall \mathbf{k} \in \mathbb{Z}^d$). For $1 \leq p < \infty$, we define the function-spaces

$$L_p(\mathbb{R}^d) = \left\{ s(\mathbf{x}) : \|s\|_{L_p} = \left| \int_{\mathbb{R}^d} |s(\mathbf{x})|^p d\mathbf{x} \right|^{\frac{1}{p}} < +\infty \right\},$$

$$\ell_p(\mathbb{Z}^d) = \left\{ \{c[\mathbf{k}]\} : \|c\|_{\ell_p} = \left| \sum_{\mathbf{k} \in \mathbb{Z}^d} |c[\mathbf{k}]|^p \right|^{\frac{1}{p}} < +\infty \right\}.$$

We also denote $\|s\|_{L_\infty} = \text{ess sup}_{\mathbf{x} \in \mathbb{R}^d} |s(\mathbf{x})| = \inf\{C \geq 0 : |s(\mathbf{x})| \leq C \text{ a.e.}\}$ (essential supremum) and $\|c\|_{\ell_\infty} = \sup_{\mathbf{k} \in \mathbb{Z}^d} |c[\mathbf{k}]|$. Correspondingly, we obtain the spaces

$$L_\infty(\mathbb{R}^d) = \{s(\mathbf{x}) : \|s\|_{L_\infty} < +\infty\}, \quad \text{and} \quad \ell_\infty(\mathbb{Z}^d) = \{\{c[\mathbf{k}]\} : \|c\|_{\ell_\infty} < +\infty\}.$$

The continuous-space Fourier transform of $f(\mathbf{x}) \in L_1(\mathbb{R}^d) \cap L_2(\mathbb{R}^d)$ is

$$\mathcal{F}\{f\}(\boldsymbol{\omega}) = \hat{f}(\boldsymbol{\omega}) = \int_{\mathbb{R}^d} f(\mathbf{x}) e^{-j\boldsymbol{\omega}^T \mathbf{x}} d\mathbf{x},$$

and the discrete-space Fourier transform of $\{c[\mathbf{k}]\} \in \ell_1(\mathbb{Z}^d) \cap \ell_2(\mathbb{Z}^d)$ is

$$C(e^{j\boldsymbol{\omega}}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] e^{-j\boldsymbol{\omega}^T \mathbf{x}},$$

where $\boldsymbol{\omega}^T \mathbf{x} = \sum_{i=1}^d \omega_i x_i$ and $e^{j\boldsymbol{\omega}} = (e^{j\omega_1 x_1}, e^{j\omega_2 x_2}, \dots, e^{j\omega_d x_d})$. For $\mathbf{x} \in \mathbb{R}^d$, we

also define the Euclidean norm $\|\mathbf{x}\|^2 = \sum_{i=1}^d x_i^2$.

We denote the L_2 -inner product between two complex functions $f(\mathbf{x})$ and $s(\mathbf{x})$ by

$$\langle f, s \rangle_{L_2} = \int_{\mathbb{R}^d} f^*(\mathbf{x}) s(\mathbf{x}) \, d\mathbf{x},$$

and the ℓ_2 -inner product between two complex sequences $c[\mathbf{k}]$ and $b[\mathbf{k}]$ by

$$\langle c, b \rangle_{\ell_2} = \sum_{\mathbf{k} \in \mathbb{Z}^d} c^*[\mathbf{k}] b[\mathbf{k}],$$

where $(\bullet)^*$ indicates the complex conjugate. Corresponding to the above inner products, we have the norms

$$\|f\|_{L_2} = \sqrt{\langle f, f \rangle_{L_2}} \quad \text{and} \quad \|c\|_{\ell_2} = \sqrt{\langle c, c \rangle_{\ell_2}}.$$

Continuous-space convolution of $f(\mathbf{x})$ and $s(\mathbf{x})$ is denoted using the “ \star ” symbol

$$(f \star s)(\mathbf{x}) = \int_{\mathbb{R}^d} f(\mathbf{u}) s(\mathbf{x} - \mathbf{u}) \, d\mathbf{u} = \langle f, s(\mathbf{x} - \bullet) \rangle_{L_2},$$

and its discrete counterpart is denoted using the “ $*$ ” symbol

$$(c * b)[\mathbf{k}] = \sum_{\mathbf{m} \in \mathbb{Z}^d} c[\mathbf{m}] b[\mathbf{k} - \mathbf{m}] = \langle c, b[\mathbf{k} - \bullet] \rangle_{\ell_2},$$

respectively.

2.2.1.1 Sobolev Spaces

The approximation theoretic results that we review in this chapter are specified with respect to the Sobolev space W_2^ρ of (possibly non-integer) order $\rho > \frac{d}{2}$

defined as

$$W_2^\rho = \left\{ s(x) : \|s\|_{W_2^\rho}^2 \stackrel{\text{def}}{=} \int_{\mathbb{R}^d} (1 + \|\boldsymbol{\omega}\|^2)^\rho |\hat{s}(\boldsymbol{\omega})|^2 d\boldsymbol{\omega} < +\infty \right\},$$

which contains finite energy functions whose $\frac{\gamma}{2}$ -iterated d -dimensional Laplacians,³ $0 \leq \gamma \leq \rho$, also have finite L_2 -norms.

2.2.1.2 Useful Identities

We will also frequently rely on the following mathematical identities in our analysis:

Parseval's identity: For $f, s \in L_2(\mathbb{R}^d)$,

$$\langle f, s \rangle_{L_2} = \frac{1}{(2\pi)^d} \langle \hat{f}^*, \hat{s} \rangle_{L_2}.$$

Similarly, for $\{b[\mathbf{k}]\}, \{c[\mathbf{k}]\} \in \ell_2(\mathbb{Z}^d)$,

$$\langle b, c \rangle_{\ell_2} = \frac{1}{(2\pi)^d} \int_{[0, 2\pi]^d} B^*(e^{j\boldsymbol{\omega}}) C(e^{j\boldsymbol{\omega}}) d\boldsymbol{\omega} = \langle B^*(e^{j\boldsymbol{\omega}}), C(e^{j\boldsymbol{\omega}}) \rangle_{L_2([0, 2\pi]^d)},$$

where $[0, 2\pi]^d = \{\boldsymbol{\omega} = (\omega_1, \omega_2, \dots, \omega_d) \mid \omega_i \in [0, 2\pi), i = 1, 2, \dots, d\}$.

Young's inequality:⁴ For any $f \in L_1(\mathbb{R}^d)$ and $g \in L_p(\mathbb{R}^d)$ [56],

$$\|f \star g\|_{L_p} \leq \|f\|_{L_1} \|g\|_{L_p}.$$

The same holds true for sequences, with $L_p(\mathbb{R}^d)$ being replaced by $\ell_p(\mathbb{Z}^d)$.

Inclusion-property of ℓ_p -spaces: For $1 \leq p_0 < p_1 \leq \infty$,

$$\ell_{p_0}(\mathbb{Z}^d) \subset \ell_{p_1}(\mathbb{Z}^d).$$

2.2.2 Shift-Invariant Spaces and Riesz Bases

Here, we first present details about a general sampling lattice and then provide conditions necessary for stable representation of functions in shift-invariant spaces.

³The Fourier transform of a $\frac{\gamma}{2}$ -iterated Laplacian of $s(\mathbf{x})$ is given by $\|\boldsymbol{\omega}\|^\gamma \hat{s}(\boldsymbol{\omega})$ (in the sense of distributions).

⁴A more general form of Young's inequality is written as $\|f \star g\|_{L_r} \leq \|f\|_{L_p} \|g\|_{L_q}$ for $f \in L_p, g \in L_q$, and $\frac{1}{r} = \frac{1}{p} + \frac{1}{q} - 1$.

2.2.2.1 Sampling Lattices

We consider sampling points chosen from the set $\mathbf{D}\mathbb{Z}^d$ (sampling lattice) specified as⁵

$$\mathbf{D}\mathbb{Z}^d = \{\mathbf{z}_1 : \mathbf{z}_1 = \mathbf{D}\mathbf{k}, \mathbf{k} \in \mathbb{Z}^d\}, \quad (2.6)$$

where \mathbf{D} is a $d \times d$ invertible matrix; $\mathbf{D}\mathbf{k}$ represents the matrix product between \mathbf{D} and the $d \times 1$ column vector \mathbf{k} , and $|\mathbf{D}|$ denotes the absolute value of the determinant of \mathbf{D} . For $\mathbf{D} = T\mathbf{I}$, we obtain the standard sampling lattice $T\mathbb{Z}^d$ that corresponds to a sampling-step $T > 0$ (along every dimension).

The dual lattice, which plays an important role in the Fourier analysis, is written as $2\pi\mathbf{D}^{-T}\mathbb{Z}^d$ and is given by the set

$$2\pi\mathbf{D}^{-T}\mathbb{Z}^d = \{\mathbf{z}_2 : \mathbf{z}_2 = 2\pi\mathbf{D}^{-T}\mathbf{k}, \mathbf{k} \in \mathbb{Z}^d\}, \quad (2.7)$$

where $\mathbf{D}^{-T} = (\mathbf{D}^{-1})^T$. Then, $\forall \mathbf{z}_1 \in \mathbf{D}\mathbb{Z}^d$ and $\mathbf{z}_2 \in 2\pi\mathbf{D}^{-T}\mathbb{Z}^d$, we have $e^{j\mathbf{z}_2^T\mathbf{z}_1} = 1$.

We also require the Fourier transform of $f(\mathbf{D}\mathbf{x})$ which is given by

$$\mathcal{F}\{f(\mathbf{D}\bullet)\}(\boldsymbol{\omega}) = \frac{1}{|\mathbf{D}|} \hat{f}(\mathbf{D}^{-T}\boldsymbol{\omega}), \quad (2.8)$$

and the Poisson's summation formula⁶ which, in the current context, is written as

$$\sum_{\mathbf{k} \in \mathbb{Z}^d} f(\mathbf{D}\mathbf{k}) e^{j(\mathbf{D}^T\boldsymbol{\omega})^T\mathbf{k}} = \frac{1}{|\mathbf{D}|} \sum_{\mathbf{m} \in \mathbb{Z}^d} \hat{f}(\boldsymbol{\omega} - 2\pi\mathbf{D}^{-T}\mathbf{m}) \quad (2.9)$$

whenever the summations in the l.h.s. and the r.h.s. in the above equation converge.

2.2.2.2 Specification of the Reconstruction Space

The reconstruction is performed in the function-space $V_{\mathbf{D}}(\varphi)$ specified as

$$V_{\mathbf{D}}(\varphi) = \left\{ s(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] \varphi(\mathbf{D}^{-1}\mathbf{x} - \mathbf{k}) : \{c[\mathbf{k}]\} \in \ell_2(\mathbb{Z}^d) \right\}, \quad (2.10)$$

⁵Detailed description of sampling lattices can be found in [57, 58].

⁶Details of the derivation of Poisson's formula for the $\mathbf{D}\mathbb{Z}^d$ sampling lattice are given in Appendix A.1.

where $\varphi(\mathbf{x})$ is called the generating function which has acceptable time and frequency localization. In using $V_{\mathbf{D}}(\varphi)$, we retain the shift-invariant flavor of Shannon's model,⁷ but now, invariance is possible only for shifts on the lattice $\mathbf{D}\mathbb{Z}^d$; that is, if $s(\mathbf{x}) \in V_{\mathbf{D}}(\varphi)$, then for any $\mathbf{m} \in \mathbb{Z}^d$, $s(\mathbf{x} - \mathbf{D}\mathbf{m}) \in V_{\mathbf{D}}(\varphi)$ as seen below:

$$\begin{aligned} s(\mathbf{x} - \mathbf{D}\mathbf{m}) &= \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] \varphi(\mathbf{D}^{-1}(\mathbf{x} - \mathbf{D}\mathbf{m}) - \mathbf{k}) \\ &= \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] \varphi(\mathbf{D}^{-1}\mathbf{x} - \mathbf{m} - \mathbf{k}) \\ &= \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k} - \mathbf{m}] \varphi(\mathbf{D}^{-1}\mathbf{x} - \mathbf{k}) \in V_{\mathbf{D}}(\varphi). \end{aligned} \quad (2.11)$$

It must be noted that the coefficients $\{c[\mathbf{k}]\}$ are not necessarily the samples of the signal and that $\varphi(\mathbf{x})$ can be different from $\text{sinc}(\mathbf{x})$. In fact, we are interested in the scenario where $\varphi(\mathbf{x})$ has a finite support, e.g., polynomial B-splines [14], which helps reduce computational load significantly during reconstruction.

2.2.2.3 Riesz Bases

In order for $s(\mathbf{x}) \in V_{\mathbf{D}}(\varphi)$ to have a unique and stable representation in terms of $\{c[\mathbf{k}]\}$, the set $\{\varphi(\mathbf{D}^{-1}\mathbf{x} - \mathbf{k})\}_{\mathbf{k} \in \mathbb{Z}^d}$ must form a Riesz basis of $V_{\mathbf{D}}(\varphi)$; that is, $\forall \{c[\mathbf{k}]\} \in \ell_2(\mathbb{Z}^d)$, there must exist constants $A_{R_{\varphi},2} > 0$ and $B_{R_{\varphi},2} < +\infty$ such that⁸

$$A_{R_{\varphi},2}^2 \|c\|_{\ell_2}^2 \leq \left\| \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] \varphi(\mathbf{D}^{-1}\mathbf{x} - \mathbf{k}) \right\|_{L_2}^2 \leq B_{R_{\varphi},2}^2 \|c\|_{\ell_2}^2. \quad (2.12)$$

Proposition 2.1. *The Riesz condition (2.12) is equivalent to having*

$$A_{R_{\varphi},2}^2 \leq A_{\varphi}(\boldsymbol{\omega}) \leq B_{R_{\varphi},2}^2 \text{ a.e.}, \quad (2.13)$$

where

$$A_{\varphi}(\boldsymbol{\omega}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2.$$

⁷Due to the interpolating nature of the sinc-function, Shannon's model satisfies the shift-invariant property for non-integer shifts also: If $s(\mathbf{x}) \in V_{\mathbf{D}}(\text{sinc})$, then $s(\mathbf{x} - \mathbf{D}\boldsymbol{\tau}) \in V_{\mathbf{D}}(\text{sinc})$ for $\boldsymbol{\tau} \in \mathbb{R}^d$.

⁸The basis is orthogonal if and only if $A_{R_{\varphi},2} = B_{R_{\varphi},2} = 1$.

While this is now a standard result in approximation and wavelet theory, it is hard to trace back its origin in the literature. Our proof of this result is based on the proof given in [51, Theorem 2] for the case of $d = 1$.

Proof. We note that

$$\mathcal{F} \left\{ \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] \varphi(\mathbf{D}^{-1} \bullet - \mathbf{k}) \right\} (\boldsymbol{\omega}) = C(e^{j\mathbf{D}^T \boldsymbol{\omega}}) |\mathbf{D}| \hat{\varphi}(\mathbf{D}^T \boldsymbol{\omega}). \quad (2.14)$$

Then using Parseval's identity, we obtain

$$\begin{aligned} \left\| \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] \varphi(\mathbf{D}^{-1} \mathbf{x} - \mathbf{k}) \right\|_{L_2}^2 &= \frac{|\mathbf{D}|^2}{(2\pi)^d} \int_{\mathbb{R}^d} |C(e^{j\mathbf{D}^T \boldsymbol{\omega}})|^2 |\hat{\varphi}(\mathbf{D}^T \boldsymbol{\omega})|^2 d\boldsymbol{\omega} \\ &= \frac{|\mathbf{D}|}{(2\pi)^d} \int_{\mathbb{R}^d} |C(e^{j\boldsymbol{\omega}})|^2 |\hat{\varphi}(\boldsymbol{\omega})|^2 d\boldsymbol{\omega} \\ &= \frac{|\mathbf{D}|}{(2\pi)^d} \sum_{\mathbf{k} \in \mathbb{Z}^d} \int_{[0, 2\pi)^d} |C(e^{j\boldsymbol{\omega}})|^2 |\hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 d\boldsymbol{\omega}, \\ &= \frac{|\mathbf{D}|}{(2\pi)^d} \int_{[0, 2\pi)^d} |C(e^{j\boldsymbol{\omega}})|^2 A_\varphi(\boldsymbol{\omega}) d\boldsymbol{\omega}, \end{aligned} \quad (2.15)$$

where we have interchanged the summation and integral (since the integrand is positive) while using the fact that $C(e^{j\boldsymbol{\omega}})$ is (2π) -periodic along each dimension. Then, since $|\mathbf{D}| > 0$ and because

$$\|c\|_{\ell_2}^2 = \frac{1}{(2\pi)^d} \int_{[0, 2\pi)^d} |C(e^{j\boldsymbol{\omega}})|^2 d\boldsymbol{\omega},$$

the result directly follows by noting that for $\left\| \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] \varphi(\mathbf{D}^{-1} \mathbf{x} - \mathbf{k}) \right\|_{L_2}^2$ to be

bounded below and above by $A_{R_\varphi, 2}^2 \|c\|_{\ell_2}^2$ and $B_{R_\varphi}^2 \|c\|_{\ell_2}^2$, respectively, $A_\varphi(\boldsymbol{\omega})$ must satisfy (2.13). The technicalities associated with this last step are provided in Appendix A.2. \blacksquare

Remark 2.1. *It must be noted that the (2π) -periodic (along each dimension) $A_\varphi(\boldsymbol{\omega})$ is nothing but the Fourier transform of the samples of the autocorrelation of $\varphi(\mathbf{x})$: $\{(\varphi \star \bar{\varphi})(\mathbf{k})\}$. Moreover, if $\{(\varphi \star \bar{\varphi})(\mathbf{k})\} \in \ell_1(\mathbb{Z}^d)$, then $A_\varphi(\boldsymbol{\omega})$ is uniformly continuous and (2.13) holds $\forall \boldsymbol{\omega} \in [0, 2\pi)^d$.*

The existence of a non-zero $A_{R_\varphi,2}$ ensures that the basis functions are linearly independent,⁹ while a finite $B_{R_\varphi,2}$ guarantees that any $s(\mathbf{x}) \in V_{\mathbf{D}}(\varphi)$ has a finite L_2 -norm which in turn implies that $V_{\mathbf{D}}(\varphi) \subset L_2(\mathbb{R}^d)$. Note that $\varphi = \text{sinc}$ is a special case: $\text{sinc}(\mathbf{x})$ automatically satisfies the Riesz condition (2.12) since $A_{\text{sinc}}(\boldsymbol{\omega}) = 1, \forall \boldsymbol{\omega} \in [0, 2\pi)^d$.

2.2.2.4 Equivalent Bases

It is important to note that the generator φ of $V_{\mathbf{D}}(\varphi)$ is not unique. In fact, any function $\varphi_{\text{eq}}(\mathbf{x})$ of the form

$$\varphi_{\text{eq}}(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} q[\mathbf{k}] \varphi(\mathbf{x} - \mathbf{k}),$$

generates a L_2 -stable Riesz basis of $V_{\mathbf{D}}(\varphi)$ —equivalent basis—provided the sequence $\{q[\mathbf{k}]\}$ fulfills certain conditions specified in the following proposition (this is the d -dimensional version of Proposition 6 in [51]).

Proposition 2.2. *When φ satisfies (2.13), the function $\varphi_{\text{eq}}(\mathbf{x})$ also generates a L_2 -stable Riesz basis of $V_{\mathbf{D}}(\varphi)$ if and only if the Fourier transform of $\{q[\mathbf{k}]\}$, $Q(e^{j\boldsymbol{\omega}})$, satisfies*

$$0 < m_q \leq |Q(e^{j\boldsymbol{\omega}})| \leq M_q < +\infty \quad a.e. \quad (2.16)$$

Proof. From the definition of φ_{eq} , we have $A_{\varphi_{\text{eq}}}(\boldsymbol{\omega}) = |Q(e^{j\boldsymbol{\omega}})|^2 A_\varphi(\boldsymbol{\omega})$. The result then follows from the fact that $A_{\varphi_{\text{eq}}}(\boldsymbol{\omega})$ fulfills (2.13) if and only if $Q(e^{j\boldsymbol{\omega}})$ satisfies (2.16). ■

2.2.2.5 Separability

In the multidimensional setting, separable basis functions play an important role since many reconstruction algorithms (e.g., spline based interpolation, etc) rely on separability to reduce the computational load. A separable generator is given by

$$\varphi_{\text{sep}}(\mathbf{x}) = \prod_{i=1}^d \varphi(x_i).$$

⁹That means to say, $s(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] \varphi(\mathbf{D}^{-1}\mathbf{x} - \mathbf{k}) = 0$, if and only if $c[\mathbf{k}] = 0, \forall \mathbf{k} \in \mathbb{Z}^d$.




Figure 2.2: Generalized sampling setup: An arbitrary finite-energy analog signal $f(\mathbf{x})$ is first prefiltered with $\frac{1}{|\mathbf{D}|}h(\mathbf{D}^{-1}\mathbf{x})$ (also called the analysis function). The obtained samples $\tilde{c}[\mathbf{k}]$ are then processed by a digital correction filter q , which provides the coefficients $\{c[\mathbf{k}]\}$ for shift-invariant reconstruction by $\varphi(\mathbf{D}^{-1}\mathbf{x})$.

Correspondingly, all expressions—integrals, summations, Fourier transforms, etc—can be decomposed into similar tensor products; this simplifies the analysis since what holds true for one dimension remains true for all other dimensions. Specifically, for the Riesz condition, one has

$$A_{\varphi_{\text{sep}}}(\boldsymbol{\omega}) = \prod_{i=1}^d A_{\varphi}(\omega_i). \quad (2.17)$$

Then, one needs to verify (2.13) for only $\hat{\varphi}(\omega)$ (the 1-D counterpart); automatically, $\hat{\varphi}_{\text{sep}}(\boldsymbol{\omega})$ also satisfies (2.13) because of separability.

2.3 Sampling in Shift-Invariant Spaces

The approximation schemes that we consider are all based on the standard three-step approach—prefiltering, sampling, and reconstruction. Figure 2.2 illustrates the generic setup where the input signal is first filtered by the prefilter $\frac{1}{|\mathbf{D}|}h(\mathbf{D}^{-1}\mathbf{x})$ which is followed by sampling represented as the product with a sequence of Dirac impulses. We also consider using a digital filter q that computes the expansion coefficients. The reconstruction is performed by applying

the synthesis filter $\varphi(\mathbf{D}^{-1}\mathbf{x})$. Mathematically, we have

$$\begin{aligned}\tilde{c}[\mathbf{k}] &= (f \star \frac{1}{|\mathbf{D}|} h(\mathbf{D}^{-1}\bullet))(\mathbf{D}\mathbf{k}) = \int_{\mathbb{R}^d} f(\boldsymbol{\xi}) \frac{1}{|\mathbf{D}|} h(\mathbf{k} - \mathbf{D}^{-1}\boldsymbol{\xi}) d\boldsymbol{\xi} \\ &= \left\langle f, \frac{1}{|\mathbf{D}|} h(\mathbf{k} - \mathbf{D}^{-1}\bullet) \right\rangle_{L_2},\end{aligned}\quad (2.18)$$

and $c[\mathbf{k}] = (\tilde{c} \star q)[\mathbf{k}]$ so that the reconstruction is given by

$$f_{\text{rec}}(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} (\tilde{c} \star q)[\mathbf{k}] \varphi(\mathbf{D}^{-1}\mathbf{x} - \mathbf{k}).\quad (2.19)$$

Accordingly, we define the approximation operator

$$\mathcal{Q}_{\mathbf{D}}\{f\}(\mathbf{x}) = \sum_{\substack{\mathbf{k} \in \mathbb{Z}^d \\ \mathbf{m} \in \mathbb{Z}^d}} \left(\int_{\mathbb{R}^d} f(\boldsymbol{\xi}) h(\mathbf{m} - \mathbf{D}^{-1}\boldsymbol{\xi}) \frac{d\boldsymbol{\xi}}{|\mathbf{D}|} \right) q[\mathbf{k} - \mathbf{m}] \varphi(\mathbf{D}^{-1}\mathbf{x} - \mathbf{k}),\quad (2.20)$$

which yields the approximation $f_{\text{rec}}(\mathbf{x}) \in V_{\mathbf{D}}(\varphi)$ of the input signal $f(\mathbf{x}) \in L_2(\mathbb{R}^d)$.

Since φ completely characterizes the approximation space $V_{\mathbf{D}}(\varphi)$, its selection is important for obtaining a good approximation of the original signal. In this context, the approximation order of φ plays a key role as will be explained in Section 2.4. Once the choice of φ has been made, h and q can be determined in accordance with the design constraints. In this context, a popular constraint [5, 53] is the biorthogonality condition¹⁰

$$\left\langle \sum_{\mathbf{k} \in \mathbb{Z}^d} q[\mathbf{k}] h(\mathbf{n} - \mathbf{k} - \bullet), \varphi(\bullet) \right\rangle_{L_2} = \delta[\mathbf{n}] \stackrel{\text{def}}{=} \begin{cases} 1, & \mathbf{n} = \mathbf{0} \\ 0, & \mathbf{n} \in \mathbb{Z}^d \setminus \{\mathbf{0}\} \end{cases} .\quad (2.21)$$

A direct consequence of (2.21) is that $\mathcal{Q}_{\mathbf{D}}$ becomes a projector; that is,

$$\mathcal{Q}_{\mathbf{D}}\{\mathcal{Q}_{\mathbf{D}}\{f\}\}(\mathbf{x}) = \mathcal{Q}_{\mathbf{D}}\{f\}(\mathbf{x}).$$

This means that if the original signal $f(\mathbf{x})$ belongs to $V_{\mathbf{D}}(\varphi)$ then it can be perfectly reconstructed. This can be thought of as a generalized version of Shannon's result for which $\mathcal{Q}_{\mathbf{D}} = P_{V_{\mathbf{D}}(\text{sinc})}$ that pertains only to bandlimited functions. In the sequel, we describe different sampling schemes and provide the corresponding prefilter and the digital correction filter in each case.

¹⁰In [53], this condition is specified for a setup without the digital correction filter q . However, since we have considered q explicitly in our formulation, the standard representation of biorthogonality translates to Equation (2.21) where we have made use of the linearity of the sampling process to associate q with the prefilter h .

2.3.1 Minimum-Error Sampling

Here, we formulate the problem in accordance with Equation (2.3), but, we look for the minimum L_2 -error approximation of $f(\mathbf{x}) \in L_2(\mathbb{R}^d)$ in $V_{\mathbf{D}}(\varphi)$:

$$f_{\text{rec}} = \arg \min_{s \in V_{\mathbf{D}}(\varphi)} \|f - s\|_{L_2}^2. \quad (2.22)$$

The solution is the orthogonal projection of f on to $V_{\mathbf{D}}(\varphi)$ given by

$$\begin{aligned} f_{\text{min}}(\mathbf{x}) &= P_{V_{\mathbf{D}}(\varphi)}\{f\}(\mathbf{x}) \\ &= \sum_{\mathbf{k} \in \mathbb{Z}^d} \left\langle f, \frac{1}{|\mathbf{D}|} \hat{\varphi}(\mathbf{D}^{-1} \bullet - \mathbf{k}) \right\rangle_{L_2} \varphi(\mathbf{D}^{-1} \mathbf{x} - \mathbf{k}), \end{aligned} \quad (2.23)$$

where $\hat{\varphi}(\mathbf{x}) \in V_{\mathbf{I}}(\varphi)$ (for $\mathbf{D} = \mathbf{I}$) is the dual function that is defined via its Fourier transform:

$$\hat{\hat{\varphi}}(\boldsymbol{\omega}) = \frac{\hat{\varphi}(\boldsymbol{\omega})}{A_{\varphi}(\boldsymbol{\omega})}. \quad (2.24)$$

This type of least-squares solution is standard in approximation and wavelet theory [44, 54, 59, 60]; our formula (2.23) is the multidimensional transcription of these results for an arbitrary sampling lattice $\mathbf{D}\mathbb{Z}^d$.

Since $A_{\varphi}(\boldsymbol{\omega})$ is non-zero (because $A_{R_{\varphi,2}} > 0$), $\hat{\varphi}$ is well-defined whenever $\varphi(\mathbf{x})$ generates a Riesz basis. It is clear that $\hat{\varphi}$ satisfies the biorthogonality relation (2.21) with $h(\mathbf{x}) = \hat{\varphi}(-\mathbf{x})$ and $q[\mathbf{k}] = \delta[\mathbf{k}]$ which confirms that $P_{V_{\mathbf{D}}(\varphi)}$ is indeed a projector. Also, the solution (2.23) represents a generalization¹¹ of Shannon's model since (2.23) reduces to (2.5) when $\varphi = \text{sinc}$ (for $d = 1$).

2.3.2 Consistent Sampling

In many practical situations, the prefilter h is often specified *a priori*. In such cases, it is preferable to obtain a reconstruction that is *consistent* in the sense that $f_{\text{rec}}(\mathbf{x})$ (see Figure 2.2) yields exactly the same measurements when fed back into the system; that is,

$$\tilde{c}[\mathbf{k}] = \left\langle f, \frac{1}{|\mathbf{D}|} h(\mathbf{k} - \mathbf{D}^{-1} \bullet) \right\rangle_{L_2} = \left\langle f_{\text{rec}}, \frac{1}{|\mathbf{D}|} h(\mathbf{k} - \mathbf{D}^{-1} \bullet) \right\rangle_{L_2}. \quad (2.25)$$

¹¹It is a well-known fact that $\text{sinc}(\mathbf{x})$ is orthogonal to its integer translates and is therefore its own dual [5, 61].

The framework for this type of reconstruction was presented by Unser *et al.* in [52] for the case of $d = 1$. Here, we extend this for arbitrary \mathbf{D} in d -dimensions: Based on [52, Theorem 1], a unique reconstruction f_{rec} satisfying (2.25) can be obtained by filtering the samples $\tilde{c}[\mathbf{k}]$ with the digital correction filter q_{con} whose frequency response is given by

$$Q_{\text{con}}(e^{j\boldsymbol{\omega}}) = \frac{1}{A_{h\varphi}(\boldsymbol{\omega})} = \frac{1}{\sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k})}. \quad (2.26)$$

Using Fourier domain calculations, it can be verified that the reconstruction

$$f_{\text{con}}(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} (\tilde{c} * q_{\text{con}})[\mathbf{k}] \varphi(\mathbf{D}^{-1}\mathbf{x} - \mathbf{k}) \quad (2.27)$$

satisfies (2.25). In fact, the above reconstruction formula provides a generalization of Shannon's model for nonideal acquisition devices.

The existence of q_{con} is guaranteed whenever $\forall \boldsymbol{\omega} \in [0, 2\pi)^d$, $|A_{h\varphi}(\boldsymbol{\omega})| \geq m_{h,\varphi}$ for some $m_{h,\varphi} > 0$. Moreover, when $|A_{h\varphi}(\boldsymbol{\omega})| \leq M_{h,\varphi} < +\infty$, $\forall \boldsymbol{\omega} \in [0, 2\pi)^d$, we have that

$$0 < \frac{1}{M_{h,\varphi}} \leq |Q_{\text{con}}(e^{j\boldsymbol{\omega}})| \leq \frac{1}{m_{h,\varphi}} < +\infty.$$

Then, we can define an equivalent generating function

$$\varphi_{\text{con}}(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} q_{\text{con}}[\mathbf{k}] \varphi(\mathbf{x} - \mathbf{k}), \quad (2.28)$$

and from Proposition 2.2, $\{\varphi_{\text{con}}(\mathbf{D}^{-1}\mathbf{x} - \mathbf{k})\}_{\mathbf{k} \in \mathbb{Z}^d}$ also forms a Riesz basis of $V_{\mathbf{D}}(\varphi)$.

Equation (2.27) can also be used as a computationally efficient approximation to the minimum-error reconstruction f_{min} whenever the inner product with the dual function becomes difficult to evaluate in (2.23). Using Fourier domain calculations, it can be verified that the biorthogonality condition (2.21) is satisfied for $q = q_{\text{con}}$. In this case, the interpretation is that f_{con} is a projection onto $V_{\mathbf{D}}(\varphi)$ “perpendicular” to $V_{\mathbf{D}}(h)$ —oblique projection [52]—where $V_{\mathbf{D}}(h)$ is the shift-invariant space obtained by replacing¹² φ by h in (2.10).

¹²It is also required that $V_{\mathbf{D}}(h) \subset L_2(\mathbb{R}^d)$ which is ensured provided h satisfies (2.12).

By extending the results of [52] to d -dimensions, it can be established that the L_2 -error of f_{con} is bounded as

$$\|f - f_{\min}\|_{L_2} \leq \|f - f_{\text{con}}\|_{L_2} \leq \frac{1}{\cos(\theta_{h\varphi})} \|f - f_{\min}\|_{L_2}, \quad (2.29)$$

where $\theta_{h\varphi}$ may be interpreted as the angle between the shift-invariant spaces generated by $h(\mathbf{x})$ and $\varphi(\mathbf{x})$ and is specified as

$$\cos(\theta_{h\varphi}) = \text{ess} \inf_{\boldsymbol{\omega} \in [0, 2\pi)^d} \left| \frac{A_{h\varphi}(\boldsymbol{\omega})}{\sqrt{A_h(\boldsymbol{\omega})} \sqrt{A_\varphi(\boldsymbol{\omega})}} \right|, \quad (2.30)$$

where the *essential infimum* reduces to *infimum* when the quantity on the r.h.s. of the above equation is continuous at all $\boldsymbol{\omega} \in [0, 2\pi)^d$.

The upper bound in (2.29), which corresponds to the worst possible case, becomes close to the lower bound—minimum error—whenever the angle between analysis and reconstruction spaces is negligible. In fact, we regain the minimum L_2 -error reconstruction if $\cos(\theta_{h\varphi}) = 1$. This happens when $h(\mathbf{x}) \in V_{\mathbf{I}}(\varphi)$ (for $\mathbf{D} = \mathbf{I}$), which refers to the case where the analysis and the synthesis spaces are identical.

2.3.3 Interpolation

When $h(\mathbf{x}) = \delta(\mathbf{x})$, the Dirac impulse, the consistent sampling technique described in the previous section turns out to be an interpolation scheme, where we reconstruct an interpolating function $f_{\text{int}}(\mathbf{x})$ which exactly fits the samples¹³ of the given signal; that is,

$$f_{\text{int}}(\mathbf{x})|_{\mathbf{x}=\mathbf{D}\mathbf{k}} = f(\mathbf{D}\mathbf{k}). \quad (2.31)$$

Interpolation is a vital ingredient of many image-processing and biomedical applications [15] and its theory originates from ancient mathematics [6], cardinal series [7–10] and approximation theory [48, 62]. Equation (2.31) is also known as the *interpolation condition* that is enforced on f_{int} by filtering $\{f(\mathbf{D}\mathbf{k})\}$ with the digital correction filter q_{int} whose frequency response is [15, 63] (obtained by setting $\hat{h}(\boldsymbol{\omega}) = 1$ in (2.26))

$$Q_{\text{int}}(e^{j\boldsymbol{\omega}}) = \left(\sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \right)^{-1}. \quad (2.32)$$

¹³In order for the samples $\{f(\mathbf{D}\mathbf{k})\}$ to be well-defined in the ℓ_2 -sense, $f(\mathbf{x}) \in W_2^p$ (see Theorem 3.1 in Chapter 3).

Table 2.1: Summary of schemes for sampling in shift-invariant spaces

Sampling Scheme	Digital		Analog	Equivalent Basis
	Analog Prefilter	Correction Filter	Reconstruction Filter	
Shannon	$\text{sinc}(\mathbf{x})$	—	$\text{sinc}(\mathbf{x})$	$\text{sinc}(\mathbf{x})$
Minimum-error	$\check{\varphi}(-\mathbf{x})$	$\delta[\mathbf{k}]$	$\varphi(\mathbf{x})$	$\varphi(\mathbf{x})$
Consistent	$h(\mathbf{x})$	q_{con}	$\varphi(\mathbf{x})$	$\varphi_{\text{con}}(\mathbf{x})$
Interpolation	$\delta(\mathbf{x})$	q_{int}	$\varphi(\mathbf{x})$	$\varphi_{\text{int}}(\mathbf{x})$

The corresponding reconstruction—the interpolating signal f_{int} —is given by

$$f_{\text{int}}(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} (\tilde{c} * q_{\text{int}})[\mathbf{k}] \varphi(\mathbf{D}^{-1}\mathbf{x} - \mathbf{k}), \quad (2.33)$$

where $\tilde{c}[\mathbf{k}] = f(\mathbf{D}\mathbf{k})$ in this case. The equivalent generating function (provided $Q_{\text{int}}(e^{j\omega})$ meets $0 < m_{q_{\text{int}}} \leq |Q_{\text{int}}(e^{j\omega})| \leq M_{q_{\text{int}}} < +\infty$) is nothing but the interpolator

$$\varphi_{\text{int}}(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} q_{\text{int}}[\mathbf{k}] \varphi(\mathbf{x} - \mathbf{k}), \quad (2.34)$$

which fulfills $\varphi_{\text{int}}(\mathbf{x})|_{\mathbf{x}=\mathbf{k}} = \delta[\mathbf{k}]$: This is the equivalent of the biorthogonality condition (2.21) for the current setting. In fact, this property is naturally satisfied by the sinc-function and therefore, (2.33) may be considered as an extension of Shannon’s model to interpolation using an arbitrary basis.

We summarize in Table 2.1 the attributes of different sampling schemes considered so far and highlight the analog prefilter, the digital correction filter and the equivalent basis in each case.

2.4 Characterizing the Approximation Error

Since projection-based approaches rely on the approximation of the original signal in some subspace $V_{\mathbf{D}}(\varphi) \subset L_2(\mathbb{R}^d)$, it is imperative to characterize the associated L_2 approximation error $\|f - \mathcal{Q}_{\mathbf{D}}\{f\}\|_{L_2}$. However, we observe that $\|f - \mathcal{Q}_{\mathbf{D}}\{f\}\|_{L_2}$ obeys the following rule (which can be derived by suitable

change of variables):

$$\|f(\bullet) - \mathcal{Q}_{\mathbf{D}}\{f\}(\bullet)\|_{L_2} = (|\mathbf{D}|)^{\frac{1}{2}} \|f(\mathbf{D}\bullet) - \mathcal{Q}_{\mathbf{I}}\{f(\mathbf{D}\bullet)\}(\bullet)\|_{L_2},$$

that is, the error of approximating $f(\mathbf{x})$ using a shift-invariant representation on the lattice $\mathbf{D}\mathbb{Z}^d$ is equal (up to a multiplicative constant) to that of approximating $f(\mathbf{D}\mathbf{x})$ by the same representation on the standard lattice \mathbb{Z}^d . Therefore, for mathematical ease, we will treat the case of

$$\mathbf{D} = T\mathbf{I},$$

where $T > 0$ is the sampling-step that is identical along every dimension. Correspondingly, the error takes the form

$$\varepsilon_f(T) = \|f - \mathcal{Q}_{T\mathbf{I}}\{f\}\|_{L_2}.$$

The following result then provides an exact Fourier domain formula for $\varepsilon_f(T)$ (this is the d -dimensional version of Theorem 1 in [54]).

Theorem 2.1. *Let the Fourier transforms $|\hat{h}(\boldsymbol{\omega})| \leq M_h < +\infty \forall \boldsymbol{\omega} \in \mathbb{R}^d$ and $|Q(e^{j\boldsymbol{\omega}})| \leq M_q < +\infty \forall \boldsymbol{\omega} \in [0, 2\pi)^d$. Then, $\forall f \in W_2^\rho$, the approximation error $\varepsilon_f(T)$ is given by*

$$\varepsilon_f(T) = \varepsilon_{\text{dom}}(f, T) + \mathcal{O}(T^\rho), \quad (2.35)$$

where¹⁴ $\varepsilon_{\text{dom}}(f, T)$ is the dominating component that can be computed exactly in the Fourier domain as

$$\varepsilon_{\text{dom}}(f, T) = \left[\frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \mathcal{E}_{h,\varphi}(\boldsymbol{\omega}T) |\hat{f}(\boldsymbol{\omega})|^2 d\boldsymbol{\omega} \right]^{\frac{1}{2}}, \quad (2.36)$$

and where $\mathcal{E}_{h,\varphi}(\boldsymbol{\omega})$ represents the Fourier domain error kernel given by

$$\begin{aligned} \mathcal{E}_{h,\varphi}(\boldsymbol{\omega}) &= \left| 1 - Q(e^{j\boldsymbol{\omega}}) \hat{h}(\boldsymbol{\omega}) \hat{\varphi}(\boldsymbol{\omega}) \right|^2 \\ &+ \sum_{\mathbf{k} \in \mathbb{Z}^d \setminus \{\mathbf{0}\}} |Q(e^{j\boldsymbol{\omega}})|^2 \left| \hat{h}(\boldsymbol{\omega}) \right|^2 |\hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2. \end{aligned} \quad (2.37)$$

Proof. The proof is given in Appendix A.3. ■

¹⁴We write $s(x) = \mathcal{O}(x^n)$ if $\limsup_{x \rightarrow 0} \left| \frac{s(x)}{x^n} \right| < +\infty$.

Since the form of the multidimensional error kernel $\mathcal{E}_{h,\varphi}(\boldsymbol{\omega})$ is similar to that of the one-dimensional one in [54, Equation (21)] the accompanying analysis in [54] also remains largely valid in our case; that is, $\varepsilon_{\text{dom}}(f, T)$ allows for the prediction of the general error behavior of a sampling scheme by simple examination of $\mathcal{E}_{h,\varphi}(\boldsymbol{\omega})$. To see this, we manipulate and rewrite the error kernel as

$$\mathcal{E}_{h,\varphi}(\boldsymbol{\omega}) = \left(1 - \frac{|\hat{\varphi}(\boldsymbol{\omega})|^2}{A_\varphi(\boldsymbol{\omega})}\right) + A_\varphi(\boldsymbol{\omega}) \left|Q^*(e^{j\boldsymbol{\omega}})\hat{h}^*(\boldsymbol{\omega}) - \hat{\varphi}(\boldsymbol{\omega})\right|^2. \quad (2.38)$$

Then, for the case of minimum L_2 -error sampling ($h(\mathbf{x}) = \hat{\varphi}(-\mathbf{x})$ and $Q(e^{j\boldsymbol{\omega}}) = 1$), the error kernel reduces to

$$\mathcal{E}_{\hat{\varphi},\varphi}(\boldsymbol{\omega}) = 1 - \frac{|\hat{\varphi}(\boldsymbol{\omega})|^2}{A_\varphi(\boldsymbol{\omega})}, \quad (2.39)$$

which is in agreement with the fact that the orthogonal projection onto $V_{T\mathbf{I}}(\varphi)$ minimizes the approximation error. For Shannon's model, orthogonal projection onto the space of bandlimited functions results in the kernel (see Figure 2.3 for the case of $d = 1$)

$$\mathcal{E}_{\text{sinc,sinc}}(\boldsymbol{\omega}) = 1 - \text{rect}\left(\frac{\boldsymbol{\omega}}{2\pi}\right), \quad (2.40)$$

which makes sense since in this case the error is entirely due to the out-of-band portion of the signal.

We plot (for $d = 1$) $\mathcal{E}_{h,\varphi}(\omega)$ in Figure 2.3 for the considered sampling schemes—minimum L_2 -error sampling, consistent sampling and interpolation—corresponding to the following two cases: (a) $h(x) = \text{rect}(x)$ and $\varphi(x) = \beta^1(x)$ (the linear B-spline or the tent function [14]), and (b) $h(x) = \beta^1(x)$ and $\varphi(x) = \beta^3(x)$ (the cubic B-spline [14]). We also show $\mathcal{E}_{\text{sinc,sinc}}(\omega)$ (in gray) corresponding to Shannon's setup.

As expected, $\mathcal{E}_{\hat{\varphi},\varphi}(\omega)$, which corresponds to the minimum L_2 -error solution, is closer to $\mathcal{E}_{\text{sinc,sinc}}(\omega)$. Note that the kernel corresponding to consistent sampling— $\mathcal{E}_{h,\varphi}(\omega)$ —is only slightly suboptimal; it converges to the minimum-error solution when the analysis filter becomes close to the synthesis space (see Figure 2.3(b)). This is also indicated by the $\frac{1}{\cos(\theta_{h,\varphi})}$ values (see Equation (2.29)) which is 1.155 for case (a) while it is only 1.030 for case (b). Finally, sampling without prefiltering ($h(x) = \delta(x)$)—interpolation—is the least favorable approach as the signal components above the Nyquist frequency contribute twice to the approximation error both due to aliasing and because the reconstruction is not perfect in $V_{T\mathbf{I}}(\varphi)$ for in-band frequencies.

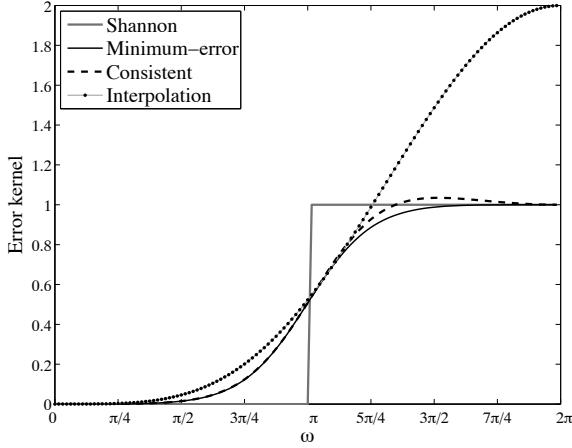
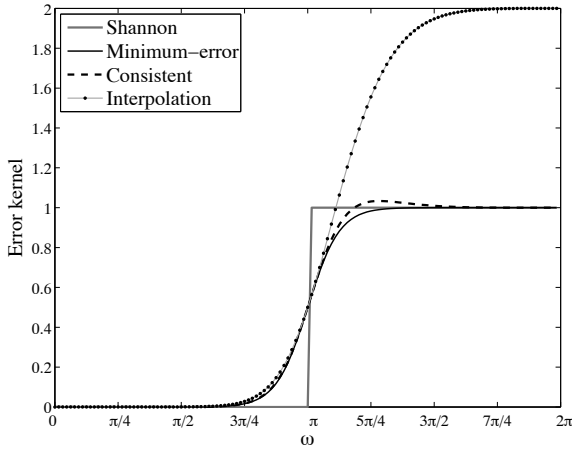
(a) $\mathcal{E}_{h,\varphi}(\omega)$ for $h(x) = \text{rect}(x)$ and $\varphi(x) = \beta^1(x)$ (b) $\mathcal{E}_{h,\varphi}(\omega)$ for $h(x) = \beta^1(x)$ and $\varphi(x) = \beta^3(x)$

Figure 2.3: Approximation error kernels corresponding to the schemes namely, Shannon sampling, minimum L_2 -error sampling, consistent sampling and interpolation, for two different setups: In both (a) and (b) the kernel $\mathcal{E}_{\hat{\varphi},\varphi}$ corresponding to minimum L_2 -error sampling is closest to the ideal one of Shannon $\mathcal{E}_{\text{sinc},\text{sinc}}$ (in gray), while $\mathcal{E}_{h,\varphi}$ (consistent sampling) closely follows $\mathcal{E}_{\hat{\varphi},\varphi}$. The kernel $\mathcal{E}_{\delta,\varphi}$ corresponding to interpolation is always higher than the other two indicating that sampling without prefiltering can be detrimental.

2.5 Approximation Order

It is reasonable to expect the approximation error $\varepsilon_f(T)$ to decrease as the sampling-step T becomes smaller. A key concept in this context is the *approximation order* which describes the rate at which $\varepsilon_f(T)$ decays; specifically, we write that $V_{T\mathbf{I}}(\varphi)$ provides approximation order L [44] if $\forall f \in W_2^\rho$, $\rho \geq L$, the approximation error corresponding to the orthogonal projection behaves as

$$\|f - P_{V_{T\mathbf{I}}(\varphi)}\{f\}\|_{L_2} \leq \text{constant} \times T^L \|f\|_{W_2^\rho}. \quad (2.41)$$

2.5.1 Strang-Fix Conditions

A popular result due to Strang and Fix [45] says that (2.41) is true if and only if the Fourier transform $\hat{\varphi}(\boldsymbol{\omega})$ of $\varphi(\mathbf{x})$ satisfies¹⁵ the following conditions:

$$\hat{\varphi}(\mathbf{0}) \neq 0, \quad \text{and} \quad \hat{\varphi}^{(|\boldsymbol{\alpha}|)}(2\pi\mathbf{k}) = 0 \quad \text{for} \quad \begin{cases} \mathbf{k} \in \mathbb{Z}^d \setminus \{\mathbf{0}\}, \\ 0 \leq |\boldsymbol{\alpha}| \leq L - 1, \end{cases} \quad (2.42)$$

where $|\boldsymbol{\alpha}| = \sum_{i=1}^d \alpha_i$ and $\hat{\varphi}^{(|\boldsymbol{\alpha}|)}(\boldsymbol{\omega}) = \frac{\partial^{|\boldsymbol{\alpha}|} \hat{\varphi}(\boldsymbol{\omega})}{\partial \omega_1^{\alpha_1} \partial \omega_2^{\alpha_2} \dots \partial \omega_d^{\alpha_d}}$. If φ satisfies (2.42), then it is said to be of order L . This is more apparent if we expand $\hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k})$, $\mathbf{k} \in \mathbb{Z}^d$ in terms of its Taylor series around $\boldsymbol{\omega} = \mathbf{0}$:

$$\hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k}) = \sum_{l=0}^{\infty} \sum_{|\boldsymbol{\alpha}|=l} \hat{\varphi}^{(|\boldsymbol{\alpha}|)}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \Big|_{\boldsymbol{\omega}=\mathbf{0}} \frac{\boldsymbol{\omega}^{\boldsymbol{\alpha}}}{\boldsymbol{\alpha}!},$$

where we have used the standard multi-index notation: $\frac{\boldsymbol{\omega}^{\boldsymbol{\alpha}}}{\boldsymbol{\alpha}!} = \prod_{i=1}^d \frac{\omega_i^{\alpha_i}}{\alpha_i!}$. Then,

from (2.42), we see that $\forall \mathbf{k} \in \mathbb{Z}^d$,

$$\hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k}) = \hat{\varphi}(\mathbf{0}) \delta[\mathbf{k}] + \mathcal{O}(\|\boldsymbol{\omega}\|^L). \quad (2.43)$$

In a similar manner, the concept of approximation order can be extended to the general projector $\mathcal{Q}_{T\mathbf{I}}$; obviously this necessitates conditions similar to (2.43) jointly on q , h and φ : Specifically, it can be shown that if $\forall \mathbf{k} \in \mathbb{Z}^d$,

$$Q(e^{j\boldsymbol{\omega}}) \hat{h}(\boldsymbol{\omega}) \hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k}) = \delta[\mathbf{k}] + \mathcal{O}(\|\boldsymbol{\omega}\|^L), \quad (2.44)$$

¹⁵The first-order ($L = 1$) Strang-Fix condition with $\hat{\varphi}(\mathbf{0}) = 1$ is equivalent to a partition of unity, namely, $\sum_{\mathbf{k} \in \mathbb{Z}^d} \varphi(\mathbf{x} - \mathbf{k}) = 1$, $\forall \mathbf{x} \in [0, 1)^d$. If φ satisfies the partition of unity condition, then $\lim_{T \rightarrow 0} \varepsilon_f(T) = 0$ [5].

then all the terms in the r.h.s. of (2.37) are $\mathcal{O}(\|\boldsymbol{\omega}\|^L)$. Using the same Taylor expansion argument, then we have that $\forall f \in W_2^\rho$, $\rho \geq L$,

$$\|f - \mathcal{Q}_{\text{TI}}\{f\}\|_{L_2} \leq \text{constant} \times T^L \|f\|_{W_2^L}, \quad (2.45)$$

which is a standard result in approximation theory [44–46, 64, 65].

2.5.2 Asymptotic Form of the Error for $d = 1$

One can obtain an explicit form for the leading constant in (2.45) in the case of $d = 1$ (1-D signals) as described in [53, 54]. For $d = 1$, (2.44) is equivalent to saying that that φ and $h_q(\bullet) = \sum_{k \in \mathbb{Z}} q[k] h(\bullet - k)$ are *quasibiorthonormal* of order L ; that is, in addition to the Strang-Fix conditions of order L on φ , the moments of $h_q(-x)$ and of $\tilde{\varphi}(x)$ are equal [54]:

$$\int_{\mathbb{R}} x^\alpha h_q(-x) dx = \int_{\mathbb{R}} x^\alpha \tilde{\varphi}(x) dx, \quad (2.46)$$

for $\alpha = 0, 1, \dots, L - 1$. Then, $\forall f \in W_2^\rho$, $\rho \geq L$, the asymptotic form of the error $\epsilon_f(T)$ can be deduced as [53, 54]:

$$\epsilon_f(T) = C_L T^L \|f^{(L)}\|_{L_2} \quad \text{as } T \rightarrow 0, \quad (2.47)$$

where $f^{(L)}$ is the L -th derivative of f and C_L is the asymptotic error constant given by

$$C_L = \sqrt{\frac{\mathcal{E}_{h,\varphi}^{(2L)}(0)}{(2L)!}}, \quad (2.48)$$

where $\mathcal{E}_{h,\varphi}^{(2L)}(T)$ is the $(2L)$ -th derivative of $\mathcal{E}_{h,\varphi}(T)$.

Equations (2.45) and (2.48) play a crucial role in the selection of the basis function since the approximation order L of φ dictates the decay rate of the error: The higher the value of L , the faster the decay and therefore, the better is the quality of the reconstruction. For a fixed φ , we can also study the performance of various approximation algorithms such as projectors (orthogonal and oblique) and interpolators by comparing the leading constant C_L corresponding to each algorithm. For instance, let us consider the case (a) of Figure 2.3. Here, for the interpolation algorithm $C_{L_{\text{int}}} = \frac{1}{2\sqrt{30}} \approx 9.1287 \times 10^{-2}$, while for the orthogonal and oblique projectors $C_{L_{\text{ortho}}} = C_{L_{\text{obl}}} = \frac{1}{12\sqrt{5}} \approx 3.7267 \times 10^{-2}$,

which is the smallest possible constant achievable in the current setting. This further substantiates the argument at the end of Section 2.4 for the improved performance of the projectors (orthogonal and oblique) over the interpolator. Case (b) is no different, where again, the projectors have a good edge over the interpolator: $C_{L_{\text{ortho}}} = C_{L_{\text{obl}}} = \frac{1}{240\sqrt{21}} \approx 9.0924 \times 10^{-4}$, while $C_{L_{\text{int}}} = \frac{1}{72\sqrt{70}} \approx 1.66 \times 10^{-3}$, respectively.

2.6 Quasi-Interpolation

The performance of interpolation can be boosted by giving up the exact fitting constraint (see Section 2.3.3) and by designing the digital correction filter $Q(e^{j\boldsymbol{\omega}})$ so as to minimize the second term¹⁶ on the r.h.s. of $\mathcal{E}_{h,\varphi}(\boldsymbol{\omega})$ in (2.38) for $h(\mathbf{x}) = \delta(\mathbf{x})$. The resulting scheme is termed *quasi-interpolation* where “quasi” refers to the fact that the exact fitting condition is satisfied only for polynomials of degree $n = 0, 1, \dots, L - 1$ [53, 54, 65].

Under the assumption that the input signal has a more or less constant spectrum over the bandlimited support, the above minimization leads to the optimal filter [54]

$$q_{\text{quasi}}[\mathbf{k}] = \frac{1}{(2\pi)^d} \int_{[-\pi, \pi]^d} \hat{\varphi}(\boldsymbol{\omega}) e^{-j\boldsymbol{\omega}^T \mathbf{k}} d\boldsymbol{\omega}. \quad (2.49)$$

Thus, $Q_{\text{quasi}}(e^{j\boldsymbol{\omega}})$ is chosen so that it matches $\hat{\varphi}(\boldsymbol{\omega})$ over $\boldsymbol{\omega} \in [-\pi, \pi]^d$, thereby attempting to mimic the minimum L_2 -error solution in this frequency band.

The filter q can also be designed based on the asymptotic form of the error $\varepsilon_f(T)$ whenever the assumption of constant spectrum does not apply. The corresponding design constraint is

$$Q^*(e^{j\boldsymbol{\omega}}) - \hat{\varphi}(\boldsymbol{\omega}) = \mathcal{O}(\|\boldsymbol{\omega}\|^N), \quad (2.50)$$

which is based on a weaker hypothesis and is therefore more relevant to practical situations.

2.7 Summary

In this chapter, we reviewed some key concepts in linear approximation theory and their application to the sampling problem. The common strategy underly-

¹⁶The first term on the r.h.s. of (2.38) is the minimum of $\mathcal{E}_{h,\varphi}(\boldsymbol{\omega})$ that is achieved by the orthogonal projector.

ing the many extensions considered here—minimum L_2 -error sampling, consistent sampling, and interpolation—is to pose sampling and reconstruction as an approximation problem in some given shift-invariant space.

We studied results pertaining to the analytical computation of the L_2 -approximation error, the dominant component of which can be determined by integrating the spectrum of the input signal against a Fourier domain error kernel $\mathcal{E}_{h,\varphi}(\omega)$. Based on quantitative analysis of the approximation error, it was seen that the approximation order of the basis function is the key parameter that controls the quality of reconstruction.

Chapter 3

Nonideal Sampling in the Presence of Noise

3.1 Introduction

The approximation-theoretic solutions discussed in the previous chapter provide a mathematically elegant framework for extending Shannon’s work. However, they rely on the assumption that the sampling process is noise-free which is usually not the case in practice. Moreover, prior to sampling, real-world signals often pass through an acquisition device whose impulse response is in general nonideal.

Digital filtering techniques such as inverse filtering [34, 52] have been proposed to partially compensate for these departures from “idealness”. In the presence of noise, the digital correction filter is sufficiently regularized for stability purposes.¹ This forms the key aspect of the work of Eldar *et al.* [66] who presented an extended formulation of projection-based approach for the problem of nonideal sampling in the presence of noise. Their idea is to compute the reconstruction of a 1-D signal $f(x)$ from noisy (nonideal) measurements in some general integer-shift-invariant space $V(\varphi)$, where $\varphi(x)$ is an arbitrary generating function specified *a priori*, by suitable digital processing of the nonideal samples. As in the classical case [34], this corresponds to a special kind of

¹This chapter is based on the article: S. Ramani, D. Van De Ville, T. Blu and M. Unser, “Nonideal Sampling and Regularization Theory,” IEEE Transactions on Signal Processing, vol. 56, no. 3, pp: 1055–1070, March 2008.

Table 3.1: Characteristics of Various Sampling Schemes

Method	Acquisition Device	Noise	Reconstruction Space	Dimension
Shannon [2]	Ideal	No	Bandlimited	1-D
Interpolation [6, 15, 19, 65]	Ideal	No	Shift-Invariant, Splines	Multi-D
Peterson <i>et al.</i> [67]	Ideal	No	Bandlimited	Multi-D
Minimum-Error [44]	Dual	No	Shift-Invariant	Multi-D
Papoulis [68]	Nonideal	No	Bandlimited	1-D
Consistent, Unser <i>et al.</i> [52]	Nonideal	No	Shift-Invariant, Splines	1-D
Eldar <i>et al.</i> [66]	Nonideal	Yes	Shift-Invariant	1-D
This work	Nonideal	Yes	“Matched”	Multi-D

(regularized-inverse, Wiener) filtering, except that the digital correction filter explicitly depends on φ [66]. The bottom-line is that the user is free to select the reconstruction space he wishes, provided of course that he optimizes the digital correction filter accordingly.

Having this added flexibility is desirable, but it also raises the important issue of the selection of a reconstruction space that is most suitable for a particular sampling set-up and/or class of input signals. This is precisely the question that we will address here while also extending some of the previous formulations to the multidimensional setting. Table 3.1 summarizes the characteristics of some well-known methods in literature in comparison to what is developed in this chapter. In contrast with the previous works (such as those listed in Table 3.1), we want to infer the “optimal” reconstruction space together with the reconstruction algorithm in a deductive fashion through a global mathematical optimization process.

We consider a realistic setting where a multidimensional signal is prefiltered prior to sampling, and the samples are corrupted by additive noise as illus-




Figure 3.1: Block diagram of the generalized sampling problem: The acquisition device specifies the prefilter $h(\mathbf{x})$ which acts in the continuous domain. The choice of the generator $\varphi(\mathbf{x})$ specifies the reconstruction space. The optimal coefficients $\{c[\mathbf{k}]\}$ are determined from the noisy samples $\{g[\mathbf{k}]\}$ as the result of a numerical linear / nonlinear optimization process.

trated in Figure 3.1. We formulate the reconstruction problem in a variational framework wherein the continuous-space solution is obtained by the minimization of a data-fidelity term subject to a continuous-space regularization constraint (Tikhonov-like functional). The data-fidelity term (which is possibly non-quadratic) ensures that the continuous-space solution is consistent. For the regularization, we choose a convex-increasing function of the L_2 -norm of a generalized derivative $L\{\bullet\}$ of the reconstruction (Tikhonov regularization). This gives rise to a general setting which is mathematically tractable and also leads to the specification of an “optimal” reconstruction space.

3.1.1 Contributions

We summarize below the main contributions of the thesis pertaining to this chapter:

- We establish through mathematical minimization of the variational criterion that the “optimal” reconstruction space is shift-invariant in nature and that the continuous-space global minimum solution has an exact analytical form.
- For quadratic data-fidelity, we prove that the optimal solution can be ob-

tained by a simple one-step hybrid-linear filtering (discrete in - continuous out) of the measurements. For the non-quadratic case, the solution can be computed by nonlinear optimization of the coefficients of its shift-invariant representation.

- We show that the generator of the “optimal” reconstruction space is tied to the regularization operator L via an optimality condition. As a direct consequence, the regularization operator completely characterizes the reconstruction space which can be interpreted as a generalization of [66].
- We also justify that the reconstruction framework is closer to spline theory than it is to the traditional view of sampling/deconvolution. A conceptual advantage is that there are no numerical approximations involved (e.g. finite differences to estimate derivatives, etc): All calculations are exact and performed analytically in the continuous domain (similar to what is done in the context of the wavelet transform).
- We present statistical arguments that suggest the same type of shift-invariant reconstruction mechanism when the input signal belongs to certain classes of stochastic processes. In fact, we uncover a functional equivalence between the variational and Wiener solutions to the problem of nonideal sampling in the presence of noise, which in turn helps us to select the most appropriate regularization functional. In this context, we also develop a new Matérn B-spline which is appropriate for sampling and interpolation of images with $\|\omega\|^{-\tau}$ -like spectra [69].

3.1.2 Organization of the Chapter

We review in Section 3.2 some mathematical results that we will use in our analysis. The generalized sampling problem is investigated in Section 3.3. We first give a precise statement of the problem and discuss the mathematical hypotheses that are essential to the analysis. We then derive the general continuous-space solution to our variational reconstruction problem and prove that it is included in some optimal shift-invariant space. Section 3.4 is dedicated to the analysis of the stability of the reconstruction; the end result is a set of relatively mild constraints on the frequency response of the regularization operator. We illustrate our results with some concrete examples in Section 3.5. In particular, we consider explicit classes of separable and isotropic regularization operators and characterize the corresponding reconstruction spaces. This also leads to

the definition of a new brand of multidimensional Matérn splines, which extend Rabut's polyharmonic family [70]. Finally, in Section 3.6, we revisit our sampling problem from an estimation theoretic point of view and derive the corresponding minimax and MMSE solutions. This allows us to draw an interesting link with the variational formulation leading to some equivalences of solutions for the deterministic and stochastic cases.

3.2 Preliminaries

For ease of notation, we use the standard sampling lattice, \mathbb{Z}^d (that corresponds to $\mathbf{D} = \mathbf{I}$ in (2.6) in Chapter 2 and normalized sampling-step $T = 1$), henceforth unless specified otherwise.² We represent the integer-shift-invariant space³ generated by $\varphi(\mathbf{x})$ as

$$V(\varphi) = \left\{ s(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] \varphi(\mathbf{x} - \mathbf{k}) : c \in \ell_2(\mathbb{Z}^d) \right\}. \quad (3.1)$$

For the expansion in (3.1) to be well defined, $\varphi(\mathbf{x})$ must satisfy the following stability conditions.

Definition 3.1. *A function $\varphi(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}$ is a stable generator of $V(\varphi)$ if and only if it satisfies the two “stable representation” (**SR**) conditions:*

$$\mathbf{SR} : \left\{ \begin{array}{l} A_{R,\varphi,2} = \operatorname{ess\,inf}_{\boldsymbol{\omega} \in [0, 2\pi)^d} \sqrt{\sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2} > 0, \\ \sup_{\mathbf{x} \in [0, 1)^d} \sum_{\mathbf{k} \in \mathbb{Z}^d} |\varphi(\mathbf{x} + \mathbf{k})| < +\infty. \end{array} \right. \quad (3.2)$$

This is more general than the L_2 -Riesz condition we saw in Chapter 2 in the sense that we are now dealing with L_p -spaces for $1 \leq p \leq \infty$, which requires the stronger second **SR** condition. The two **SR** conditions together ensure that the set of functions $\{\varphi(\mathbf{x} - \mathbf{k})\}$ forms a L_p -stable Riesz basis for all $1 \leq p \leq \infty$ [71];

²The results can be easily extended to handle an arbitrary sampling lattice $\mathbf{D}\mathbb{Z}^d$ using the Fourier transform relation (2.8) and the suitably modified Poisson's summation formula (2.9) of Chapter 2.

³We drop the subscript T in $V_T(\varphi)$ for ease of notation since we have set $T = 1$ here.

that is, $\forall \{c[\mathbf{k}]\} \in \ell_p(\mathbb{Z}^d)$,

$$A_{R,\varphi,p} \|c\|_{\ell_p} \leq \left\| \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] \varphi(\mathbf{x} - \mathbf{k}) \right\|_{L_p} \leq B_{R,\varphi,p} \|c\|_{\ell_p}, \quad (3.3)$$

where $A_{R,\varphi,p} > 0$ and $B_{R,\varphi,p} < +\infty$ are appropriate constants. The norm equivalence (3.3) implies that $V(\varphi) \subset L_p(\mathbb{R}^d)$ for all $1 \leq p \leq \infty$.

We will also take advantage of the following Young-type inequality which asserts that the second **SR** condition is preserved through convolution.

Proposition 3.1. *Let $\varphi_1(\mathbf{x})$ satisfy the second **SR** condition and let $\varphi_2(\mathbf{x}) = (f \star \varphi_1)(\mathbf{x})$ for some $f(\mathbf{x}) \in L_1(\mathbb{R}^d)$. Then $\varphi_2(\mathbf{x})$ also satisfies the second **SR** condition; that is,*

$$\sup_{x \in [0, 1)^d} \sum_{\mathbf{k} \in \mathbb{Z}^d} |\varphi_2(\mathbf{x} + \mathbf{k})| < +\infty. \quad (3.4)$$

Proof. Consider the series

$$\begin{aligned} \sum_{\mathbf{k} \in \mathbb{Z}^d} |\varphi_2(\mathbf{x} + \mathbf{k})| &= \sum_{\mathbf{k} \in \mathbb{Z}^d} \left| \int_{\mathbb{R}^d} f(\mathbf{x}') \varphi_1(\mathbf{x} + \mathbf{k} - \mathbf{x}') \, d\mathbf{x}' \right| \\ &\leq \int_{\mathbb{R}^d} |f(\mathbf{x}')| \left(\sum_{\mathbf{k} \in \mathbb{Z}^d} |\varphi_1(\mathbf{x} + \mathbf{k} - \mathbf{x}')| \right) \, d\mathbf{x}' \\ &\leq \|f\|_{L_1} \left(\sup_{\mathbf{x}' \in [0, 1)^d} \sum_{\mathbf{k} \in \mathbb{Z}^d} |\varphi_1(\mathbf{x}' + \mathbf{k})| \right) < +\infty, \end{aligned}$$

where the above inequality holds $\forall \mathbf{x} \in [0, 1)^d$, including the value \mathbf{x}_0 for which the l.h.s. achieves its maximum. \blacksquare

Remark 3.1. *It should be noted that $(f \star \varphi_1)(\mathbf{k}) \in \ell_1(\mathbb{Z}^d)$, which is obvious when we set $\mathbf{x} = \mathbf{0}$ in (3.4).*

Proposition 3.2. *Let φ satisfy the two **SR** conditions. Then, the function*

$$\varphi_{\text{eq}}(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} q[\mathbf{k}] \varphi(\mathbf{x} - \mathbf{k})$$

satisfies the first **SR** condition, if and only if the Fourier transform of the sequence $\{q[\mathbf{k}]\}$, $Q(e^{j\boldsymbol{\omega}})$, obeys

$$|Q(e^{j\boldsymbol{\omega}})| \geq A_Q > 0 \quad \text{a.e.} \quad (3.5)$$

A sufficient condition for $\varphi_{\text{eq}}(\mathbf{x})$ to satisfy the second **SR** condition is that $\|q\|_{\ell_1} < +\infty$.

Proof. The first part directly follows from Proposition 2.2 in Chapter 2. The second part follows from Proposition 3.1 above. ■

Remark 3.2. A direct consequence of Proposition 3.2 is that φ_{eq} also generates a L_p -stable Riesz basis of $V(\varphi)$.

3.3 Nonideal Sampling and Regularization

3.3.1 Problem Specification

The problem we consider is the recovery of a continuous-space real-valued signal $f(\mathbf{x})$ given some equally spaced, noisy measurements $\{g[\mathbf{k}]\}$. The deterministic signal $f(\mathbf{x})$ is convolved with a prefilter $h(\mathbf{x})$ prior to sampling and the generalized samples are corrupted by additive noise. Specifically, the measurement model associated with Figure 3.1 is

$$g[\mathbf{k}] = (h \star f)(\mathbf{k}) + n[\mathbf{k}] = y[\mathbf{k}] + n[\mathbf{k}], \quad (3.6)$$

where $\{g[\mathbf{k}]\}$ are the measurements, $y[\mathbf{k}] = (h \star f)(\mathbf{k})$, $\mathbf{k} \in \mathbb{Z}^d$, are the generalized samples (samples of the prefiltered signal) and $n[\mathbf{k}]$ is a discrete additive, zero-mean noise component.

Throughout this chapter, we assume that the analog prefilter $h(\mathbf{x})$, which represents the point spread function (PSF) of the nonideal acquisition device, is of either forms below:

(i) PSF defined over a continuum: $h(\mathbf{x}) \in L_1(\mathbb{R}^d)$,

(ii) Sampled PSF (or discrete filter): $h(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} h_d[\mathbf{k}] \delta(\mathbf{x} - \mathbf{k})$, where $h_d[\mathbf{k}] \in \ell_1(\mathbb{Z}^d)$.

A necessary condition for the sampling problem to be well defined is that the prefilter be BIBO (Bounded-Input Bounded-Output) stable,⁴ as implied by (i) or (ii). The PSF in (ii) is not rigorously in $L_1(\mathbb{R}^d)$; however, it can be shown that it satisfies all corresponding Young-type inequalities. Note that (ii) also includes the identity filter $h(\mathbf{x}) = \delta(\mathbf{x})$ (ideal sampling) as a special case.

We adopt a variational approach and formulate the reconstruction problem as an optimization task. The solution is obtained by minimizing an error criterion denoted by $J(g, u)$ which depends on the input measurements $\{g[\mathbf{k}]\}$ and the continuous-space reconstruction $u(\mathbf{x})$. The specification of the cost-functional $J(g, u)$ is based on the following two key points:

- (a) The reconstruction should be sufficiently constrained (e.g., smooth or slowly varying) to make up for the fact that we are missing information in between pixels and to counterbalance the effect of noise, and
- (b) The generalized samples corresponding to the reconstruction should be close to the given measurements to ensure some level of consistency.

Specifically, among all continuously-defined functions $u(\mathbf{x}) \in L_2(\mathbb{R}^d)$, we are seeking the optimal signal reconstruction,

$$f_{\text{rec}}(\mathbf{x}) = \arg \min_{u \in L_2(\mathbb{R}^d)} J(g, u), \quad (3.7)$$

where the cost functional $J(g, u)$ is given by

$$J(g, u) = \underbrace{\sum_{\mathbf{k} \in \mathbb{Z}^d} |g[\mathbf{k}] - (h \star u)(\mathbf{k})|^p}_{\text{Data-Fidelity Term}} + \lambda \underbrace{\Theta(\|L\{u\}\|_{L_2}^2)}_{\text{Regularization}}. \quad (3.8)$$

L is a suitable shift-invariant differential (or regularization) operator, λ a positive real number, and $\Theta(\bullet)$ is a convex increasing function. The L_2 -norm in the regularization term is a measure of the “roughness” of the reconstruction. Minimization of $J(g, u)$ therefore ensures sufficient smoothness in the reconstruction while the data-fidelity term constrains the generalized samples of the reconstruction to be “close” to the measurements in the ℓ_p -sense. The parameter $\lambda \geq 0$ controls the amount of regularization imposed on the reconstruction. We will discuss the issue of selecting an appropriate value for this parameter in Section 3.6 and in a more general context in Chapters 5 and 6.

⁴Correspondingly, we have $|\hat{h}(\boldsymbol{\omega})| < +\infty, \forall \boldsymbol{\omega} \in \mathbb{R}^d$.

The above cost functional reduces to the well known Tikhonov criterion when $p = 2$ and $\Theta(x) = x$. Moreover, we will show in Section 3.6 that the solution corresponding to the Tikhonov criterion is functionally equivalent to the ones obtained for the minimax estimation and the stochastic (or Wiener) formulation of the generalized sampling problem.

At this point, it is important to note that $J(g, u)$ is a hybrid criterion that has a discrete part—the data-fidelity term—and an analog one—the regularization functional—that imposes smoothness constraints on the continuous-domain solution. It is this latter term together with the extent of the search space⁵ that differentiates our problem from a more traditional deconvolution task which is usually entirely formulated in the discrete domain. Here, we are attempting to solve the deconvolution and interpolation problems simultaneously and hoping that the criterion will dictate an “optimal” discretization procedure.

We believe that the present cost-functional is the most general one that can be solved analytically in the continuous-space domain. It allows for a non-quadratic data-fidelity term in the spirit of Fu *et al.* [72] and Nikolova [73], but it excludes some popular non-quadratic regularization such as total-variation (TV) [74, 75], which are not mathematically tractable in the present continuous-space framework. We must admit that this restriction constitutes a limitation of our formulation, but it is also clear that the generalized sampling problem is more difficult than the classical deconvolution problem: We are not only trying to get the optimal solution at the sample locations, but also in-between pixels, which adds another level of ill-posedness.

As we shall see, it is the presence of the L_2 -norm of $L\{\bullet\}$ in the regularization term that makes the derivation of the continuous-space solution feasible. Indeed, we will show that the continuous-space solution is well-defined and that it belongs to an “optimal” subspace, $V(\varphi_{\text{opt}})$, that is shift-invariant and independent upon the input signal. Moreover, for $p = 2$, there is a direct solution that can be computed elegantly by simple digital filtering of the discrete measurements. This is illustrated in Figure 3.1 in the reconstruction setup where the measurements $\{g[\mathbf{k}]\}$ are first compensated for their nonideality using linear/nonlinear optimization techniques (that can be implemented efficiently via digital filtering); the reconstruction then is performed in the shift-invariant space generated by $\varphi = \varphi_{\text{opt}}$.

⁵We are optimizing the criterion over $L_2(\mathbb{R}^d)$, which is considerably larger than the subspace of bandlimited functions. To make an analogy, $L_2(\mathbb{R}^d)$ is to $V(\text{sinc})$ —or, equivalently, $\ell_2(\mathbb{Z}^d)$ —what the real numbers are to the integers.

3.3.2 Consistent Sampling in $V(\varphi)$

In the noise-free case, it is reasonable to decrease the weight of the regularization ($\lambda \rightarrow 0$) and to seek a signal reconstruction that is consistent with the measurements. This corresponds to the case where the samples of the continuous-space function are equal to the measurements (so that the data-fidelity term is zero). This is the approach to the generalized sampling problem that was developed initially in [51] for a signal reconstruction in some predefined shift-invariant space $V(\varphi)$ (see Section 2.3.2 in Chapter 2). We will review this solution here, keeping in mind that it is not necessarily optimal because of the restriction on the search spaces (that is, $V(\varphi)$ instead of $L_2(\mathbb{R}^d)$). Under suitable conditions [51, Theorem 1], the consistent signal reconstruction in $V(\varphi)$ is unique and is given by

$$f_{\text{con}}(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} (g * r_0)[\mathbf{k}] \varphi(\mathbf{x} - \mathbf{k}),$$

where r_0 is the digital restoration filter whose frequency response is [51]

$$R_0(e^{j\boldsymbol{\omega}}) = \frac{1}{\sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k})}. \quad (3.9)$$

Note that this filter corresponds to the convolution inverse of the sequence $\{(h * \varphi)(\mathbf{k})\}$. We check for the consistency of $(h * f_{\text{con}})(\mathbf{x})$ by sampling it at the integers:

$$(h * f_{\text{con}})(\mathbf{x})|_{\mathbf{x}=\mathbf{n} \in \mathbb{Z}^d} = \sum_{\mathbf{k} \in \mathbb{Z}^d} (g * r_0)[\mathbf{k}] (h * \varphi)(\mathbf{n} - \mathbf{k}) = g[\mathbf{n}].$$

This holds true because $(r_0 * (h * \varphi))[\mathbf{n}] = \delta[\mathbf{n}]$ and the interchange of the sum and integral is justified using Lebesgue's dominated convergence theorem in combination with Cauchy-Schwarz's inequality. The condition for the existence (and unicity) of the consistent sampling (CS) solution is that the denominator of (3.9) is non-vanishing; that is, $\forall \boldsymbol{\omega} \in [0, 2\pi)^d$,

$$\text{CS} : \left| \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \right| > 0, \quad (3.10)$$

which imposes a joint constraint on $\varphi(\mathbf{x})$ and $h(\mathbf{x})$. This guarantees the BIBO stability of the reconstruction filter r_0 . Indeed, the L_p -stability of φ (SR condition) implies that $\{(h * \varphi)(\mathbf{k})\} \in \ell_1(\mathbb{Z}^d)$ (see Remark 3.1 at the beginning of

this chapter). The results then follow from Wiener's lemma [76], which ensures that the inverse filter $\{r_0[\mathbf{k}]\} \in \ell_1(\mathbb{Z}^d)$. The argument also holds in the case where h is a sampled PSF.

3.3.3 Space of Admissible Solutions

We now go back to our initial minimization problem (3.7) with $\lambda > 0$. While searching for a global solution in $L_2(\mathbb{R}^d)$, we must exclude all potential candidates for which the cost is infinite. We therefore say that the function $u(\mathbf{x}) \in L_2(\mathbb{R}^d)$ is admissible with respect to the criterion (3.7) if and only if $J(g, u)$ is finite. In particular, this implies that

$$\|\mathbf{L}\{u\}\|_{L_2}^2 = \int_{\mathbf{x} \in \mathbb{R}^d} |\mathbf{L}\{u(\mathbf{x})\}|^2 d\mathbf{x} \leq \mu_0^2 < +\infty. \quad (3.11)$$

This together with the fact that we are looking for a solution with finite L_2 -norm gets translated into $u(\mathbf{x}) \in W_2^{\mathbf{L}}$, where

$$W_2^{\mathbf{L}} = \left\{ f(\mathbf{x}) : \int_{\mathbb{R}^d} |\hat{f}(\boldsymbol{\omega})|^2 (1 + |\hat{\mathbf{L}}(\boldsymbol{\omega})|^2) d\boldsymbol{\omega} < +\infty \right\}$$

is the generalized Sobolev space associated with the operator \mathbf{L} . Thus, the problem can be restated as

$$f_{\text{rec}}(\mathbf{x}) = \arg \min_{u \in W_2^{\mathbf{L}}} J(g, u). \quad (3.12)$$

Additionally, for the data-fidelity term in $J(g, u)$ to be finite (in the $p = 2$ case), the samples of the solution $f_{\text{rec}} \in W_2^{\mathbf{L}}$ should be well-defined in the ℓ_2 -sense (we examine the cases where $p \neq 2$ when we actually present the solution to the minimization problem). This is ensured provided that \mathbf{L} acts as a differential operator and enforces sufficient smoothness on the solution. In multiple dimensions, a classical choice for \mathbf{L} is the $\frac{\gamma}{2}$ -iterated multidimensional Laplacian operator ($\gamma > 0$), whose transfer function⁶ is $\hat{\mathbf{L}}(\boldsymbol{\omega}) = \|\boldsymbol{\omega}\|^\gamma$ [77–79]. This leads to the Sobolev space W_2^γ of order γ that contains finite energy functions whose Laplacians up to order $\frac{\gamma}{2}$ have finite L_2 -norm [70, 79]. Among other properties of W_2^γ , the one that is of interest to us is for real $\gamma > \frac{d}{2}$,

$$f(\mathbf{x}) \in W_2^\gamma \implies \{f(\mathbf{k})\}_{\mathbf{k} \in \mathbb{Z}^d} \in \ell_2(\mathbb{Z}^d).$$

⁶The transfer function of a linear differential operator is defined in the sense of distributions: $\mathbf{L}\{\psi\} \xrightarrow{\mathcal{F}} \hat{\mathbf{L}}(\boldsymbol{\omega}) \hat{\psi}(\boldsymbol{\omega})$, where ψ is a valid test function. For instance, in 1-D, $\frac{d}{dx} \xrightarrow{\mathcal{F}} j\omega$.

For a proof of this result (in 1-D), see [54, Appendix C].

In the present work, we propose to extend this result to a larger class of multidimensional differential operator L . To that end, we impose the following admissibility condition on L which guarantees the minimum required degree of smoothness.

Definition 3.2. L is an admissible multidimensional differential operator if and only if

$$\sum_{\mathbf{k} \in \mathbb{Z}^d} \frac{1}{1 + |\hat{L}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2} \leq C_0 < +\infty. \quad (3.13)$$

The above condition implicitly controls the minimal “growth” rate of $\hat{L}(\boldsymbol{\omega})$, because the above series converges only when $\hat{L}(\boldsymbol{\omega})$ grows faster than $\|\boldsymbol{\omega}\|^{\frac{d}{2}}$ which is the limit case. If the generalized differential operator L is admissible, then the following theorem ensures that the associated generalized Sobolev space W_2^L has the properties that we demand.

Theorem 3.1. Let W_2^L be the generalized Sobolev space associated with the admissible regularization operator L . Then

$$f(\mathbf{x}) \in W_2^L \implies \{f(\mathbf{k})\} \in \ell_2(\mathbb{Z}^d),$$

and the Poisson’s summation formula holds:

$$\sum_{\mathbf{k} \in \mathbb{Z}^d} f(\mathbf{k}) e^{-j\boldsymbol{\omega}^T \mathbf{k}} = \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{f}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \quad a.e.$$

Proof. The proof is given in Appendix B.1. ■

For the present context, we extend the scope of the above theorem to functions of type $(h \star f)(\mathbf{x})$ using the following proposition.

Corollary 3.1. If $f(\mathbf{x}) \in W_2^L$ and $h(\mathbf{x}) \in L_1(\mathbb{R}^d)$, then $(h \star f)(\mathbf{x}) \in W_2^L$ and $(h \star f)(\mathbf{k}) \in \ell_2(\mathbb{Z}^d)$.

Proof. Since $h(\mathbf{x}) \in L_1(\mathbb{R}^d)$, $|\hat{h}(\boldsymbol{\omega})| \leq \|h\|_{L_1} < +\infty$, and the result is obvious from the definition of W_2^L and the application of Theorem 3.1. The same holds for the sampled PSF case too. ■

It is also clear that the search for a suboptimal solution in some subspace $V(\varphi)$ only makes sense if $V(\varphi) \subset W_2^L$. This is equivalent to requiring that the following upper Riesz bound be finite:

$$B_{R_{L\varphi},2} = \text{ess sup}_{\boldsymbol{\omega} \in [0, 2\pi)^d} \sqrt{\sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{L}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2} < +\infty.$$

3.3.4 Solution to the Variational Problem

To minimize the cost functional $J(g, u)$, we first observe that the data-fidelity term depends on the generalized samples of the solution exclusively, which suggests a two-stage optimization strategy. We first deal with the regularization part of the problem. Let us denote by $u_0(\mathbf{x})$ the optimal solution to our problem. The main part of the argument will be the construction of a function $f_{\text{con}}(\mathbf{x})$ which belongs to some “optimal” shift-invariant space $V(\varphi_{\text{opt}})$, that is consistent with $u_0(\mathbf{x})$ in the sense that $\forall \mathbf{k} \in \mathbb{Z}^d$,

$$y[\mathbf{k}] = (h \star u_0)(\mathbf{k}) = (h \star f_{\text{con}})(\mathbf{k}),$$

and which has the least $\Theta(\|L\{\bullet\}\|_{L_2}^2)$ among all consistent functions in W_2^L (see Section 3.3.4.1). This leads us to conclude that there is a solution of our global optimization problem that belongs to the space $V(\varphi_{\text{opt}})$ (see Sections 3.3.4.2 and 3.3.4.3). Once the optimal space is known, we only need to determine the coefficients $\{c[\mathbf{k}]\}$ of the shift-invariant representation which yields the discretized version of the problem presented in Section 3.3.4.4. The quadratic case, which can be solved explicitly, is dealt with in Section 3.3.4.5.

3.3.4.1 Consistent, Shift-Invariant Solution

The optimal generating function for our problem will be denoted by $\varphi_{\text{opt}}(\mathbf{x})$. We also introduce a sufficient condition for optimality which will be justified by Theorem 3.2 below.

Definition 3.3. *Let $\varphi_{\text{opt}}(\mathbf{x})$ be a generating function that satisfies the **SR** and **CS** conditions for a given prefilter $h(x)$. Then, $\varphi_{\text{opt}}(\mathbf{x})$ is said to be optimal with respect to the problem (3.7) if there exists a sequence $\{q[\mathbf{k}]\} \in \ell_1(\mathbb{Z}^d)$ such that*

$$\text{Optimality Condition OC : } L^*L\{\varphi_{\text{opt}}(\mathbf{x})\} = \sum_{\mathbf{k} \in \mathbb{Z}^d} (q * \bar{q})[\mathbf{k}] \bar{h}(\mathbf{x} - \mathbf{k}), \quad (3.14)$$

where L^*L is the self-adjoint operator whose transfer function is $|\hat{L}(\boldsymbol{\omega})|^2$.

We now construct a function $f_{\text{con}}(\mathbf{x}) \in V(\varphi_{\text{opt}})$ that is consistent with the optimal solution $u_0(\mathbf{x})$ and therefore yields the same data-fidelity term. This can always be done, as stated below.

Proposition 3.3. *Let $\varphi_{\text{opt}}(\mathbf{x})$ be optimal as in Definition 3.3. Then, if the samples $\{y[\mathbf{k}] = (h \star u_0)(\mathbf{k})\} \in \ell_2(\mathbb{Z}^d)$, there exists a unique consistent function $f_{\text{con}}(\mathbf{x}) \in V(\varphi_{\text{opt}})$ such that $(h \star u_0)(\mathbf{k}) = (h \star f_{\text{con}})(\mathbf{k})$, $\forall \mathbf{k} \in \mathbb{Z}^d$. It is given by*

$$f_{\text{con}}(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} \underbrace{(y \star r_0)[\mathbf{k}]}_{c[\mathbf{k}]} \varphi_{\text{opt}}(\mathbf{x} - \mathbf{k}), \quad (3.15)$$

where the digital reconstruction filter r_0 is specified by (3.9) with $\varphi = \varphi_{\text{opt}}$. Moreover, we have the equivalence $\{(h \star u_0)(\mathbf{k})\} \in \ell_2(\mathbb{Z}^d) \iff f_{\text{con}}(\mathbf{x}) \in W_2^L$.

Proof. The first part of the proposition is a direct application of the consistent sampling solution presented in Section 3.3.2. Since $\{r_0[\mathbf{k}]\} \in \ell_1(\mathbb{Z}^d)$ (**CS** Condition) and $\{y[\mathbf{k}] = (h \star u_0)(\mathbf{k})\} \in \ell_2(\mathbb{Z}^d)$, we have $\{c[\mathbf{k}] = (y \star r_0)[\mathbf{k}]\} \in \ell_2(\mathbb{Z}^d)$ by Young's inequality. This implies that $\|f_{\text{con}}(\mathbf{x})\|_{L_2} < +\infty$, because φ_{opt} generates a Riesz basis. To further prove that $\|L\{f_{\text{con}}\}\|_{L_2} < +\infty$, we need to show that $L\{\varphi_{\text{opt}}\}$ has an upper Riesz bound: By Proposition 3.1, we have $(h \star \varphi_{\text{opt}})(\mathbf{k}) \in \ell_1(\mathbb{Z}^d)$. This together with the fact that $|Q(e^{j\boldsymbol{\omega}})|$ is bounded from above yields

$$\begin{aligned} & \sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{L}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{\varphi}_{\text{opt}}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 \\ &= |Q(e^{j\boldsymbol{\omega}})|^2 \left(\sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{h}^*(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{\varphi}_{\text{opt}}^*(\boldsymbol{\omega} + 2\pi\mathbf{k}) \right) \\ &< +\infty, \end{aligned}$$

which ensures that $B_{R_{L\varphi_{\text{opt}}}, 2}$ is finite. We have thus established the forward implication: $\{(h \star u_0)(\mathbf{k})\} \in \ell_2(\mathbb{Z}^d) \implies f_{\text{con}}(\mathbf{x}) \in W_2^L$. The converse is a consequence of Theorem 3.1. \blacksquare

An interesting property that will be used later on is that the frequency response $R_0(e^{j\boldsymbol{\omega}})$ of the corresponding optimal reconstruction filter is always

strictly positive. This can be seen by writing (3.14) in the Fourier domain as:

$$\begin{aligned}\hat{\varphi}_{\text{opt}}(\boldsymbol{\omega}) &= \hat{h}^*(\boldsymbol{\omega}) \frac{|Q(e^{j\boldsymbol{\omega}})|^2}{|\hat{L}(\boldsymbol{\omega})|^2} \\ \implies \hat{h}(\boldsymbol{\omega}) \hat{\varphi}_{\text{opt}}(\boldsymbol{\omega}) &= |\hat{h}(\boldsymbol{\omega})|^2 \frac{|Q(e^{j\boldsymbol{\omega}})|^2}{|\hat{L}(\boldsymbol{\omega})|^2} \geq 0,\end{aligned}\quad (3.16)$$

which implies $\sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{\varphi}_{\text{opt}}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \geq 0$, $\forall \boldsymbol{\omega} \in [0, 2\pi)^d$. In fact, the latter is a strict inequality because of the **CS** condition.

3.3.4.2 Global Optimality of the Shift-Invariant Solution

Now that we have constructed the consistent function $f_{\text{con}}(\mathbf{x})$, it only remains to show that it is the one with the least $\Theta(\|\mathbb{L}\{\bullet\}\|_{L_2}^2)$ value. To do this, we establish a norm identity that decomposes $\|\mathbb{L}\{u\}\|_{L_2}^2$ into two orthogonal components.

Theorem 3.2. *Let $V(\varphi_{\text{opt}}) \subset W_2^L$ be the function space generated by φ_{opt} as in Definition 3.3. Then, the following orthogonality property holds:*

$$\forall u(\mathbf{x}) \in W_2^L, \quad \|\mathbb{L}\{u\}\|_{L_2}^2 = \|\mathbb{L}\{u - f_{\text{con}}\}\|_{L_2}^2 + \|\mathbb{L}\{f_{\text{con}}\}\|_{L_2}^2, \quad (3.17)$$

where $f_{\text{con}}(\mathbf{x})$ is the unique consistent function in $V(\varphi_{\text{opt}})$ such that

$$(h \star f_{\text{con}})(\mathbf{x})|_{\mathbf{x}=\mathbf{k}} = (h \star u)(\mathbf{k}), \quad \forall \mathbf{k} \in \mathbb{Z}^d,$$

as specified in Proposition 3.3.

Proof. The proof is given in Appendix B.2. ■

Setting $u = f_{\text{con}}$ in (3.17) minimizes the function norm $\|\mathbb{L}\{\bullet\}\|_{L_2}$. Thus, $f_{\text{con}}(\mathbf{x}) \in V(\varphi_{\text{opt}})$ minimizes $\|\mathbb{L}\{u\}\|_{L_2}^2$ among all functions in W_2^L having consistent samples.

3.3.4.3 General Solution

Putting the pieces together, we get our main theorem, which states that the solution to our optimization problem has a unique representation in the integer-shift-invariant space $V(\varphi_{\text{opt}})$.

Theorem 3.3. *Given the measurements $\{g[\mathbf{k}]\} \in \ell_p(\mathbb{Z}^d)$ and the “optimal” subspace of W_2^L , $V(\varphi_{\text{opt}}) = \text{span}\{\varphi_{\text{opt}}(\mathbf{x} - \mathbf{k})\}_{\mathbf{k} \in \mathbb{Z}^d}$ with φ_{opt} as in Definition 3.3, one has the following problem equivalence*

$$\min_{u \in W_2^L} J(g, u) \equiv \min_{f \in V(\varphi_{\text{opt}})} J(g, f), \quad (3.18)$$

for any general cost function of the form (3.8). This implies that the optimal signal reconstruction can be written as

$$f_{\text{rec}}(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] \varphi_{\text{opt}}(\mathbf{x} - \mathbf{k}),$$

and the optimization performed over the discrete set of coefficients $\{c[\mathbf{k}]\} \in \ell_2(\mathbb{Z}^d)$, which narrows down the search considerably.

Proof. We first concentrate on the case $p \geq 2$. Since we are looking for a reconstruction $u(\mathbf{x}) \in W_2^L$, the generalized samples $\{(h \star u)(\mathbf{k})\} \in \ell_2(\mathbb{Z}^d) \subset \ell_p(\mathbb{Z}^d)$ and the cost functional $J(g, u)$ is always well defined. The cost (3.8) is convex over W_2^L because it is the sum of two (strictly) convex sub-functionals. We are therefore guaranteed that the left hand side of the problem in (3.18) has a global minimum associated with the solution $u_0(\mathbf{x})$. The corresponding consistent reconstruction in $V(\varphi_{\text{opt}})$ is denoted by $f_{\text{con},0}(\mathbf{x})$ and is such that $(h \star u_0)(\mathbf{k}) = (h \star f_{\text{con},0})(\mathbf{k})$, $\forall \mathbf{k} \in \mathbb{Z}^d$. Applying Theorem 3.2 to the regularization part of the criterion, we have

$$\begin{aligned} J(g, u) &= \sum_{\mathbf{k} \in \mathbb{Z}^d} |g[\mathbf{k}] - (h \star u)(\mathbf{k})|^p \\ &\quad + \lambda \Theta (\|L\{f_{\text{con},0}\}\|_{L_2}^2 + \|L\{u - f_{\text{con},0}\}\|_{L_2}^2), \end{aligned} \quad (3.19)$$

where the data-fidelity term and the consistent function $f_{\text{con},0} \in V(\varphi_{\text{opt}})$ are fixed and uniquely tied to u . The optimality of u_0 implies that $J(g, u_0) \leq J(g, f_{\text{con},0})$. On the other hand, the comparison of (3.19) for $u = u_0$ and $u = f_{\text{con},0}$ indicates that $J(g, u_0) \geq J(g, f_{\text{con},0})$ because the data-fidelity term is the same in both cases and the function Θ is convex increasing. Thus, the conclusion is that $J(g, u_0) = J(g, f_{\text{con},0})$, which proves our assertion.

For $p < 2$, the situation is more restricted because the cost $J(g, u)$ explodes if the samples $\{(h \star u_0)(\mathbf{k})\} \in \ell_2(\mathbb{Z}^d) \setminus \ell_p(\mathbb{Z}^d)$. The minimization of $J(g, u)$ therefore automatically confines the samples $\{(h \star u_0)(\mathbf{k})\}$ to lie in $\ell_p(\mathbb{Z}^d)$. Since $\ell_p(\mathbb{Z}^d) \subset \ell_2(\mathbb{Z}^d)$ for $1 \leq p < 2$, the shift-invariant solution $f_{\text{con},0}(\mathbf{x}) \in V(\varphi_{\text{opt}})$ is still valid by Proposition 3.3 and Theorem 3.2. ■

The above result is conceptually pleasing because the continuous-space optimization problem (3.12) does not make any *a priori* assumption on the form of the reconstruction. The shift-invariant structure of the solution comes out as a result of the mathematical optimization. The generator of the optimal reconstruction space is specified by the optimality condition **OC**; that is, the operator L determines the generator $\varphi_{\text{opt}}(\mathbf{x})$. This simply means that the reconstruction space should be “matched” to the regularization operator.

Note that the fundamental solution of (3.14)—that is, $q[\mathbf{k}] = \delta[\mathbf{k}]$ —is $\varphi_{\text{opt}}(\mathbf{x}) = (\bar{h} \star \rho)(\mathbf{x})$, where $\rho(\mathbf{x})$ is a Green’s function of the self-adjoint operator L^*L , and this generator is generally not bandlimited.

3.3.4.4 Optimal Discretization of the Problem

Once the reconstruction space $V(\varphi_{\text{opt}})$ is specified, we only need to search for the expansion coefficients $\{c[\mathbf{k}]\} \in \ell_2(\mathbb{Z}^d)$ of the solution. To do this, we write $(h \star f_{\text{con},0})(\mathbf{k}) = (c \star \xi)(\mathbf{k})$, where $\xi[\mathbf{k}] = (h \star \varphi_{\text{opt}})(\mathbf{k})$. Using (3.14) for the second term, the cost can be rewritten in terms of the signal coefficients $\{c[\mathbf{k}]\}$:

$$\begin{aligned} J(g, f_{\text{rec}}) &= J(g; c) \\ &= \|g - (c \star \xi)\|_{\ell_p}^p + \lambda \Theta(\langle q \star \bar{q} \star c, \xi \star c \rangle_{\ell_2}), \end{aligned} \quad (3.20)$$

where we have used the fact that

$$\|L\{f_{\text{rec}}\}\|_{L_2}^2 = \langle L^*L\{f_{\text{rec}}\}, f_{\text{rec}} \rangle_{L_2} = \langle q \star \bar{q} \star c, \xi \star c \rangle_{\ell_2}.$$

We are therefore faced with a nonlinear optimization problem. Even though the problem does generally not have an explicit analytical solution, the good news is that $J(g; c)$ is a convex function of the coefficients c , which ensures that any local minimum automatically yields a global solution. The minimization can therefore be done by using any standard gradient-based nonlinear optimization technique [80].

3.3.4.5 Solution of the Quadratic/Tikhonov Problem

When $p = 2$ and $\Theta(x) = x$, the cost functional (3.20) is quadratic in $\{c[\mathbf{k}]\}$ (Tikhonov criterion [35]) and the derivation of the optimal solution can be carried out analytically. This yields an efficient digital filtering reconstruction algorithm.

Corollary 3.2. *When $p = 2$, $\Theta(x) = x$ and $\{g[\mathbf{k}]\} \in \ell_2(\mathbb{Z}^d)$, the global minimum of the cost functional $J(g, u)$ is achieved by*

$$f_{\text{rec}}(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} (r_\lambda * g)[\mathbf{k}] \varphi_{\text{opt}}(\mathbf{x} - \mathbf{k}), \quad (3.21)$$

where the frequency response of the optimal restoration (digital correction) filter r_λ is given by

$$R_\lambda(e^{j\boldsymbol{\omega}}) = \frac{1}{\sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{\varphi}_{\text{opt}}(\boldsymbol{\omega} + 2\pi\mathbf{k}) + \lambda |Q(e^{j\boldsymbol{\omega}})|^2}. \quad (3.22)$$

Proof. Setting $p = 2$ and $\Theta(x) = x$ in (3.20) and equating the partial derivative of $J(g; c)$ with respect to $c[\mathbf{k}]$ to zero yields,

$$(\bar{\xi} * \xi * c)[\mathbf{k}] + \lambda(\bar{\xi} * q * \bar{q} * c)[\mathbf{k}] = (\bar{\xi} * g)[\mathbf{k}].$$

Rewriting the above equality in the Fourier domain, we obtain the desired result:

$$C(e^{j\boldsymbol{\omega}}) = \frac{G(e^{j\boldsymbol{\omega}})}{\sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{\varphi}_{\text{opt}}(\boldsymbol{\omega} + 2\pi\mathbf{k}) + \lambda |Q(e^{j\boldsymbol{\omega}})|^2} = R_\lambda(e^{j\boldsymbol{\omega}}) G(e^{j\boldsymbol{\omega}}),$$

where $G(e^{j\boldsymbol{\omega}})$ is the Fourier transform of the given samples $\{g[\mathbf{k}]\}$. ■

Since $R_0(e^{j\boldsymbol{\omega}})$ is strictly positive and because $Q(e^{j\boldsymbol{\omega}})$ is bounded, the filter $R_\lambda(e^{j\boldsymbol{\omega}})$ is strictly positive and bounded as well. Therefore, by writing the solution (3.21) in the Fourier domain, we obtain

$$\hat{f}_{\text{rec}}(\boldsymbol{\omega}) = \hat{\varphi}_{\text{eq}}(\boldsymbol{\omega}) G(e^{j\boldsymbol{\omega}}), \quad (3.23)$$

where $\hat{\varphi}_{\text{rec}}(\boldsymbol{\omega}) = R_\lambda(e^{j\boldsymbol{\omega}}) \hat{\varphi}_{\text{opt}}(\boldsymbol{\omega})$ is the equivalent basis function (see Remark 3.2) that needs to be applied to the measurements $G(e^{j\boldsymbol{\omega}})$ to produce the continuous-space signal reconstruction, as illustrated in Figure 3.2. Indeed, we are ensured that $\varphi_{\text{rec}}(\mathbf{x})$ generates a L_p -stable Riesz basis because of the boundedness of $R_\lambda(e^{j\boldsymbol{\omega}})$ (see Proposition 3.2).

3.4 Stability of the Reconstruction

In this section, we examine φ_{opt} from a spline-theoretic point of view and provide guidelines for selecting L so that the problem is well posed. To this end, we




Figure 3.2: Reconstruction setup for the case of $p = 2$: Linear Filtering of measurements $\{g[\mathbf{k}]\}$.

rewrite the optimality condition (3.16) in the Fourier domain as $\hat{\varphi}_{\text{opt}}(\boldsymbol{\omega}) = \hat{h}^*(\boldsymbol{\omega}) |\hat{\beta}(\boldsymbol{\omega})|^2$, where

$$\hat{\beta}(\boldsymbol{\omega}) = \frac{Q(e^{j\boldsymbol{\omega}})}{\hat{L}(\boldsymbol{\omega})} \quad (3.24)$$

is interpreted as the Fourier transform of the generalized B-spline associated with the operator L [70, 81]. The construction of a B-spline is typically achieved by selecting a filter $Q(e^{j\boldsymbol{\omega}})$ that cancels the zeros of $\hat{L}(\boldsymbol{\omega})$ in order to produce a frequency response $\hat{\beta}(\boldsymbol{\omega})$ that is bounded. In the sequel, we will only consider “spline admissible” operators for which the corresponding generalized B-spline generates an L_p -stable Riesz basis.

Definition 3.4. *The operator L is said to be spline-admissible with B-spline $\beta(\mathbf{x})$ if and only if*

$$L\{\beta\}(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} q[\mathbf{k}] \delta(\mathbf{x} - \mathbf{k}), \quad (3.25)$$

where $\{q[\mathbf{k}]\} \in \ell_1(\mathbb{Z}^d)$, and $\beta(\mathbf{x})$ satisfies the **SR** condition.

Note that the above equation is just a restatement of (3.24) in the signal domain. The non-trivial aspect here is the existence of the lower Riesz bound $A_{R,\beta,2}$, which needs to be checked on a case-by-case basis.

The optimal generating function can therefore be regarded a compound generalized B-spline given by

$$\varphi_{\text{opt}}(\mathbf{x}) = (\bar{h} \star \beta \star \bar{\beta})(\mathbf{x}). \quad (3.26)$$

We will now see that this interpretation of φ_{opt} greatly simplifies our task of making sure that the conditions **SR** and **CS** for the well-posedness of our reconstruction problem are met. Interestingly, both conditions are tightly linked when we are working with the optimal basis.

Proposition 3.4. *Let φ_{opt} be the optimal generator. If $h(\mathbf{x}) \in L_1(\mathbb{R}^d)$ and L is spline-admissible, then we have the following equivalence:*

φ_{opt} satisfies the **SR** conditions $\iff \varphi_{\text{opt}}$ satisfies the **CS** condition.

Proof. Let us start from the left. Since φ_{opt} satisfies the first **SR** condition, we have a strictly positive lower bound for the series

$$\sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{\varphi}_{\text{opt}}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 = \sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 |\hat{\beta}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^4 \geq A_{R, \varphi_{\text{opt}}, 2}^2 > 0,$$

$\forall \boldsymbol{\omega} \in [0, 2\pi)^d$. Since all individual terms are positive, continuous functions⁷ of $\boldsymbol{\omega}$ (because $h(\mathbf{x}) \in L_1(\mathbb{R}^d)$ and $\beta(\mathbf{x}) \in L_1(\mathbb{R}^d)$ since L is spline-admissible), this means that for each $\boldsymbol{\omega}_0 \in [0, 2\pi)^d$, there is at least one $\mathbf{k} \in \mathbb{Z}^d$ such that $\hat{h}(\boldsymbol{\omega}_0 + 2\pi\mathbf{k}) \hat{\beta}(\boldsymbol{\omega}_0 + 2\pi\mathbf{k}) \neq 0$. Thus, $\forall \boldsymbol{\omega} \in [0, 2\pi)^d$, the sum

$$\sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{\varphi}_{\text{opt}}(\boldsymbol{\omega} + 2\pi\mathbf{k}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 |\hat{\beta}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 > 0, \quad (3.27)$$

is strictly positive so that φ_{opt} satisfies the **CS** condition. For the converse implication, we see right away that the second **SR** condition is satisfied by φ_{opt} because of (3.26) and Proposition 3.1. Then, from Remark 3.1, we have $\{(h \star \varphi_{\text{opt}})(\mathbf{k})\} \in \ell_1(\mathbb{Z}^d)$. Therefore, the l.h.s. of (3.27) is continuous and since it is also strictly positive, there exists $m > 0$, such that

$$\sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 |\hat{\beta}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^4 \geq m > 0,$$

which proves that the first **SR** condition is met as well. ■

⁷If a function $f(\mathbf{x}) \in L_1(\mathbb{R}^d)$, then its Fourier transform $\hat{f}(\boldsymbol{\omega})$ is uniformly continuous. The same is true for sequences: if $\{c[\mathbf{k}]\} \in \ell_1(\mathbb{Z}^d)$, then $C(e^{j\boldsymbol{\omega}})$ is uniformly continuous.

The key to the above equivalence is the optimality of $\varphi_{\text{opt}}(\mathbf{x})$ which leads to series of positive terms in both the **SR** and the **CS** conditions. However, the question still remains as how to ensure the existence of at least one non-vanishing product $\hat{h}(\boldsymbol{\omega}_0 + 2\pi\mathbf{k})\hat{\beta}(\boldsymbol{\omega}_0 + 2\pi\mathbf{k})$ for each $\boldsymbol{\omega}_0 \in [0, 2\pi)^d$ and some $\mathbf{k} = \mathbf{k}(\boldsymbol{\omega}_0) \in \mathbb{Z}^d$. Unfortunately, this may be quite tedious to check directly. What we propose here is an easy alternative where the user only has to worry about the placement of the zeros of $\hat{L}(\boldsymbol{\omega})$ in relation to those of $\hat{h}(\boldsymbol{\omega})$.

Theorem 3.4. *Let the following be true:*

- (i) L be spline-admissible,
- (ii) $h(\mathbf{x}) \in L_1(\mathbb{R}^d)$ be non-pathological in the sense that there exists

$$A_{\mathbb{R}^n, 2}^2 = \inf_{\boldsymbol{\omega} \in [0, 2\pi)^d} \sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 > 0, \text{ and,} \quad (3.28)$$

- (iii) $\hat{L}(\boldsymbol{\omega})$ and $\hat{h}(\boldsymbol{\omega})$ have no common zeros.

Then, φ_{opt} satisfies **SR** and **CS** conditions.

Proof. First we note that $\hat{h}(\boldsymbol{\omega})$ and $\hat{\beta}(\boldsymbol{\omega})$ are continuous and bounded because $h(\mathbf{x}) \in L_1(\mathbb{R}^d)$ and L is spline-admissible. Suppose that $\hat{L}(\boldsymbol{\omega}_n) = 0$ for $\boldsymbol{\omega}_n \in \mathbb{R}^d$, $\mathbf{n} \in \mathcal{I}$ where $\mathcal{I} \subset \mathbb{Z}^d$ is a finite index set. These zeros must necessarily be cancelled by some corresponding zeros of $Q(e^{j\boldsymbol{\omega}})$ which is 2π -periodic. This implies that

$$\hat{\beta}(\boldsymbol{\omega}_n + 2\pi\mathbf{k}) = \begin{cases} \mathcal{C}_{\beta, \mathbf{n}}, & \mathbf{k} = \mathbf{0}, \\ \mathbf{0}, & \text{otherwise,} \end{cases}$$

where $\mathcal{C}_{\beta, \mathbf{n}} \neq 0$ is a real constant. Moreover, β satisfies the **SR** conditions because of the spline admissibility of L . Therefore, the only zeros of $\hat{\beta}(\boldsymbol{\omega})$ are $\{\boldsymbol{\omega}_n + 2\pi\mathbf{k}, \mathbf{k} \in \mathbb{Z}^d \setminus \{\mathbf{0}\}, \mathbf{n} \in \mathcal{I}\}$. Since $\hat{h}(\boldsymbol{\omega}_n) \neq 0$ ($\hat{L}(\boldsymbol{\omega})$ and $\hat{h}(\boldsymbol{\omega})$ have no common zeros), we see that

$$\sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{h}(\boldsymbol{\omega}_n + 2\pi\mathbf{k})|^2 |\hat{\beta}(\boldsymbol{\omega}_n + 2\pi\mathbf{k})|^2 = |\hat{h}(\boldsymbol{\omega}_n)|^2 \mathcal{C}_{\beta, \mathbf{n}}^2 > 0,$$

for all the zeros of $\hat{\beta}(\boldsymbol{\omega})$. Therefore, the only way $\hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k})\hat{\beta}(\boldsymbol{\omega} + 2\pi\mathbf{k})$ can vanish is when $\hat{h}(\boldsymbol{\omega})$ has a 2π -periodic zero, which cannot be the case because of (3.28). Thus, we are ensured that for each $\boldsymbol{\omega}_0 \in [0, 2\pi)^d$, there is at least one $\mathbf{k} \in \mathbb{Z}^d$ such that $\hat{h}(\boldsymbol{\omega}_0 + 2\pi\mathbf{k})\hat{\beta}(\boldsymbol{\omega}_0 + 2\pi\mathbf{k}) \neq 0$. Following the argument used in Proposition 3.4, this implies that **SR** and **CS** conditions are both satisfied by $\varphi_{\text{opt}}(\mathbf{x})$. ■

The placement of the zeros $\hat{L}(\boldsymbol{\omega})$ has some important effect on the reconstruction. Qualitatively, the regularization will be the least at the frequencies where $|\hat{L}(\boldsymbol{\omega})|$ is minimum. In the limit when $\hat{L}(\boldsymbol{\omega}_n) = 0$, the restoration filter simplifies to

$$\begin{aligned} R_\lambda(e^{j\boldsymbol{\omega}_n}) &= \frac{1}{\sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{h}(\boldsymbol{\omega}_n + 2\pi\mathbf{k})|^2 |\hat{\beta}(\boldsymbol{\omega}_n + 2\pi\mathbf{k})|^2} \\ &= \frac{1}{\hat{h}(\boldsymbol{\omega}_n)\hat{\varphi}_{\text{opt}}(\boldsymbol{\omega}_n)} < +\infty, \end{aligned} \quad (3.29)$$

because $Q(e^{j\boldsymbol{\omega}})|_{\boldsymbol{\omega}=\boldsymbol{\omega}_n} = 0$ and $\hat{\beta}(\boldsymbol{\omega}_n + 2\pi\mathbf{k}) = 0$, $\mathbf{k} \in \mathbb{Z}^d \setminus \{\mathbf{0}\}$, by construction. Interestingly, this is the same response as that of a classical (non-regularized) inverse filter.

In light of this observation and Theorem 3.4, it makes good sense to place the zeroes (or minimal values) of $\hat{L}(\boldsymbol{\omega})$, where $\hat{h}(\boldsymbol{\omega})$ takes its maximum (typically, $\boldsymbol{\omega} = 0$) and vice versa. We will now consider some concrete examples and specify families of regularization operators that are well suited for lowpass systems. We will then return to the issue of the selection of the “best” regularization operator in Section 3.6, where we present an alternative stochastic formulation of the generalized sampling problem.

3.5 Case Illustrations

In this section, we examine a few special cases of the generalized sampling setup. The key point is that the reconstruction space is derived from the regularization operator via the optimality condition **OC**. We illustrate this connection with concrete examples that are relevant to image processing.

3.5.1 Ideal versus Nonideal Sampling

The prefilter $h(\mathbf{x})$ in the generalized sampling setup in Figure 3.1 models the PSF of the acquisition device. The simplest, idealized case is $h(\mathbf{x}) = \delta(\mathbf{x})$, which corresponds to a perfect sampling of the signal. The problem described by (3.12) then reduces to the multidimensional version of the smoothing spline problem investigated in [78, 81, 82]. In that case, we can simply ignore $\hat{h}(\boldsymbol{\omega})$ in all formulas and invoke Theorem 3.4 to show that the problem is well posed, provided that the regularization operator L is spline-admissible.

Likewise, we can account for the physical effect of an optical system. In the paraxial approximation of optics, the system is shift-invariant with a general lowpass behavior; the optical PSF is in $L_1(\mathbb{R}^3)$ and does not exhibit structured set of zeros (2π -periodic). Hence, stability is usually not a problem.

By contrast, some man-made sensors, such as CCD cameras whose impulse responses are indicator functions, do exhibit zeroes on a regular grid in the Fourier domain; e.g., $h_{\text{CCD}}(x_1, x_2) = \text{rect}(x_1, x_2) \Rightarrow \hat{h}_{\text{CCD}}(\boldsymbol{\omega}) = 0$, $\boldsymbol{\omega} = 2\pi\mathbf{k}$, $\mathbf{k} \neq \mathbf{0}$. Fortunately, the gain at zero frequency is non-zero (lowpass behavior) so that condition (3.28) is generally satisfied.

3.5.2 Regularization Operators and Reconstruction Spaces

Since most natural images are predominantly lowpass, it is desirable to reconstruct the lower part of the spectral content with minimum distortion which can be achieved through a judicious placement of the zeros of $\hat{L}(\boldsymbol{\omega})$ near the origin (see Equation (3.29)). This strategy is also justified by the lowpass behavior of most PSFs and our desire to minimize instabilities by having $\hat{L}(\boldsymbol{\omega})$ small where $\hat{h}(\boldsymbol{\omega})$ is large, and vice versa. In what follows, we consider examples of multidimensional regularization operators that are associated with the most prominent families of spline functions: Tensor-product polynomial splines [48, 62], and polyharmonic splines [70, 83], the latter being the non-separable counterpart of the former. We also introduce a generalized class of isotropic operators that lead to a new brand of ‘‘Matérn’’ splines, the relevance of which will be further justified in Section 3.6.

3.5.2.1 Separable Operators

We first study the separable case, where the Fourier transform of the multidimensional operator $\hat{L}(\boldsymbol{\omega})$ can be decomposed into a product of simple monomials. The prototypical example is a succession of n_k -th order derivatives along each spatial coordinate leading to

$$\hat{L}_{\text{Sep}}(\boldsymbol{\omega}) = \prod_{k=1}^d (j\omega_k)^{n_k}, \quad n_k \geq 1. \quad (3.30)$$

Since $\hat{L}_{\text{Sep}}(\boldsymbol{\omega})$ has multiple zeros at $\boldsymbol{\omega} = \mathbf{0}$, we must choose $Q(e^{j\boldsymbol{\omega}})$ in (3.24) to cancel these out. The canonical choice is

$$Q_{\text{Sep}}(e^{j\boldsymbol{\omega}}) = \prod_{k=1}^d (1 - e^{-j\omega_k})^{n_k},$$

which yields a multidimensional B-spline whose Fourier transform is

$$\hat{\beta}_{\text{Sep}}(\boldsymbol{\omega}) = \prod_{k=1}^d \left(\frac{1 - e^{-j\omega_k}}{j\omega_k} \right)^{n_k} = \prod_{k=1}^d e^{-\frac{j\omega_k n_k}{2}} \left(\frac{2 \sin(\omega_k/2)}{\omega_k} \right)^{n_k}. \quad (3.31)$$

This expression is separable and can be inverted in a coordinate-wise fashion. The end result is a tensor-product B-spline:

$$\beta_{\text{Sep}}(\mathbf{x}) = \prod_{k=1}^d \beta^{n_k-1} \left(x_k - \frac{n_k}{2} \right), \quad (3.32)$$

where $\beta^{n-1}(x)$ is Schoenberg's symmetric polynomial B-spline of degree $n-1$ (or order n) [62]. Since the 1-D B-splines are compactly supported and generate 1-D L_p -stable Riesz bases, the same holds true for $\beta_{\text{Sep}}(\mathbf{x})$ in higher dimensions.

Let us now look at two concrete examples of sampling configurations. The first is $h(\mathbf{x}) = \delta(\mathbf{x})$ (ideal sampling) with a second derivative regularization $n_k = 2$ for all k . In that case,

$$\varphi_{\text{opt}}(\mathbf{x}) = (\beta_{\text{Sep}} \star \bar{\beta}_{\text{Sep}})(\mathbf{x}) = \prod_{k=1}^d \beta^3(x_k),$$

which is a tensor-product symmetrical cubic B-spline. Thus, we may rely on the present variational formalism to justify the use of cubic interpolation which is quite popular in image processing applications; this corresponds to the limiting case $\lambda \rightarrow 0$ where the data-fidelity term vanishes (perfect fit).

The second example is $h(\mathbf{x}) = \text{rect}(\mathbf{x})$ (CCD camera) with a first derivative regularization $n_k = 1$ for all k . In that case,

$$\varphi_{\text{opt}}(\mathbf{x}) = (\bar{h} \star \beta_{\text{Sep}} \star \bar{\beta}_{\text{Sep}})(\mathbf{x}) = \prod_{k=1}^d \beta^2(x_k),$$

which is a quadratic B-spline. Note that without the prefilter, we would have ended up with a linear spline solution which is the poor man's solution to the interpolation problem. As far as image analysis is concerned, we believe that a quadratic spline model is preferable, one of the reason being that it is better suited for the evaluation of image differentials because of its higher order continuity [84]. Thanks to the present formulation and the fact that most images are captured using a pixel-integration device, we can invoke the present variational argument to support the use of quadratic splines in imaging applications. We are not aware of any previous justification in that direction.

3.5.2.2 Laplace Operators

The prime example of a non-separable, isotropic operator is the $\frac{\gamma}{2}$ -iterated (or fractional) Laplacian whose Fourier transform is given by

$$\hat{L}_{\text{Lap}}(\boldsymbol{\omega}) = \|\boldsymbol{\omega}\|^\gamma, \quad \gamma > \frac{d}{2}. \quad (3.33)$$

For $\gamma = 2$, we recover the classical Laplacian which is a popular, local operator.

Similar to the 1-D separable case, the choice of an appropriate localization filter

$$Q_{\text{Lap}}(e^{j\boldsymbol{\omega}}) = \left\| 2 \sin\left(\frac{\boldsymbol{\omega}}{2}\right) \right\|^\gamma$$

that cancels out the multiple zeros at $\boldsymbol{\omega} = \mathbf{0}$ yields the polyharmonic B-spline function of order γ for $\beta(\mathbf{x})$ [70, 85]; these can be written in the Fourier domain as

$$\hat{\beta}_\gamma(\boldsymbol{\omega}) = \frac{\left\| 2 \sin\left(\frac{\boldsymbol{\omega}}{2}\right) \right\|^\gamma}{\|\boldsymbol{\omega}\|^\gamma}, \quad (3.34)$$

where $\|\boldsymbol{\omega}\| = \sqrt{\sum_{k=1}^d \omega_k^2}$ and $\|\sin(\boldsymbol{\omega})\| = \sqrt{\sum_{k=1}^d \sin^2(\omega_k)}$. The polyharmonic B-splines fulfill the **SR** condition as well as the convolution property:

$$\beta_{\gamma_1} \star \beta_{\gamma_2} = \beta_{\gamma_1 + \gamma_2}.$$

If we now consider the case of an ideal sampling device together with a $\frac{\gamma}{2}$ -iterated Laplacian regularizer, we end up with the optimal generator

$$\varphi_{\text{opt}}(\mathbf{x}) = (\beta_\gamma \star \bar{\beta}_\gamma)(\mathbf{x}) = \beta_{2\gamma}(\mathbf{x}),$$

which is a polyharmonic B-spline of order 2γ . The corresponding solution to the quadratic problem in Section 3.3.4.5 is then the equivalent of the polyharmonic smoothing spline estimator investigated by Tirosh *et al.* [78].

3.5.2.3 Generalized Isotropic Operators

We now introduce a generalization of the fractional Laplacian operator that has the important advantage of being associated with an extended family of stationary processes, including the so-called Matérn class [40, 86]. This stochastic link

will be made explicit in Sections 3.6.2 and 3.7 in connection with the MMSE (or Wiener) estimator. We note that there is also a stochastic interpretation of the polyharmonic spline estimator [69, 78], but it is much more intricate because it involves fractional Brownian motion fields which are non-stationary processes [69, 78]. The idea here is to keep the isotropy property by having $\hat{L}(\boldsymbol{\omega})$ be a function of $\|\boldsymbol{\omega}\|$, but to displace the zeros from the origin by considering an operator whose Fourier transform is

$$\hat{L}_{\text{Iso}}(\boldsymbol{\omega}) = \hat{L}_{\text{Iso}}(\|\boldsymbol{\omega}\|) = \prod_{m=1}^N (\alpha_m + \|\boldsymbol{\omega}\|^2)^{\gamma_m}, \quad (3.35)$$

with $\alpha_m > 0, \gamma_m > \frac{d}{2}$, $m = 1, 2, \dots, N$.

Since $\hat{L}_{\text{Iso}}(\boldsymbol{\omega})$ does not vanish, there is actually no need for localization. So we simply set $Q_{\text{Iso}}(e^{j\boldsymbol{\omega}}) = 1$ and define the generalized Matérn B-spline with parameters $\{(\alpha_m, \gamma_m)\}_{m=1,2,\dots,N}$ as

$$\hat{\beta}_{\text{Iso}}(\boldsymbol{\omega}) = \prod_{m=1}^N \hat{\beta}_{\alpha_m, \gamma_m}(\boldsymbol{\omega}),$$

where $\hat{\beta}_{\alpha_m, \gamma_m}(\boldsymbol{\omega}) = (\alpha_m + \|\boldsymbol{\omega}\|^2)^{-\gamma_m}$. Interestingly, we are able to compute the inverse Fourier transform of $\hat{\beta}_{\alpha_m, \gamma_m}(\boldsymbol{\omega})$ which is given by (see Appendix B.3)

$$\beta_{\alpha_m, \gamma_m}(\mathbf{x}) = \left(\frac{\|\mathbf{x}\|}{2\sqrt{\alpha_m}} \right)^{\nu_m} \frac{\mathcal{K}_{\nu_m}(\sqrt{\alpha_m} \|\mathbf{x}\|)}{2^{d-1} \sqrt{\pi^d} \Gamma(\nu_m + \frac{d}{2})}, \quad (3.36)$$

with $\nu_m = \gamma_m - d/2$; $\mathcal{K}_\nu(x)$ is the modified Bessel function of the second kind which rapidly decays (faster than polynomial decay) for increasing x and is positive for $\nu > 0$ and $x > 0$ [87]. This result constitutes the multidimensional extension of the 2-D formula given in [86].

Therefore, $\beta_{\alpha_m, \gamma_m}(\mathbf{x})$ is positive for $\gamma_m > d/2$, and consequently $\beta_{\text{Iso}}(\mathbf{x})$, which is a N -fold convolution product,

$$\beta_{\text{Iso}}(\mathbf{x}) = \beta_{\alpha_1, \gamma_1}(\mathbf{x}) \star \beta_{\alpha_2, \gamma_2}(\mathbf{x}) \star \dots \star \beta_{\alpha_N, \gamma_N}(\mathbf{x}), \quad (3.37)$$

is positive as well. Moreover, since $\hat{\beta}_{\text{Iso}}(\boldsymbol{\omega}) \in L_1(\mathbb{R}^d)$ [87], $\beta_{\text{Iso}}(\mathbf{x})$ is uniformly continuous. The following proposition then ensures that $\hat{L}_{\text{Iso}}(\boldsymbol{\omega})$ is spline-admissible.

Proposition 3.5. *For $\gamma_m > \frac{d}{2}$, $m = 1, 2, \dots, N$, the operator $\hat{L}_{\text{Iso}}(\boldsymbol{\omega})$ is spline-admissible; that is, the function $\beta_{\text{Iso}}(\mathbf{x})$ given by (3.37) satisfies the **SR** conditions.*

Proof. Consider the periodic function $F(\boldsymbol{\omega})$ defined as

$$\begin{aligned} F(\boldsymbol{\omega}) &= \sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{\beta}_{\text{Iso}}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 = \sum_{\mathbf{k} \in \mathbb{Z}^d} \prod_{m=1}^N \frac{1}{(\alpha_m + \|\boldsymbol{\omega} + 2\pi\mathbf{k}\|^2)^{2\gamma_m}} \\ &> \prod_{m=1}^N \frac{1}{(\alpha_m + 4\pi^2 d)^{2\gamma_m}} > 0, \end{aligned}$$

which involves strictly positive terms only. Thus, we have the lower Riesz bound inequality satisfied:

$$A_{\beta_{\text{Iso},2}} \geq \prod_{m=1}^N (\alpha_m + 4\pi^2 d)^{-\gamma_m} > 0.$$

To prove the second **SR** condition, we observe that $\beta_{\text{Iso}}(\mathbf{x}) > 0, \forall \mathbf{x} \in \mathbb{R}^d$. Therefore, the task boils down to showing that

$$\sum_{\mathbf{k} \in \mathbb{Z}^d} \beta_{\text{Iso}}(\mathbf{x} + \mathbf{k}) < +\infty, \forall \mathbf{x} \in [0, 1)^d.$$

We first prove that this condition is satisfied by each $\beta_{\alpha_m, \gamma_m}(\mathbf{x})$. Consider the series $\sum_{\mathbf{n} \in \mathbb{Z}^d} \hat{\beta}_{\alpha_m, \gamma_m}(2\pi\mathbf{n})$ which can be bounded from above

$$\begin{aligned} \sum_{\mathbf{n} \in \mathbb{Z}^d} \hat{\beta}_{\alpha_m, \gamma_m}(2\pi\mathbf{n}) &= \sum_{\mathbf{n} \in \mathbb{Z}^d} \frac{1}{(\alpha_m + \|2\pi\mathbf{n}\|^2)^{\gamma_m}} \\ &< \frac{1}{\alpha_m^{\gamma_m}} + \sum_{\mathbf{n} \in \mathbb{Z}^d \setminus \{\mathbf{0}\}} \frac{1}{\|2\pi\mathbf{n}\|^{2\gamma_m}} < +\infty, \end{aligned}$$

where the series on the r.h.s. is known to converge whenever $\gamma_m > \frac{d}{2}$. We then use Poisson's formula to obtain

$$\begin{aligned} \sum_{\mathbf{k} \in \mathbb{Z}^d} \beta_{\alpha_m, \gamma_m}(\mathbf{x} + \mathbf{k}) &= (2\pi)^d \sum_{\mathbf{n} \in \mathbb{Z}^d} \hat{\beta}_{\alpha_m, \gamma_m}(2\pi\mathbf{n}) e^{j2\pi\mathbf{n}^T \mathbf{x}} \\ &\leq (2\pi)^d \sum_{\mathbf{n} \in \mathbb{Z}^d} \hat{\beta}_{\alpha_m, \gamma_m}(2\pi\mathbf{n}) < +\infty. \end{aligned}$$

The result then follows from the repeated application of Proposition 3.1. ■

3.6 Stochastic Formulations

The subject matter of this chapter until now has been the theoretical analysis of a variational approach to the nonideal sampling problem in the presence of noise. The signal and measurements were treated as deterministic entities, and the solution to the problem was obtained by finding the global minimum of a cost functional. However, there are also alternative formulations for the situations where the measurements are of stochastic nature and where we have some *a priori* knowledge on the class of signals. The reconstruction problem can then be posed as an estimation problem where the continuous-space reconstruction is estimated from the given measurement so as to minimize the mean square error (MSE). Following [66] we consider two formulations. In the first, the signal to be reconstructed is treated as deterministic, while the additive noise is considered to be of stochastic nature. In this case, we extend the results of [66] to multiple dimensions and perform a minimax estimation to obtain the reconstruction. In the second case, which we call the Wiener formulation, both the signal (continuous) and the noise (discrete) are modeled as stationary processes. The solution is obtained by minimizing the MSE of the reconstruction at each point in \mathbb{R}^d . Interestingly, this brings out a direct connection with the variational problem for the $p = 2$ case (the Tikhonov criterion) and leads to identical reconstruction algorithms for some particular choice of L and λ .

3.6.1 Connection with the Minimax Estimator

When the signal is deterministic and the noise is stationary, the reconstruction problem can be posed as a minimax estimation problem [66]. In this case, we assume that the signal $s(\mathbf{x}) \in W_2^L$ with $\|L\{s\}\|_{L_2} \leq \mu_0$ where the upper bound expresses our *a priori* knowledge on the class of input signals. The additive noise is modeled as a discrete, zero-mean stationary process with power spectral density (PSD) $C_n(e^{j\omega})$.

The criterion we minimize is the worst case projected MSE which eliminates explicit dependence of the solution on the unknown signal [66]. The projection is made onto some reconstruction space $V(\varphi)$ generated by $\varphi(\mathbf{x})$ which is decided *a priori* (and not necessarily optimal). In that framework, the minimization of the criterion over the signal coefficients $\{c[\mathbf{k}]\}$ yields a digital correction filter $r_{\text{MX}}[\mathbf{k}]$ which, when applied on the measurements together with the basis function, gives the reconstruction, as illustrated in Figure 3.2. Thus, the reconstruction

at $\mathbf{x}_0 \in \mathbb{R}^d$ is obtained as

$$f_{\text{rec}}(\mathbf{x}_0) = \arg \min_{s \in V(\varphi)} \max_{\|L\{f\}\|_{L_2} \leq \mu_0} E\{|P_{V(\varphi)}\{f\}(\mathbf{x}_0) - s(\mathbf{x}_0)|^2\}, \quad (3.38)$$

where $E\{\bullet\}$ represents the expectation operation with respect to the random vector corresponding to the noise and $P_{V(\varphi)}\{f\}(\mathbf{x})$ is the orthogonal projection of the deterministic input signal $f(\mathbf{x})$ onto the shift-invariant space $V(\varphi)$. The minimax reconstruction in $V(\varphi)$ is then given by

$$f_{\text{rec}}(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} (r_{\text{MX}} * g)[\mathbf{k}] \varphi(\mathbf{x} - \mathbf{k}). \quad (3.39)$$

By extending the argumentation of [66] to multiple dimensions, we can derive the frequency response of digital correction filter $r_{\text{MX}}[\mathbf{k}]$ that minimizes the minimax criterion (3.38); it is given by

$$R_{\text{MX}}(e^{j\boldsymbol{\omega}}) = \frac{\left(\sum_{\mathbf{k} \in \mathbb{Z}^d} \frac{\hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k})\hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k})}{|\hat{L}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2} \right)}{\left(\sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 \right) \left(\frac{C_n(e^{j\boldsymbol{\omega}})}{\mu_0^2} + \sum_{\mathbf{k} \in \mathbb{Z}^d} \frac{|\hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2}{|\hat{L}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2} \right)}. \quad (3.40)$$

The key point for our purpose is that, in the case of additive white noise, the minimax and Tikhonov reconstruction filters can be made rigorously equivalent, as stated in the following proposition.

Proposition 3.6. *The reconstruction filters for the above minimax estimation problem and the Tikhonov problem in Section 3.3.4.5 are identical provided that:*

- (1) *The discrete additive noise is white; that is, $C_n(e^{j\boldsymbol{\omega}}) = \sigma^2$,*
- (2) *The generator $\varphi(\mathbf{x})$ for the minimax method satisfies the optimality condition **OC**; that is, $\varphi(\mathbf{x}) = \varphi_{\text{opt}}(\mathbf{x})$,*

- (3) *The regularization parameter is set to its optimal value $\lambda_{\text{eq}} = \frac{\sigma^2}{\mu_0^2}$.*

Proof. By substituting $C_n(e^{j\boldsymbol{\omega}}) = \sigma^2$ and $\hat{\varphi}(\boldsymbol{\omega}) = \hat{\varphi}_{\text{opt}}(\boldsymbol{\omega}) = \hat{h}^*(\boldsymbol{\omega})\hat{\phi}(\boldsymbol{\omega})$ where $\hat{\phi}(\boldsymbol{\omega}) = \frac{|Q(e^{j\boldsymbol{\omega}})|^2}{|\hat{L}(\boldsymbol{\omega})|^2}$ and multiplying and dividing by $|Q(e^{j\boldsymbol{\omega}})|^2$ in the r.h.s. of

(3.40), we obtain

$$\begin{aligned}
 R_{\text{MX}}(e^{j\boldsymbol{\omega}}) &= \left(\sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{\varphi}_{\text{opt}}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{\phi}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \right) \\
 &\quad \times \left(\sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{\varphi}_{\text{opt}}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 \right)^{-1} \\
 &\quad \times \left(\frac{\sigma^2}{\mu_0^2} |Q(e^{j\boldsymbol{\omega}})|^2 + \sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 \hat{\phi}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \right)^{-1} \\
 &= \frac{1}{\left(\lambda_{\text{eq}} |Q(e^{j\boldsymbol{\omega}})|^2 + \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{\varphi}_{\text{opt}}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \right)} = R_{\lambda_{\text{eq}}}(e^{j\boldsymbol{\omega}}).
 \end{aligned}$$

Thus, (3.39) is equivalent to (3.21) for $\lambda_{\text{eq}} = \frac{\sigma^2}{\mu_0^2}$. ■

It is interesting to note that this equivalence only holds when we are considering the optimal generator φ_{opt} . Otherwise, the Tikhonov and minimax solution are generally different, as discussed in [66].

3.6.2 Unification with the Wiener Formulation

We now move one step further and consider that the input signal is a stochastic entity as well. Specifically, we assume that $s(\mathbf{x})$ is a realization of a continuous-space zero-mean stationary process with PSD $\hat{c}_s(\boldsymbol{\omega}) = \mathcal{F}\{c_f(\bullet)\}(\boldsymbol{\omega})$, where $c_f(\mathbf{x}) \in L_2(\mathbb{R}^d)$ is the corresponding autocovariance function. The measurements are the same as in (3.6), except that the additive noise is now modeled as a zero-mean discrete stationary process whose autocorrelation and PSD are denoted by $c_n[\mathbf{k}]$ and $C_n(e^{j\boldsymbol{\omega}})$, respectively. It is further assumed that the signal and noise are uncorrelated.

The solution in this formulation is obtained by minimizing the MSE of the reconstruction at any fixed location $\mathbf{x}_0 \in \mathbb{R}^d$,

$$\min E\{|f(\mathbf{x}_0) - f_{\text{rec}}(\mathbf{x}_0)|^2\}. \quad (3.41)$$

It should be noted that we minimize the MSE of the reconstruction and not the projected MSE, as done in [66]. Hence, our solution is the global minimum of

the MSE. Keeping in line with the previous reconstruction framework, we will see that the optimal solution belongs once again to an integer-shift-invariant space that is generated by some optimal function $\varphi_w(\mathbf{x})$ where the subscript stands for Wiener.

We want to emphasize the fact that the minimization of (3.41) is a hybrid version (discrete input - continuous output) of the standard Wiener problem, which is usually either stated in the discrete or purely continuous domain [88]. This said, the method of proof⁸ and solution are quite similar to the traditional ones and have already been deployed in the context of generalized sampling [89]. Therefore, in what follows, we simply present the hybrid Wiener solution (without proof) in a form that suits our notation.

Proposition 3.7. *Consider the measurements $\mathcal{G} = \{g[\mathbf{k}] = (h \star f)(\mathbf{k}) + n[\mathbf{k}]\}_{\mathbf{k} \in \mathbb{Z}^d}$, where $s(\mathbf{x})$ is a realization of a continuous-space stationary process with autocorrelation function $c_s(\mathbf{x}) \in L_2(\mathbb{R}^d)$ and $\{n[\mathbf{k}]\}$ is a discrete stationary noise component with PSD $C_n(e^{j\boldsymbol{\omega}})$. Then, the linear minimum mean square error (LMMSE) estimator of the signal $f(\mathbf{x}_0)$ given \mathcal{G} can be written as*

$$f_{\text{rec}}(\mathbf{x}_0) = \sum_{\mathbf{k} \in \mathbb{Z}^d} (r_w \star g)[\mathbf{k}] \varphi_w(\mathbf{x}_0 - \mathbf{k}), \quad (3.42)$$

where the optimal Wiener generator is $\varphi_w(\mathbf{x}) = (\bar{h} \star c_f)(\mathbf{x})$, and where the frequency response of the optimal restoration (digital correction) filter r_w is given by

$$R_w(e^{j\boldsymbol{\omega}}) = \frac{1}{C_n(e^{j\boldsymbol{\omega}}) + \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{\varphi}_w(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k})}. \quad \blacksquare \quad (3.43)$$

It is important to note that this LMMSE formula is valid for any $\mathbf{x}_0 \in \mathbb{R}^d$ and that the complete estimator $f_{\text{rec}}(\mathbf{x})$, $\mathbf{x} \in \mathbb{R}^d$ is included in the integer-shift-invariant subspace space $V(\varphi_w)$. It is also well known from estimation theory that the Wiener solution is optimal among all estimators (not just among the linear ones) when both the signal and noise are Gaussian distributed.

We observe that the PSD of the measurements is given by

$$C_g(e^{j\boldsymbol{\omega}}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{c}_s(\boldsymbol{\omega} + 2\pi\mathbf{k}) |\hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 + C_n(e^{j\boldsymbol{\omega}}),$$

⁸Here, we are interested in the linear minimum MSE estimate. The proof is then based on setting up a Yule-Walker-like system which when solved results in the integer-shift-invariant structure for the solution in Equation (3.42).

which coincides with the denominator of (3.43). Thus, a necessary and sufficient condition for the Wiener filter to be well defined (that is, bounded) is that this quantity be non-vanishing (a sufficient condition is $C_n(e^{j\boldsymbol{\omega}}) > 0$). Also, if we take the Fourier transform of (3.42), we see that the solution is of the same form as (3.23) where the Fourier transform of the equivalent basis function is given by

$$\hat{\varphi}_{w, \text{eq}}(\boldsymbol{\omega}) = \frac{\hat{\varphi}_w(\boldsymbol{\omega})}{C_g(e^{j\boldsymbol{\omega}})} = \frac{\hat{h}^*(\boldsymbol{\omega}) \hat{c}_s(\boldsymbol{\omega})}{C_g(e^{j\boldsymbol{\omega}})}. \quad (3.44)$$

Once again, the optimal reconstruction space is generally not bandlimited, unless either $h(\mathbf{x})$ or $s(\mathbf{x})$ are bandlimited to start with. Here too, the reconstruction filters for the deterministic and stochastic cases can be made equivalent. This equivalence helps us not only associate the optimal generating function to the autocorrelation function of the signal, but also allows us to choose the best regularization operator L and the best regularization parameter λ_{opt} for the Tikhonov solution in Corollary 3.2. Comparing the solutions (3.21) and (3.42) for the deterministic and stochastic cases respectively, we obtain the equivalence which corresponds to the case $Q(e^{j\boldsymbol{\omega}}) = 1$ and $C_n(e^{j\boldsymbol{\omega}}) = \sigma^2$; that is, the Tikhonov and Wiener reconstruction algorithms become equivalent when the discrete additive noise is white and the regularization operator is chosen such that it *whitens* the input signal.

Proposition 3.8. *The Tikhonov and Wiener solutions are functionally equivalent provided that:*

- (1) L is the whitening operator of the underlying continuous-space stochastic process; that is, $L^*L\{c_f(\mathbf{x})\} = \sigma_0^2 \delta(\mathbf{x})$,
- (2) The measurement noise is white with variance σ^2 , and
- (3) The regularization parameter is set to its optimal value $\lambda_{\text{opt}} = \frac{\sigma^2}{\sigma_0^2}$.

Proof. The proof is obvious once conditions (1) and (2) are substituted in equations (3.14) and (3.22). ■

We now conclude the section by illustrating this equivalence for the generalized isotropic operator introduced in Section 3.5.2.3. Consider a stationary processes (for instance, the Matérn process) whose PSD is given by

$$\hat{c}_s(\boldsymbol{\omega}) = \sigma_0^2 (\alpha^2 + \|\boldsymbol{\omega}\|^2)^{-2\gamma}. \quad (3.45)$$

It is readily seen that the corresponding whitening operator is

$$\hat{L}_{\text{Iso}}(\boldsymbol{\omega}) = (\alpha^2 + \|\boldsymbol{\omega}\|^2)^\gamma, \quad (3.46)$$

which is a special instance of (3.35) with $\alpha_m = \alpha^2$, $\gamma_m = \gamma$ and $N = 1$. Therefore the optimal estimate for such a process lies in the space generated by

$$\varphi_{\text{w}}(\mathbf{x}) = (\bar{h} \star \beta_{\text{Iso}} \star \bar{\beta}_{\text{Iso}})(\mathbf{x}),$$

and the optimal regularization parameter is given by $\lambda_{\text{opt}} = \frac{\sigma^2}{\sigma_0^2}$. In the case where $h(\mathbf{x}) = \delta(\mathbf{x})$ (ideal sampling), the optimal generator happens to be a Matérn B-spline with parameters $(\alpha, 2\gamma)$ whose explicit form is given by (3.36). Thus, the conclusion is that the Matérn splines are the optimal basis functions for approximation of such signals from their (noisy) samples. In the following section, we provide some experimental justification for the use of this model for sampling and interpolation of natural images.

Note that the equivalence established in the above corollary is meaningful only for operators of the type (3.35) and is not applicable to the Laplacian because the corresponding PSD is not defined in the classical sense.

3.7 Matérn Model for the Nonideal Sampling Problem

The Matérn class of functions—those of the type (3.36)—derive from the Matérn stochastic model that is of practical value in geostatistics for the prediction/estimation—also referred to as kriging—of spatial data [40]. The Matérn class can be viewed as a (fractional) generalization of the Markov Random Field model which is often used for modeling images [90]. Also relevant is the observation that many natural images exhibit a $\|\boldsymbol{\omega}\|^{-\tau}$ -like spectral decay [69], which is compatible with the Matérn PSD given by (3.45).

Our purpose here is to present further arguments in favor of the Matérn class and to demonstrate its usefulness for stochastic image modeling and processing. We will concentrate on the 2-D case and consider the generalized sampling problem that is summarized in Figure 3.2. Specifically, we assume that the signal $f(\mathbf{x})$ that we are sampling is a realization of a Matérn stochastic process (with PSD as given by Equation (3.45)). We also propose an extended version of the Matérn B-spline that can account for image anisotropies in 2-D—anisotropic Matérn B-spline β_{Ani} —and demonstrate its suitability for natural images. We

then compute the Wiener solution (3.42) keeping in mind that it is equivalent to the Tikhonov solution that is graphically portrayed in Figure 3.2 when the operator \mathbf{L} whitens the process $f(\mathbf{x})$. To make the method practical, we develop an estimation procedure for identifying the Matérn B-spline parameters (both of isotropic and anisotropic) and the noise variance from the noisy measurements $\{g[\mathbf{k}]\}$.

3.7.1 Anisotropic Matérn B-spline

Here, we introduce an anisotropic version of the Matérn B-spline in 2-D by applying a similarity transform (scaling and rotation) to the coordinate system; that is, we construct the new frequency variable

$$\tilde{\boldsymbol{\omega}} = \mathbf{R}_\theta \mathbf{K} \boldsymbol{\omega}, \quad (3.47)$$

where \mathbf{R}_θ is the matrix corresponding to the counter-clockwise rotation of the \mathbb{R}^2 -plane by θ and \mathbf{K} represents the scaling operation along the ω_2 axis by $(1 - \kappa)^{-1}$:

$$\mathbf{R}_\theta = \begin{bmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{bmatrix}, \quad \text{and} \quad \mathbf{K} = \begin{bmatrix} 1 & 0 \\ 0 & \left(\frac{1}{1-\kappa}\right) \end{bmatrix}, \quad (3.48)$$

where $0 \leq \kappa < 1$ is called as the ellipticity of the anisotropy. Then, the anisotropic Matérn B-spline is defined via its Fourier transform as

$$\hat{\beta}_{\text{Ani}}(\boldsymbol{\omega}) = \frac{\sigma_0^2}{(\alpha^2 + \omega_e^2(\boldsymbol{\omega}))^\gamma}, \quad (3.49)$$

where $\omega_e(\boldsymbol{\omega})$ is the distance from the origin to a point $\boldsymbol{\omega}$ whose locus is an ellipse: Starting from (3.47), we have

$$\begin{aligned} \omega_e^2(\boldsymbol{\omega}) &= \|\tilde{\boldsymbol{\omega}}\|^2 = \tilde{\boldsymbol{\omega}}^T \tilde{\boldsymbol{\omega}} = \boldsymbol{\omega}^T \mathbf{R}_{-\theta} \mathbf{K}^2 \mathbf{R}_\theta \boldsymbol{\omega} \\ &= \left(\cos^2(\theta) + \frac{\sin^2(\theta)}{(1-\kappa)^2} \right) \omega_1^2 + \left(\frac{\cos^2(\theta)}{(1-\kappa)^2} + \sin^2(\theta) \right) \omega_2^2 \\ &\quad + \sin(2\theta) \left(1 - \frac{1}{(1-\kappa)^2} \right) \omega_1 \omega_2. \end{aligned} \quad (3.50)$$

When $\kappa = 0$, there is no anisotropy and (3.49) reduces to the isotropic case

$$\hat{\beta}_{\text{Iso}}(\boldsymbol{\omega}) = \frac{\sigma_0^2}{(\alpha^2 + \|\boldsymbol{\omega}\|^2)^\gamma},$$

irrespective of the value of θ . Examples of isotropic and anisotropic Matérn B-splines are shown in Figures 3.3 and 3.4, respectively, for $\sigma_0 = 1$, $\alpha = 0.25$, $\gamma = 1.5$, $\theta = \frac{\pi}{4}$ and $\kappa = 0.75$. The anisotropic Matérn function, $\beta_{\text{Ani}}(\mathbf{x})$ also generates a L_p -stable Riesz basis as shown below.

Proposition 3.9. *For all $\gamma > 1$ (for $d = 2$), the function $\beta_{\text{Ani}}(\mathbf{x})$ satisfies both the **SR** conditions.*

Proof. We rewrite (3.50) as

$$\omega_e^2(\boldsymbol{\omega}) = (\omega_1 \cos(\theta) + \omega_2 \sin(\theta))^2 + \left(\frac{-\omega_1 \sin(\theta) + \omega_2 \cos(\theta)}{1 - \kappa} \right)^2. \quad (3.51)$$

Then, $\forall 0 \leq \kappa < 1$, we have $1 \leq (1 - \kappa)^{-1} < +\infty$ and so, we can bound $\omega_e^2(\boldsymbol{\omega})$ as

$$\|\boldsymbol{\omega}\|^2 \leq \omega_e^2(\boldsymbol{\omega}) \leq \frac{1}{(1 - \kappa)^2} \|\boldsymbol{\omega}\|^2, \quad (3.52)$$

where we made use of the fact that

$$\|\boldsymbol{\omega}\|^2 = (\omega_1 \cos(\theta) + \omega_2 \sin(\theta))^2 + (-\omega_1 \sin(\theta) + \omega_2 \cos(\theta))^2.$$

Correspondingly, for $\gamma > 1$, we have the following inequality from (3.52):

$$\frac{\sigma_0^2(1 - \kappa)^2}{(\alpha^2(1 - \kappa)^2 + \|\boldsymbol{\omega}\|^2)^\gamma} \leq \hat{\beta}_{\text{Ani}}(\boldsymbol{\omega}) \leq \frac{\sigma_0^2}{(\alpha^2 + \|\boldsymbol{\omega}\|^2)^\gamma}. \quad (3.53)$$

The result then follows from the application of Proposition 3.5 to the (isotropic) lower and upper bounds in (3.53). \blacksquare

3.7.2 Reconstruction using the Matérn Model

We propose to use the generating function $\varphi_w(\mathbf{x}) = (\bar{h} \star \beta_{\text{Ani}} \star \bar{\beta}_{\text{Ani}})(\mathbf{x})$ to specify the optimal reconstruction space which corresponds to the whitening operator with the Fourier transform

$$\hat{L}_{\text{Ani}}(\boldsymbol{\omega}) = (\alpha^2 + \omega_e(\boldsymbol{\omega}))^\gamma.$$

For illustration purposes, we consider a texture (Brodatz-D37) and a natural image (Aerial View). The gray dots in Figures 3.5 and 3.6 indicate the values




Figure 3.3: Plot of the logarithm of Isotropic Matérn B-spline in 2-D.




Figure 3.4: Plot of the logarithm of Anisotropic Matérn B-spline in 2-D.




Figure 3.5: Aerial View image: Approximately Isotropic Matérn B-spline, $\kappa = 0.098$, $\theta = 96.77^\circ$, $\sigma_0 = 56.604$, $\alpha = 0.021$, $\gamma = 1.410$.




Figure 3.6: D37 Brodatz texture: Anisotropic PSD, $\kappa = 0.407$, $\theta = 94.68^\circ$, $\sigma_0 = 46.811$, $\alpha = 0.083$, $\gamma = 1.320$.

of the periodogram of the noise-free images as a function of ω_e (see insets). The solid line in each figure is the generalized Matérn model as a function of ω_e fit to the periodogram of the image. It is clearly seen that the generalized Matérn model captures the trend of the periodograms while accounting for the anisotropy in D37 texture via the ellipticity parameter κ .

3.7.2.1 Estimation of Matérn parameters

We estimate the parameters in β_{Ani} viz-a-viz, ellipticity κ , angle θ and the Matérn parameters $(\sigma_0, \alpha, \gamma)$, and also the noise variance σ^2 by performing a non-linear least squares minimization of the following cost function,

$$\aleph(\sigma_0, \alpha, \gamma, \kappa, \theta, \tilde{\sigma}) = \sum_{\boldsymbol{\omega} \in (0, 2\pi)^2} \left[\log \left(\tilde{\Upsilon}(\boldsymbol{\omega}) + \tilde{\sigma}^2 \right) - \log \left(\hat{\Upsilon}_{np}(\boldsymbol{\omega}) \right) - \gamma_e \right]^2,$$

which measures the “distance” between the periodogram of the noisy measurements $\hat{\Upsilon}_{np}$ and that of the reconstruction (when it is fed back into the noisy sampling system) $\tilde{\Upsilon}(\boldsymbol{\omega})$ given by

$$\tilde{\Upsilon}(\boldsymbol{\omega}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 |\hat{\beta}_{\text{Ani}}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2.$$

The constant $\gamma_e \approx 0.57721$ is the Euler-Mascheroni number, which is used here to compensate for the bias due the log function.⁹ We also exclude $\boldsymbol{\omega} = 0$ in the summation to avoid spurious entries for the DC component in $\hat{\Upsilon}_{np}$.

3.7.2.2 Simulation Results

In our experiments, the reference scale for the signal and its reconstruction was $T = 1$. The signal was blurred with a standardized Gaussian prefilter and downsampled by a factor of two. The measurements were corrupted by various levels of white Gaussian noise.

The noise variance and the parameters of the anisotropic Matérn PSD were estimated from these noisy measurements, as described in Section 3.7.2.1. The reconstruction of the images at original resolution was performed according to the procedure outlined in Proposition 3.7 using $c_f = \beta_{\text{Ani}} \star \beta_{\text{Ani}}$ with the estimated parameters. This was compared with

⁹The random variable $\hat{\Upsilon}_{np}$ is χ -squared distributed when the noise is Gaussian. The details of the connection between χ -squared random variables and the Euler-Mascheroni number can be found in [78].




Figure 3.7: Performance of reconstruction methods: SNR Variation for Aerial view.

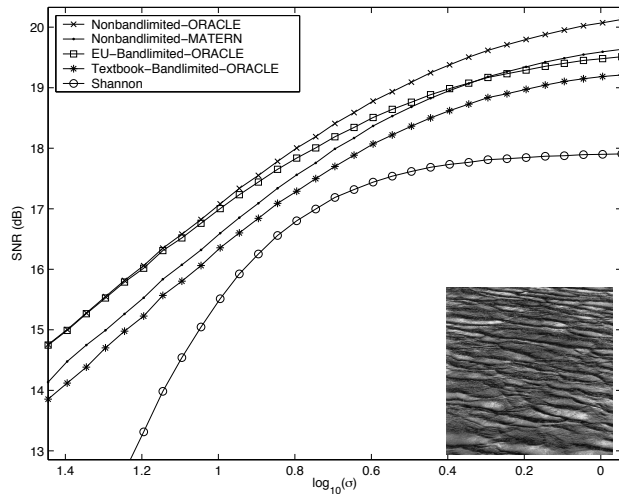


Figure 3.8: Performance of reconstruction methods: SNR Variation for D37 Brodatz texture.

- the multidimensional version of the optimally-compensated bandlimited reconstruction of Eldar *et al.* [66] (represented in Figures 3.7 and 3.8 as EU-bandlimited),
- the textbook implementation of the Wiener filter associated with a bandlimited version of the $\hat{\varphi}_w$ (that is, using $\text{rect}\left(\frac{\omega}{2\pi}\right) \hat{\varphi}_w(\omega)$), and
- sinc-interpolation without compensation (Shannon’s method).

Some of these are referred to as *oracle* solutions since they used the PSD of the noise-free image at the original resolution to derive the best possible reconstruction filter. Figures 3.7 and 3.8 show the signal-to-noise ratio (SNR) of the reconstruction for various noise levels at the sampling stage. We see that the proposed algorithm performs consistently better than the oracle textbook-bandlimited and sinc-interpolation methods. It also outperforms the oracle EU-bandlimited method at lower noise levels.

3.8 An Approximation Error Formula for the Noisy Scenario

In line with Section 2.4 of Chapter 2, we derive here a formula for the approximation error corresponding to the nonideal sampling problem. Specifically, we consider the setup of Figure 3.2, where we assume the additive noise component to be zero-mean with PSD $C_n(e^{j\omega})$ (the corresponding autocorrelation sequence is $\{c_n[\mathbf{k}]\}$). We also restrict ourselves to the case where the input signal $f(\mathbf{x})$ is a realization of a stationary stochastic process¹⁰ with PSD $\hat{c}_f(\omega)$ (the autocovariance function is $c_f(\mathbf{x})$). We also assume that f and the noise are mutually independent.

Similar to Equation (2.20) (see Chapter 2), we specify the approximation operator $\mathcal{Q}_{T\mathbf{I}}\{f\}$ corresponding to the setup of Figure 3.2 for the sampling lattice $\mathbf{D}\mathbb{Z}^d = T\mathbf{I}\mathbb{Z}^d$ (see Sections 2.2.2.1 and 2.4 of Chapter 2), where $T > 0$ is

¹⁰The case where the signal is deterministic and the noise is random is very difficult to treat technically and is not considered in this thesis: On close examination of the problem, it can be found that neither the L_2 -norm-based error (such as that considered in Section 2.4 of Chapter 2) that is appropriate for deterministic signals, nor the statistical mean-squared error measure (appropriate for random signals), nor a combination of both serve as a stable measure for quantifying the error of approximation for this case.

the sampling step, as follows:

$$f_{\text{rec}}(\mathbf{x}) = \mathcal{Q}_{T\mathbf{I}}\{f\}(\mathbf{x}) = \sum_{\mathbf{k}, \mathbf{m} \in \mathbb{Z}^d} g[\mathbf{m}] r[\mathbf{k} - \mathbf{m}] \varphi\left(\frac{\mathbf{x}}{T} - \mathbf{k}\right),$$

where $\left(\frac{\mathbf{x}}{T}\right) = \left(\frac{x_1}{T}, \frac{x_2}{T}, \dots, \frac{x_d}{T}\right)$, $\{g[\mathbf{k}]\}$ are the noisy measurements whose \mathbf{k} -th sample is

$$g[\mathbf{k}] = \int_{\mathbb{R}^d} f(\boldsymbol{\xi}) \frac{1}{T^d} h\left(\mathbf{k} - \frac{\boldsymbol{\xi}}{T}\right) d\boldsymbol{\xi} + n[\mathbf{k}],$$

$\frac{1}{T^d} h\left(\frac{\bullet}{T}\right)$ and $\varphi\left(\frac{\bullet}{T}\right)$ are the accordingly-scaled analog prefilter and reconstruction filter, respectively, and $\{n[\mathbf{k}]\}$ represents the additive noise component.

For quantifying the approximation error, we consider the statistical expectation of the “time”-averaged error [53, 66] given by

$$\varepsilon_{f, \text{noisy}}(T) = E \left\{ \lim_{\tau \rightarrow +\infty} \frac{1}{(2\tau)^d} \int_{[-\tau, \tau]^d} |f(\mathbf{x}) - \mathcal{Q}_{T\mathbf{I}}\{f\}(\mathbf{x})|^2 d\mathbf{x} \right\},$$

where the expectation is taken jointly over f and the noise. Since the integrand in the above equation is non-negative, we can move the expectation inside the integral (Lebesgue’s monotone-convergence theorem). Then, as $E\{|f(\mathbf{x}) - \mathcal{Q}_{T\mathbf{I}}\{f\}(\mathbf{x})|^2\}$ is T -periodic¹¹ (it is T -periodic along every dimension), $\varepsilon_{f, \text{noisy}}(T)$ reduces to [53]

$$\varepsilon_{f, \text{noisy}}(T) = \frac{1}{T^d} \int_{[0, T]^d} E \left\{ |f(\mathbf{x}) - \mathcal{Q}_{T\mathbf{I}}\{f\}(\mathbf{x})|^2 \right\} d\mathbf{x}.$$

The following result then yields an explicit formula for $\varepsilon_{f, \text{noisy}}(T)$ that can be computed in the Fourier domain similar to Equation (2.36) in Chapter 2.

Proposition 3.10. *Let $h(\mathbf{x}) \in L_1(\mathbb{R}^d)$, $\{r[\mathbf{k}]\} \in \ell_1(\mathbb{Z}^d)$ and let φ satisfy the **SR** conditions. Then, provided $c_f \in W_2^L$ and $\{c_n[\mathbf{k}]\} \in \ell_2(\mathbb{Z}^d)$, we have*

$$\varepsilon_{f, \text{noisy}}(T) = \varepsilon_f(T) + \varepsilon_{\text{noise}}(r, \varphi), \tag{3.54}$$

where

$$\varepsilon_f(T) = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \mathcal{E}_{h, \varphi}(\boldsymbol{\omega}T) \hat{c}_f(\boldsymbol{\omega}) d\boldsymbol{\omega}.$$

¹¹This can be easily verified by expanding $E\{|f(\mathbf{x}) - \mathcal{Q}_{T\mathbf{I}}\{f\}(\mathbf{x})|^2\}$.

$\mathcal{E}_{h,\varphi}(\boldsymbol{\omega})$ is the Fourier domain error kernel given by

$$\mathcal{E}_{h,\varphi}(\boldsymbol{\omega}) = 1 - \frac{|\hat{\varphi}(\boldsymbol{\omega})|^2}{A_\varphi(\boldsymbol{\omega})} + A_\varphi(\boldsymbol{\omega}) \left| R^*(e^{j\boldsymbol{\omega}}) \hat{h}^*(\boldsymbol{\omega}) - \hat{\varphi}(\boldsymbol{\omega}) \right|^2, \quad (3.55)$$

and

$$\varepsilon_{\text{noise}}(r, \varphi) = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} C_n(e^{j\boldsymbol{\omega}}) |R(e^{j\boldsymbol{\omega}})|^2 |\hat{\varphi}(\boldsymbol{\omega})|^2 d\boldsymbol{\omega},$$

where $R(e^{j\boldsymbol{\omega}})$ is the Fourier transform of the sequence $\{r[\mathbf{k}]\}$.

Proof. The proof is given in Appendix B.4. ■

We make the following observations about the error $\varepsilon_{f,\text{noisy}}(T)$:

- It is composed of two components: The first one, $\varepsilon_f(T)$, corresponds to the error due to approximation of f by f_{rec} when there is no noise; this is equivalent to the $\varepsilon_{\text{dom}}(f, T)$ component in (2.36) in Chapter 2 with $|\hat{f}(\boldsymbol{\omega})|^2$ replaced by $\hat{c}_f(\boldsymbol{\omega})$ (this is also the d -dimensional version of [53, Equation (18)]).
- The second term, $\varepsilon_{\text{noise}}(r, \varphi)$, is the contribution of the noise component exclusively; that is, it corresponds to the energy of reconstructing the noise only when there is no signal.

$\varepsilon_{\text{noise}}(r, \varphi)$ depends only on the choice of the digital correction filter r and the generating function φ and is independent of f and T which is quite expected since, in the current setting, the noise has “no knowledge” of f . Moreover, making the sampling-step size small has no noise-reducing effect. Therefore, as $T \rightarrow 0$, the first term will progressively disappear, while the second term becomes predominant.

- When the noise is white, $C_n(e^{j\boldsymbol{\omega}}) = \sigma^2$, and $\varepsilon_{\text{noise}}(r, \varphi)$ simply reduces to $\varepsilon_{\text{noise}}(r, \varphi) = \sigma^2 \|\varphi_{\text{eq}}\|_{L_2}^2$, where φ_{eq} is the equivalent basis function (see Figure 3.2). Therefore, for a given σ , $\varepsilon_{f,\text{noisy}}(T)$ represents a trade-off between the error of approximation of the noise-free signal and the (signal-independent) contribution due to noise.

When $\sigma \rightarrow 0$, we get back the result (d -dimensional version) corresponding the noise-free case of [53]: $\varepsilon_f(T)$ plays the dominant role and the approximation theoretic results of Chapter 2 (see Section 2.4) are applicable.

Conversely, for large σ , $\varepsilon_{\text{noise}}(r, \varphi)$ becomes the dominant term; then the reasoning using approximation order (see Section 2.5 in Chapter 2) may no longer be valid in this setting. We present an illustration in Figure 4.1 in Chapter 4 that highlights this observation.

The only way $\varepsilon_{\text{noise}}(r, \varphi)$ can be reduced is by using a digital correction filter r that makes $\|\varphi_{\text{eq}}\|_{L_2}^2$ small. This is precisely achieved by the regularized filter r_λ (with Fourier transform $R_\lambda(e^{j\omega})$ given by Equation (3.22)), as illustrated in Figures 3.9 and 3.10 (in the 1-D case) for $\lambda = 0$ and $\lambda = \frac{1}{2}$, respectively. The prefilter and the reconstruction filter are the same as those considered in case (a) of the illustration in Section 2.4 of Chapter 2: $h(x) = \text{rect}(x)$ and $\varphi(x) = \beta^1(x)$, the linear B-spline.

For $\lambda = 0$, r_λ is not regularized (see Figure 3.9(c) where $R_\lambda(e^{j\omega})$ enhances the higher frequencies predominantly) and the situation corresponds to consistent reconstruction or exact fitting of the noisy measurements. Correspondingly, the contribution of noise to the error is $\varepsilon_{\text{noise}}(r, \varphi) \propto \|\varphi_{\text{eq}}\|_{L_2}^2 \approx 297.224$; the shaded area under the curve in Figure 3.9(b) spans frequencies $|\omega| \leq \pi$. However, for $\lambda = \frac{1}{2}$, r_λ is sufficiently regularized (Figure 3.10(c) where $R_\lambda(e^{j\omega})$ has a low-pass behavior) and consequently, the spectrum of φ_{eq} has significantly shrunk (see Figure 3.10(b)); the corresponding value of $\|\varphi_{\text{eq}}\|_{L_2}^2$ is ≈ 52.408 which is lesser than in the case of $\lambda = 0$. This clearly illustrates the effectiveness of regularization in tackling the noise problem.

However, the error kernel $\mathcal{E}_{h,\varphi}$ is also dependent on r_λ : Selecting a value $\lambda > 0$ implies that we are altering the approximation characteristics of the reconstruction setup. This is seen in Figures 3.9(a) and 3.10(a): For the case of $\lambda = 0$, the error kernel $\mathcal{E}_{h,\varphi}$ is closer to zero for $|\omega| \leq \frac{\pi}{2}$ indicating that the (noise-free) approximation error— ε_f in (3.54)—is less for these frequencies, while for $\lambda = \frac{1}{2}$ this range has shrunk to $|\omega| \lesssim \frac{\pi}{4}$. Thus, for a stochastic signal whose PSD is predominantly concentrated in the low-pass region (*e.g.*, Matérn class of PSDs), it can be easily verified that the two components of $\varepsilon_{f,\text{noisy}}$ have conflicting behaviors; that is, an increase in the value of λ reduces $\varepsilon_{\text{noise}}(r, \varphi)$ but at the cost of increasing the (noise-free) approximation error ε_f and vice versa. This is precisely what we meant by the term trade-off earlier in this context. Therefore, λ must be selected carefully so as to reduce $\varepsilon_{\text{noise}}(r, \varphi)$ while not increasing ε_f too much: This observation is coherent with the argument in Section 3.3.1 that is in general applicable to many variational problems. We take up the problem of choosing an appropriate value for λ in Chapters 5 and 6.




Figure 3.9: Plots of (a) the error kernel $\mathcal{E}_{h,\varphi}(\omega)$, (b) $|\varphi_{\text{eq}}(\omega)|^2$, and (c) $R_\lambda(e^{j\omega})$, for $h(x) = \text{rect}(x)$, $\varphi(x) = \beta^1(x)$ and $\lambda = 0$. The shaded region in (b) denotes the area under $|\varphi_{\text{eq}}(\omega)|^2$ which is nothing but $\|\varphi_{\text{eq}}\|_{L_2}^2 \approx 297.224$.




Figure 3.10: Plots of (a) the error kernel $\mathcal{E}_{h,\varphi}(\omega)$, (b) $|\varphi_{\text{eq}}(\omega)|^2$, and (c) $R_\lambda(e^{j\omega})$, for $h(x) = \text{rect}(x)$, $\varphi(x) = \beta^1(x)$ and $\lambda = \frac{1}{2}$. The shaded region in (b) denotes the area under $|\varphi_{\text{eq}}(\omega)|^2$ which is nothing but $\|\varphi_{\text{eq}}\|_{L_2}^2 \approx 52.408$.

3.9 Summary

In this chapter, we formulated nonideal sampling in the presence of noise as a variational problem where the continuous-space solution is obtained by the minimization of a data-fidelity term subject to a continuous-space regularization constraint (Tikhonov-like functional) based on generalized differential operator L . In this formulation the signal and noise are treated as deterministic entities. By carrying out the mathematical optimization, we showed that the global minimum of the criterion lies in an integer-shift-invariant space whose generator φ_{opt} is matched to the operator L through an optimality condition (**OC**) that gives rise to a B-spline interpretation of $\varphi_{\text{opt}}(\mathbf{x})$ and ensures the L_p -stability of the representation.

Once a reconstruction space is specified, the expansion coefficients of the solution can be obtained by means of a non-linear optimization process. In the particular case where $p = 2$ and $\Theta(x) = x$ (Tikhonov criterion), the solution has an explicit analytical form and can be computed by a one-step hybrid-linear filtering (discrete in - continuous out) of the measurements.

The proposed formulation is quite general and extends most of the solutions to the sampling problem that have been proposed so far (see Chapter 2); it also encompasses the various classes of smoothing spline estimators [78,81]. We further justified this fact by presenting illustrations of various regularization operators and the corresponding reconstruction spaces.

Finally, we also presented a stochastic formulation of the generalized sampling problem, providing explicit minimax and LMMSE solutions. We showed that, for appropriate regularization operator L and the regularization parameter λ , the solution corresponding to the deterministic Tikhonov criterion is functionally equivalent to the ones obtained in the stochastic signal processing framework. The bottomline is that the regularization operator should be matched to the spectral behavior of the signal (whitening operator) and the regularization strength set inversely proportional to the signal-to-noise ratio.

In the stochastic context, we also developed an exact formula for the approximation error $\varepsilon_{f,\text{noisy}}$ for the noisy scenario and showed that it is composed of two components namely, the error of approximation of the noise-free signal ε_f and a signal-independent contribution due to the noise $\varepsilon_{\text{noise}}(r, \varphi)$ that depends only on the reconstruction setup. We also demonstrated that the use of regularization (r_λ for $\lambda > 0$) reduces $\varepsilon_{\text{noise}}(r, \varphi)$ but at the cost of increasing ε_f . This serves as another argument that strongly indicates the need for careful selection of the regularization parameter λ .

Chapter 4

Non-quadratic Regularization for Interpolation of Noisy Data

4.1 Introduction

Interpolation is the process of fitting discrete data samples with a continuously-defined model that is usually represented as a weighted sum of shifted basis functions [15, 19]. The standard approach is to specify the model so as to fit the data exactly [6]. In the context of sampling, interpolation is the reconstruction procedure that is employed when the acquisition device is ideal (see Section 2.3.3 in Chapter 2). When the basis functions of the interpolation model take a unit value at the origin and are zero at all other integer locations—an interpolator—the weights are given by the data samples themselves. Otherwise, they are determined by imposing an exact-fitting requirement [63] (see Equation 2.31 in Chapter 2). The standard approach, though meaningful in the noise-free scenario, is less appropriate when data samples are corrupted by noise since the model is forced to fit noise also.

To deal with the problem of noisy data,¹ we extend the notions that we developed in the previous chapter—variational approach and the use of regularization—

¹This chapter is based on the article: S. Ramani, P. Thévenaz and M. Unser, “Regularized Interpolation for Noisy Images,” submitted to IEEE Transactions on Medical Imaging.

for the problem of interpolation in the presence of noise. The resulting scheme is termed *regularized interpolation* where the objective is to obtain the solution by minimizing a variational criterion that jointly measures the data-fitting error and the regularity of the continuous-domain solution. Various researchers have formulated interpolation in the variational framework to accommodate regularization constraints [36–39, 66, 81, 91–95]. These works can be broadly classified into digital-domain or analog-domain approaches: The former refer to the case where the solution is a discrete entity defined on a grid that is finer than that of the data—these methods specifically cater to the image-upsampling problem [92–95]. In the latter case, a continuously-defined solution—typically, a smoothing-spline—is obtained by minimizing the L_2 -norm of some scalar derivative of the solution [36–39, 66, 81, 91] subject to certain data-fitting requirements. In fact, the latter corresponds to the special cases that we discussed in Section 3.5 of the previous chapter.

In this chapter, we propose a regularized interpolation scheme that extends the continuous approach so as to include a much wider-class of regularizations based on non-quadratic criteria—the motivation is to overcome the shortcoming of the Tikhonov-like quadratic regularization which is known for its over-smoothing behavior that smears important signal features such as edges in images.

4.1.1 Contributions

- Keeping in mind the desirable feature of affine invariance (that is, invariance to translation, rotation and scaling operations) of the solution, we show that the L_p -norm of an appropriate vector derivative is the most suitable choice of regularization for our problem. This includes the edge-preserving total-variation (TV) regularization ($p = 1$) which we propose to use in this work.

We consider a shift-invariant model for our regularized interpolation scheme and obtain the weights by minimizing the L_p -norm subject to a data-domain constraint that measures the statistical infidelity of the solution to the given data (in terms of a negative log-likelihood function).

- We handle the problem in a numerical optimization framework and present monotonically convergent algorithms based on the *majorize-minimize* (MM) strategy [96] for performing the corresponding minimization.

We validate our method by carrying out 2-D rotation experiments on noisy medical images and illustrate its superior performance over standard

interpolation. We also numerically justify that the use of TV-like non-quadratic regularization brings about further improvement over quadratic regularization.

4.1.2 Organization of the Chapter

For the sake of consistency, we briefly review the standard interpolation technique [63] (described in detail in Section 2.3.3 in Chapter 2) in Section 4.2 following which we present a numerical example in Section 4.2.1 to demonstrate its poor performance when applied to noisy data. In Section 4.3, we elaborate the proposed regularized interpolation scheme. We first specify the problem mathematically and provide hypotheses related to the data-fidelity and regularization terms. Then, we present algorithms to carry out the corresponding optimization based on an appropriate discretization of the problem. Section 4.4 is dedicated to experimental results where we numerically verify the superiority of the proposed approach over other methods in the literature. In Section 4.5, we provide evidence that the proposed scheme achieves a reasonable tradeoff between computational cost and performance improvement and also discuss issues related to the selection of the regularization parameter. Finally, we draw our conclusions in Section 4.6.

4.2 Standard Interpolation

Standard interpolation is the process of computing a continuously-defined function $f_{\text{int}}(\mathbf{x})$, which exactly fits an unknown analog signal $f(\mathbf{x})$ at the given sample points $\{f(\mathbf{k})\}_{\mathbf{k} \in \mathbb{Z}^d}$. Typically, f_{int} is constructed as

$$f_{\text{int}}(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} f(\mathbf{k}) \varphi_{\text{int}}(\mathbf{x} - \mathbf{k}), \quad (4.1)$$

where $\varphi_{\text{int}}(\mathbf{x})$ is an interpolating function; that is, $\varphi_{\text{int}}(\mathbf{0}) = 1$ and $\varphi_{\text{int}}(\mathbf{x})|_{\mathbf{x}=\mathbf{k} \in \mathbb{Z}^d \setminus \{\mathbf{0}\}} = 0$. Popular examples of φ_{int} are the linear B-spline [14] and the sinc-function which perform linear and bandlimited interpolation, respectively.

Thévenaz *et al.* [63] provided an equivalent formulation of (4.1) for an arbitrary non-interpolating function $\varphi(\mathbf{x})$ by considering the integer-shift-invariant model

$$f_{\text{int}}(\mathbf{x}) = \sum_{\mathbf{m} \in \mathbb{Z}^d} c[\mathbf{m}] \varphi(\mathbf{x} - \mathbf{m}), \quad (4.2)$$

where the coefficients $\{c[\mathbf{m}]\}$ are determined by solving the set of linear equations²

$$f_{\text{int}}(\mathbf{x})|_{\mathbf{x}=\mathbf{k}} = \sum_{\mathbf{m} \in \mathbb{Z}^d} c[\mathbf{m}] \varphi(\mathbf{k} - \mathbf{m}) = f(\mathbf{k}), \quad \mathbf{k} \in \mathbb{Z}^d,$$

which ensures perfect fitting of the given samples. In this chapter, we propose to use the shift-invariant model (4.2) for our regularized interpolation scheme, but here, we are going to specify $\{c[\mathbf{m}]\}$ based on certain regularization requirements on the interpolation model in addition to the data-fitting constraint (see Section 4.3.2).

4.2.1 Standard Interpolation in the Presence of Noise

While standard interpolation (exact fitting) is desirable in the noise-free scenario, it can lead to unfavorable results when applied on noise-corrupted data. To demonstrate this, we consider the following experiment: First, we rotate some noise-free input image by a random angle using a high-quality interpolator. Then, we add zero-mean white Gaussian noise of variance σ^2 to obtain a prescribed signal-to-noise ratio (SNR). Finally, we rotate back the noisy image with nearest-neighbour (using the rect-function), linear, and cubic B-spline based interpolation. We repeat this for a fixed number of realizations and average the SNR of the output image over all realizations. We show in Figure 4.1 the plot of the (averaged) SNR of the output image for a range of input SNRs. We observe that piecewise linear interpolation outperforms cubic splines at high noise levels which contradicts the usual behavior.³

This can be explained as follows: Since standard interpolation is a linear operation (see Section 2.3.3 in Chapter 2) and because noise is zero-mean and uncorrelated with the image (which we assume is a realization of a stationary process), we can use the approximation error formula $\varepsilon_{f,\text{noisy}}$ (see Equation (3.54)) that we derived in Section 3.8 of Chapter 3 to this setting:⁴ $\varepsilon_{f,\text{noisy}}$ is composed of the noise-free error component ε_f and the contribution $\varepsilon_{\text{noise}}$ that is entirely due to noise.

²Solving the linear system is equivalent to filtering the given samples with an IIR interpolation filter (see Section 2.3.3 in Chapter 2)). However, in practice, the interpolation filter is implemented using computationally efficient recursive filtering techniques [63].

³For the noise-free case, we saw in Chapter 2 (based on approximation-theoretic results) that cubic-spline interpolation provides better quality compared to piecewise linear interpolation.

⁴In the context of interpolation, $h(\mathbf{x}) = \delta(\mathbf{x})$ and $\varphi_{\text{eq}} = \varphi_{\text{int}}$, the interpolator.




Figure 4.1: Rotation experiment in the presence of noise: Piecewise linear interpolation performs better than cubic splines at high noise levels (input SNR < 24 dB) contradicting the expected behavior. We show in the inset the central region of the image which was used for computing SNR.

From Equation (3.54) (see Chapter 3), it is seen that ε_f is completely characterized by the approximation order L [53] (see Section 2.5 in Chapter 2); that is, the ability of the model to reproduce polynomials of degree $n = 0, 1, \dots, L - 1$. It is known from approximation theory [53] that the higher the value of L , the lower the ε_f error. Since $L_{\text{Cubic}} > L_{\text{Linear}} > L_{\text{Rect}}$, we have that

$$\varepsilon_{f_{\text{Cubic}}} < \varepsilon_{f_{\text{Linear}}} < \varepsilon_{f_{\text{Rect}}}. \quad (4.3)$$

As given in Proposition 3.10 in Chapter 3, the contribution due to noise is nothing but $\varepsilon_{\text{noise}}(\varphi) = \sigma^2 \|\varphi_{\text{int}}\|_{L_2}^2$, where σ^2 is the noise variance. Computing it for $\varphi = \text{rect}$ ($\varepsilon_{\text{noise}}(\text{rect}) = \sigma^2$), $\varphi = \beta^1$, the linear B-spline ($\varepsilon_{\text{noise}}(\beta^1) = \sigma^2 \frac{2}{3}$), and $\varphi = \beta^3$, the cubic B-spline ($\varepsilon_{\text{noise}}(\beta^3) \approx 0.874 \sigma^2$), we observe that

$$\varepsilon_{\text{noise}}(\beta^1) < \varepsilon_{\text{noise}}(\beta^3) < \varepsilon_{\text{noise}}(\text{rect}). \quad (4.4)$$

Therefore, in Figure 4.1, at high noise levels (input SNRs in the range 0 to 24

dB) where the effect of $\varepsilon_{\text{noise}}$ is dominant, piecewise linear interpolation does better than cubic splines because of (4.4). At low noise levels (input SNR > 24 dB), ε_f becomes effective and the cubic splines take the lead due to (4.3). Nearest-neighbour interpolation has the poorest performance at all noise levels since both ε_f and $\varepsilon_{\text{noise}}$ are higher for the rect function than for linear and cubic B-splines.

4.3 Regularized Interpolation

When the input data is noisy, it is meaningful to adopt a variational approach to enforce regularity constraints on the interpolation model that counterbalance the effect of noise. This leads to a regularized interpolation scheme that searches for a “smooth” function $f(\mathbf{x})$ which provides a reasonably good fit of the input data $\{g[\mathbf{k}]\}$.

4.3.1 Problem Formulation

We develop the problem in a penalized-likelihood setting in the spirit of [97–99] where the solution is obtained by minimizing a cost functional $J\{g, f\}$ composed of a negative log-likelihood term (also called the data-fidelity term) $\mathcal{L}\{g, f\}$ and a continuous-space regularization functional $\Psi\{f\}$: The log-likelihood measures the goodness-of-fit between $\{g[\mathbf{k}]\}$ and $\{f(\mathbf{k})\}$ in a statistical sense, while the regularization penalizes heavy oscillations in the solution.

Mathematically, this is written as

$$f_\lambda = \arg \min_f J\{g, f\}, \quad (4.5)$$

where

$$J\{g, f\} = \mathcal{L}\{g, f\} + \lambda \Psi\{f\},$$

and where $\lambda > 0$ is the regularization parameter that governs the tradeoff between goodness-of-fit and smoothness of f_λ . We shall take up the problem of choosing an appropriate λ in Chapters 5 and 6 where we propose a practical scheme that provides the “best” fit within the given class of solutions.

4.3.1.1 Data-Fidelity Term

In the penalized-likelihood framework, we have that $\mathcal{L}\{g, f\} = -\log(\wp(g|f))$ where \wp is the probability density of $\{g[\mathbf{k}]\}$ given $\{f(\mathbf{k})\}$, or, equivalently, the

probability density of the random vector characterizing the noise in the data. Here, $\mathcal{L}\{g, f\}$ is always a discrete-domain entity since it measures the statistical infidelity of the samples of f to the data. We consider that the noise is statistically independent at different sample locations. Then, the joint-probability density can be written as $\wp(g|f) = \prod_{\mathbf{k}} \wp_{\mathbf{k}}(g[\mathbf{k}]|f(\mathbf{k}))$, where $\wp_{\mathbf{k}}$ is the marginal density, so that⁵

$$\mathcal{L}\{g, f\} = - \sum_{\mathbf{k}} \log(\wp_{\mathbf{k}}(g[\mathbf{k}]|f(\mathbf{k}))). \quad (4.6)$$

In this chapter, we deal with two specific instances of (4.6). The first is i.i.d. zero-mean additive white Gaussian noise (AWGN), where

$$\wp_{\mathbf{k}}(g[\mathbf{k}]|f(\mathbf{k})) \propto e^{-\kappa(g[\mathbf{k}]-f(\mathbf{k}))^2}.$$

This leads to the negative log-likelihood

$$\mathcal{L}_Q\{g, f\} = \sum_{\mathbf{k}} (g[\mathbf{k}] - f(\mathbf{k}))^2. \quad (4.7)$$

From a signal-processing perspective, the AWGN model is often preferred for mathematical ease as the quadratic nature of \mathcal{L}_Q simplifies the optimization process.

When the noise is non-Gaussian in nature the corresponding log-likelihood is non-quadratic. As an example of this case, we consider the signal-dependent Poisson model (non-stationary) that is appropriate for imaging applications such as fluorescence microscopy [100] and emission tomography [101]. Here, the sample $g[\mathbf{k}]$ represents the detector counts at the \mathbf{k} -th pixel and

$$\wp_{\mathbf{k}}(g[\mathbf{k}]|f(\mathbf{k})) = (g[\mathbf{k}]!)^{-1} e^{-f(\mathbf{k})} (f(\mathbf{k}))^{g[\mathbf{k}]}$$

The corresponding negative log-likelihood is given by (up to the additive constant $\log(g[\mathbf{k}]!)$ which is irrelevant for optimization purposes)

$$\mathcal{L}_{\text{Poisson}}\{g, f\} = - \sum_{\mathbf{k} \in \mathbb{Z}^d} g[\mathbf{k}] \log(f(\mathbf{k})) + \sum_{\mathbf{k} \in \mathbb{Z}^d} f(\mathbf{k}). \quad (4.8)$$

For numerical tractability, we only use the log-likelihood of the pure Poisson component neglecting other factors like the detector gain and variance of the background radiation since these can be modeled as being part of $\{f(\mathbf{k})\}$.

⁵In this context, the ℓ_p data-fidelity in Equation (3.8) in the previous chapter can be interpreted as the log-likelihood corresponding to a zero-mean additive generalized-Gaussian noise model.

4.3.1.2 Regularization

We consider the following general form as our regularization functional:

$$\Psi\{f(\bullet)\} = \int_{\mathbb{R}^d} \Phi(\|\mathbf{L}\{f(\bullet)\}(\mathbf{x})\|) \, d\mathbf{x}, \tag{4.9}$$

where \mathbf{L} is a vector composed of linear differential operators $\{L_m\}_{m=1}^{d_{\mathbf{L}}}$ which measures the “smoothness” of $f(\mathbf{x})$ in terms of

$$\|\mathbf{L}\{f\}(\mathbf{x})\| = \left(\sum_{m=1}^{d_{\mathbf{L}}} (L_m\{f\}(\mathbf{x}))^2 \right)^{\frac{1}{2}},$$

and where Φ is called the potential function that defines the penalty associated with $\|\mathbf{L}\{f\}(\mathbf{x})\|$.

The main difference from the regularization $\Theta(\|\mathbf{L}\{\bullet\}\|_{L_2}^2)$ in (3.8) (see Chapter 3) is that we now use a vector differential operator \mathbf{L} in place of a scalar one (L in Chapter 3) and most importantly, we explore the possibility of a non-quadratic regularization, in the sense that we do not restrict ourselves to the case of the quadratic norm $\|\mathbf{L}\{\bullet\}\|_{L_2}^2$ in (3.8). Due to the positivity of $\|\mathbf{L}\{f\}\|$, Φ need only be specified on the set of non-negative real numbers.

Definition 4.1. *The one-sided potential function Φ is said to be appropriate for the purpose of regularization if it is non-negative, $\Phi(x) \geq 0$, and increasing $\Phi'(x) > 0, \forall x \geq 0$. ■*

This definition is consistent with the minimization in (4.5) since we wish to increase the penalty whenever $\|\mathbf{L}\{f\}\|$ increases. In order for (4.9) to be beneficial for the interpolation problem, we additionally require that $\Psi\{f\}$ be invariant to translation, rotation, and dilation. Then, it is guaranteed that the solution is invariant to such transformations of the given data, thereby becoming independent of the data-grid. Mathematically, the invariance requirements are prescribed as follows: We want the value of $\Psi\{f\}$ to remain unchanged (up to a multiplicative constant) when f is

1. shifted by $\mathbf{x}_0 \in \mathbb{R}^d$ (translation-invariance)

$$\Psi\{f(\bullet - \mathbf{x}_0)\} = \Psi\{f(\bullet)\}, \tag{4.10}$$

2. rotated about the origin by an arbitrary angle θ (rotation-invariance)

$$\Psi\{f(\mathbf{R}_\theta \bullet)\} = \Psi\{f(\bullet)\}, \text{ and} \tag{4.11}$$

3. dilated by $\tau > 0$ (scale-invariance)

$$\Psi \left\{ f \left(\frac{\bullet}{\tau} \right) \right\} = C(\tau) \Psi \{ f(\bullet) \}, \quad (4.12)$$

where $C(\tau) > 0$ is an appropriate scalar that is differentiable with respect to τ ; its role is to balance the regularization against a change of scale at which it is calculated.

Since $\Psi \{ f(\bullet) \}$ is specified via the norm $\| \mathbf{L} \{ f(\bullet) \}(\mathbf{x}) \|$, affine invariance of $\Psi \{ f(\bullet) \}$ (4.10)-(4.12) necessitates that $\| \mathbf{L} \{ f(\bullet) \}(\mathbf{x}) \|$ be preserved under translation, rotation, and dilation of f , up to the Lebesgue measure in the integral (4.9). This curtails the choice of \mathbf{L} to those that are shift-, rotation-, and scale-invariant in nature.

Definition 4.2. *The vector-differential operator \mathbf{L} is said to be a shift-, rotation-, and scale-invariant operator if $\| \mathbf{L} \{ f(\bullet) \}(\mathbf{x}) \|$ commutes with translation*

$$\| \mathbf{L} \{ f(\bullet - \mathbf{x}_0) \}(\mathbf{x}) \| = \| \mathbf{L} \{ f(\bullet) \}(\mathbf{x} - \mathbf{x}_0) \|, \quad (4.13)$$

with rotation

$$\| \mathbf{L} \{ f(\mathbf{R}_\theta \bullet) \}(\mathbf{x}) \| = \| \mathbf{L} \{ f(\bullet) \}(\mathbf{R}_\theta \mathbf{x}) \|, \quad (4.14)$$

and with dilation

$$\left\| \mathbf{L} \left\{ f \left(\frac{\bullet}{\tau} \right) \right\}(\mathbf{x}) \right\| = \rho(\tau) \left\| \mathbf{L} \{ f(\bullet) \} \left(\frac{\mathbf{x}}{\tau} \right) \right\|, \quad (4.15)$$

where $\rho(\bullet) > 0$ is a differentiable function that captures the response of \mathbf{L} to a scaling of the coordinate system.

Interestingly, common multivariate differential operators such as the gradient ($\mathbf{L} = \nabla$) and the Laplacian ($\mathbf{L} = \Delta$) turn out to be shift-, rotation-, and scale-invariant in nature. In the case of the gradient operator, we have $L_m = \frac{\partial}{\partial x_m}$, $m = 1, 2, \dots, d$. Obviously, the relation

$$\frac{\partial f(\mathbf{x}/\tau)}{\partial x_m} = \tau^{-1} \frac{\partial f(\mathbf{x})}{\partial x_m} \Big|_{\mathbf{x}=\mathbf{x}/\tau}$$

implies that the response of ∇ to the dilation operation is $\rho(\tau) = \tau^{-1}$. Similarly,

for the case of the Laplacian $\mathbf{L} = \Delta = \sum_{m=1}^d \frac{\partial^2}{\partial x_m^2}$ (scalar operator), we see that $\rho(\tau) = \tau^{-2}$ in (4.15).

Going back to (4.9), it may seem that Φ can be arbitrarily chosen. This is true with respect to translation and rotation invariance of $\Psi\{f\}$ since (4.13) and (4.14) ensure that without the need for specifying an explicit functional form for Φ . However, invariance of $\Psi\{f\}$ to dilation calls for special attention as it couples the scale invariance of \mathbf{L} and the effect of dilation on the potential function Φ . In fact, this connection together with (4.12) narrows down the choice of Φ as shown in the following theorem.

Theorem 4.1. *Let \mathbf{L} be a linear, scale-, rotation-, and shift-invariant differential operator and the potential function Φ be as defined in Definition 4.1. Then, $\Psi\{f\}$ is invariant to scaling of the coordinates if and only if $\Phi(x) = x^p, \forall x \geq 0$ with $p > 0$.*

Proof: Writing down (4.12) explicitly in terms of integrals yields

$$\int_{\mathbb{R}^d} \Phi \left(\left\| \mathbf{L} \left\{ f \left(\frac{\cdot}{\tau} \right) \right\} (\mathbf{x}) \right\| \right) d\mathbf{x} = \zeta(\tau) \int_{\mathbb{R}^d} \Phi(\|\mathbf{L}\{f\}(\mathbf{x})\|) d\mathbf{x}. \quad (4.16)$$

We start from the l.h.s. of (4.16) and use the fact that \mathbf{L} is scale-invariant (4.15) to obtain

$$\begin{aligned} \int_{\mathbb{R}^d} \Phi \left(\left\| \mathbf{L} \left\{ f \left(\frac{\cdot}{\tau} \right) \right\} (\mathbf{x}) \right\| \right) d\mathbf{x} &= \int_{\mathbb{R}^d} \Phi \left(\rho(\tau) \left\| \mathbf{L} \{f\} \left(\frac{\mathbf{x}}{\tau} \right) \right\| \right) d\mathbf{x} \\ &= \tau^d \int_{\mathbb{R}^d} \Phi(\rho(\tau) \|\mathbf{L}\{f\}(\mathbf{x})\|) d\mathbf{x}. \end{aligned} \quad (4.17)$$

Then, comparing the r.h.s. of (4.17) and (4.16), we infer that Φ must necessarily satisfy

$$\Phi(\rho(\tau)x) = \vartheta(\tau)\Phi(x) \quad \forall x \in \mathbb{R}, \quad (4.18)$$

where $\vartheta(\tau) = \tau^{-d}\zeta(\tau)$. Differentiating (4.18) with respect to τ and setting $\tau = 1$, we get

$$x\Phi'(x) = p\Phi(x), \quad (4.19)$$

where we have used the fact that $\rho(1) = 1$ (there is no scaling for $\tau = 1$ in (4.15)) and $p = \frac{\vartheta'(1)}{\rho'(1)}$ is a real number. The general solution to (4.19) is of the form⁶ $\Phi(x) = \gamma_1 x_+^p + \gamma_2 (-x)_+^p$, where γ_1 and γ_2 are arbitrary constants, and

$$x_+^p = \begin{cases} x^p, & \text{if } x \geq 0, \\ 0, & \text{otherwise.} \end{cases}$$

⁶The general solution may contain distributions for negative integer values of p [102]. However, in the present context, we would like the solution to be a true function of x , which leads to the given form for Φ .

Using the hypothesis that Φ is defined only for $x \geq 0$ and is non-negative and increasing, we see from (4.19) that $p > 0$, which leads to the desired result (up to a multiplicative real constant) $\Phi(x) = x^p, \forall x \geq 0$ with $p > 0$.

Conversely, it is verified that the form $\Phi(x) = x^p$ ensures scale invariance of $\Psi\{f\}$. ■

As a direct consequence of Theorem 4.1, we see that the following L_p -norm is the only choice of regularization with respect to (4.12): Substituting $\Phi(x) = x^p$ in (4.9), we get

$$\Psi\{f\} = \int_{\mathbb{R}^d} \|\mathbf{L}\{f\}(\mathbf{x})\|^p d\mathbf{x}. \quad (4.20)$$

In this work, we shall focus on the convex class of regularization functionals in (4.20) which precludes $p < 1$. Therefore, the practical range of interest of the p -values is $1 \leq p \leq 2$. Some popular instances of convex $\Psi\{f\}$ in (4.20) that can be found in regularization or spline literature are

1. Total-variation regularization [74, 103, 104] where $p = 1$, $\mathbf{L} = \nabla$, and $\zeta(\tau) = \tau^{d-1}$,
2. Quadratic regularization: Set $p = 2$
 - (a) Laplacian semi-norm [105] with $\mathbf{L} = \Delta$ and $\zeta(\tau) = \tau^{d-4}$,
 - (b) Duchon's semi-norm of order M [106] where \mathbf{L} is a vector composed of every possible M -th order partial derivative operator and $\zeta(\tau) = \tau^{d-2M}$.

4.3.2 Discretization of the Problem

As we are dealing with an interpolation problem, we seek a solution of the form

$$f_\lambda(\mathbf{x}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} c_\lambda[\mathbf{k}] \varphi(\mathbf{x} - \mathbf{k}) \quad (4.21)$$

for our regularized scheme. Then, the original problem (4.5) can be posed as the discrete-domain optimization problem

$$\{c_\lambda[\mathbf{k}]\}_{\mathbf{k} \in \mathbb{Z}^d} = \arg \min_{\{c[\mathbf{k}]\}_{\mathbf{k} \in \mathbb{Z}^d}} J_\lambda\{g, c\}. \quad (4.22)$$

The associated cost J_λ is a function of $\{c[\mathbf{k}]\}$ given by

$$J_\lambda\{g, c\} = \mathcal{L}\{g, (c * b)\} + \lambda \int_{\mathbb{R}^d} \left(\sum_{m=1}^{d_L} \left(\sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] L_m\{\varphi\}(\mathbf{x} - \mathbf{k}) \right)^2 \right)^{\frac{p}{2}} d\mathbf{x}, \quad (4.23)$$

where we have used $f(\mathbf{k}) = \sum_{\mathbf{m} \in \mathbb{Z}^d} c[\mathbf{m}] b[\mathbf{k} - \mathbf{m}]$ with $b[\mathbf{k}] = \varphi(\mathbf{k})|_{\mathbf{k} \in \mathbb{Z}^d}$ in the data-fidelity term. Thus, while we consider the same continuous-domain model in (4.2) and (4.21), standard and regularized schemes differ in the way the coefficients are obtained. However, when $\lambda \rightarrow 0$ in (4.23), the regularized scheme (4.22) converges to the standard case (4.2) since we only minimize $\mathcal{L}\{g, (c * b)\}$ which leads to close-fitting of data; at the other extreme, as $\lambda \rightarrow \infty$, (4.22) results in a maximum-likelihood estimate within the null-space of L_m , $m = 1, \dots, d_L$.

4.3.2.1 Quadratic Regularization

When $p = 2$, the integrand in the r.h.s. of (4.23) is a quadratic term. Up to technical details related⁷ to φ and $\{L_m\}_{m=1}^{d_L}$ [66, 91], the integral can be evaluated analytically to yield

$$\sum_{m=1}^{d_L} \int_{\mathbb{R}^d} \left(\sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] L_m\{\varphi\}(\mathbf{x} - \mathbf{k}) \right)^2 d\mathbf{x} = \langle c * q, c \rangle_{\ell_2} = \Psi_Q\{c\},$$

where $\{q[\mathbf{k}]\}_{\mathbf{k} \in \mathbb{Z}^d}$ is the discrete sequence given by

$$q[\mathbf{k}] = \sum_{m=1}^{d_L} q_m[\mathbf{k}] \quad (4.24)$$

$$= \sum_{m=1}^{d_L} \int_{\mathbb{R}^d} L_m\{\varphi\}(\mathbf{x}) L_m\{\varphi\}(\mathbf{x} - \mathbf{k}) d\mathbf{x}. \quad (4.25)$$

Thus, in the quadratic case ($p = 2$), the discretization of $\Psi\{f\}$ implicitly follows from (4.25) and leads to a quadratic function $\Psi_Q\{c\}$ of the coefficients.

⁷As discussed at the end of Section 3.3.3 in Chapter 3, we can argue here that the summation and the integral can be exchanged whenever $L_m\{\varphi\}$ satisfies the first **SR** condition in (3.2), $\forall m = 1, 2, \dots, d_L$.

4.3.2.2 Non-Quadratic Regularization

In this case, the problem (4.22) can be handled only in a numerical optimization framework. For the purpose of numerical tractability, we replace the integral in (4.23) by a Riemann sum which leads to the discrete non-quadratic regularization

$$\Psi_{\text{NQ}}\{c\} = \sum_{\mathbf{k} \in \mathbb{Z}^d} \left(\sum_{m=1}^{d_{\mathbf{L}}} ((c * \eta_m)[\mathbf{k}])^2 \right)^{\frac{p}{2}}, \quad (4.26)$$

where $\eta_m[\mathbf{k}] = L_m\{\varphi\}(\mathbf{k})$ represents the discretized version of the differential operator L_m . The use of $\Psi_{\text{NQ}}\{c\}$ for regularized interpolation distinguishes this work from those in the literature [36, 37, 39, 66, 81, 91] which primarily deal with quadratic regularization.

From (4.26), it is seen that the numerical characterization of $\Psi\{f\}$ for $p \neq 2$ still accommodates the following important analog-domain feature: The discretized derivatives η_m are obtained by sampling the corresponding continuous-domain derivatives $L_m\{\varphi\}$, $m = 1, 2, \dots, d_{\mathbf{L}}$. We illustrate this connection in Section 4.3.4 where we present η_m , $m = 1, \dots, d_{\mathbf{L}}$, for the case of the gradient operator and polynomial B-splines.

4.3.3 Algorithms for Regularized Interpolation

When the cost J_λ is quadratic (Gaussian likelihood \mathcal{L}_Q and $p = 2$), the optimization can be performed mathematically [66, 91] which leads to an explicit closed-form solution that is related to $\{g[\mathbf{k}]\}$ in a linear fashion (as was seen in Section 3.3.4.5 in Chapter 3). Otherwise, the optimization has to be performed numerically and the corresponding solution depends on $\{g[\mathbf{k}]\}$ in a non-linear fashion. The same is true whenever the data-fidelity is dictated by a non-Gaussian likelihood irrespective of whether the regularization is quadratic or not. This is summarized in Table 4.1 where we present the characteristics of different regularized interpolation schemes.

In the sequel, we describe in detail the minimization procedure that we adopt for the Gaussian-likelihood setting (first two rows of Table 4.1). We also briefly mention some conventional optimization strategies for a general non-Gaussian likelihood and give a detailed exposition for the Poisson-likelihood case.

Table 4.1: Characteristics of different Penalized-likelihood interpolation schemes

Likelihood Model	Regularization	Optimization	Solution	Type of Algorithm
Gaussian	Quadratic	Analytical	Explicit, Closed-Form	Linear, One-Step
Gaussian	Non-quadratic	Numerical	Numerical	Non-Linear, Iterative
Non-Gaussian	Quadratic, Non-quadratic	Numerical	Numerical	Non-Linear, Iterative

4.3.3.1 Gaussian Likelihood with Quadratic Regularization

The cost to be minimized is composed of \mathcal{L}_Q and Ψ_Q ; that is,

$$J_{\text{GLQ}}\{g, c\} = \sum_{\mathbf{k} \in \mathbb{Z}^d} (g[\mathbf{k}] - (c * b)[\mathbf{k}])^2 + \lambda \langle c * q, c \rangle_{\ell_2}. \quad (4.27)$$

Since $J_{\text{GLQ}}\{g, c\}$ is quadratic in $\{c[\mathbf{k}]\}$, the minimization is performed (as done in Section 3.3.4.5 in Chapter 3) by setting its derivative with respect to $c[\mathbf{k}]$ to zero which yields

$$(c * b * \bar{b} + \lambda q * c)[\mathbf{k}] = (g * \bar{b})[\mathbf{k}], \quad (4.28)$$

where $\bar{b}[\mathbf{k}] = b[-\mathbf{k}]$.

Equation (4.28) can be solved in the discrete-space Fourier domain to obtain

$$C_\lambda(e^{j\boldsymbol{\omega}}) = R_\lambda(e^{j\boldsymbol{\omega}}) G(e^{j\boldsymbol{\omega}}) \quad (4.29)$$

$$= \left(\frac{B^*(e^{j\boldsymbol{\omega}})}{|B(e^{j\boldsymbol{\omega}})|^2 + \lambda Q(e^{j\boldsymbol{\omega}})} \right) G(e^{j\boldsymbol{\omega}}), \quad (4.30)$$

where $C_\lambda(e^{j\boldsymbol{\omega}})$, $G(e^{j\boldsymbol{\omega}})$, $B(e^{j\boldsymbol{\omega}})$, and $Q(e^{j\boldsymbol{\omega}})$ are the Fourier transforms⁸ of the solution $\{c_\lambda[\mathbf{k}]\}$, the data $\{g[\mathbf{k}]\}$, and the sequences $\{b[\mathbf{k}]\}$ and $\{r[\mathbf{k}]\}$, respectively.

⁸Conditions necessary for the existence and stability of $R_\lambda(e^{j\boldsymbol{\omega}})$ have to be checked on a case-by-case basis and can be done in a fashion similar to that described in Section 3.4 in Chapter 3.

Equation (4.30) amounts to writing $c_\lambda[\mathbf{k}] = (r_\lambda * g)[\mathbf{k}]$ where r_λ is the digital-correction filter whose frequency response is $R_\lambda(e^{j\omega})$. The solution is therefore linear, in the sense that $f_\lambda(\mathbf{x})$ is related to $\{g[\mathbf{k}]\}$ in a linear fashion. The implementation of (4.30) is straightforward and can be done via Fast-Fourier-Transform (FFT). We also note that $R_\lambda(e^{j\omega}) \rightarrow \frac{1}{B(e^{j\omega})}$ for $\lambda \rightarrow 0$, in which case f_λ performs standard interpolation of $\{g[\mathbf{k}]\}$ [63]. However, while $B(e^{j\omega})$ can be separable, this is never the case for $R_\lambda(e^{j\omega})$ when $\lambda > 0$.

Algorithm 4.1. *Algorithm for minimizing J_{GLQ} for a given λ .*

- (1) Precompute and store $B(e^{j\omega})$, $Q(e^{j\omega})$, and $G(e^{j\omega})$
- (2) Construct $R_\lambda(e^{j\omega})$ and compute $C_\lambda(e^{j\omega})$ using (4.30) for the given λ
- (3) Perform inverse Fourier transform of $C_\lambda(e^{j\omega})$ to obtain $\{c_\lambda[\mathbf{k}]\}$

4.3.3.2 Gaussian Likelihood with Non-Quadratic Regularization

In this case, we maintain \mathcal{L}_Q in the cost while we plug in Ψ_{NQ} for the regularization which results in

$$J_{\text{GLNQ}}\{g, c\} = \sum_{\mathbf{k} \in \mathbb{Z}^d} (g[\mathbf{k}] - (c * b)[\mathbf{k}])^2 + \lambda \Psi_{\text{NQ}}\{c\}. \quad (4.31)$$

The presence of Ψ_{NQ} in (4.31) makes $\mathcal{J}_{\text{GLNQ}}$ non-quadratic. Nevertheless, we note that J_{GLNQ} is convex in $\{c[\mathbf{k}]\}$ for the considered range of p . Several methodologies such as the majorize-minimize (MM) approach [96] (or, equivalently, bounded-optimization) and the half-quadratic method [107–109] have been developed in the past to handle non-quadratic convex costs. More recently, Nikolova *et al.* showed that the half-quadratic and bounded-optimization approaches are equivalent for cost functions of the type J_{GLNQ} [110]. In the present work, we resort to the MM approach as it is easy to comprehend: The idea is to replace the original difficult task by several easy-to-optimize problems that will guarantee monotonic decrease of the cost criterion [96]. Below, we briefly summarize the mathematical details underlying the MM philosophy, after which we apply it to minimize J_{GLNQ} .

In the MM setting, we construct an auxiliary cost function $J_{\text{AUX}}\{c|c_t\}$ at the

current estimate $\{c_t[\mathbf{k}]\}$ such that

$$\begin{aligned} J_{\text{GLNQ}}\{g, c_t\} &= J_{\text{AUX}}\{g, c_t|c_t\}, \\ J_{\text{GLNQ}}\{g, c\} &< J_{\text{AUX}}\{g, c|c_t\}, \quad c \neq c_t. \end{aligned} \quad (4.32)$$

Then, the idea is to find the estimate c_{t+1} at the next iteration such that

$$J_{\text{AUX}}\{g, c_{t+1}|c_t\} < J_{\text{AUX}}\{g, c_t|c_t\}, \quad (4.33)$$

which leads to a definite decrease of J_{GLNQ} as shown below:

$$\begin{aligned} J_{\text{GLNQ}}\{g, c_{t+1}\} &= J_{\text{AUX}}\{g, c_{t+1}|c_t\} + \underbrace{J_{\text{GLNQ}}\{g, c_{t+1}\} - J_{\text{AUX}}\{g, c_{t+1}|c_t\}}_{< 0 \text{ using (4.32)}} \\ &< J_{\text{AUX}}\{g, c_t|c_t\} \quad (\text{using (4.33)}) \\ &= J_{\text{GLNQ}}\{g, c_t\}. \end{aligned}$$

Our construction of the auxiliary cost J_{AUX} is based on the inequality [96]

$$|x|^p \leq \frac{2-p}{2}|x_0|^p + \frac{p}{2}x^2|x_0|^{p-2},$$

for $1 \leq p < 2$ and some $x_0 \neq 0$, where the equality holds only when $p = 2$ or $x = x_0$. The inequality is preserved under summation and multiplication by $\lambda > 0$ which leads to

$$\lambda \Psi_{\text{NQ}}\{c\} \leq \lambda \frac{p}{2} \sum_{\mathbf{k} \in \mathbb{Z}^d} |\chi_t[\mathbf{k}]|^{p-2} \sum_{m=1}^{d_{\mathbf{L}}} ((c * \eta_m)[\mathbf{k}])^2 + \mathcal{A}_{\chi_t}, \quad (4.34)$$

where

$$\chi_t[\mathbf{k}] = \sqrt{\sum_{m=1}^{d_{\mathbf{L}}} ((c_t * \eta_m)[\mathbf{k}])^2} \quad \text{and} \quad \mathcal{A}_{\chi_t} = \lambda \left(\frac{2-p}{2} \right) \Psi_{\text{NQ}}\{c_t\} \quad (4.35)$$

is a constant independent of $\{c[\mathbf{k}]\}$. Thus, we obtain

$$\begin{aligned} J_{\text{AUX}}\{g, c|c_t\} &= \sum_{\mathbf{k} \in \mathbb{Z}^d} (g[\mathbf{k}] - (c * b)[\mathbf{k}])^2 \\ &\quad + \lambda \frac{p}{2} \sum_{\mathbf{k} \in \mathbb{Z}^d} |\chi_t[\mathbf{k}]|^{p-2} \sum_{m=1}^{d_{\mathbf{L}}} ((c * \eta_m)[\mathbf{k}])^2 \\ &\quad + \mathcal{A}_{\chi_t}, \end{aligned}$$

which satisfies (4.32). Moreover, we note that J_{AUX} is quadratic in $\{c[\mathbf{k}]\}$ —we minimize it by setting its derivative to zero. This results in the following system of linear equations: $\forall \mathbf{k} \in \mathbb{Z}^d$

$$\left. \begin{aligned} & (c * b * \bar{b})[\mathbf{k}] \\ & + \\ & \lambda \frac{p}{2} \sum_{m=1}^{d_L} \sum_{\mathbf{m} \in \mathbb{Z}^d} |\chi_t[\mathbf{m}]|^{p-2} (c * \eta_m)[\mathbf{m}] \eta_m[\mathbf{m} - \mathbf{k}] \end{aligned} \right\} = (g * \bar{b})[\mathbf{k}]. \quad (4.36)$$

The presence of $\{\chi_t[\mathbf{k}]\}$ in (4.36) prohibits the use of Fourier-domain techniques such as those used in Section 4.3.3.1. Nevertheless, the MM technique requires only a decrease of J_{AUX} ; there is no need to minimize it completely. Therefore, efficient numerical schemes such as the conjugate-gradient (CG) method can be utilized to solve (4.36) partially. This leads to Algorithm 4.2.

Algorithm 4.2. *MM algorithm for minimizing J_{GLNq} for a given λ and p .*

- (1) *Precompute and store $\{(\bar{b} * b)[\mathbf{k}]\}$ and $\{(g * \bar{b})[\mathbf{k}]\}$*
- (2) *Initial estimate = $\{c_0[\mathbf{k}]\}_{\mathbf{k} \in \mathbb{Z}^d}$; $t = 0$*

Repeat Steps 3, 4 and 5 until Stop Criterion is met
- (3) *Compute $\{\chi_t[\mathbf{k}]\}_{\mathbf{k} \in \mathbb{Z}^d}$ using $\{c_t[\mathbf{k}]\}_{\mathbf{k} \in \mathbb{Z}^d}$*
- (4) *Apply CG iterations to partially solve for $\{c_{t+1}\}$ in (4.36)*
- (5) *Set $t = t + 1$*

4.3.3.3 Non-Gaussian Likelihood

Non-Gaussian likelihood functions invariably require iterative techniques for minimization of the associated cost. If the negative log-likelihood $\mathcal{L}\{g, (c * b)\}$ is a non-convex function of $\{c[\mathbf{k}]\}$, conventional techniques such as the non-linear conjugate gradient and BFGS methods [111] may be used for determining the local minima of the associated cost J_λ . The basic idea underlying these methods is to find a numerical solution by iteratively reducing the norm of the gradient of J_λ .

However, if $\mathcal{L}\{g, (c * b)\}$ is a convex function, the MM-strategy (4.32) can be used for developing elegant algorithms that guarantee monotonic decrease of J_λ .

Moreover, it is well known that the minimization of a convex J_λ always leads to a solution that is the global-minimum of J_λ . In this chapter, we consider the Poisson likelihood as a prototypical example of a convex non-quadratic data-fidelity term. The corresponding optimization procedure (Algorithm C.1) is given in Appendix C for both quadratic (Ψ_Q) and non-quadratic regularization (Ψ_{NQ}).

4.3.4 Spline-Based Regularized Interpolation

Here, we make explicit the link between the sequences $\{\eta_m\}_{m=1}^{d_L}$, $\{q_m\}_{m=1}^{d_L}$, and φ and \mathbf{L} for spline-based interpolation which is well-suited for imaging problems [14, 63]. We consider a separable generating function given by

$$\varphi(\mathbf{x}) = \varphi(x_1) \varphi(x_2) \cdots \varphi(x_d), \quad (4.37)$$

where $\varphi(x)$ is the univariate generating function. For \mathbf{L} , we select the gradient operator: Then, the proposed regularization includes the total-variation functional (for $p = 1$) which is of particular interest to us because of its edge-preserving characteristics [74].

The separable nature of $\varphi(\mathbf{x})$ in (4.37) leads to the separability of the corresponding discrete sequences $\{b[\mathbf{k}]\}$, $\{\eta_m[\mathbf{k}]\}_{m=1}^{d_L}$, and $\{q_m[\mathbf{k}]\}_{m=1}^{d_L}$, as written below:

$$b[\mathbf{k}] = b[k_1] b[k_2] \cdots b[k_d],$$

where $b[k] = \varphi(x)|_{x=k}$. For the gradient operator, we have $d_L = d$, which yields for $m = 1, \dots, d$,

$$\eta_m[\mathbf{k}] = \prod_{\substack{l=1 \\ l \neq m}}^d b[k_l] b^{(1)}[k_m],$$

where $b^{(1)}[k] = \varphi'(x)|_{x=k}$ and

$$q_m[\mathbf{k}] = \prod_{\substack{l=1 \\ l \neq m}}^d a[k_l] a^{(2)}[k_m],$$

where $a[k] = (\bar{\varphi} \star \varphi)(x)|_{x=k}$, $\bar{\varphi}(x) = \varphi(-x)$, and $a^{(2)}[k] = (\bar{\varphi}' \star \varphi')(x)|_{x=k}$. We list out the sequences $\{b[k]\}$, $\{b^{(1)}[k]\}$, and $\{a^{(2)}[k]\}$ (in terms of their transfer functions $B(z)$, $B^{(1)}(z)$, and $A^{(2)}(z)$, respectively) in Tables 4.2 and 4.3, for

Table 4.2: Transfer functions $B(z)$ and $B^{(1)}(z)$ corresponding to various spline degrees

n	$B(z)$	$B^{(1)}(z)$
1	1	$\frac{1}{2}(z^{-1} - z)$
2	$\frac{3}{4} + \frac{1}{8}(z^{-1} + z)$	$\frac{1}{2}(z^{-1} - z)$
3	$\frac{4}{6} + \frac{1}{6}(z^{-1} + z)$	$\frac{1}{2}(z^{-1} - z)$
4	$\frac{115}{192} + \frac{19}{96}(z^{-1} + z) + \frac{1}{384}(z^{-2} + z^2)$	$\frac{11}{24}(z^{-1} - z) + \frac{1}{48}(z^{-2} - z^2)$
5	$\frac{11}{20} + \frac{13}{60}(z^{-1} + z) + \frac{1}{120}(z^{-2} + z^2)$	$\frac{5}{6}(z^{-1} - z) + \frac{1}{24}(z^{-2} - z^2)$

$\varphi(x) = \beta^n(x)$, which is the symmetric polynomial B-spline of degree $n \geq 1$ [14]. We see that $\{b^{(1)}[k]\}$ and $\{a^{(2)}[k]\}$ are very different from the simple finite-difference filters (pixel-to-pixel difference) that are typically used in the literature for TV-based image restoration.

4.4 Results

We validate the proposed regularized interpolation scheme by carrying out 2-D rotation experiments in the presence of noise similar to those described in Section 4.2.1. We adopt the polynomial B-spline model in (4.37) for all implementations and use the gradient operator for regularization. Standard interpolation is performed as described⁹ in [63], while, in the regularized case, we implement Algorithm 4.1 using FFTs; the various steps of the non-linear algorithms (*cf.* Algorithm 4.2 and Algorithm C.1—see Appendix C) are executed via discrete convolutions. In the experiments, we set $p = 1$ for the non-quadratic regularization (NQR). The stopping criterion for Algorithm 4.2 and Algorithm C.1 is $\|c_{t+1} - c_t\|^2 < 10^{-4} N^{-1} \|g\|_{\ell_2}$, where N is the size of $\{g[\mathbf{k}]\}$, while the maximum

⁹The coefficients $\{c[\mathbf{k}]\}$ are obtained by prefiltering the samples with the interpolation filter (q_{int} in (2.32) in Chapter 2). Then, the interpolation is performed by computing the sum of shifted φ 's weighted by the coefficients.

Table 4.3: Transfer function $A^{(2)}(z)$ corresponding to various spline degrees

n	$A^{(2)}(z)$
1	$2 - (z^{-1} + z)$
2	$1 - \frac{1}{3}(z^{-1} + z) - \frac{1}{6}(z^{-2} + z^2)$
3	$\frac{2}{3} - \frac{1}{8}(z^{-1} + z) - \frac{1}{5}(z^{-2} + z^2) - \frac{1}{120}(z^{-3} + z^3)$
4	$\frac{35}{72} - \frac{11}{360}(z^{-1} + z) - \frac{17}{90}(z^{-2} + z^2) - \frac{59}{2520}(z^{-3} + z^3) - \frac{1}{5040}(z^{-4} + z^4)$
5	$\frac{809}{2160} + \frac{1}{64}(z^{-1} + z) - \frac{31}{1890}(z^{-2} + z^2) - \frac{907}{24192}(z^{-3} + z^3) - \frac{25}{18144}(z^{-4} + z^4) - \frac{1}{362880}(z^{-5} + z^5)$

number of allowed iterations is set to 25 (with 20 CG iterations in each main loop) for Algorithm 4.2 and it is set to 300 for Algorithm C.1.

The performance of all methods is quantified by the signal-to-noise ratio (SNR) computed as

$$\text{SNR} = 10 \log_{10} \left(\frac{\sum_{\mathbf{k} \in \Omega} f^2(\mathbf{k})}{\sum_{\mathbf{k} \in \Omega} (f(\mathbf{k}) - f_\lambda(\mathbf{k}))^2} \right), \tag{4.38}$$

where $\{f(\mathbf{k})\}$ and $\{f_\lambda(\mathbf{k})\}$ are the values of the continuous-domain noise-free signal and regularized output, respectively, sampled on the grid of data. The SNR is evaluated inside a circular region Ω concentric with the image so as to avoid boundary effects.

Our main aim in this chapter is to characterize the best-possible performance of the proposed regularized interpolation methods. For this purpose, we conduct oracle-based experiments; that is, we set λ so as to obtain f_λ that yields the highest SNR for a given realization of the noisy data $\{g[\mathbf{k}]\}$. In Chapters 5 and 6, we take up the problem of developing a data-driven scheme for obtaining the “optimal” λ directly from $\{g[\mathbf{k}]\}$. Specifically, we develop a Monte-Carlo

method in Chapter 6 for estimating the SNR corresponding to an arbitrary interpolation/denoising algorithm and apply it for selecting the “optimal” λ for our regularized interpolation problem.

4.4.1 Regularized versus Standard Interpolation

We compare the proposed regularized scheme against standard interpolation by carrying out two sets of 2-D rotation experiments where the setup is exactly similar to that described in Section 4.2.1. For the first experiment, we use a 512×512 image of a CT slice (central portion shown in the inset of Figure 4.2). For the second, we consider a stack of clean MRI images [112] (T1, T2, and PB images that are zero-padded to size 300×300 so as to provide enough margins for the rotations) where we use different slices picked randomly from the MRI stack for different realizations. We perform standard interpolation (non-regularized) using nearest-neighbour, linear, and cubic B-spline, while, for regularized methods, we consider quadratic regularized (QR) cubic B-spline (Algorithm 4.1) and non-quadratic regularized (NQR) cubic B-spline (Algorithm 4.2), respectively.

4.4.1.1 SNR Comparison

We plot the output SNR for each of these methods in Figures 4.2 and 4.3 corresponding to the CT slice and the MRI stack, respectively. In both figures, we observe the following: The nearest-neighbour method always performs poorly, which is expected. For low input SNRs (< 24 dB in Figure 4.2 and < 32 dB in Figure 4.3), piecewise linear interpolation does better than cubic B-spline: This behavior was already explained in Section 4.2.1. The most important observation is that both regularized interpolation methods (QR and NQR) perform far better than the non-regularized ones at high noise levels. The consistently superior trend exhibited by regularized interpolation in both figures clearly illustrate its robustness against noise. Finally, the output SNR of the two regularized methods (for cubic B-spline) converges to that of the non-regularized cubic B-spline for relatively high input SNRs since the effect of regularization becomes negligible under very low noise.

4.4.1.2 Visual Comparison

We present in Figure 4.4 output images for one realization of the experiment in Section 4.4.1 (with MRI stack). Since standard cubic B-spline interpolation is not regularized, the corresponding output (Figures 4.4(c)) is still noisy. On

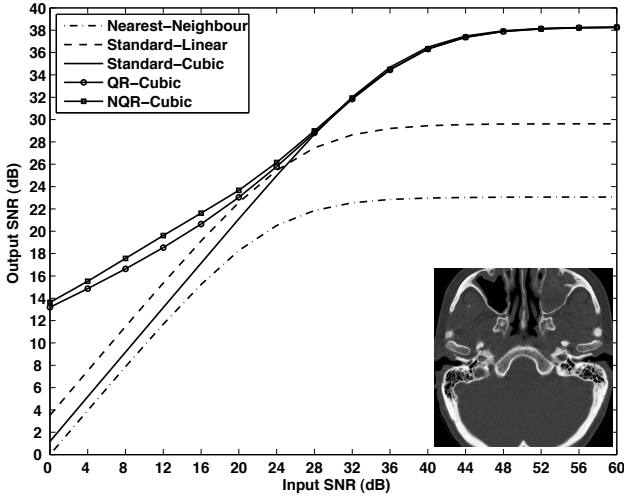


Figure 4.2: Rotation experiments on noisy versions of a slice of a CT image: Comparison of performance of non-regularized and regularized interpolation.

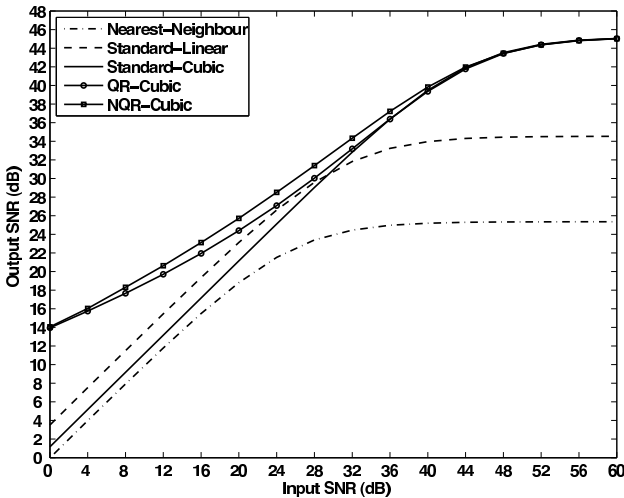


Figure 4.3: Rotation experiments on noisy versions of MRI slices: Comparison of performance of non-regularized and regularized interpolation.




Figure 4.4: Visual Comparison: (a) Noise-free image; (b) Noisy data, rotation angle $\approx -60^\circ$, SNR ≈ 16 dB; (c) Non-regularized cubic B-spline interpolation result, SNR = 17.12 dB; (d) Quadratic regularized (QR) cubic interpolation result ($p = 2$, optimal λ in Algorithm 4.1) SNR = 19.49 dB; (e) Non-quadratic regularized (NQR) cubic interpolation result ($p = 1$, optimal λ in Algorithm 4.2) SNR = 21.04 dB. The reported SNR values were calculated over the support of the noise-free image displayed here.

the contrary, regularized methods lead to significant noise reduction as seen in Figures 4.4(d) and 4.4(e) corresponding to QR-cubic and NQR-cubic outputs, respectively.

4.4.2 Regularized Interpolation: Varying Spline Degree

To study the effect of the basis function φ on the discretization of the non-quadratic regularization (NQR) in (4.26), we repeat the experiment in Section 4.4.1 (for MRI images) with B-splines of integer degree varying from 1 to 5 and focus on NQR interpolation (Algorithm 4.2). We show in Figure 4.5 the performance of the NQR interpolation based on linear ($n = 1$), quadratic ($n = 2$), cubic ($n = 3$), quartic ($n = 4$), and pentic ($n = 5$) B-splines, respectively.

The output SNR consistently increases with the degree of the B-spline over the entire range of input SNRs indicating that higher-degree B-splines yield better performance with NQR interpolation. Particularly, there is a notable improvement going from linear to a higher-degree B-spline. This is probably because, for the linear B-spline, the discretization does not adequately capture the features of the corresponding continuous-domain model, while the situation improves when $n \geq 2$. For higher-order B-splines ($n \geq 4$), we only observe marginal increments in the output SNR that tends to saturate. This is to be expected since the cardinal splines (corresponding to these B-splines) rapidly converge to the sinc-function with increasing n [14].

4.4.3 Experiments with Poisson Noise

Until now, we have presented results that demonstrate the superior performance of regularized interpolation (using Algorithm 4.1 and Algorithm 4.2) for data corrupted by additive Gaussian noise. In this section, we investigate the proposed approach using Algorithm C.1 (in Appendix C) by performing rotations in the presence of signal-dependent Poisson noise (results not shown). The outcome was very similar to that exhibited in Figures 4.2 and 4.3, and we summarize our findings here by stating again that the proposed regularized interpolation scheme outperforms standard methods by a wide margin.

Our concern here is rather to investigate whether or not the choice of the data-fidelity term based on the likelihood model is crucial for regularized interpolation. Specifically, we propose to study the performance of quadratic data fidelity \mathcal{L}_Q (which corresponds to a Gaussian likelihood) when applied to signal-dependent Poisson noise and compare it with $\mathcal{L}_{\text{Poisson}}$ which is statistically the most appropriate data-fidelity term for this type of noise. For this, we repeat




Figure 4.5: Non-quadratic regularized (NQR) interpolation for varying spline-degree.

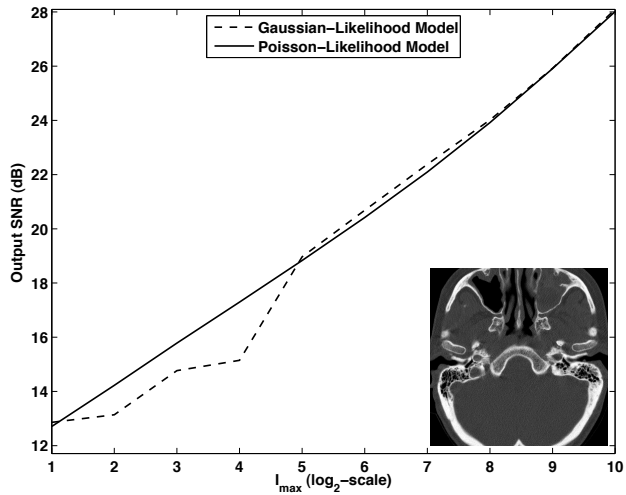


Figure 4.6: Comparison of the effect of likelihood on non-quadratic regularized (NQR) interpolation.

the experiment (with CT slice) described in Section 4.4.1, but, here, we consider data corrupted by signal-dependent Poisson noise. To generate Poisson noise, we make use of the RandomJ plugin [113] for the ImageJ freeware. The intensity value of the noise-free image at a given pixel characterizes the Poisson noise at that pixel (mean of the Poisson random variable). So we vary the maximum-allowed intensity I_{\max} in the test image to simulate Poisson noise of varying strengths. For instance, a low I_{\max} value indicates a low-count situation (*i.e.*, low input SNR). We perform NQR-cubic interpolation using Algorithm 4.2 (Gaussian-likelihood) and Algorithm C.1 (Poisson likelihood). The respective output SNRs are plotted in Figure 4.6 as a function of I_{\max} (in \log_2 -scale).

For $I_{\max} \leq 2^5$, which corresponds to an average of six or less counts per pixel for the CT slice image in Figure 4.6, Algorithm C.1 yields substantial improvement (≥ 1 dB) over Algorithm 4.2 pointing to the superiority of the Poisson-likelihood model for low-count situations. However, when $I_{\max} > 2^5$, Algorithm 4.2 and Algorithm C.1 yield more-or-less similar performances indicating that the effect of the likelihood model is no longer significant.

4.5 Discussion

4.5.1 Quadratic versus Non-Quadratic Regularization

Among the regularized schemes investigated in this work, NQR interpolation (based on TV regularization) performs significantly better than QR interpolation both in terms of visual quality and SNR: In Figure 4.4 the NQR result (Figure 4.4(e)) is sharper and less noisy than the QR result (Figure 4.4(d)), while in Figures 4.2 and 4.3 we see that NQR brings about a consistent SNR improvement over QR interpolation at high noise levels. These results are a direct consequence of the fact that TV-like NQRs have a good ability to preserve important signal features (such as edges) while Tikhonov-like QRs tend to blur them thus compromising the quality. Therefore, from a performance point of view, it is better to employ TV-like NQR for regularized interpolation.

4.5.2 Influence of the Likelihood Model

It is commonly believed that the likelihood term plays a significant role in variational problems. While this is true in some situations, it often turns out that, in the presence of a strong regularization, the solution is dictated more by the regularization than by the likelihood. We presented an example in Section 4.4.3

to demonstrate this for the case of Poisson noise: The Poisson likelihood $\mathcal{L}_{\text{Poisson}}$ outperforms the Gaussian likelihood \mathcal{L}_Q for low-count Poisson data, while they yield almost the same performance in the high-count case. Thus, when the choice of the data-fidelity term is not crucial, algorithms can be designed to reduce the amount of computations. For instance, since Algorithm 4.2 uses \mathcal{L}_Q for the data fidelity, it only requires the linearization of the regularization—the corresponding optimization is simple as it only amounts to solving a set of linear equations. This is difficult to accomplish under the Poisson-likelihood model—the logarithm in $\mathcal{L}_{\text{Poisson}}$ requires careful handling of the problem as positivity of the solution often becomes a harsh constraint. In most practical situations, the number of counts is sufficiently large that the deviation from the Poisson likelihood model is negligible. In such cases, \mathcal{L}_Q can be used in place of $\mathcal{L}_{\text{Poisson}}$ without compromising on the performance.

4.5.3 Computational Cost

In our analysis so far, we ranked the various regularized interpolation algorithms purely based on performance gain. However, we must also consider the computation cost associated with these algorithms. In [63], Thévenaz *et al.* performed a thorough cost-performance analysis for standard interpolation; for a given quality measure, their emphasis was on reducing the cost of evaluating $\varphi(\mathbf{x} - \mathbf{k})$ for many arguments $(\mathbf{x} - \mathbf{k})$. In the context of regularized interpolation, since we are concerned with obtaining $\{c[\mathbf{k}]\}$, we only consider the cost of computing the coefficients for a given φ . In the sequel, we provide a rough measure of this cost for the various (regularized and non-regularized) algorithms presented in this work.

Since, for standard interpolation, the $\{c[\mathbf{k}]\}$ are computed by linear filtering of the data [63], it can be achieved with $O(N)$ complexity. In the context of regularized interpolation, Algorithm 4.1 is easily implemented in the Fourier domain (via FFTs)—it requires about $O(N \log(N))$ operations. Thus, with only a slightly larger computation load (an extra $\log(N)$ factor), Algorithm 4.1 yields significant improvement in quality and is generally preferable to standard interpolation.

Algorithm 4.2 is iterative in nature and requires the evaluation of several convolutions in its intermediate steps. The convolutions are always performed on the estimate $\{c_t[\mathbf{k}]\}$ with $\{b[\mathbf{k}]\}$, $\{\eta_m[\mathbf{k}]\}_{m=1}^{d_L}$. To simplify the process, we disregard the size of these sequences and associate the same cost C to the computation of $\{(b * \bar{b} * c_t)[\mathbf{k}]\}$ and $\{(\eta_m * c_t)[\mathbf{k}]\}$. Since the CG-solver is the predominant step in Algorithm 4.2, we wish to analyze the complexity associated

with one CG iteration by taking the following into consideration:

- (1) The cost of evaluating the l.h.s. of (4.36),
- (2) The cost of calculating $\{\chi[\mathbf{k}]\}$ is averaged over the total number Q of CG iterations (the $\{\chi[\mathbf{k}]\}$ are calculated outside the CG-loop),
- (3) The convolution in the r.h.s. of (4.36) is ignored as it can be precomputed.

Thus, we require

$$\frac{d_{\mathbf{L}}C}{Q} + (2d_{\mathbf{L}} + 1)C + d_{\mathbf{L}}N \quad (4.39)$$

operations for one CG-iteration in Algorithm 4.2.

Similarly, for Algorithm C.1, at each iteration, we need to evaluate the constants $\{\mathcal{A}_{1\mathbf{m}}\}$, $\{\mathcal{A}_{2\mathbf{m}}\}$, and $\{\mathcal{A}_{3\mathbf{m}}\}$ (see Appendix C) which are of the same size as $\{c_t[\mathbf{k}]\}$. Going through the corresponding expressions, it can be deduced that the associated cost is

$$C(4d_{\mathbf{L}} + 2) + N(d_{\mathbf{L}} + 1) \quad (4.40)$$

for one iteration of Algorithm C.1. From (4.39) and (4.40), it is clear that Algorithm C.1 is more demanding than Algorithm 4.2 which is coherent with our discussion in Section 4.5.2. Similarly, from a quality point of view, Figures 4.2 and 4.3 indicate that Algorithm 4.2 is preferable to Algorithm 4.1 for strong Gaussian noise. Therefore, Algorithm 4.2 is best-suited to carry out regularized interpolation in a general setting. However, at low noise levels, the effect of regularization becomes negligible (the output SNR curves in Figures 4.2 and 4.3 eventually meet); then, it may be desirable to use Algorithm 4.1 as it has the lowest computational complexity.

4.6 Summary

Standard interpolation performs exact fitting of the given data. In the presence of noise, we have shown that this can have a detrimental influence on the interpolation quality. To interpolate noisy data, we have developed a regularized scheme that counterbalances the effect of noise by imposing smoothness constraints (variational approach) on the resulting continuous-domain solution. We have adopted an integer shift-invariant signal model for interpolation where the model parameters (coefficients of the integer-shift-invariant expansion) are

obtained by minimizing the statistical infidelity of the solution to the given data (negative log-likelihood data term) subject to a regularization.

Based on affine invariance of the solution, we showed that the L_p -norm of a vector derivative is the most suitable choice of regularization for our purpose. In addition to Tikhonov-like quadratic functional (which leads to smoothing-spline-like interpolants in Chapter 3), this includes non-quadratic ones such as the edge-preserving total-variation (TV) regularization. We have presented algorithms based on the MM (majorize-minimize) strategy to perform the optimization of the coefficients for non-quadratic cost criteria.

We have carried out 2-D rotation experiments in the presence of noise and numerically quantified the performance of standard (non-regularized) and regularized interpolation. We observed that regularized interpolation always yields a significant improvement in quality over standard interpolation. Furthermore, among the regularized methods, we could note that non-quadratic regularized (NQR) interpolation consistently outperforms the quadratic regularized (QR) one; we associate this phenomenon to the fact that NQR interpolation preserves edges and achieves better noise reduction than QR interpolation. We conclude from these observations that regularized interpolation, specifically, NQR interpolation, can be of potential interest in medical-imaging applications.

Chapter 5

Data-driven Methods for Risk Estimation

5.1 Introduction

The successful application of regularization in the variational formulation of image reconstruction (discussed in the previous two chapters) entails the selection of an appropriate value for the regularization parameter λ . The same problem is encountered in the more general context of inverse problems in imaging¹ (e.g., deconvolution, and the less-ill-posed denoising problem) where parameters of a reconstruction algorithm—represented as the vector $\boldsymbol{\lambda}$ that includes e.g., number of iterations, step-size of the algorithm, regularization parameter, etc—must be adjusted “correctly” in order to obtain satisfactory results. To accomplish this, researchers usually resort to empirical methods: The most-common techniques include the use of the discrepancy principle [114], the L-curve methods [115–117] and generalized cross-validation (GCV) [39, 114, 118–126]. Alternatively, the problem can also be formulated in a Bayesian framework by imposing model-based constraints as prior knowledge on the noise-free signal [127–130].

¹This chapter is based on the following articles:
S. Ramani, T. Blu and M. Unser, “Monte-Carlo SURE: A Black-Box Optimization of Regularization Parameters for General Denoising Algorithms,” *IEEE Transactions on Image Processing*, vol. 17, no. 9, pp. 1540–1554, 2008, and
S. Ramani, C. Vonesch and M. Unser, “Deconvolution of 3D Fluorescence Micrographs with Automatic Risk Minimization,” *ISBI’08*, Paris, France, pp. 732–735, May 14–17, 2008.

In the context of denoising and deconvolution, however, the mean-squared error (MSE)—also called the *risk*—of the reconstruction is the preferred measure of quality to optimize λ . This is also the case with the regularized interpolation algorithm (see Chapter 4) where we use the SNR measure which is directly dependent on the MSE.² Unfortunately, MSE depends on the noise-free signal (as is the case with Equation (4.38) in Chapter 4) which is generally unavailable or unknown *a priori*. A practical approach, therefore, is to replace the true MSE by a reliable estimate in the scheme of things. This forms the main theme of this chapter where we review and derive analytical formulas for risk estimates (that is, estimates of MSE): The key idea is to exploit the statistical properties of the noise which helps us get away with the dependence of MSE on the original unknown signal. Accordingly, we have to specify a suitable noise model for our imaging/sampling process—this is not difficult since, in many practical applications, noise is more easy to characterize or observe than the original signal. The corresponding risk estimate will then depend on the parameters of the noise model which can be also be estimated from the given data.

5.2 Contributions and Chapter Organization

In this chapter, we focus on two types of noise models: The first is the additive Gaussian model which is the most prominent one in the signal-processing literature. Here, we review a well-known result in the statistical [131] and image-denoising literature [132–136] that goes by the name *Stein’s unbiased risk estimate*—SURE. We present details of the derivation of SURE for the denoising problem in Section 5.3.

Next, we consider a Poisson-Gaussian mixture model that is appropriate for fluorescence microscopy imaging [100, 137, 138]. In this case, we derive a new risk estimate for the problem of deconvolution of 3-D fluorescence micrographs in Section 5.4.1. We also propose a simple technique for estimating the noise-model parameters from the given data.

We provide simulation results in Sections 5.3.3 and 5.4.4 which justify the applicability of risk estimates for tuning λ in denoising and deconvolution problems. We also present experimental results on real-world fluorescence microscopy data. Specifically, we focus on linear algorithms³ in this chapter in order to illustrate the concepts clearly. The case of non-linear algorithms is treated in Chapter 6.

²In fact, the denominator of the SNR measure in Equation (4.38) is nothing but MSE.

³Linear algorithms are those whose outputs depend on the input in a linear fashion.

5.3 Denoising Using Stein’s Unbiased Risk Estimate

5.3.1 Notation and Problem formulation

We slightly digress from the continuous-domain notations introduced in Chapter 2 and work with discrete-domain quantities to accurately formulate what we said earlier in this chapter. We adopt the standard vector formulation of a denoising problem: We observe the noisy data $\mathbf{y} \in \mathbb{R}^N$ given by

$$\mathbf{y} = \mathbf{x} + \mathbf{b}, \quad (5.1)$$

where $\mathbf{x} \in \mathbb{R}^N$ represents the vector containing the samples of the unknown deterministic noise-free signal and $\mathbf{b} \in \mathbb{R}^N$ denotes the vector containing zero-mean white Gaussian noise of variance σ^2 , respectively. We are given a denoising algorithm which is represented by the operator $\mathbf{f}_{\boldsymbol{\lambda}} : \mathbb{R}^N \rightarrow \mathbb{R}^N$ that maps the input data \mathbf{y} onto the signal estimate:

$$\tilde{\mathbf{x}} = \mathbf{f}_{\boldsymbol{\lambda}}(\mathbf{y}), \quad (5.2)$$

where $\boldsymbol{\lambda}$ represents the set of parameters characterizing $\mathbf{f}_{\boldsymbol{\lambda}}$; these should be adjusted appropriately to yield the best estimate of the signal.

Here, we are interested in the optimization of $\boldsymbol{\lambda}$ so as to minimize MSE given by

$$\text{MSE}(\boldsymbol{\lambda}) = \frac{1}{N} \|\mathbf{x} - \mathbf{f}_{\boldsymbol{\lambda}}(\mathbf{y})\|^2 = \frac{1}{N} \sum_{k=1}^N (x_k - f_{\boldsymbol{\lambda}k}(\mathbf{y}))^2, \quad (5.3)$$

where $f_{\boldsymbol{\lambda}k}(\mathbf{y})$ represents the k -th component of the vector $\mathbf{f}_{\boldsymbol{\lambda}}(\mathbf{y})$. By expanding the r.h.s. of (5.3), we get

$$\text{MSE}(\boldsymbol{\lambda}) = \frac{1}{N} \|\mathbf{x}\|^2 + \frac{1}{N} \|\mathbf{f}_{\boldsymbol{\lambda}}(\mathbf{y})\|^2 - \frac{2}{N} \mathbf{x}^T \mathbf{f}_{\boldsymbol{\lambda}}(\mathbf{y}).$$

The first term on the r.h.s. of the above equation, $\|\mathbf{x}\|^2$, is an irrelevant constant that does not depend on $\boldsymbol{\lambda}$. The term $\|\mathbf{f}_{\boldsymbol{\lambda}}(\mathbf{y})\|^2$ is easily calculated as it depends on the output $\mathbf{f}_{\boldsymbol{\lambda}}(\mathbf{y})$. The only term that prevents us from using $\text{MSE}(\boldsymbol{\lambda})$ is $\mathbf{x}^T \mathbf{f}_{\boldsymbol{\lambda}}(\mathbf{y})$ since this term cannot be computed as \mathbf{x} is usually unknown *a priori*. In order to circumvent this difficulty, we propose to estimate $\text{MSE}(\boldsymbol{\lambda})$ (that is, compute SURE) purely from \mathbf{y} and $\mathbf{f}_{\boldsymbol{\lambda}}(\mathbf{y})$, as illustrated by the “MSE estimation” box in Figure 5.1.




Figure 5.1: The signal estimate $\tilde{\mathbf{x}}$ is obtained by applying the λ -dependent denoising algorithm on the observed data \mathbf{y} . The MSE box then computes the estimate $\text{SURE}(\lambda)$ of the MSE between the noise-free \mathbf{x} and the denoised $\tilde{\mathbf{x}}$ as a function of λ , knowing only \mathbf{y} and $\mathbf{f}_\lambda(\mathbf{y})$. The best estimate of the signal is obtained by finding that λ which minimizes the surrogate mean-squared error.

5.3.2 Stein’s Unbiased Risk Estimate—SURE

In his hallmark paper [131], Stein established the framework for unbiased estimation of the risk (or MSE) of an arbitrary estimator in the presence of Gaussian noise. SURE is a well-established technique in the statistical literature and is beginning to gain attention within the signal-processing community. Notably, it has been used in the context of (orthonormal) wavelet denoising [132, 133] where the SURE strategy has proven to be quite powerful and has been incorporated in some state-of-the-art algorithms [134–136]. In what follows, we briefly review the theory of SURE and then illustrate the concept in the simpler case of a linear algorithm, which also yields a closed-form solution.

5.3.2.1 Theoretical Background

In the sequel, we assume that \mathbf{f}_λ is a continuous and bounded operator (that is, the input-output mapping is continuous and a small perturbation of the input necessarily results in a small perturbation of the output). We also require that

the divergence of \mathbf{f}_λ with respect to the data \mathbf{y} given by

$$\operatorname{div}_{\mathbf{y}}\{\mathbf{f}_\lambda(\mathbf{y})\} = \sum_{k=1}^N \frac{\partial f_{\lambda k}(\mathbf{y})}{\partial y_k} \quad (5.4)$$

is well defined in the weak sense.⁴ Then, the following theorem, which is a version of Stein's result in [131], provides us with an expression for an unbiased estimate—SURE—of MSE.

Theorem 5.1. *Given \mathbf{y} as in (5.1), SURE corresponding⁵ to $\mathbf{f}_\lambda(\mathbf{y})$ is a random variable $\eta : \mathbb{R}^N \rightarrow \mathbb{R}$, specified as*

$$\eta(\lambda) = \frac{1}{N} \|\mathbf{y} - \mathbf{f}_\lambda(\mathbf{y})\|^2 - \sigma^2 + \frac{2\sigma^2}{N} \operatorname{div}_{\mathbf{y}}\{\mathbf{f}_\lambda(\mathbf{y})\}. \quad (5.5)$$

Moreover, $\eta(\lambda)$ is an unbiased estimator of $\operatorname{MSE}(\lambda)$; that is,

$$E_{\mathbf{b}} \left\{ \frac{1}{N} \|\mathbf{x} - \mathbf{f}_\lambda(\mathbf{y})\|^2 \right\} = E_{\mathbf{b}} \{ \eta(\lambda) \}, \quad (5.6)$$

where $E_{\mathbf{b}}\{\bullet\}$ represents the expectation with respect to \mathbf{b} .

Proof. The proof is given in Appendix D.1. ■

In the SURE formulation, $\operatorname{MSE}(\lambda)$ is estimated purely based on the input data \mathbf{y} , the divergence of $\mathbf{f}_\lambda(\mathbf{y})$, and the noise statistics (here, it depends on the noise variance σ^2); it requires no knowledge whatsoever of the noise free signal \mathbf{x} . The basis for the approach is that there are many more data points⁶ than unknown parameters λ . Therefore, thanks to the law of large numbers, both $\frac{1}{N} \|\mathbf{x} - \mathbf{f}_\lambda(\mathbf{y})\|^2$ and $\operatorname{div}_{\mathbf{y}}\{\mathbf{f}_\lambda(\mathbf{y})\}$ are quite stable estimates of $E_{\mathbf{b}}\{\frac{1}{N} \|\mathbf{x} - \mathbf{f}_\lambda(\mathbf{y})\|^2\}$ and $E_{\mathbf{b}}\{\operatorname{div}_{\mathbf{y}}\{\mathbf{f}_\lambda(\mathbf{y})\}\}$, respectively, meaning that SURE provides a

⁴That is to say, $\langle \operatorname{div}_{\mathbf{y}}\{\mathbf{f}_\lambda(\mathbf{y})\}, \psi(\mathbf{y}) \rangle = -\langle \mathbf{f}_\lambda(\mathbf{y}), \nabla_{\mathbf{y}}\psi(\mathbf{y}) \rangle$, $\forall \psi \in \mathcal{S}$ where $\nabla_{\mathbf{y}}\psi(\mathbf{y})$ is the gradient of $\psi(\mathbf{y})$ with respect to \mathbf{y} and \mathcal{S} is the space of rapidly decaying test functions that are infinitely differentiable. We also require (as explain in Appendix D.1) $E_{\mathbf{b}} \left\{ \left| \frac{\partial f_{\lambda k}(\mathbf{y})}{\partial y_k} \right| \right\} < \infty$, $k = 1, 2, \dots, N$ [131, 139], which is satisfied by most algorithms encountered in practice.

⁵It can be easily shown that when the additive Gaussian noise has mean $\boldsymbol{\mu} \in \mathbb{R}^N$ and is colored with covariance matrix \mathbf{C} , $\eta(\lambda) = \frac{1}{N} \|\mathbf{y} - \mathbf{f}_\lambda(\mathbf{y})\|^2 - \sigma^2 + \frac{2}{N} \boldsymbol{\mu}^T \mathbf{f}_\lambda(\mathbf{y}) + \frac{2}{N} \operatorname{div}_{\mathbf{y}}\{\mathbf{C} \mathbf{f}_\lambda(\mathbf{y})\}$. This is slightly more general in that it reduces to the expression in the r.h.s. of (5.5) when $\boldsymbol{\mu} = \mathbf{0}$ and $\mathbf{C} = \sigma^2 \mathbf{I}$ (corresponding to the zero-mean white noise scenario).

⁶For image denoising applications, N represents the number of pixels and is usually very large; typically, $N \geq 256^2$.

fairly accurate proxy for the true MSE. Hence, it can be applied for data-driven optimization of a wide range of denoising problems.

The divergence $\text{div}_{\mathbf{y}}\{\mathbf{f}_{\lambda}(\mathbf{y})\}$ requires the evaluation of N partial derivatives of \mathbf{f}_{λ} . However, when \mathbf{f}_{λ} is a linear algorithm, the divergence evaluation simplifies significantly as demonstrated next. For the case of non-linear algorithms (including those for which \mathbf{f}_{λ} does not have an explicit functional form), we develop a novel Monte-Carlo procedure for divergence estimation in Chapter 6.

5.3.2.2 Application to Denoising and Regularized Interpolation Based on Linear Algorithms

Linear algorithms are usually associated with quadratic cost functions: The better-known examples are Tikhonov filters [117, 122] and smoothing splines [70, 78, 81, 140] and quadratic-regularized (QR) interpolation (see Chapter 4) in the variational setting, MAP estimators under the Gaussian prior [129], and Wiener filter [89, 122] in the stochastic setting. Such estimators can be described by the following matrix transformation:

$$\mathbf{f}_{\lambda}(\mathbf{y}) = \mathbf{F}_{\lambda}\mathbf{y}, \quad (5.7)$$

where \mathbf{F}_{λ} is a $N \times N$ matrix that depends on λ . Thus, the desired divergence term is explicitly evaluated as

$$\text{div}_{\mathbf{y}}\{\mathbf{f}_{\lambda}(\mathbf{y})\} = \text{div}_{\mathbf{y}}\{\mathbf{F}_{\lambda}\mathbf{y}\} = \text{trace}\{\mathbf{F}_{\lambda}\}, \quad (5.8)$$

which yields an explicit expression for SURE.

In this context, smoothing splines and quadratic-regularized (QR) interpolation algorithms deserve a special mention because \mathbf{F}_{λ} corresponding to their discretized-versions is circulant and so its structure can be exploited for efficient computation of the trace. If we re-sample the smoothing splines solution (Section 3.3.4.5 in Chapter 3 with $h(\mathbf{x}) = \delta(\mathbf{x})$) or the quadratic-regularized solution (Section 4.3.3.1 in Chapter 4), it yields a discretized signal estimate of the form

$$\mathbf{f}_{\lambda}(y)[\mathbf{k}] = \sum_{\mathbf{m} \in \mathbb{Z}^d} (y * r_{\lambda})[\mathbf{m}] \varphi(\mathbf{k} - \mathbf{m}) = (y * r_{\lambda} * b)[\mathbf{k}], \quad (5.9)$$

where $\{y[\mathbf{k}]\}$ represents the infinite sequence of noise-corrupted input, $\mathbf{f}_{\lambda}(y)[\mathbf{k}]$ is the \mathbf{k} -th component of the infinite-dimensional output $\mathbf{f}_{\lambda}(y)$, $b[\mathbf{k}] = \varphi(\mathbf{x})|_{\mathbf{x}=\mathbf{k}}$,

and $\boldsymbol{\lambda} = \lambda$ is the regularization parameter. The required divergence is $\text{div}_{\mathbf{y}}\{\mathbf{f}_{\boldsymbol{\lambda}}(y)\}$ whose \mathbf{k} -th component is given by

$$\frac{\partial \mathbf{f}_{\boldsymbol{\lambda}}(y)[\mathbf{k}]}{\partial y[\mathbf{k}]} = (r_{\lambda} * b)[\mathbf{0}]. \quad (5.10)$$

It is independent of \mathbf{k} and can be computed in the Fourier domain as

$$(r_{\lambda} * b)[\mathbf{0}] = \frac{1}{(2\pi)^d} \int_{[0, 2\pi)^d} \underbrace{R_{\lambda}(e^{j\boldsymbol{\omega}})}_{F_{\lambda}(e^{j\boldsymbol{\omega}})} \left(\sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \right) d\boldsymbol{\omega}, \quad (5.11)$$

where $R_{\lambda}(e^{j\boldsymbol{\omega}})$ is the frequency response of r_{λ} .

In the finite dimensional case, the smoothing spline denoised output can be obtained using (5.7) where \mathbf{F}_{λ} is the block-circulant matrix formed from the filter taps $(r_{\lambda} * b)[\mathbf{k}]$ and is diagonalized by the Fast-Fourier-Transform (FFT) matrix. Its eigenvalues are nothing but the samples of the frequency response $F_{\lambda}(e^{j\boldsymbol{\omega}})$ whose sum yields the desired trace.

5.3.3 Results

The simulations presented in this section are based on the data model (5.1): We consider a variety of test images (including some standard ones; see Figure 5.2) corrupted by white Gaussian noise whose standard deviation σ is assumed to be known.⁷ We consider two types of problems: The first is denoising using polyharmonic smoothing splines⁸ (PSS) [78] and the second is quadratic-regularized interpolation (QR-cubic) discussed in Section 4.4.1 of Chapter 4. For denoising experiments, we quantify the performance of the methods using

$$\text{SNR} = 10 \log_{10} \left(\frac{\|\mathbf{x}\|^2}{\|\mathbf{x} - \mathbf{f}_{\boldsymbol{\lambda}}(\mathbf{y})\|^2} \right), \quad (5.12)$$

while for interpolation experiments, we compute SNR given by (4.38).

In both problems, we consider the task of selecting the regularization parameter λ . Our approach is to do this by minimizing SURE ($\eta(\lambda)$ given by (5.5)) which requires the knowledge⁷ of σ . We also compare the SURE-based

⁷ σ can be estimated reliably in practice using the median estimator of Donoho *et al.* [132].

⁸In doing this, we only wish to exemplify the use of SURE and do not contend with state-of-the-art denoising methods.




Figure 5.2: Standard noise-free test images: (a) Boats (512×512), (b) Barbara (512×512), (c) Peppers (256×256), (d) MRI (256×256), and (e) the Shepp-Logan phantom (256×256). In (f), we consider a realization of fractional Brownian motion (fBm) with Hurst exponent 0.5 on a uniform grid of size 256×256 for which the first-degree PSS estimator is optimal in the MSE sense [78].

Table 5.1: Comparison of Performances of GCV- and SURE-based Denoising by First-degree Polyharmonic Smoothing Splines (PSS)

		Input SNR (dB)	4	8	12	16	20
Image	Measure						
Boats	Oracle	11.83	13.69	15.81	18.27	21.20	
	SURE	11.83	13.69	15.81	18.27	21.20	
	GCV	11.76	13.36	14.80	16.04	20.02	
Barbara	Oracle	9.76	11.63	14.08	17.08	20.51	
	SURE	9.76	11.63	14.08	17.08	20.51	
	GCV	9.55	9.55	12.04	16.02	20.00	
MRI	Oracle	12.40	14.63	17.02	19.59	22.35	
	SURE	12.40	14.63	17.02	19.59	22.35	
	GCV	11.98	13.57	14.77	16.07	20.04	
Peppers	Oracle	10.74	12.47	14.70	17.44	20.68	
	SURE	10.74	12.47	14.70	17.44	20.68	
	GCV	10.74	12.42	12.10	16.04	20.01	
Shepp-Logan	Oracle	9.91	11.79	14.13	17.06	20.45	
	SURE	9.91	11.79	14.13	17.06	20.45	
	GCV	9.88	11.78	14.12	17.00	20.28	
fBm	Oracle	15.29	16.87	18.64	20.59	22.85	
	SURE	15.29	16.87	18.64	20.59	22.85	
	GCV	15.29	16.85	18.52	20.12	21.23	

methodology with GCV-based approach [39] (that is, where λ is obtained by minimizing the GCV measure): GCV is a popular alternative to SURE as it does not require the knowledge of noise variance (unlike SURE) and is asymptotically optimal; that is, λ_{GCV} minimizes the MSE as $N \rightarrow \infty$. However, its applicability is rather restricted to linear algorithms for which it is defined as [39]

$$\text{GCV}(\lambda) = \frac{N^{-1} \|\mathbf{y} - \mathbf{F}_\lambda \mathbf{y}\|^2}{(1 - N^{-1} \text{trace}\{\mathbf{F}_\lambda\})^2}. \quad (5.13)$$

5.3.3.1 Denoising Using Polyharmonic Smoothing Splines (PSS)

We implement the smoothing spline estimator of [78] using first-degree polyharmonic B-splines. All SNR values reported in this section were computed by averaging over three independent simulations. The output SNR⁹ obtained by adjusting λ based on SURE and generalized cross-validation (GCV) is tabulated for various input SNRs and test images in Table 5.1. Also listed are the oracle output SNR values corresponding to the case where λ was selected by minimizing $\text{MSE}(\lambda)$.

As seen from the table, the output SNR values corresponding to SURE is equal or very close to the oracle values. While the performance of GCV is similar to that of the SURE at high noise levels, it becomes steadily poorer with decreasing noise level. This may be due to the fact that GCV does only perform optimally under special conditions [141, Proposition 3.1] which are probably not fulfilled in the present experiments.

This is also in tune with the plots of $\text{MSE}(\lambda)$ and $\text{SURE}(\lambda)$ versus λ in Figures 5.3(a) and 5.3(b) where we see that $\text{SURE}(\lambda)$ faithfully reproduces the trend of the oracle $\text{MSE}(\lambda)$ in both cases. Also λ_{SURE} obtained by minimizing $\text{SURE}(\lambda)$ is very close to λ_{Oracle} that achieves the minimum of $\text{MSE}(\lambda)$. However, λ_{GCV} (minimum of $\text{GCV}(\lambda)$) is far from λ_{Oracle} and correspondingly, the performance of the GCV-based approach is not optimal in the MSE-sense compared to the SURE-based procedure.

5.3.3.2 Quadratic-Regularized (QR) Interpolation

Since MSE for regularized interpolation algorithms¹⁰ (the denominator of Equation (4.38) in Chapter 4) is exactly the same as (5.3) up to notation, the SURE formula for denoising (*cf.* Equation (5.5)) is also applicable to the regularized interpolation problem discussed in the previous chapter.¹¹ Specifically, we concentrate on the case of QR-cubic algorithm as it is linear (*cf.* Algorithm 4.1 in Chapter 4); the procedure for computing the desired divergence outlined in Section 5.3.2.2 is therefore applicable to this case too.

We repeat now the experiment (with MRI stack) described in Section 4.4.1 (see Chapter 4) where, for each noisy realization, we compute $\text{GCV}(\lambda)$ and

⁹The input SNR (*cf.* Table 5.1) is computed in all experiments by replacing the denominator in (5.12) by $N\sigma^2$.

¹⁰This supplements the results presented in Section 4.4.1 of Chapter 4.

¹¹But, it does not apply to the problem of reconstruction from nonideal samples discussed in Chapter 3 since that also involves deconvolution implicitly. We shall take up the problem of risk estimation for deconvolution in Section 5.4.1.




Figure 5.3: Plots of $MSE(\lambda)$ and $SURE(\lambda)$ versus λ for a realization (input SNR = 8 dB) of the denoising experiment using polyharmonic smoothing splines (PSS) on (a) Boats image, and (b) MRI image. In both cases, it is clearly seen that $SURE(\lambda)$ captures the trend of the oracle $MSE(\lambda)$. Moreover, λ_{SURE} is very close to λ_{Oracle} , while λ_{GCV} is far from λ_{Oracle} which yields a sub-optimal performance in terms of MSE.




Figure 5.4: Plots of $\text{MSE}(\lambda)$ and $\text{SURE}(\lambda)$ versus λ for two realizations (input SNR = 0 dB) of the rotation experiment (see Section 4.4.1 in Chapter 4) with QR-cubic interpolation. Again, in both cases, $\text{SURE}(\lambda)$ captures the trend of $\text{MSE}(\lambda)$ and unlike λ_{GCV} , λ_{SURE} is close to λ_{Oracle} .

SURE(λ) along with MSE(λ). The noise variance was assumed to be known⁷ in all experiments for computing SURE(λ). We plot SURE(λ) and MSE(λ) as functions of λ in Figures 5.4(a) and 5.4(b) for two realizations of this experiment at 0 dB input SNR (test images shown in the inset). It is clearly seen that the SURE curves capture the trend of the corresponding MSE curves. Moreover, they also yield very close estimates of the optimal λ in both illustrations, while GCV fails to achieve this. This indicates that a SURE-based tuning procedure is more reliable than GCV even at such low levels of input SNR.

We further quantify the performance of GCV and SURE in Table 5.2 where we list out the output SNR (averaged over all the realizations of the MRI-stack experiment corresponding to an input SNR) obtained by minimizing GCV and SURE for the QR-cubic interpolation algorithm. Also provided are the oracle values of the output SNR (obtained by minimizing the corresponding MSE). Again, we see that the GCV-based result is far from the oracle for input SNRs in the range $8 \leq \text{SNR} \leq 28$ dB. This behavior is similar to that observed in Table 5.1 and the same explanation applies here too. But SURE-based tuning consistently yields SNR values very close to the oracle.

One must also note the following important difference between the denoising and the interpolation experiment: Unlike the former (see Section 5.3.3.1) where the reported SNR accounts only for denoising, the output SNR values for the latter take into account both the effect of denoising and the detrimental action of interpolation. Thus, in Table 5.2, we observe a saturation effect of output SNR which is quite expected: At lower noise levels, the effect of interpolation is more pronounced than that of denoising and therefore, the output SNR due to GCV, SURE and even the oracle MSE, converge to the same value corresponding to the non-regularized cubic B-spline interpolation (see Figures 4.2 and 4.3 in Chapter 4).

The results from Sections 5.3.3.1 and 5.3.3.2 confirm our claim that SURE provides a reliable means of estimating the MSE. Moreover, we demonstrated that SURE is more robust than GCV and that SURE-based adjustment of parameters yields a performance that is close to the oracle-optimum. These observations indicate that SURE can be reliably employed for data-driven adjustment of parameters of (linear) denoising and regularized interpolation algorithms. The method can also be extended for non-linear algorithms, as shown in Chapter 6.

Table 5.2: Output SNR for QR-cubic interpolation obtained by MSE-based (Oracle), SURE-based and GCV-based tuning of λ

Input SNR (dB)	Output SNR (dB)		
	Oracle	SURE	GCV
0	13.95	13.95	13.82
4	15.75	15.75	15.43
8	17.66	17.66	16.94
12	19.71	19.71	18.28
16	21.95	21.95	19.67
20	24.40	24.40	21.39
24	27.09	27.09	25.20
28	30.03	30.03	29.09
32	33.19	33.19	32.86
36	36.39	36.39	36.39
40	39.35	39.35	39.23
44	41.78	41.78	41.88
48	43.43	43.42	43.40
52	44.38	44.38	44.37
56	44.83	44.82	44.80
60	45.04	45.04	44.99

5.4 Deconvolution of 3-D Fluorescence Micrographs

Deconvolution is widely used to enhance 3-D fluorescence microscopy images, both to facilitate their visual inspection and to improve the results of subsequent computer-assisted analysis steps [100]. Although many deconvolution techniques are available [142], they seldom yield agreeable results without human intervention—they are not completely automatic. As explained earlier, the oft encountered problem here is the selection of the appropriate regularization parameter λ . Here, we extend the concept of risk estimation (discussed in Section 5.3) to the deconvolution problem and propose to apply it for adjusting λ .

Since, in practical applications, fluorescence images are obtained via a CCD detector, the additive Gaussian noise model alone is inadequate; one must also take into account the signal-dependent Poisson noise model [100]. To this end,

we consider a Poisson-Gaussian mixture model where the Poisson component accounts for photon-counting noise and the Gaussian component describes the read-out noise of the CCD detector [137]; these two components describe the CCD characteristics satisfactorily [138]. We then derive a new unbiased risk estimate (URE) for this data model. To the best of our knowledge, this has neither been documented before, nor applied for deconvolution of 3-D fluorescence microscopy images. The computation of the proposed URE requires the knowledge of the gain of the Poisson component and the mean and the variance of Gaussian read-out noise. These quantities depend on the calibration of the microscope and are either unspecified or difficult to measure in practice. We therefore also propose a simple technique to estimate these parameters from the given data.

In what follows, we briefly describe the image formation model and derive the formula for the new URE that corresponds to this model. We then show how to compute this for Tikhonov deconvolution (a linear algorithm) and validate it experimentally by presenting numerical results on simulated data. Following that we apply our method to real 3-D microscopy data.

5.4.1 New Unbiased Risk Estimate for a Poisson-Gaussian Mixture Model

5.4.1.1 Image Formation

The image formation of a wide-field microscope is classically described by the following 3-D convolution model [100]: $Y(\bullet) = (h * \chi)(\bullet)$ where $\chi(\bullet)$ is the 3-D object of interest (we have neglected the effect of the background), $Y(\bullet)$ is the intensity map of the blurred object and $h(\bullet)$ (under the paraxial approximation) is the shift-invariant blur function of the microscope [100]. Typically, the blurred image Y is digitized by a CCD detector array which, in the process, introduces measurement noise. For numerical tractability, we apply a suitable discretization of the continuous convolution ($h * \chi$) which leads to the following matrix formulation of the model:

$$\mathbf{y} = \alpha \mathcal{P}\{\mathbf{H}\mathbf{x}\} + \mathbf{b}, \quad (5.14)$$

where α is the gain factor and the operator $\mathcal{P}\{\bullet\}$ represents the effect of Poisson noise; the mean and variance of the Gaussian read-out noise \mathbf{b} are denoted by μ and σ^2 , respectively and the matrix \mathbf{H} , which approximates the continuously defined h , is treated as block-circulant. Additionally, we assume independence

of the individual components of the random variable \mathbf{y} and that of the photon-counting process and the read-out noise.

5.4.1.2 Unbiased Risk Estimate—URE

We use the notation introduced in Section 5.3.1: Here, $\mathbf{f}_\lambda : \mathbb{R}^N \rightarrow \mathbb{R}^N$ represents a deconvolution algorithm that operates on the $N \times 1$ data vector \mathbf{y} to yield the deconvolved signal estimate $\tilde{\mathbf{x}}$; that is, $\tilde{\mathbf{x}} = \mathbf{f}_\lambda(\mathbf{y})$, where $\lambda = \lambda$ is the regularization parameter. We also assume that \mathbf{f}_λ can be decomposed as

$$\mathbf{f}_\lambda(\mathbf{y}) = \mathbf{H}^T \bar{\mathbf{f}}_\lambda(\mathbf{y}),$$

where $\bar{\mathbf{f}}_\lambda$ is a bounded and continuous operator with a well-defined divergence¹² (see Equation (5.4)). Our interest here is to derive a formula for an unbiased risk estimate (URE) of $\tilde{\mathbf{x}}$.

Definition 5.1. *The risk estimate that corresponds to the Poisson-Gaussian (PG) mixture data model (5.14) is a random variable $\eta : \mathbb{R}^N \rightarrow \mathbb{R}$ given by*

$$\eta_{\text{PG}}(\boldsymbol{\lambda}) = \frac{1}{N} \|\mathbf{x}\|^2 + \frac{1}{N} \|\mathbf{f}_\lambda(\mathbf{y})\|^2 + \frac{2\sigma^2}{N\alpha} \text{div}_{\mathbf{y}}\{\bar{\mathbf{g}}_\lambda(\mathbf{y})\} - \frac{2}{N\alpha} (\mathbf{y} - \mu\mathbf{1})^T \bar{\mathbf{g}}_\lambda(\mathbf{y}),$$

where $\mathbf{1}$ is a $N \times 1$ column vector of 1's, \mathbf{e}_k is a $N \times 1$ vector whose components are all zero except for the k -th one which is unity and the components of $\bar{\mathbf{g}}_\lambda$ are such that

$$\bar{\mathbf{g}}_{\lambda k}(\mathbf{y}) = \bar{\mathbf{f}}_{\lambda k}(\mathbf{y} - \alpha \mathbf{e}_k).$$

The following theorem then ensures that $\eta_{\text{PG}}(\boldsymbol{\lambda})$ is indeed unbiased.

Theorem 5.2. *Let \mathbf{f}_λ be such that $E_{\mathbf{y}}\{|\bar{\mathbf{f}}_{\lambda k}(\mathbf{y})|\} < \infty$, and*

$$E_{\mathbf{y}} \left\{ \left| \frac{\partial \bar{\mathbf{f}}_{\lambda k}(\mathbf{y})}{\partial y_k} \right| \right\} < \infty,$$

for $k = 1, 2, \dots, N$. Then random variable $\eta_{\text{PG}}(\boldsymbol{\lambda})$ is an unbiased estimate of $\text{MSE}(\boldsymbol{\lambda})$; that is,

$$E_{\mathbf{y}}\{\text{MSE}(\boldsymbol{\lambda})\} = E_{\mathbf{y}}\{\eta_{\text{PG}}(\boldsymbol{\lambda})\},$$

where $E_{\mathbf{y}}\{\bullet\}$ represents the expectation with respect to the random variable \mathbf{y} .

¹²More precisely, we require that (as explained in Appendix D.2) $E_{\mathbf{y}}\{|\bar{\mathbf{f}}_{\lambda k}(\mathbf{y})|\} < \infty$ and $E_{\mathbf{y}} \left\{ \left| \frac{\partial \bar{\mathbf{f}}_{\lambda k}(\mathbf{y})}{\partial y_k} \right| \right\} < \infty$, $k = 1, 2, \dots, N$.

Proof. The proof of this result is given in Appendix D.2. ■

The term $(\mathbf{y} - \mu\mathbf{1})^T \bar{\mathbf{g}}_\lambda(\mathbf{y})$ in $\eta_{\text{PG}}(\lambda)$ corresponds to the contribution of the Poisson component of the noise, while $\text{div}_{\mathbf{y}}\{\bar{\mathbf{g}}_\lambda(\mathbf{y})\}$ is due to the Gaussian component (as exemplified in the proof in Appendix D.2), respectively. We may then interpret $\eta_{\text{PG}}(\lambda)$ as a more general version of Equation (5.5) in the sense that it includes the Gaussian-based risk estimate for deconvolution as a special case: If the Poisson component of the noise were absent, then, $\eta_{\text{PG}}(\lambda)$ reduces to

$$\eta_{\text{G}}(\lambda) = \frac{1}{N}\|\mathbf{x}\|^2 + \frac{1}{N}\|\mathbf{f}_\lambda(\mathbf{y})\|^2 + \frac{2\sigma^2}{N}\text{div}_{\mathbf{y}}\{\bar{\mathbf{f}}_\lambda(\mathbf{y})\} - \frac{2}{N}(\mathbf{y} - \mu\mathbf{1})^T \bar{\mathbf{f}}_\lambda(\mathbf{y}),$$

which is the risk estimate corresponding to a pure additive Gaussian noise model (with mean μ and covariance matrix $\sigma^2\mathbf{I}$) for the deconvolution problem.

The unbiasedness of η_{PG} (*cf.* Theorem 5.2) indicates the equality of η_{PG} and MSE in the event $N \rightarrow \infty$, where N is the number of pixels. For practical purposes, therefore, η_{PG} can be used as a reliable substitute for MSE for very large N (especially image stacks). Moreover, we note that, except for the first term $\|\mathbf{x}\|^2$ in η_{PG} (which is a non-relevant constant for our purpose), all the other terms are purely derived from \mathbf{y} and therefore computable. In this work, we propose to minimize $\eta_{\text{PG}}(\lambda)$ for finding the optimal λ .

5.4.2 Deconvolution Algorithm

The URE η_{PG} has a closed form expression that is in principle computable for certain classes of estimators. In the sequel, we restrict ourselves to the linear setting which is adequate for our purpose and completely tractable mathematically. We consider the Tikhonov-regularized algorithm [142] which yields an explicit expression for \mathbf{f}_λ . The signal estimate is obtained by minimizing the Tikhonov criterion \mathcal{J}_{Tik} (which is based on discrete entities):

$$\begin{aligned} \mathbf{f}_\lambda(\mathbf{y}) &= \arg \min_{\mathbf{u}} J_{\text{Tik}}\{\mathbf{u}\}, \\ \mathcal{J}_{\text{Tik}}\{\mathbf{u}\} &= \|\mathbf{y} - \mathbf{H}\mathbf{u}\|^2 + \lambda\|\mathbf{P}\mathbf{u}\|^2, \end{aligned} \quad (5.15)$$

where \mathbf{P} is the matrix (typically block-circulant) that represents a suitable regularization operator (discretized laplacian or gradient) and λ is regularization parameter. The solution of (5.15) takes the classical form

$$\mathbf{f}_\lambda(\mathbf{y}) = \underbrace{(\mathbf{H}^T\mathbf{H} + \lambda\mathbf{P}^T\mathbf{P})^{-1}}_{\mathbf{F}_\lambda} \mathbf{H}^T \mathbf{y}. \quad (5.16)$$

When \mathbf{H} and \mathbf{P} are block-circulant, so is \mathbf{F}_λ . Then the above matrix solution is efficiently implemented via FFTs. Moreover, commutativity of \mathbf{F}_λ and \mathbf{H}^T (circulant matrices) yields $\tilde{\mathbf{f}}_\lambda(\mathbf{y}) = \mathbf{F}_\lambda \mathbf{y}$. Thus the URE (without the constant $\|\mathbf{x}\|^2$) can be deduced as

$$\eta_{\text{PG, Tik}}(\lambda) = \frac{1}{N} \|\mathbf{F}_\lambda \mathbf{H}^T \mathbf{y}\|^2 + \frac{2}{N\alpha} \begin{pmatrix} (\sigma^2 - \mu\alpha) \text{Trace}\{\mathbf{F}_\lambda\} \\ -(\mathbf{y} - \mu \mathbf{1})^T \mathbf{F}_\lambda \mathbf{y} \\ +\alpha \text{Trace}\{\mathbf{D}_y \mathbf{F}_\lambda\} \end{pmatrix}, \quad (5.17)$$

where $\mathbf{D}_y = \text{diag}\{\mathbf{y}\}$. The circulant nature of \mathbf{F}_λ facilitates the computation of $\text{Trace}\{\mathbf{F}_\lambda\}$ and $\text{Trace}\{\mathbf{D}_y \mathbf{F}_\lambda\}$. The former is easily computed in the Fourier domain: It is the sum of the corresponding DFT coefficients, while the latter simplifies to $\text{Trace}\{\mathbf{D}_y \mathbf{F}_\lambda\} = \frac{1}{N} \text{Trace}\{\mathbf{F}_\lambda\} \mathbf{1}^T \mathbf{y}$.

5.4.3 Estimation of CCD parameters

The use of $\eta_{\text{PG, Tik}}$ requires the knowledge of the CCD dependent parameters (α, μ, σ^2) . While μ and σ^2 (of the read-out noise) may be measured by running separate calibration experiments [137], α is not totally synonymous with the gain that is typically provided in microscopes and must therefore be estimated from the given data. For this purpose, we develop a simple procedure that is essentially based on the following identities:

$$\mu_y[k] \stackrel{\text{def}}{=} E_y\{y_k\} = \alpha(\mathbf{H}\mathbf{x})_k + \mu, \quad (5.18)$$

$$\sigma_y^2[k] \stackrel{\text{def}}{=} \text{Var}_y\{y_k\} = \alpha^2(\mathbf{H}\mathbf{x})_k + \sigma^2, \quad (5.19)$$

where $(\mathbf{H}\mathbf{x})_k$ is the k -th component of $\mathbf{H}\mathbf{x}$. This leads to the relationship

$$\sigma_y^2[k] = \alpha \mu_y[k] + (\sigma^2 - \mu\alpha).$$

Then, the idea is to estimate $\mu_y[k]$ and $\sigma_y^2[k]$, $k = 1, 2, \dots, K$ for some $K < N$ and perform a linear regression analysis on $(\mu_y[k], \sigma_y^2[k])_{k=1}^K$: The slope yields the gain α and the intercept yields the constant $\kappa = (\sigma^2 - \mu\alpha)$.

To estimate $(\mu_y[k], \sigma_y^2[k])_{k=1}^K$, we first segment \mathbf{y} into K non-overlapping regions wherein the underlying blurred signal $\mathbf{H}\mathbf{x}$ is approximately constant in each region. This is accomplished by running a very few number of iterations (typically < 4) of heavily regularized 2-D total-variation denoising (TVD) (see Section 6.4.2 of Chapter 6) on each slice of \mathbf{y} —the advantage of using TVD

is that it not only reduces the noise in \mathbf{y} but automatically yields a piecewise constant approximation \mathbf{y}_{PC} of $\mathbf{H}\mathbf{x}$. We then define level values

$$\epsilon_k = \min\{\mathbf{y}_{\text{PC}}\} + (k - 1)\Delta,$$

and corresponding segments

$$\Omega_k = \left\{ 1 \leq i \leq N : \epsilon_k - \frac{\Delta}{2} \leq y_{\text{PC}i} < \epsilon_k + \frac{\Delta}{2} \right\},$$

where $k = 1, 2, \dots, K$ and Δ is a predefined width of each level. This basically ensures a non-redundant partitioning of the image-stack \mathbf{y} : $\Omega_k \cap \Omega_l = \emptyset$, $k \neq l$ and $\bigcup_{k=1}^K \Omega_k = \mathcal{I}_N$, where $\mathcal{I}_N = \{1, 2, \dots, N\}$.

Estimates of $\mu_{\mathbf{y}}[k]$ and $\sigma_{\mathbf{y}}^2[k]$ are computed for each segment as the corresponding sample mean and sample variance,¹³ respectively:

$$\begin{aligned} \hat{\mu}_{\mathbf{y}}[k] &= \frac{1}{N_k} \sum_{i \in \Omega_k} y_i, \\ \hat{\sigma}_{\mathbf{y}}^2[k] &= \frac{1}{N_k - 1} \sum_{i \in \Omega_k} (y_i - \hat{\mu}_{\mathbf{y}}[k])^2, \end{aligned}$$

where N_k is the cardinality of the set Ω_k . Since most of the noise is sufficiently reduced in \mathbf{y}_{PC} , assuming that $\min\{\mathbf{H}\mathbf{x}\} = 0$ (which is mostly the case with biological images), μ is simply estimated as $\min\{\mathbf{y}_{\text{PC}}\}$. Thus, knowing α , κ and μ , we easily obtain a formula for estimating σ^2 .

5.4.4 Results

5.4.4.1 Simulations

First, we numerically validate the proposed URE $\eta_{\text{PG, Tik}}$ (*cf.* Equation (5.17)) and the estimation scheme described in Section 5.4.3. For this, we performed the following simulation: We used the 2-D uniform blur kernel of size 9×9 and blurred the image show in the inset of Figure 5.5 according to (5.14). The maximum intensity in the noise-free blurred signal ($\mathbf{H}\mathbf{x}$) was set so as to simulate an average of 8 photons per pixel. Then, we generated a noisy realization of the blurred version for the following parameter setting: $(\alpha, \mu, \sigma^2) = (10, 20, 10)$.

¹³The more sophisticated median-of-absolute-deviations (MAD) approach can also be used for estimating the variance. However, due to a large N , sample-variance estimates are sufficiently robust and so we choose to use this for simplicity.

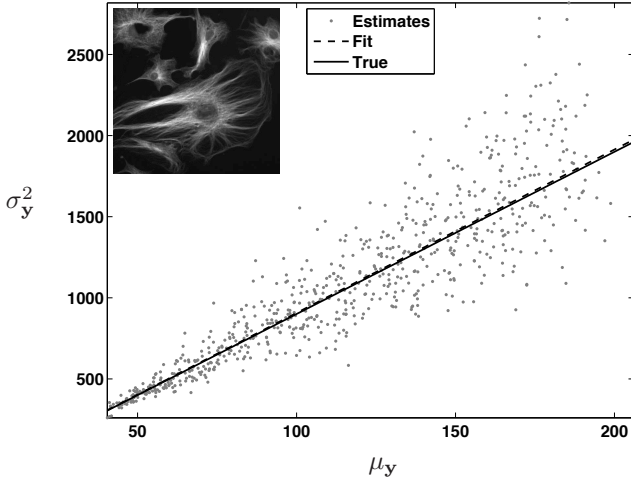


Figure 5.5: Plot of $(\hat{\mu}_y, \hat{\sigma}_y^2)$ and corresponding fit for a noisy realization $(\alpha, \mu, \sigma^2) = (10, 20, 10)$ of the image shown in the inset. Estimated parameters are $(10.04, 18.54, 9.3)$.




Figure 5.6: Plot of $\text{MSE}(\lambda)$, and the URE $\eta_{\text{PG, Tik}}(\lambda)$ computed using estimated parameters $(\hat{\alpha}, \hat{\mu}, \hat{\sigma}) = (10.07, 18.54, 9.3)$. The URE captures the trend of MSE and yields λ_{URE} that is close to the oracle optimum λ_{MSE} unlike λ_{GCV} : Correspondingly, $\text{MSE}(\lambda_{\text{URE}}) < \text{MSE}(\lambda_{\text{GCV}})$ indicating that URE-based approach yields a better result.




Figure 5.7: Zoomed in version of (a) Original unblurred noise-free image, (b) Blurred noisy image, and Deconvolution results obtained by minimizing (c) MSE (Oracle), (d) GCV, and (e) URE. The URE-based result is visually close to the oracle and is also less noisy compared to GCV-based one

From the blurred noisy image, we estimated the CCD parameters using the method described in Section 5.4.3 with $\Delta = 0.2$. Figure 5.5 shows the corresponding regression plot. The estimated parameter values are $\hat{\alpha} = 10.07$, $\hat{\mu} = 18.54$, and $\hat{\sigma} = 9.3$: It is clearly seen that the fit (dashed line) almost overlaps with the true relationship (solid line) between $\mu_{\mathbf{y}}$ and $\sigma_{\mathbf{y}}^2$ indicating the reliability of our estimation procedure.

We then applied the deconvolution algorithm (5.16) using a Laplacian operator for \mathbf{P} . The URE $\eta_{\text{PG,Tik}}$ was computed using the estimated CCD parameter values and plotted in Figure 5.6 in comparison with the MSE (*cf.* Equation (5.3)) without the $\|\mathbf{x}\|^2$ term. It is clearly seen that the URE correctly predicts the shape of MSE curve for the entire range of λ . The slight offset is due to the multiplicative nature of the gain α : Even a slight difference from the true value shifts the URE vertically. But this does not move the optimal value of λ , as seen in Figure 5.6 where the URE closely determines the minimum of the MSE.

For comparison, we also tested the performance of GCV ((see Equation 5.13)): λ_{GCV} yields a sub-optimal performance compared to the λ_{URE} ($\lambda_{\text{GCV}} = 0.0548$ is far from the oracle optimum $\lambda_{\text{Oracle}} = 0.1692$ and is less than $\lambda_{\text{URE}} = 0.1841 \approx \lambda_{\text{Oracle}}$). Correspondingly, the GCV-based result in Figure 5.7(d) is under-regularized [122] and is therefore more noisy compared to the URE-based result in Figure 5.7(e). This is also reflected in the MSE values:

$$\begin{aligned} \text{MSE}(\lambda_{\text{GCV}}) &= 5.81, \\ \text{MSE}(\lambda_{\text{URE}}) &= 2.89, \\ \text{MSE}(\lambda_{\text{Oracle}}) &= 2.87. \end{aligned}$$

Thus, the proposed URE achieves superior performance; the lesser performance of GCV may be due to the fact that it is linked to MSE in the measurement domain (see Equation 5.13) which is suitable for denoising, but less appropriate for deconvolution.

5.4.4.2 Real-World Data

We now present results for automatic Tikhonov deconvolution of real 3-D wide-field fluorescence micrographs. The data was acquired on a Leica DM 5500B system equipped with a cooled CCD camera, at the EPFL BioImaging and Optics Platform.¹⁴ Figure 5.9(a) shows the maximum-intensity projection (MIP) of the acquired 3-D image-stack of a *C. Elegans* embryo ($420 \times 320 \times 270$ voxels ≈ 70 MB in TIFF format). The sample was labeled with three fluorophores: Alexa568

¹⁴Courtesy: Prof. Pierre Gönczy, EPFL/SV/ISREC.




Figure 5.8: Plot of $(\hat{\mu}_y, \hat{\sigma}_y^2)$ for the (a) Red, (b) Green, and (c) Blue channels of real data and corresponding fits.




Figure 5.9: Real Data Set 1: Maximum-intensity projections (MIP) of (a) Noisy real-world data, (b) GCV-based deconvolution result, and (c) URE-based deconvolution result.




Figure 5.10: Real Data Set 2: Maximum-intensity projections (MIP) of (a) Noisy real-world data and (b) URE-based deconvolution result

(vesicles, red channel), Alexa488 (microtubules, green channel), and Hoechst (chromosomes, blue channel). The blurring matrix \mathbf{H} to be used in the deconvolution algorithm (see Equation 5.16)) was generated mathematically based on a standard diffraction-limited PSF model [100] using the manufacturer-provided imaging parameters (magnification: 63 \times ; numerical aperture: 1.4; immersion oil refractive index: 1.518).

Figure 5.8 documents the linear regression on (μ_y, σ_y^2) estimated for data in the red, green, and the blue channel, respectively. The estimated parameters are

$$\begin{aligned} \text{Red channel:} & \quad (\hat{\alpha}, \hat{\mu}, \hat{\sigma}) = (2.15, 217.65, 29.41), \\ \text{Green channel:} & \quad (\hat{\alpha}, \hat{\mu}, \hat{\sigma}) = (2.13, 143.81, 26.51), \\ \text{Blue channel:} & \quad (\hat{\alpha}, \hat{\mu}, \hat{\sigma}) = (1.08, 25.27, 11.09). \end{aligned}$$

For this data set, the MATLAB (FFT-based) implementation of the Tikhonov deconvolution algorithm required approximately 3 GB of memory and took about 20 seconds to produce a deconvolved signal estimate for a single λ on a 2.66 GHz (dual core) Intel-Mac with 6 GB RAM. The three channels were processed separately using the method described above for λ that minimized the GCV measure and $\eta_{\text{PG,Tik}}$. The MIP of the deconvolved data set is shown in Figures 5.9(b) and 5.9(c) corresponding to GCV- and URE-based approaches, respectively. It is seen that the URE-based result has a better contrast and is slightly less noisy (in the Green channel) compared to the GCV-based one. Also, the URE-based result captures a lot more of the filament structure of the microtubules than does the GCV-based one indicating the superior performance of the proposed URE over GCV.

Another set of deconvolution results is presented in Figure 5.10 corresponding to a different *C. Elegans* embryo data set ($464 \times 320 \times 270$ voxels ≈ 76 MB in TIFF format). Here, the reduction of typical wide-field haze in the Red channel is prominent in the URE-based deconvolution result (see Figure 5.10(b)): The red light coming from the vesicles has been reassigned to a more concentrated and brighter volume; this is in accordance with the fact that vesicles are relatively small biological structures. Also, the filament structure of the microtubules is visible with significantly better contrast.

5.5 Summary

In this chapter, we focussed on the problem of data-driven optimization of algorithm parameters λ for some inverse problems (denoising and deconvolution) in

imaging. Our approach was to do this so as to achieve a minimum-MSE performance. In order to circumvent the explicit dependence of MSE on the unknown noise-free signal, we resorted to the concept of risk estimation wherein the MSE is estimated purely from the given data and the mathematical description of the algorithm. We first studied a well-known result due to Stein [131] (Stein's unbiased risk estimate)—SURE—that is applicable whenever the noise is additive white Gaussian. We showed how to compute SURE for linear denoising and regularized interpolation algorithms and demonstrated experimentally that SURE-based selection of the regularization parameter for these two problems yields optimal results in terms of MSE.

Next, we considered the problem of deconvolution of 3-D fluorescence microscopy images. Based on a practical model for CCD data that takes into account Poisson photon-counting noise and Gaussian read-out noise, we derived the formula for a new unbiased risk estimate (URE) for an arbitrary deconvolution algorithm. We also devised a scheme for estimating the CCD parameters (the detector gain, mean and variance of the read-out noise) necessary for computing the proposed URE from the given data. We showed how to compute the URE for a linear deconvolution algorithm and validated it using phantom data. We also presented deconvolution results for real 3D wide-field microscopy images. Based on these results (denoising, regularized interpolation and deconvolution), we come to the conclusion that risk estimation provides a robust framework for data-driven tuning of algorithm parameters based on MSE.

Chapter 6

Monte-Carlo Risk Minimization for Non-linear Algorithms

In this chapter,¹ we focus on the problem of computing SURE² (Stein’s unbiased risk estimate) for MSE-based data-driven adjustment of the parameters ($\boldsymbol{\lambda}$) of an arbitrary non-linear algorithm $\mathbf{f}_{\boldsymbol{\lambda}}$. For ease of analysis, we only deal with the denoising problem, while the technique introduced here can, in principle, also be extended to other inverse problems such as deconvolution [143]. We also show that the technique developed in this chapter is applicable to the non-quadratic regularized (NQR) interpolation problem of Chapter 4.

6.1 Problem Statement

As explained in Section 5.3.2.1 of the previous chapter, the divergence $\operatorname{div}_{\mathbf{y}}\{\mathbf{f}_{\boldsymbol{\lambda}}(\mathbf{y})\}$ of the denoising operator $\mathbf{f}_{\boldsymbol{\lambda}}$ is the key ingredient of SURE. It can be computed analytically only in some special cases such as when the denoising op-

¹This chapter is based on the article: S. Ramani, T. Blu and M. Unser, “Monte-Carlo SURE: A Black-Box Optimization of Regularization Parameters for General Denoising Algorithms,” *IEEE Transactions on Image Processing*, vol. 17, no. 9, pp. 1540–1554, 2008.

²Throughout this chapter, we restrict ourselves to the additive white Gaussian noise model. Nevertheless, we discuss the possibility of extending the current approach to a signal-dependent Poisson noise model in the next chapter.

erator performs a coordinate-wise non-linear mapping, when the signal estimate is obtained by a linear transformation of the noisy data (linear filtering [122]), or when both are combined in a specific way (e.g., wavelet thresholding [132, 133, 135, 136]). For linear algorithms, the desired divergence reduces to the trace of the corresponding matrix transformation (see Section 5.3.2.2 of Chapter 5). However, in a general setting, the explicit evaluation of the divergence is often out of reach. Especially challenging are cases where the functional form of the denoising operator is not known, for example when the denoised output is the result of an iterative optimization procedure. Since most of the variational and Bayesian methods fall into this category, there are many key algorithms for which the evaluation of the required divergence term is neither tractable mathematically nor even feasible numerically.³ In this chapter, we address this limitation by proposing a novel and computationally efficient scheme that is applicable for a general denoising scenario.

6.2 Contributions and Chapter Organization

Our method is based on Monte-Carlo simulation: The denoising algorithm is probed with additive noise and the response signal is manipulated to estimate the desired divergence. This leads to a black-box interpretation of the proposed technique—it completely relies on the output of the denoising operator and does not need any information about its functional form. We also provide sufficiency conditions on the denoising operator so that estimation procedure is stable. We validate the proposed scheme by presenting numerical results for a variety of popular denoising methods—total-variation (TV) denoising, redundant-wavelet soft-thresholding (for which the desired divergence evaluation is known to be a non-trivial task), and some classical ones such as orthonormal-wavelet soft-thresholding.

The chapter is structured as follows: In Section 6.3, we present Monte-Carlo strategies for estimating the divergence of an arbitrary denoising algorithm. First, we review a simple scheme for the special case of arbitrary linear algorithms⁴ and then proceed to describe a new method for arbitrary non-linear operators. In Section 6.4, we briefly describe the working of some denoising

³The divergence computation requires the evaluation of N partial derivatives of the denoising operator where N (typically, $\geq 256^2$) is the number of pixels in image-processing applications. Therefore, for a general denoising operator, this task can be quite overwhelming from a computational point of view.

⁴It must be noted we had concentrated only on the case of linear algorithms characterized by circulant matrices in the previous chapter (see Section 5.3.2.2).

algorithms that we consider in this work following which we numerically verify the validity of the proposed Monte-Carlo procedure in Section 6.5. We present experimental results in Section 6.6 and demonstrate numerically that SURE, computed using the new Monte-Carlo strategy, faithfully imitates the behavior of MSE. We show that it is always capable of uncovering the optimal value of the parameter (regularization parameter for the variational methods and soft-threshold value for the wavelet-based methods). Additionally, we illustrate that the new Monte-Carlo scheme is suitable for denoising methods characterized by multiple parameters. In the process, we compare the performance of these denoising algorithms in terms of visual quality and signal-to-noise ratio (SNR). For the sake of completeness, we also present experimental results to exemplify the applicability of the proposed method for the non-quadratic regularized (NQR) interpolation (see Chapter 4).

6.3 Monte-Carlo Estimation of the Divergence

We now investigate Monte-Carlo techniques for the computation of $\text{div}_{\mathbf{y}}\{\mathbf{f}_{\lambda}(\mathbf{y})\}$. We start by revisiting a method that is valid in the linear case only [144, 145], but which can be very useful when the matrix \mathbf{F}_{λ} is not available explicitly. Following that, we introduce a more general technique that is applicable for arbitrary (non-linear) algorithms.

6.3.1 Linear Algorithm with Unstructured System Matrices

In many practical situations, especially with large data-sets, the matrix \mathbf{F}_{λ} is not available explicitly; instead Equation (5.7) (see Section 5.3.2.2 in Chapter 5) is implemented iteratively by using some suitable numerical solver (e.g., conjugate gradient, multigrid technique [146, 147]). It follows that the trace is not directly accessible. There are matrix methods (such as the power method) that can produce an estimate of $\text{trace}\{\mathbf{F}_{\lambda}\}$ in an iterative fashion starting from (5.7), but they tend to be memory- and computation-intensive. To tackle this difficulty, we propose the use of the following Monte-Carlo algorithm which estimates the required trace stochastically with $O(N)$ computational cost (up to the complexity of realizing (5.7)). It is implemented by applying the estimator to noise only, as described next.

Algorithm 6.1. Monte-Carlo algorithm for estimating $\frac{1}{N}\text{trace}\{\mathbf{F}_\lambda\}$.

- Generate a zero-mean i.i.d. random vector \mathbf{b}' of unit variance.
- For a given $\lambda = \lambda_0$ do the following:
 1. Evaluate $\tilde{\mathbf{b}} = \mathbf{F}_\lambda \mathbf{b}'$ for $\lambda = \lambda_0$
 2. Compute the estimate of $\frac{1}{N}\text{trace}\{\mathbf{F}_\lambda\}$ as $\frac{1}{N}\mathbf{b}'^T \tilde{\mathbf{b}}$

Algorithm 6.1 is a standard procedure in the literature [144,145] and has a twofold advantage over the iterative matrix methods mentioned before: Firstly, it is memory-efficient because, at any given point, it only stores $\mathbf{F}_\lambda \mathbf{b}'$ and not \mathbf{F}_λ itself. Secondly, from a computation point of view, the method is as good as the initial algorithm itself since we can simply apply it to noise. The validity of the algorithm is guaranteed by the fact that the random variable $\mathbf{b}'^T \mathbf{F}_\lambda \mathbf{b}'$ is an unbiased estimator of $\text{trace}\{\mathbf{F}_\lambda\}$, which is a well-established result in the literature [144,145,148,149].

Proposition 6.1. Let \mathbf{b}' be a zero-mean i.i.d. random vector with unit variance and $\hat{t} = \frac{1}{N}\mathbf{b}'^T \mathbf{F}_\lambda \mathbf{b}'$, where the factor $\frac{1}{N}$ accounts for the averaging of MSE (see Equation (5.3) in Section 5.3.1 of Chapter 5) over all samples. Then,⁵

$$E_{\mathbf{b}'}\{\hat{t}\} = \frac{1}{N}\text{trace}\{\mathbf{F}_\lambda\}. \quad \blacksquare \quad (6.1)$$

For image-processing applications, it is reasonable to believe that a single realization of \mathbf{b}' will yield a sufficiently low variance estimate [144,150]. This is because, in practice, most denoising algorithms operate only “locally” (that is, \mathbf{F}_λ is more or less diagonal with rapidly decaying off-diagonal elements). Qualitatively speaking, the components $\{\tilde{b}_i\}_{i=1}^N$ of $\tilde{\mathbf{b}}$ are therefore “nearly” independent. Since N is large for images (typically $N \geq 256^2$), by law of large numbers $\hat{t} - E_{\mathbf{b}'}\{\hat{t}\}$ does not fluctuate more than $\frac{1}{\sqrt{N}}$; this eliminates any necessity for additional algorithm evaluations. A more quantitative argument can be made by computing the variance of \hat{t} which is given by

$$\text{Var}_{\mathbf{b}'}\{\hat{t}\} = \frac{1}{N^2} \left(\text{trace}\{\mathbf{F}_\lambda^T \mathbf{F}_\lambda\} + \text{trace}\{\mathbf{F}_\lambda^2\} + (E_{\mathbf{b}'}\{b'^4\} - 3) \sum_{k=1}^N F_{kk}^2 \right),$$

⁵The proof of this result is straightforward and can be found in [144].

where F_{kk} is the k -th diagonal element of \mathbf{F}_λ and $E_{b'}\{b'^4\}$ is the fourth-order moment of the random variable b' . Again, since \mathbf{F}_λ is “approximately” diagonal, the quantities $\text{trace}\{\mathbf{F}_\lambda^T \mathbf{F}_\lambda\}$ and $\text{trace}\{\mathbf{F}_\lambda^2\}$ are of the order of N . The variance is then bounded as

$$\text{Var}_{\mathbf{b}'}\{\hat{t}\} \leq \frac{\text{constant}}{N}.$$

Thus, in principle, \hat{t} asymptotically converges to $N^{-1}\text{trace}\{\mathbf{F}_\lambda\}$ in the mean-squared-error sense. A further option is to reduce $\text{Var}_{\mathbf{b}'}\{\hat{t}\}$ by selecting a \mathbf{b}' that has small a fourth-order moment. For instance, it has been suggested to choose \mathbf{b}' such that its components are either $+1$ or -1 with probability $\frac{1}{2}$ [145, 148, 149]; for such a \mathbf{b}' , the variance is lower than that obtained using a Gaussian \mathbf{b}' [145, 149].

6.3.2 General Algorithm for Non-Linear Problems

Similar to the technique described above, our strategy for a non-linear \mathbf{f}_λ is essentially based on probing the system with noise, but is slightly more involved because of the nonlinearity of \mathbf{f}_λ . Specifically, we propose to investigate $\mathbf{f}_\lambda(\mathbf{y} + \varepsilon\mathbf{b}')$ which may be thought of as a random perturbation around the operating point of the algorithm. The output is then compared with $\mathbf{f}_\lambda(\mathbf{y})$ which yields a differential response of \mathbf{f}_λ evaluated at \mathbf{y} . The following theorem states that this differential response yields the desired divergence as ε decreases.

Theorem 6.1. *Let \mathbf{b}' be a zero-mean i.i.d. random vector (that is independent of \mathbf{y}) with unit variance and bounded higher order moments. Then,*

$$\text{div}_{\mathbf{y}}\{\mathbf{f}_\lambda(\mathbf{y})\} = \lim_{\varepsilon \rightarrow 0} E_{\mathbf{b}'} \left\{ \mathbf{b}'^T \left(\frac{\mathbf{f}_\lambda(\mathbf{y} + \varepsilon\mathbf{b}') - \mathbf{f}_\lambda(\mathbf{y})}{\varepsilon} \right) \right\}, \quad (6.2)$$

provided that \mathbf{f}_λ admits a well-defined second-order Taylor expansion. Otherwise, the expression is still valid in the weak sense (sufficient to apply Theorem 6.1) provided that

$$\|\mathbf{f}_\lambda(\mathbf{y})\| \leq C_0(1 + \|\mathbf{y}\|^{n_0}), \quad (6.3)$$

for some $n_0 > 1$ and $C_0 > 0$ (that is, \mathbf{f}_λ is tempered).

Proof. We write the second-order Taylor expansion of $\mathbf{f}_\lambda(\mathbf{y} + \varepsilon\mathbf{b}')$ as

$$\mathbf{f}_\lambda(\mathbf{y} + \varepsilon\mathbf{b}') = \mathbf{f}_\lambda(\mathbf{y}) + \varepsilon \mathbf{J}_{\mathbf{f}_\lambda}(\mathbf{y}) \mathbf{b}' + \varepsilon^2 \mathbf{r}_{\mathbf{f}_\lambda}, \quad (6.4)$$

where $\mathbf{J}_{\mathbf{f}_\lambda(\mathbf{y})}$ is the Jacobian matrix of \mathbf{f}_λ evaluated at \mathbf{y} and $\mathbf{r}_{\mathbf{f}_\lambda}$ represents the vector containing the (Lagrange) remainder terms corresponding to each component of \mathbf{f}_λ . In this case, the components $r_{\mathbf{f}_\lambda k}$ of $\mathbf{r}_{\mathbf{f}_\lambda}$ are bounded in the expectation sense; that is, $E_{\mathbf{b}'}\{|r_{\mathbf{f}_\lambda k}|\} < +\infty$, $k = 1, 2, \dots, N$.

Then, subtracting $\mathbf{f}_\lambda(\mathbf{y})$ from (6.4) and multiplying by \mathbf{b}'^T from the left yields

$$\begin{aligned} E_{\mathbf{b}'}\{\mathbf{b}'^T(\mathbf{f}_\lambda(\mathbf{y} + \varepsilon\mathbf{b}') - \mathbf{f}_\lambda(\mathbf{y}))\} &= \varepsilon E_{\mathbf{b}'}\{\mathbf{b}'^T \mathbf{J}_{\mathbf{f}_\lambda}(\mathbf{y}) \mathbf{b}'\} + \varepsilon^2 E_{\mathbf{b}'}\{\mathbf{b}'^T \mathbf{r}_{\mathbf{f}_\lambda}\} \\ &= \varepsilon \text{trace}\{\mathbf{J}_{\mathbf{f}_\lambda}(\mathbf{y})\} + C_2 \varepsilon^2, \end{aligned}$$

where $E_{\mathbf{b}'}\{\mathbf{b}'^T \mathbf{r}_{\mathbf{f}_\lambda}\} = C_2$ and $|C_2| < +\infty$ because $\{E_{\mathbf{b}'}\{|r_{\mathbf{f}_\lambda k}|\}\} < +\infty\}_{k=1}^N$ and \mathbf{b}' has bounded higher-order moments. When $\varepsilon \rightarrow 0$, we immediately see that

$$\lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} E_{\mathbf{b}'}\{\mathbf{b}'^T(\mathbf{f}_\lambda(\mathbf{y} + \varepsilon\mathbf{b}') - \mathbf{f}_\lambda(\mathbf{y}))\} = \text{trace}\{\mathbf{J}_{\mathbf{f}_\lambda}(\mathbf{y})\} = \text{div}_{\mathbf{y}}\{\mathbf{f}_\lambda(\mathbf{y})\},$$

which yields the desired result.

We could also obtain the proof of the weak form of the result (when the second derivatives are not necessarily well-defined), but it is more technical. It involves standard but tedious usage of mathematical tools of measure theory such as the Fubini theorem and the Lebesgue's dominated convergence theorem and is given in Appendix E.2. ■

Theorem 6.1 is a powerful result since nowhere did we have to express the functional form of \mathbf{f}_λ explicitly, thus making (6.2) suitable for a wide variety of algorithms. The important point is that \mathbf{f}_λ is treated as a black box, meaning that we only need access to the output of the operator, irrespective of how it is implemented. From a calculus point of view, it can be regarded as the stochastic definition of the divergence of a vector field in multiple dimensions where $\mathbf{f}_\lambda(\mathbf{y} + \varepsilon\mathbf{b}') - \mathbf{f}_\lambda(\mathbf{y})$ may be understood as the first-order (random) difference of \mathbf{f}_λ . It may also be thought of as a generalization of a result due to Wahba [151, 152] and Girard [153] developed in the context of RGCV which is only applicable for “mildly” non-linear problems, in the sense that $\mathbf{J}_{\mathbf{f}_\lambda}(\mathbf{y}) \approx \mathbf{J}_{\mathbf{f}_\lambda}(\mathbf{x})$. We discuss this further in Section 6.6.1.1.

Equation (6.2) (including the limit) forms the basis of our Monte-Carlo approach for computing SURE for a non-linear \mathbf{f}_λ . Since, in practice, the limit in (6.2) cannot be implemented due to finite machine precision, we propose the following approximation:

$$\frac{1}{N\varepsilon} \text{div}_{\mathbf{y}}\{\mathbf{f}_\lambda(\mathbf{y})\} \approx \frac{1}{N\varepsilon} \mathbf{b}'^T(\mathbf{f}_\lambda(\mathbf{y} + \varepsilon\mathbf{b}') - \mathbf{f}_\lambda(\mathbf{y})), \quad (6.5)$$

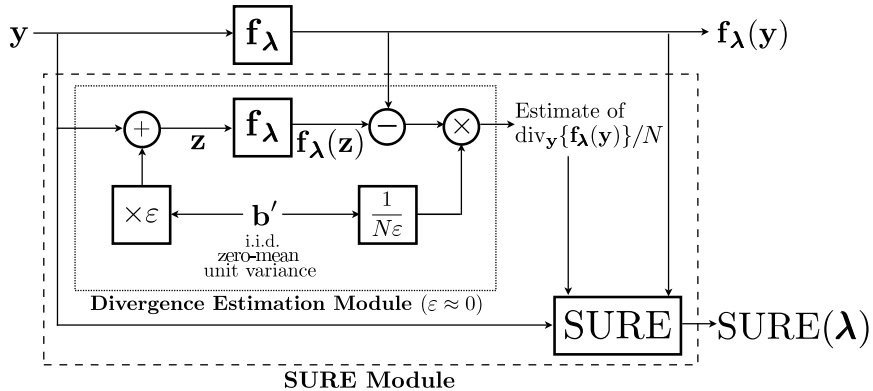


Figure 6.1: The dotted box depicts the module that estimates $\frac{1}{N} \text{div}_{\mathbf{y}}\{f_\lambda(\mathbf{y})\}$ according to (6.5). The dashed box represents the SURE module (depicted as MSE estimation box in Figure 5.1 of Chapter 5) which computes the SURE according to Equation (5.5) (see Chapter 5).

where the factor $\frac{1}{N}$ accounts for the averaging (of SURE) over all the pixels. The r.h.s. of (6.5) amounts to adding a small amount of noise (of variance ϵ^2) to \mathbf{y} and evaluate $f_\lambda(\mathbf{y} + \epsilon \mathbf{b}')$. The difference $f_\lambda(\mathbf{y} + \epsilon \mathbf{b}') - f_\lambda(\mathbf{y})$ is then used to obtain an estimate of the divergence. The schematics of implementing (6.5) is illustrated in Figure 6.1. The validity of the approximation in (6.5) depends on how small ϵ can be made. In practice, we must select a ϵ small enough to mimic the limit, but still large enough so as to avoid round-off errors in $f_\lambda(\mathbf{y} + \epsilon \mathbf{b}')$. As demonstrated in Section 6.5, the admissible range of ϵ covers several decades, so that the choice of this parameter is not critical.

We now give an algorithm for Monte-Carlo divergence estimation (and SURE) which is quite straightforward and easy to implement. It assumes that a “suitably” small ϵ has been selected and a zero-mean unit variance i.i.d. random vector \mathbf{b}' has been generated. It also uses only one realization of \mathbf{b}' for the same reason given in Section 6.3.1: The law of large numbers is applicable to $\frac{1}{N} \mathbf{b}'^T f_\lambda(\mathbf{y} + \epsilon \mathbf{b}')$ whenever $f_{\lambda k}(\mathbf{y} + \epsilon \mathbf{b}')$ is “approximately” independent for different k . This assumption is quite valid in practice because f_λ mostly performs “local” operations (for instance, finite-length wavelet filters and coordinate-wise thresholding are used in wavelet-based methods and finite-difference filters are

used in TV denoising). We present experimental results in Section 6.6 that support this claim.

Algorithm 6.2. Algorithm for estimating $\frac{1}{N} \operatorname{div}_{\mathbf{y}} \{\mathbf{f}_{\lambda}(\mathbf{y})\}$ and $SURE(\lambda)$ for a given $\lambda = \lambda_0$.

1. For $\lambda = \lambda_0$, evaluate $\mathbf{f}_{\lambda}(\mathbf{y})$
2. Build $\mathbf{z} = \mathbf{y} + \varepsilon \mathbf{b}'$. Evaluate $\mathbf{f}_{\lambda}(\mathbf{z})$ for $\lambda = \lambda_0$
3. Compute $\operatorname{div} = \frac{1}{N\varepsilon} \mathbf{b}'^T (\mathbf{f}_{\lambda}(\mathbf{z}) - \mathbf{f}_{\lambda}(\mathbf{y}))$ and $SURE(\lambda_0)$ using (5.5).

Another significant observation is that whenever \mathbf{f}_{λ} is linear, the two Monte-Carlo algorithms discussed in this work turn out to be rigorously equivalent. This is formally stated in the following proposition.

Proposition 6.2. Let \mathbf{f}_{λ} be linear as in Equation (5.7) (see Chapter 5) and \mathbf{b}' be a zero-mean i.i.d. random vector with unit variance. Then, without the limit, the r.h.s. of (6.2) reduces to that of (6.1), independent of ε .

Proof. Starting from the main term in the r.h.s. of (6.2) and using (5.7) and (6.1) we obtain the desired result:

$$\begin{aligned} \frac{1}{\varepsilon} E_{\mathbf{b}'} \{ \mathbf{b}'^T (\mathbf{f}_{\lambda}(\mathbf{y} + \varepsilon \mathbf{b}') - \mathbf{f}_{\lambda}(\mathbf{y})) \} &= \frac{1}{\varepsilon} E_{\mathbf{b}'} \{ \mathbf{b}'^T \mathbf{F}_{\lambda} (\mathbf{y} + \varepsilon \mathbf{b}' - \mathbf{y}) \} \\ &= E_{\mathbf{b}'} \{ \mathbf{b}'^T \mathbf{F}_{\lambda} \mathbf{b}' \} \\ &= \operatorname{Trace} \{ \mathbf{F}_{\lambda} \}. \end{aligned}$$

■

6.4 Description of Some Popular Denoising Algorithms

Now that we have a practical means of estimating $\operatorname{div}_{\mathbf{y}} \{\mathbf{f}_{\lambda}(\mathbf{y})\}$ for an arbitrary \mathbf{f}_{λ} , we demonstrate the applicability of Monte-Carlo SURE for some popular denoising algorithms such as total-variation denoising (TVD) and redundant scale-dependent wavelet soft-thresholding (RSWST). Also included in the evaluation are orthonormal scale-dependent wavelet soft-thresholding and smoothing

splines⁶ for which SURE takes an explicit form. For the variational methods (TVD and smoothing splines), the parameter $\boldsymbol{\lambda} = \lambda$ represents the regularization tradeoff, while for the wavelet-based methods, $\boldsymbol{\lambda}$ controls the scale-dependent thresholds. In the forthcoming sections, we first describe each algorithm along with its associated characteristics.⁷ We then discuss numerical issues related to choice of ε to be used in Algorithm 6.2. Finally, we present experimental results that validate our arguments.

6.4.1 Scale-Dependent Wavelet Soft-Thresholding

6.4.1.1 Orthonormal Wavelets (OSWST)

If \mathbf{W} is the matrix corresponding to an orthonormal wavelet transform, the OSWST denoised signal is given by $\mathbf{f}_\lambda(\mathbf{y}) = \mathbf{W}^T \tilde{\mathbf{c}}$, where

$$\tilde{\mathbf{c}} = \arg \min_{\mathbf{c}} \underbrace{\left\{ \|\mathbf{y} - \mathbf{W}^T \mathbf{c}\|^2 + \sum_{i,k} \lambda_{i,s,q} |c_k^i|^q \right\}}_{J_{\mathbf{W}}\{\mathbf{y}, \mathbf{c}\}}. \quad (6.6)$$

The second term in the r.h.s. of the above equation is equivalent to the Besov norm of the corresponding continuously defined signal estimate [154]. The quantity c_k^i is the k -th wavelet coefficient in the i -th sub-vector of \mathbf{c} (corresponding to the i -th sub-band) and

$$\lambda_{i,s,q} = 2^{-iq(s + \frac{d}{2} - \frac{d}{q})} \lambda$$

is the scale-dependent regularization parameter for s , $\lambda \in \mathbb{R}^+$; the dimension of the data is d , while q corresponds to the ℓ_q -norm of the coefficient vector. For our experiments, we set $d = 2$ and $q = 1$ (for image denoising with ℓ_1 constraint on the wavelet coefficients), which yields the scale-dependent regularization parameter

$$\lambda_{i,s} = 2^{-i(s-1)} \lambda. \quad (6.7)$$

The advantage of selecting an orthogonal transform is that it decouples $J_{\mathbf{W}}$ so that (6.6) is equivalent to independently minimizing scalar cost functions on a

⁶We again present the results for denoising by polyharmonic smoothing splines (PSS, see Section 5.3.3.1 of Chapter 5) for comparison and completeness.

⁷For the description of SURE computation for smoothing splines, see Section 5.3.2.2 in Chapter 5.

coefficient-by-coefficient basis. The minimization of scalar cost corresponding to \tilde{c}_k^i is then simply achieved by a soft-thresholding operation [154] with the threshold $\frac{\lambda_{i,s}}{2}$ so that

$$\tilde{c}_k^i = \mathcal{T}_{\lambda_{i,s}}(c_k^i) = \begin{cases} c_k^i - \frac{\lambda_{i,s}}{2} \text{sign}(c_k^i) & \text{if } |c_k^i| > \frac{\lambda_{i,s}}{2}, \\ 0 & \text{if } |c_k^i| \leq \frac{\lambda_{i,s}}{2}, \end{cases} \quad (6.8)$$

where c_k^i is the k -th wavelet coefficient in the i -th sub-band of the wavelet transform $\mathbf{c} = \mathbf{W}\mathbf{y}$. Due to the orthonormality of \mathbf{W} , MSE (and hence SURE) is invariant under the transform (Parseval equivalence). Therefore \mathbf{c} replaces \mathbf{y} , while $\mathcal{T}_{\lambda_{i,s}}$ replaces \mathbf{f}_λ in Equation (5.5) (see Chapter 5). The required divergence is then simply computed to be $\sigma^2 \sum_{i,k} \mathbb{1}_{\mathcal{A}}(c_k^i)$, where $\mathcal{A} = \{c_k^i : |c_k^i| > \frac{\lambda_{i,s}}{2} \forall i, k\}$ and $\mathbb{1}\{\bullet\}$ is the indicator function.

The OSWST is akin to the *SureShrink* algorithm of Donoho *et al.* [132] in that they both apply soft-thresholding in an orthonormal (wavelet) transform domain. However, the two methods significantly differ from each other in the way they select the threshold levels: While *SureShrink* assigns a threshold value to each sub-band by independent sub-band minimization of SURE, OSWST optimizes the threshold parameters (λ, s) (that characterize the sub-band dependent threshold value in Equation (6.7)) by minimization of SURE computed over all the sub-bands (entire wavelet decomposition).

6.4.1.2 Redundant Wavelets (RSWST)

Redundant discrete wavelet transforms are over-complete representations that are advantageous for denoising, mainly due to their better shift-invariant properties [155–157]. We consider the undecimated wavelet transform (UWT) with an orthonormal filter pair in the redundant paradigm (tight-frame). Our denoising function is again the scale dependent soft-thresholding operator $\mathcal{T}_{\lambda_{i,s}}$ but now applied on the UWT coefficients. For $s = 1$ in (6.7), $\lambda_{i,s} = \lambda$ yields the same threshold level for all sub-bands i in which case both OSWST and RSWST perform universal soft-thresholding of the corresponding wavelet coefficients. However, unlike OSWST, there is no cost function associated with RSWST. Moreover, as shown in [136], Parseval's equivalence is no longer valid in the redundant wavelet domain which forces us to evaluate SURE in the signal domain.

Writing $\mathbf{f}_\lambda(\mathbf{y}) = \mathbf{W}^T \mathcal{T}(\mathbf{W}\mathbf{y})$, where \mathbf{W} is a UWT matrix and \mathcal{T} the vector containing the soft-thresholding operators $\mathcal{T}_{\lambda_{i,s}}$ (see Equation (6.8)), it is immediately clear that evaluating $\text{div}_{\mathbf{y}}\{\mathbf{f}_\lambda(\mathbf{y})\}$ is arduous because the output of \mathcal{T}

depends on $\mathbf{W}\mathbf{y}$ thus demanding explicit access to each element of \mathbf{W} . However, since the soft-thresholding operator is continuous and weakly-differentiable [133], RSWST (and OSWST included) satisfies the weaker hypotheses of Theorem 6.1 and therefore qualifies for Monte-Carlo estimation procedure described in Section 6.3.2. In fact, RSWST constitutes a good demonstration example for illustrating the signal-domain computation of SURE using Algorithm 6.2 to perform a combined optimization of the two threshold parameters $\boldsymbol{\lambda} = (\lambda, s)$.

6.4.2 Total-Variation Denoising (TVD)

While wavelet-based denoising forms an active research area in its own right, other denoising procedures that have flourished in the literature, include variational and PDE based methods of which the most popular is TV denoising [74]. The idea behind TVD is to minimize the total-variation of an image that is constrained to be “close” to the given noisy data. The problem has been formulated in both continuous and discrete domains [74, 158]. The solution is either found by evolving a PDE derived from the Euler-Lagrange equation or by performing some kind of iterative optimization (e.g., bounded optimization using Majorization-Minimization (MM) [96] or half-quadratic [159] optimization).

Here, we consider the discrete domain formulation of Figueiredo *et al.* [158] where the TV denoised image is obtained by minimizing the cost functional

$$J_{\text{TV}}(\mathbf{y}, \mathbf{u}) = \|\mathbf{y} - \mathbf{u}\|^2 + \lambda \text{TV}(\mathbf{u}), \quad (6.9)$$

where

$$\text{TV}(\mathbf{u}) = \sum_k \sqrt{(\mathbf{D}_h \mathbf{u})[k]^2 + (\mathbf{D}_v \mathbf{u})[k]^2}$$

is the discrete 2-D total-variation norm and \mathbf{D}_h and \mathbf{D}_v are matrices corresponding to the first-order finite difference in the horizontal and vertical directions, respectively. J_{TV} is convex and can be minimized using an iterative MM algorithm [158]. Then, starting from the update equation, it can be established in a straightforward (but tedious) manner that $\mathbf{f}_{\boldsymbol{\lambda}}$ for TVD admits at least a second-order Taylor expansion.⁸ TVD is a typical example where SURE cannot be evaluated analytically while our Monte-Carlo method circumvents the difficulty.

⁸The derivation of this result is given in Appendix E.1.

6.5 Range of Validity of the Proposed Monte-Carlo SURE

The two main conditions for Algorithm 6.2 to work are that \mathbf{f}_λ satisfies the hypotheses of Theorem 6.1 and ε is “small”. Ideally, we would like to let ε tend towards zero in (6.5) as dictated by (6.2), but this cannot be realized exactly in practice due to finite machine precision. When ε is too small, numerical round-off errors become more prominent because \mathbf{f}_λ becomes insensitive to changes in ε . In effect, this phenomenon fixes a lower bound for ε which may vary depending on the sensitivity of \mathbf{f}_λ . To elucidate this, we selected the following non-linear algorithms: TVD and RSWST (with threshold value $\frac{\lambda}{2}$) which satisfy at least one of the hypotheses of Theorem 6.1 and found, based on numerical experiments with Java that $\varepsilon \geq 10^{-12}$ was admissible for these algorithms. We then applied Algorithm 6.2 with Gaussian \mathbf{b}' for each of these methods for different values of ε and a wide range of λ for the Boats test image with input SNR⁹ 4 dB.

We observed that when ε was decreased from $\varepsilon = 1$ down to 10^{-12} , Algorithm 6.2 yielded SURE values which not only captured the trend of MSE over a wide range of λ but also yielded very good estimates of the optimal λ for the TVD and RSWST methods, in agreement with Theorem 6.1. We illustrate this in Figures 6.2(a) and 6.2(b) for the cases of $\varepsilon = 0.1$ and $\varepsilon = 0.01$ for TVD and RSWST, respectively, where the corresponding curves nearly overlap and are also close to the MSE curve over the entire range of λ . At the other end, as soon as $\varepsilon \gtrsim 2$, we started to observe significant bias (*cf.* uppermost curves in Figure 6.2 corresponding to $\varepsilon = 10$) which indicates that large ε is not desirable for non-linear problems. We therefore conclude that whenever the assumptions of Theorem 6.1 are valid, the proposed estimation procedure is quite robust with respect to ε (when $\varepsilon \rightarrow 0$) and it yields meaningful results when ε is made “small”.

Next, to investigate the relevance of the underlying differentiability hypotheses in Theorem 6.1, we applied Algorithm 6.2 to RSWHT which performs hard-thresholding with the threshold value $\frac{\lambda}{2}$. Since the hard-thresholding operator is neither continuous nor weakly-differentiable [160], RSWHT violates the hypotheses of Theorem 6.1. Numerically, this is reflected in the increasing instability of the SURE curves as ε decreases in Figure 6.2(c). In this case, violating the hypotheses of Theorem 6.1 leads to a variance of Monte-Carlo SURE that

⁹The input SNR in all experiments is computed using (5.12) in Chapter 5 by replacing the denominator with $N\sigma^2$.




Figure 6.2: Plots of $MSE(\lambda)$ and (Monte-Carlo) $SURE(\lambda)$ for different ϵ : (a) TVD, (b) Haar-RSWST with threshold value $\frac{\lambda}{2}$, and (c) Haar-RSWHT with threshold value $\frac{\lambda}{2}$ for noisy Boats image with $SNR = 4$ dB ($\sigma = 29.45$).

increases without bound with decreasing ε .

It must be noted that the hard-thresholding function is quite an extreme case and has been considered here purely to illustrate the sharpness of the hypotheses of Theorem 6.1 to certify whether or not a denoising algorithm is suitable for the proposed Monte-Carlo SURE. Fortunately for us, most common algorithms encountered in practice satisfy the required differentiability hypothesis and can be optimized with Algorithm 6.2 as demonstrated next.

6.6 Results

6.6.1 Results with One-Parameter Optimization

We now present numerical results for SURE-based optimization of a single parameter (only λ) for the denoising methods discussed in Section 6.4. In doing this, we exemplify the use of SURE, but do not contend with state-of-the-art denoising methods. For our experiments, we consider different categories of test images in Figure 5.2 (see Chapter 5). We choose the Haar wavelet transform for the wavelet based methods to match the wavelet filter with the first-order finite difference filter employed in TVD. We used $J = 4$ levels of decomposition in all cases and did not perform any thresholding on the coarse-scale projection of the signal. We also considered the polyharmonic smoothing splines (PSS) method (see Section 5.3.3.1 of Chapter 5) for the sake of completeness.

The performance of the methods is quantified by the SNR of the output $\mathbf{f}_\lambda(\mathbf{y})$ given by (5.12) (see Chapter 5). All SNR values reported here were obtained by averaging over three independent simulations. We consider images corrupted by white Gaussian noise whose standard deviation σ is again assumed to be known.¹⁰ In all experiments, the value of σ is set to achieve the desired input SNR computed by replacing $\|\mathbf{x} - \mathbf{f}_\lambda(\mathbf{y})\|^2$ with $N\sigma^2$ in (5.12). Besides, in the implementation of the denoising methods, periodic boundary conditions were used when required. For PSS and OSWST, SURE was computed analytically, while for TVD and RSWST, the proposed Monte-Carlo method (Algorithm 6.2) was used with zero-mean i.i.d. Gaussian random vectors of standard deviation $\varepsilon = 0.1$.

¹⁰In practice, σ can be reliably estimated using the median estimator of Donoho *et al.* [132].

Table 6.1: Comparison of RGCV- and SURE-based Performances of Denoising by RSWST

	Input SNR (dB)	4	8	12	16	20
Image	Measure					
Boats	Oracle	11.87	14.07	16.49	19.07	21.91
	SURE	11.87	14.06	16.49	19.07	21.90
	RGCV _{NL}	9.42	11.45	13.05	16.56	20.03
	RGCV* _{NL}	11.65	13.97	15.19	18.83	20.60
MRI	Oracle	12.20	14.64	17.26	20.08	23.08
	SURE	12.19	14.64	17.26	20.07	23.08
	RGCV _{NL}	9.63	12.04	13.90	17.82	21.18
	RGCV* _{NL}	12.11	14.44	16.97	19.99	22.87
Peppers	Oracle	12.05	14.57	17.28	20.04	22.88
	SURE	12.04	14.56	17.28	20.04	22.88
	RGCV _{NL}	9.34	11.96	13.93	17.86	20.86
	RGCV* _{NL}	11.94	14.27	16.18	19.98	22.82
Shepp-Logan	Oracle	13.98	17.59	21.28	25.02	28.82
	SURE	13.98	17.58	21.26	25.00	28.81
	RGCV _{NL}	10.82	14.51	17.19	22.04	25.93
	RGCV* _{NL}	13.66	16.90	19.89	24.27	28.32

6.6.1.1 Comparison with Randomized GCV

Here, we compare the performances of SURE and the randomized generalized cross-validation (RGCV) for the RSWST ($s = 1$) algorithm in terms of SNR improvement. The RGCV measure is computed using the Monte-Carlo version (for non-linear algorithms) proposed by Girard [153]:

$$\text{RGCV}_{\text{NL}}(\lambda) = \frac{N^{-1} \|\mathbf{y} - \mathbf{f}_\lambda(\mathbf{y})\|^2}{(1 - N^{-1} \varepsilon^{-1} \mathbf{b}^T [\mathbf{f}_\lambda(\mathbf{y} + \varepsilon \mathbf{b}') - \mathbf{f}_\lambda(\mathbf{y})])^2}, \quad (6.10)$$

where $\varepsilon = 0.9\sigma$ is used in (6.10) as recommended in [153]. The output SNR obtained by adjusting λ based on SURE and RGCV_{NL} is tabulated¹¹ for various input SNR and test images in Table 6.1.

¹¹Table 6.1 supplements the results of Table 5.1 in Chapter 5 for the non-linear denoising algorithm RSWST.

It was observed that the λ selected by $\text{RGCV}_{\text{NL}}(\lambda)$ was far from the optimum value in all cases: This can be attributed to the bias originating from the recommended value of ε and the fact that RSWST does not probably satisfy the “mild” non-linearity assumption of Wahba and Girard [151–153]. As a result, the performance of RGCV_{NL} is poor at all noise levels.

Following the philosophy underlying (6.2) and the argumentation in Section 6.5, we therefore decided to inspect another version of RGCV_{NL} , denoted by $\text{RGCV}_{\text{NL}}^*$, which utilized a small value: $\varepsilon = 0.1$. It is observed that $\varepsilon = 0.1$ dramatically improves the performance as reflected in the output SNR values corresponding to $\text{RGCV}_{\text{NL}}^*$: This demonstrates the validity of the proposed Monte-Carlo procedure for estimating the divergence for algorithms with “arbitrary” non-linearities. However, it should be noted that the performance of $\text{RGCV}_{\text{NL}}^*$ is still not on par with SURE, which consistently imitates the oracle (that is, MSE) for all noise levels and considered test images. This indicates that GCV-like measures, though having the advantage of not requiring σ^2 , may not always yield optimal performance for all denoising algorithms.

6.6.1.2 MSE-SURE Comparison

A series of relevant graphs ($\text{SURE}(\lambda)$, $\text{MSE}(\lambda)$ versus λ) for four denoising methods are shown in Figures 6.3 and 6.4. It is observed that SURE follows the MSE curve remarkably well in all the cases thereby leading to accurate estimates of the optimal λ . We observed the same trend for all test images and input SNRs which confirms the consistency of our method. The agreement is somewhat better in the case of larger images (Boats, Barbara) as compared to the Peppers image which is probably due to the fact that we have 4 times more pixels to estimate MSE (law of large numbers).

These results demonstrate the validity of the approximation in (6.5). The RSWST method is a borderline case for which the formula (6.2) is only true in the weak sense because the second derivative of the soft-thresholding operator is not well-defined for the two critical values $\pm \frac{\lambda}{2}$. Yet, Algorithm 6.2 still performs well in accordance with the second part of Theorem 6.1.

It should also be noted that this type of extensive estimation over a wide range of λ (as shown in Figures 6.3 and 6.4) has been done purely for the purpose of illustration. In practice, we can rely on bracketing methods (golden-mean search) which do not use any derivative information in order to find the minimum of SURE in a much smaller number of steps (typically 10 steps).




Figure 6.3: MSE(λ) and SURE(λ) for (a) Haar-OSWST ($s = 1$), (b) TVD ($\epsilon = 0.1$), (c) Haar-RSWST ($s = 1, \epsilon = 0.1$) for noisy Peppers image with SNR = 4 dB ($\sigma = 33.54$).




Figure 6.4: MSE(λ) and SURE(λ) for (a) Haar-OSWST ($s = 1$), (b) TVD ($\epsilon = 0.1$), and (c) Haar-RSWST ($s = 1, \epsilon = 0.1$) for noisy Boats image with input SNR = 4 dB ($\sigma = 29.45$).

6.6.1.3 Visual Comparison

To highlight the different characteristics of the denoising methods it is best to compare the results visually. Figure 6.5 shows the denoised outputs of four algorithms with optimized parameters. The smoothing spline estimator, as its name suggests, attempts to smooth the noisy fluctuations during the denoising process. But in doing so, it also smoothes the underlying image leading to smudged edges (as seen in Figure 6.5(c)), which is the main disadvantage of this approach.

The Haar-OSWST ($s = 1$) preserves some edge information but produces a blocky output (see Figure 6.5(d)) because small detail coefficients are set to zero by the univariate soft-thresholding operator. There is a loss of image details and the reconstructed output exhibits artifacts corresponding to the footprints of the basis function (Haar wavelet). The Haar wavelet is at the low end of what can be achieved with an orthonormal wavelet transform; the use of a wavelet with better regularization properties (symlets, higher order spline wavelets, etc) yields better results—typically +0.5 dB additional gain (results not shown).

The TV denoised image appears significantly better than the earlier two (see Figure 6.5(e)). Yet, although the edges are preserved as per the TV constraint, the output exhibits some artificial blockiness due to the fact that the algorithm tends to favor piecewise constant solutions.

The Haar-RSWST ($s = 1$) yields the best visual output, which correlates with the higher SNR value (11.90 dB). This can be attributed to the redundant nature of the underlying transform. Interestingly enough, the result is not penalized by the lower order of the Haar transform (piecewise-constant approximation), in fact, it is quite the contrary (as was also noticed in [136]). This is in contrast with the non-redundant case where higher order wavelets yield better results, but nothing that comes close to the result in Figure 6.5(f).

6.6.1.4 Computational Cost

Two main aspects of any denoising algorithm are the associated computational cost and the yielded SNR improvement. In general, these two aspects are conflicting in nature and the user must strike a good balance between them. In terms of computational efficiency, the four methods can be ranked as follows:

- (i) The Haar-OSWST method ($J = 4$ levels), which requires of the order of $2 \times 4N$ operations, while it uses the same amount of storage as the image itself.




Figure 6.5: Visual comparison of SURE-optimized denoising results for the Boats image (zoomed); (a) Noise-free image (b) Noisy observations ($\sigma = 29.45$; SNR = 4 dB); (c) First-degree PSS result (SNR = 11.84 dB); (d) Haar-OSWST ($s = 1$) result (SNR = 10.33 dB); (e) TVD result (SNR = 11.02 dB); (f) Haar-RSWST result ($s = 1$, SNR = 11.90 dB)

-
- (ii) Polyharmonic smoothing splines; these are implemented efficiently using the FFT and therefore require $O(N \log_2 N) + N$ operations while storage-wise, it is equivalent to the Haar-OSWST method.
 - (iii) The Haar-RSWST method; it is implemented using the *algorithm à trous* [156] which, for $J = 4$, requires a total of $13 \times 2 \times 4N$ computations. It should be noted that the performance improvement yielded by the redundancy of the transform is at the cost of requiring $13N$ storage locations which is probably one potential downside of this method.
 - (iv) TVD; the MM algorithm of [158] required an average of 13 main iterations. At any given iteration, the method uses few- N locations (typically $< 4N$) for storing intermediate iteration variables. Additionally, for each main iteration, we performed 20 conjugate-gradient iterations to solve an associated linear system. This leads to a total of $260N$ operations to obtain a single denoised signal estimate implying that TVD is the costliest of all the considered methods.

6.6.1.5 SNR Improvement

We now make a quantitative comparison of the methods in terms of SNR improvement. For the sake of comparison, the SNR is computed for outputs obtained by setting λ based on both MSE and SURE. This is tabulated in Table 6.2 where the first value in each cell gives the SNR obtained by choosing λ based on MSE (oracle SNR), while the second corresponds to that obtained by Monte-Carlo SURE optimization. The maximum of the SNRs with respect to all the methods is indicated in bold-face font for each image and noise variance. Several observations are in order:

- The first and the most important one is that the SNR obtained based on MSE and SURE are either equal or different only in the second decimal place for all tested cases. This indicates the reliability and robustness of our Monte-Carlo SURE optimization procedure.
- Haar-OSWST ($s = 1$) performs poorly, especially at high noise levels. This is due to the inflexible nature of the soft-thresholding operator and blocky-reconstruction of the Haar wavelet. However, as noted earlier, one may be able to boost the performance slightly by using a higher order wavelet (typically + 0.5 dB additional gain).

Table 6.2: Performance of considered methods in terms of SNR[#]

Image	Input SNR (dB)	4	8	12	16	20
Boats	PSS [†]	(11.83, 11.83)	(13.69, 13.69)	(15.81, 15.81)	(18.27, 18.27)	(21.20, 21.20)
	Haar-OSWST [†]	(10.33, 10.32)	(12.63, 12.63)	(15.22, 15.22)	(18.09, 18.08)	(21.23, 21.23)
	Haar-RSWST [†]	(11.87, 11.87)	(14.07, 14.06)	(16.49, 16.49)	(19.07, 19.07)	(21.91, 21.90)
Barbara	TV D	(10.98, 10.98)	(13.13, 13.13)	(15.61, 15.61)	(18.36, 18.36)	(21.42, 21.42)
	PSS [†]	(9.76, 9.76)	(11.63, 11.63)	(14.08, 14.08)	(17.08, 17.08)	(20.51, 20.51)
	Haar-OSWST [†]	(9.29, 9.29)	(11.71, 11.71)	(14.59, 14.58)	(17.78, 17.78)	(21.24, 21.24)
fBm	Haar-RSWST [†]	(10.56, 10.55)	(12.87, 12.86)	(15.58, 15.58)	(18.61, 18.61)	(21.89, 21.89)
	TV D	(9.45, 9.45)	(11.66, 11.66)	(14.48, 14.48)	(17.70, 17.70)	(21.15, 21.15)
	PSS [†]	(15.29, 15.29)	(16.87, 16.87)	(18.64, 18.64)	(20.59, 20.59)	(22.85, 22.85)
MRI	Haar-OSWST [†]	(11.15, 11.14)	(13.12, 13.11)	(15.38, 15.37)	(18.01, 18.00)	(21.05, 21.04)
	Haar-RSWST [†]	(13.24, 13.23)	(14.95, 14.95)	(16.90, 16.90)	(19.14, 19.14)	(21.78, 21.78)
	TV D	(12.40, 12.40)	(14.03, 14.03)	(15.93, 15.93)	(18.26, 18.26)	(21.14, 21.14)
Peppers	PSS [†]	(12.40, 12.40)	(14.63, 14.63)	(17.02, 17.02)	(19.59, 19.59)	(22.35, 22.35)
	Haar-OSWST [†]	(10.29, 10.29)	(12.86, 12.84)	(15.66, 15.66)	(18.69, 18.68)	(21.90, 21.89)
	Haar-RSWST [†]	(12.20, 12.19)	(14.64, 14.64)	(17.26, 17.26)	(20.08, 20.07)	(23.08, 23.08)
Shepp-Logan	TV D	(11.40, 11.39)	(13.70, 13.70)	(16.24, 16.24)	(19.09, 19.08)	(22.19, 22.19)
	PSS [†]	(10.74, 10.74)	(12.47, 12.47)	(14.70, 14.70)	(17.44, 17.44)	(20.68, 20.68)
	Haar-OSWST [†]	(10.07, 10.07)	(12.66, 12.64)	(15.53, 15.52)	(18.53, 18.52)	(21.68, 21.68)
TV D	Haar-RSWST [†]	(12.05, 12.04)	(14.57, 14.56)	(17.28, 17.28)	(20.04, 20.04)	(22.88, 22.88)
	TV D	(11.22, 11.22)	(13.67, 13.67)	(16.35, 16.35)	(19.17, 19.17)	(22.18, 22.18)
	PSS [†]	(9.91, 9.91)	(11.79, 11.79)	(14.13, 14.13)	(17.06, 17.06)	(20.45, 20.45)
TV D	Haar-OSWST [†]	(11.94, 11.93)	(15.47, 15.46)	(19.10, 19.09)	(22.82, 22.81)	(26.62, 26.61)
	Haar-RSWST [†]	(13.98, 13.98)	(17.58, 17.58)	(21.28, 21.26)	(25.01, 25.00)	(28.82, 28.81)
	TV D	(15.33, 15.32)	(18.92, 18.91)	(22.66, 22.66)	(26.38, 26.37)	(30.13, 30.13)

[#]Each cell is formatted as (Oracle value, Estimated value), [†]First-degree PSS; [‡]s = 1 for OSWST and RSWST.

- The linear smoothing spline technique is among the least effective method for natural image denoising. It is seemingly better than Haar-OSWST ($s = 1$) at high noise levels for almost all images due to the fact that it smoothes the noisy image thereby strongly reducing the harsh effect of noisy fluctuations. But, it also smoothes the underlying signal making it the least-preferred method for images with rich texture (for instance, the Barbara image).

However, the first-degree polyharmonic smoothing spline outperforms all the other methods for the fBm image, which is in agreement with the theory [78]. This also strengthens the fact that smoothing splines are ideal whenever the underlying image fits the statistical model. A similar behavior is observed for the MRI image which may be due to the fact that MRI images are mostly fractal-like [161] and their power spectrum can be well approximated by the $\|\omega\|^{-\alpha}$ spectral law [69].

- As expected, the use of redundant transform improves the denoising quality compared to Haar-OSWST ($s = 1$). The Haar-RSWST (with $s = 1$) method provides a gain of more than 2 dB compared to Haar-OSWST ($s = 1$) at large levels of noise. Notably, it is also the best method for all the images with the exception of fBm and the Shepp-Logan phantom.
- TVD performs better than PSS and Haar-OSWST ($s = 1$) (and even (λ, s) -optimized Haar-OSWST, see the following subsection for details), whenever the images are smooth without strong textures (for instance, the Peppers image and the Shepp-Logan phantom). This shows that TVD is competitive or even better than classical wavelet denoising methods [158] for images that fall well within the piecewise-constant category. The Shepp-Logan phantom is noteworthy in this context as it is a good example of a piecewise constant image. Unsurprisingly, TVD performs better than all the considered methods for this particular image, as indicated in Table 6.2.

In the presence of rich texture (the Barbara image), however, TVD performs worse than all wavelet based methods, which is quite expected because the TV prior is not well-suited for such images. In fact, any texture is considered part of the noise and is annihilated by TVD.

To conclude, we infer that, of the considered methods, some are better suited than others for certain type of images: While overall Haar-RSWST yields the best results for natural images, smoothing splines are well adapted to fractal-like processes and TVD does best for piecewise-constant images.

6.6.1.6 Non-Quadratic Regularized (NQR) Interpolation

This section supplements the results presented in Section 4.4.1 of Chapter 4 and Section 5.3.3.2 of Chapter 5. Here, we test the applicability of the Monte-Carlo SURE method for selecting the regularization parameter for non-quadratic regularized (NQR) interpolation which is based on a non-linear algorithm (*cf.* Algorithm 4.2, Chapter 4).

We consider the multiple-image experiment described in Section 4.4.1 (with MRI stack) where, for each noisy realization, we compute $\text{SURE}(\lambda)$ for NQR-cubic regularized interpolation scheme along with the corresponding $\text{MSE}(\lambda)$. The noise variance was assumed to be known in all experiments for computing $\text{SURE}(\lambda)$. We plot $\text{MSE}(\lambda)$ and $\text{SURE}(\lambda)$ in Figures 6.6(a) and 6.6(b) for two realization of this experiment at 0 dB input SNR (test images shown in the inset). It is clearly seen that the SURE curves capture the trend of the corresponding MSE curves. Moreover, it also yields very close estimates of the optimal λ in both cases.

Table 6.3 lists the output SNR (averaged over all the realizations of the MRI-stack experiment corresponding to an input SNR) corresponding to SURE for the NQR-cubic interpolation algorithm. Also provided are the oracle values of the output SNR (obtained by minimizing MSE). We see that SURE-based tuning consistently yields SNR values very close to the oracle indicating that it can be reliably employed for data-driven adjustment of λ for NQR interpolation.

6.6.2 Results with Multi-Parameter Optimization

So far we have only provided results for SURE-based one-parameter optimization (specifically, for the regularization parameter). However, there is no major difficulty in applying our method for multi-parameter optimization as well. The brute force approach would be to perform an exhaustive search in multiple dimensions to find the best parameter values that minimize SURE. A better way is to perform the search by applying derivative-free optimization. The Powell-Brent algorithm, which uses bracketing and parabolic interpolation for line-search and takes about $n(n+1)/2$ iterations to converge for n set of parameters, is well-suited for our problem as long as the number of parameters stays reasonably small (typically $n < 10$).

Here, we test the concept with the optimization of $\boldsymbol{\lambda} = (\lambda, s)$ for the PSS, Haar-OSWST and Haar-RSWST methods. For the PSS method, s matches the order of the spline to the Hurst exponent of the underlying noise-free signal. This fact has been applied in [78] where the optimal (λ, s) is obtained by fitting




Figure 6.6: Plots of $\text{MSE}(\lambda)$ and $\text{SURE}(\lambda)$ versus λ for two realizations (input SNR = 0 dB) of the rotation experiment (see Section 4.4.1 in Chapter 4) with NQR-cubic interpolation. In both cases, $\text{SURE}(\lambda)$ captures the trend of $\text{MSE}(\lambda)$ and also yields a value of the regularization parameter that is close to oracle optimum which minimizes MSE.

Table 6.3: Output SNR for NQR-cubic interpolation obtained by MSE-based (Oracle) and SURE-based tuning[†] of λ

Input SNR (dB)	Output SNR (dB)			
	QR-Cubic		NQR-Cubic	
	Oracle	SURE	Oracle	SURE
0	13.95	13.95	14.13	14.11
4	15.75	15.75	16.10	16.10
8	17.66	17.66	18.22	18.22
12	19.71	19.71	20.45	20.44
16	21.95	21.95	22.85	22.85
20	24.40	24.40	25.44	25.43
24	27.09	27.09	28.25	28.24
28	30.03	30.03	31.19	31.19
32	33.19	33.19	34.23	34.22
36	36.39	36.39	37.20	37.19
40	39.35	39.35	39.86	39.85
44	41.78	41.78	42.02	42.02
48	43.43	43.42	43.51	43.51
52	44.38	44.38	44.38	44.38
56	44.83	44.82	44.83	44.83
60	45.04	45.04	45.02	45.02

[†]Also presented are the results for QR-cubic interpolation (for comparison and completeness).

a fractal-like model to the power spectrum of the noisy image. However, in our approach this is not required as λ and s are optimized together using SURE. For the wavelet methods, adjusting s changes the threshold value in each sub-band according to (6.7) and our understanding is that this yields better denoising performance than universal soft-thresholding. In all our experiments, we observed that the 2-D Powell optimization of the respective methods took no more than 4 iterations at various noise levels for all the test images. The results are tabulated in Table 6.4.

With PSS, the combined optimization does not yield any significant improvement for the fBm since a degree 1 spline is theoretically the best in the MSE sense (Wiener solution) [78]. As expected PSS still performs the best of all the methods for the fBm image. The improvement for Boats, Barbara, Peppers and the Shepp-Logan phantom is also less significant because these images are not

Table 6.4: Comparison of (λ, s) -optimized methods[‡]

Image	Input SNR (dB)	4	8	12	16	20
Boats	PSS	(11.85, 11.85)	(13.76, 13.76)	(15.92, 15.91)	(18.38, 18.38)	(21.29, 21.29)
	Haar-OSWST	(11.07, 11.06)	(13.08, 13.06)	(15.48, 15.48)	(18.25, 18.23)	(21.31, 21.31)
	Haar-RSWST	(12.87, 12.87)	(14.92, 14.92)	(17.16, 17.16)	(19.53, 19.52)	(22.15, 22.15)
Barbara	PSS	(9.85, 9.85)	(11.63, 11.63)	(14.23, 14.23)	(17.29, 17.29)	(20.66, 20.66)
	Haar-OSWST	(9.63, 9.62)	(11.92, 11.91)	(14.74, 14.71)	(17.89, 17.88)	(21.33, 21.32)
	Haar-RSWST	(10.86, 10.79)	(13.03, 13.01)	(15.70, 15.70)	(18.84, 18.80)	(22.12, 22.12)
fBm	PSS	(15.30, 15.29)	(16.89, 16.89)	(18.67, 18.67)	(20.66, 20.66)	(22.95, 22.95)
	Haar-OSWST	(12.85, 12.84)	(14.43, 14.42)	(16.24, 16.21)	(18.42, 18.40)	(21.21, 21.21)
	Haar-RSWST	(15.04, 15.00)	(16.68, 16.67)	(18.49, 18.49)	(20.37, 20.36)	(22.16, 22.12)
MRI	PSS	(12.70, 12.70)	(15.19, 15.18)	(17.85, 17.85)	(20.65, 20.65)	(23.51, 23.51)
	Haar-OSWST	(11.09, 11.08)	(13.39, 13.36)	(16.06, 16.01)	(18.95, 18.92)	(22.07, 22.06)
	Haar-RSWST	(13.73, 13.72)	(16.05, 16.04)	(18.56, 18.56)	(21.21, 21.17)	(23.98, 23.95)
Peppers	PSS	(10.74, 10.74)	(12.47, 12.47)	(14.71, 14.71)	(17.51, 17.51)	(20.77, 20.77)
	Haar-OSWST	(10.85, 10.84)	(13.24, 13.23)	(15.97, 15.94)	(18.84, 18.82)	(21.89, 21.88)
	Haar-RSWST	(12.95, 12.94)	(15.42, 15.41)	(18.07, 18.06)	(20.71, 20.71)	(23.40, 23.40)
Shepp-Logan	PSS	(9.92, 9.92)	(11.80, 11.80)	(14.13, 14.13)	(17.06, 17.06)	(20.45, 20.45)
	Haar-OSWST	(12.40, 12.31)	(15.91, 15.85)	(19.51, 19.48)	(23.26, 23.21)	(27.04, 26.99)
	Haar-RSWST	(14.36, 14.24)	(17.90, 17.84)	(21.57, 21.54)	(25.26, 25.26)	(29.05, 29.02)

‡Each cell is formatted as (Oracle value, Estimated value)

very fractal-like. In contrast, there is a significant improvement (≈ 1 dB at high input SNR) for the MRI image which provides further support for the claim that MRI images are fractal-like and the order s must be matched to the fractal dimension to obtain best results.

As noted in Table 6.4, this combined optimization is shown to produce a consistent SNR increase for both Haar-OSWST and Haar-RSWST methods. In fact, in the redundant case it leads to an increase of about +1 dB for smooth images like Peppers, Boats and fBm at high noise levels. Thus the optimized Haar-RSWST performs the best of all the considered methods for all natural images which exemplifies the fact that redundant transforms make a powerful denoising tool.

However, it must be emphasized that the results provided in this section are purely for the purpose of illustrating multi-parameter optimization of SURE computed by the proposed Monte-Carlo scheme. In our experiments, we considered a set of popular denoising algorithms with adjustable parameters without making any specific claim concerning their overall optimality. In fact, we have intentionally chosen some test images which favor one or the other algorithm to illustrate that the issue of finding a “best” algorithm is not so clear-cut.

The reader who is interested in state-of-the-art methods is referred to the relevant literature; in particular the BiShrink (dual tree complex wavelet decomposition) [162], BLS-GSM (full steerable pyramidal decomposition) [163], ProbShrink (undecimated Daubechies symlets) [164], and SURE-LET (with redundant Haar transform) [136]. Depending on the type of image these more-advanced techniques can yield a further SNR improvement of the order of 1 dB. In some cases such as SURE-LET, they already take full advantage of the possibility of automatic SURE-based parameter adjustment, with the important difference that the underlying solution is explicit as opposed to our black-box approach where it is obtained numerically. The benefit with the latter scheme is that it requires no hypothesis concerning the analytical form of the solution and therefore has a wider range of applicability.

6.7 Summary

Computation and application of SURE for denoising problems demands the evaluation of the divergence of the denoising operator with respect to the given noisy data. The calculation of this divergence for a general denoising problem may turn out to be non-trivial, especially if the operator does not have explicit analytical form as is the case with iterative algorithms (variational, PDE-based

and Bayesian methods). In this chapter, we introduced a Monte-Carlo technique that circumvents this difficulty and makes SURE viable for an arbitrary denoising scenario, especially when the computation of the associated divergence is mathematically intractable or numerically infeasible. By adding a perturbation to the signal, our method essentially implements a random first-order difference estimator of the divergence of the denoising operator. From a calculus point of view, this can be related to a stochastic definition of the divergence of a vector field. The final outcome is a black-box scheme which yields SURE numerically using only the output of the denoising algorithm without the need for any knowledge of its internal working.

We demonstrated the applicability of our method by performing Monte-Carlo SURE optimization of some popular denoising algorithms in the wavelet (both orthonormal and redundant) and variational (linear and non-linear) settings. We found that SURE computed using the proposed technique perfectly predicts MSE in all considered cases, thereby yielding correct values for the optimal threshold and the regularization parameter for the respective problems. We also substantiated this argument for NQR (non-quadratic regularized) interpolation algorithms. Further, we provided (experimental) justification that the proposed method is applicable for optimization of multiple parameters by performing SURE-based optimization of the thresholds for denoising by scale-dependent wavelet soft-thresholding. We showed that the SNR obtained by SURE-based optimization is in almost perfect agreement with the oracle solution (minimum MSE) for all considered cases. This suggests that Monte-Carlo SURE can be reliably employed for data-driven adjustment of parameters in a large variety of denoising and NQR interpolation problems involving additive Gaussian noise.

Chapter 7

Summary and Conclusions

In this thesis, we investigated the problem of nonideal sampling and interpolation in the presence of noise. Our approach was based on a variational formulation of the problem where reconstruction/interpolation is achieved by imposing suitable continuous-domain regularization constraints to counterbalance the detrimental effect of noise.

We started with a general Tikhonov-like regularization criterion (Chapter 3) that is specified as a function of the L_2 -norm of an arbitrary scalar differential operator L . Then, we showed based on theoretical reasoning that the globally-optimal solution of the variational problem belongs to an integer-shift-invariant space $V(\varphi_{\text{opt}})$ (for the case of uniformly-spaced data samples) that is completely characterized by the operator L —the generator φ_{opt} and hence $V(\varphi_{\text{opt}})$ are together specified via an optimality condition that links L to φ_{opt} . This type of shift-invariant reconstruction mechanism closely relates to spline theory and therefore lends a smoothing-spline flavor to the solution.

Once $V(\varphi_{\text{opt}})$ is specified, the coefficients of the shift-invariant reconstruction are found by solving a suitable linear/non-linear optimization problem that involves the discrete data samples. The main point here is that, although the problem is originally formulated in the analog domain, the resulting solution can be efficiently implemented by digital filtering.

We also proposed stochastic formulations of the nonideal sampling problem—min-max and (analog-domain) minimum mean-squared error (MMSE)/Wiener estimation—and established a functional equivalence between the variational and Wiener solutions: This happens whenever L whitens the noise-free signal (that is, the application of L to the signal results in continuous-domain white

noise). This is one of the key findings of the study since it offers a way of selecting the ideal regularization operator for MMSE-optimal reconstruction of (stochastic) signals from their noisy samples. In this context, we investigated the Matérn stochastic model that describes the power spectral density of signals with $(\alpha + \|\boldsymbol{\omega}\|)^{-\gamma}$ -like spectra. We also constructed a new class of Matérn B-splines that generate the optimal reconstruction spaces for the class of Matérn stochastic processes.

In the second part of the thesis (Chapter 4), we focussed on non-quadratic regularization (NQR) prescribed as the L_p -norm ($p \neq 2$) of some general vector derivative \mathbf{L} . This particular choice was dictated by affine invariance (shift, rotation and scale invariance) of the solution to geometric transformation of the coordinate system. Interesting, it includes the popular total-variation (TV) regularization which is appropriate for signals with strong features (e.g., images with edges). We then used the L_p -norm-based NQR for the problem of regularized interpolation of noisy data and presented *majorize-minimize* (MM) algorithms for carrying out the optimization of the corresponding non-quadratic cost criterion. We found from 2-D rotation experiments on noisy medical images that the use of L_p -norm-like NQR provides better interpolation performance than quadratic regularized and non-regularized (standard) interpolation schemes.

In the third and final part of the thesis (Chapters 5 and 6), we developed novel data-driven methods for finding the optimal (in the MSE sense) regularization parameter for the variational problems considered in this work. Our approach was based on the concept of risk estimation which provides an unbiased means of estimating the MSE purely from the given data. Here, we reviewed Stein's unbiased risk estimate—SURE—for the case of additive Gaussian noise and also developed a new risk estimate that is applicable for a Poisson-Gaussian mixture model. Furthermore, we introduced a novel and efficient Monte-Carlo technique for computing SURE for arbitrary non-linear algorithms. In all our experiments we found that the regularization parameter selected by minimizing the risk estimate yields a performance that was very close or equivalent to the oracle optimum (minimum MSE), indicating the reliability of our approach for practical problems.

We conclude from our study that the proposed regularized interpolation technique together with the data-driven selection of the optimal regularization parameter provides a completely automatic and practical way of performing robust interpolation of noisy data and can be of potential interest to biomedical-imaging practitioners. Moreover, our Monte-Carlo risk estimation has the potential of being applicable to a variety of variational problems such as denoising

and deconvolution.

Future Research Directions

We now discuss some possible directions along which the present work can be extended.

Multilevel Algorithms

The current drawback of Algorithm 4.2 (in the context of non-quadratic regularized (NQR) interpolation) is its high computational cost. Multilevel algorithms may be devised to circumvent this difficulty. The general principle underlying a multilevel algorithm is to successively break down a problem prescribed at a finer scale (level or grid) into *essentially similar* tasks defined at coarser scales, solve them and combine their outputs coherently to achieve the original objective [146]. Since the problem is solved at coarser scales, a multilevel method provides a computationally cheap alternative to conventional methods that operate only at the finest scale. This principle can therefore be applied to our NQR interpolation problem so as to speed up Algorithm 4.2. Furthermore, spline-based interpolation offers multilevel methods an additional advantage in this context as they can take advantage of the inherent multiresolution properties of B-splines [14].

Risk Estimation for a Poisson Noise Model

Extensions Based on Monte-Carlo SURE

In Chapter 6, we specifically focussed on risk estimation for non-linear algorithms—Monte-Carlo SURE—for an additive Gaussian noise model. However, it is not difficult to extend this to a signal-dependent Poisson noise model (such as that considered in Section 5.4.1.1 of Chapter 5) that is appropriate for biomedical imaging [100].

A signal-dependent Poisson noise model can be thought of (up to second-order statistics) as additive Gaussian with signal-dependent variance. Therefore, we can use Monte-Carlo SURE for this scenario provided we estimate the variance at each pixel. But, since the mean and the variance are the same for a Poisson random variable, the given noisy data itself can be used as an

estimate of the variance map provided the Poisson counts at each pixel are sufficiently high. In practice, this last requirement is often ensured and therefore, our Monte-Carlo scheme can be reliably employed for risk estimation.

New Monte-Carlo Poisson Risk Estimator

For low Poisson-count-data (e.g., tomography), the Gaussian approximation described above may no longer hold and a rigorous risk estimator based on Poisson statistics (such as that described in [139]) needs to be used. For instance, for the denoising scenario, this requires the computation of the following sum

$$\sum_{k=1}^N y_k [\mathbf{f}_{\lambda k}(\mathbf{y}) - \mathbf{f}_{\lambda k}(\mathbf{y} - \mathbf{e}_k)], \quad (7.1)$$

where N is the number of pixels, $\mathbf{y} = \mathcal{P}\{\mathbf{x}\}$ is the $N \times 1$ noisy data vector, \mathbf{x} is the $N \times 1$ noise-free signal vector, $\mathbf{f}_{\lambda} : \mathbb{R}^N \rightarrow \mathbb{R}^N$ is mapping that describes the denoising algorithm, $\mathbf{f}_{\lambda k}(\mathbf{y})$ represents the k -th component of the denoised output vector $\mathbf{f}_{\lambda}(\mathbf{y})$ and \mathbf{e}_k is a column vector whose components are all zero except the k -th one which is unity.

A term-by-term evaluation becomes difficult to perform numerically for large data sizes. Approximations based on the Taylor series may then be used [165], but these may become unreliable especially for low-count data. This motivates the need for Monte-Carlo methods—such as that involving the response of \mathbf{f}_{λ} to a perturbation of \mathbf{y} by a small amount of additive noise—for the estimation of (7.1) for arbitrary \mathbf{f}_{λ} .

Hypothesis Testing for Estimation of the Noise Model

Since risk estimates by construction depend on the parameters of the noise model, estimating the model parameters forms a key component of the entire scheme. Specifically, the gain of the Poisson component in Equation (5.14) in Chapter 5 plays a vital role in bioimaging: This quantity is signal-dependent and therefore cannot be obtained from calibration experiments. The technique we devised in Section 5.4.3 of Chapter 5 is empirical in that it uses a certain denoising strategy to isolate patches in the data that are approximately constant. We believe that computationally-efficient techniques based on hypothesis testing may also be devised for the same problem. Such methods would offer a more robust means of identifying patches in the data that are statistically justified to be constant.

Appendix A

Details of Some Results of Chapter 2

A.1 Sampling in the $\mathbf{D}\mathbb{Z}^d$ lattice

The sampling of $f(\mathbf{x})$ in the lattice $\mathbf{D}\mathbb{Z}^d$ is modeled as a product of $f(\mathbf{x})$ with the series of Dirac masses $\sum_{\mathbf{k} \in \mathbb{Z}^d} \delta(\mathbf{x} - \mathbf{D}\mathbf{k})$:

$$f_{\text{samp}}(\mathbf{x}) = f(\mathbf{x}) \sum_{\mathbf{k} \in \mathbb{Z}^d} \delta(\mathbf{x} - \mathbf{D}\mathbf{k}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} f(\mathbf{D}\mathbf{k}) \delta(\mathbf{x} - \mathbf{D}\mathbf{k}).$$

Now, since

$$\begin{aligned} \delta(\mathbf{x} - \mathbf{D}\mathbf{k}) &\stackrel{\mathcal{F}}{\longleftrightarrow} e^{-j(\mathbf{D}^T \boldsymbol{\omega})^T \mathbf{k}}, \\ \sum_{\mathbf{k} \in \mathbb{Z}^d} \delta(\mathbf{x} - \mathbf{D}\mathbf{k}) &\stackrel{\mathcal{F}}{\longleftrightarrow} \frac{(2\pi)^d}{|\mathbf{D}|} \sum_{\mathbf{m} \in \mathbb{Z}^d} \delta(\boldsymbol{\omega} - 2\pi \mathbf{D}^{-T} \mathbf{m}), \end{aligned}$$

where the above Fourier transform has been performed in the sense of distributions, we have

$$\mathcal{F}\{f_{\text{samp}}(\bullet)\}(\boldsymbol{\omega}) = \mathcal{F}\left\{f(\bullet) \sum_{\mathbf{k} \in \mathbb{Z}^d} \delta(\bullet - \mathbf{D}\mathbf{k})\right\}(\boldsymbol{\omega}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} f(\mathbf{D}\mathbf{k}) e^{-j(\mathbf{D}^T \boldsymbol{\omega})^T \mathbf{k}}.$$

But,

$$\mathcal{F} \left\{ f(\bullet) \sum_{\mathbf{k} \in \mathbb{Z}^d} \delta(\bullet - \mathbf{D}\mathbf{k}) \right\}(\boldsymbol{\omega}) = \frac{1}{(2\pi)^d} \left(\hat{f}(\bullet) \star \frac{(2\pi)^d}{|\mathbf{D}|} \sum_{\mathbf{m} \in \mathbb{Z}^d} \delta(\bullet - 2\pi\mathbf{D}^{-\text{T}}\mathbf{m}) \right)(\boldsymbol{\omega}),$$

where \star represents the continuous-domain convolution operation. Evaluating the r.h.s. of the above equation, we obtain the Poisson's summation formula

$$\sum_{\mathbf{k} \in \mathbb{Z}^d} f(\mathbf{D}\mathbf{k}) e^{-j(\mathbf{D}^{\text{T}}\boldsymbol{\omega})^{\text{T}}\mathbf{k}} = \frac{1}{|\mathbf{D}|} \sum_{\mathbf{m} \in \mathbb{Z}^d} \hat{f}(\boldsymbol{\omega} - 2\pi\mathbf{D}^{-\text{T}}\mathbf{m}),$$

provided both the summations converge.

A.2 Condition for the Existence of Riesz Bases

Here, we provide a detailed proof of Proposition 2.1 in Chapter 2. First, we establish that $A_\varphi(\boldsymbol{\omega}) \geq A_{\text{R},\varphi,2}^2$ by contradiction. Suppose there exists a set

$$\mathfrak{C} = \left\{ \boldsymbol{\omega} \in [0, 2\pi)^d : A_\varphi(\boldsymbol{\omega}) < A_{\text{R},\varphi,2}^2 \right\},$$

whose measure is not zero: $\int_{\mathfrak{C}} d\boldsymbol{\omega} > 0$. Then, correspondingly, let us consider a sequence $\{c_1[\mathbf{k}]\}$ such that its Fourier transform

$$C_1(e^{j\boldsymbol{\omega}}) = \begin{cases} 1, & \boldsymbol{\omega} \in \mathfrak{C}, \\ 0, & \boldsymbol{\omega} \in [0, 2\pi)^d \setminus \mathfrak{C}. \end{cases}$$

Obviously, $\{c_1[\mathbf{k}]\} \in \ell_2(\mathbb{Z}^d)$: (using Parseval's identity)

$$\|c_1\|_{\ell_2}^2 = \frac{1}{(2\pi)^d} \int_{[0, 2\pi)^d} |C_1(\boldsymbol{\omega})|^2 d\boldsymbol{\omega} = \frac{1}{(2\pi)^d} \int_{\mathfrak{C}} d\boldsymbol{\omega} < \frac{1}{(2\pi)^d} \int_{[0, 2\pi)^d} d\boldsymbol{\omega} = 1. \quad (\text{A.1})$$

Then, we have

$$\int_{[0, 2\pi)^d} |C_1(e^{j\boldsymbol{\omega}})|^2 (A_\varphi(\boldsymbol{\omega}) - A_{\text{R},\varphi,2}^2) d\boldsymbol{\omega} = \int_{\mathfrak{C}} (A_\varphi(\boldsymbol{\omega}) - A_{\text{R},\varphi,2}^2) d\boldsymbol{\omega} < 0,$$

which contradicts the left-side inequality in (2.12) for $c = c_1$. Therefore, \mathfrak{C} must necessarily be of measure zero so that $A_\varphi(\boldsymbol{\omega}) \geq A_{\text{R},\varphi,2}^2$ a.e.

Similarly, for the upper bound, if there exists

$$\mathfrak{D} = \left\{ \boldsymbol{\omega} \in [0, 2\pi)^d : A_\varphi(\boldsymbol{\omega}) > B_{R_\varphi, 2}^2 \right\}$$

that has a non-zero measure, we consider a sequence $\{c_2[\mathbf{k}]\}$ whose Fourier transform $C_2(e^{j\boldsymbol{\omega}})$ is such that

$$C_2(e^{j\boldsymbol{\omega}}) = \begin{cases} 1, & \boldsymbol{\omega} \in \mathfrak{D}, \\ 0, & \boldsymbol{\omega} \in [0, 2\pi)^d \setminus \mathfrak{D}. \end{cases}$$

Then, similar to (A.1), we have $\{c_2[\mathbf{k}]\} \in \ell_2(\mathbb{Z}^d)$ and

$$\int_{[0, 2\pi)^d} |C_2(e^{j\boldsymbol{\omega}})|^2 (A_\varphi(\boldsymbol{\omega}) - B_{R_\varphi, 2}^2) d\boldsymbol{\omega} = \int_{\mathfrak{D}} (A_\varphi(\boldsymbol{\omega}) - B_{R_\varphi, 2}^2) d\boldsymbol{\omega} > 0,$$

which contradicts the right-side inequality in (2.12) for $c = c_2$. Therefore, we must have that $A_\varphi(\boldsymbol{\omega}) \leq B_{R_\varphi, 2}^2$ a.e.

Conversely, if $A_{R_\varphi, 2}^2 \leq A_\varphi(\boldsymbol{\omega}) \leq B_{R_\varphi, 2}^2$ a.e., we start from

$$A_{R_\varphi, 2}^2 \|c\|_{\ell_2}^2 \leq \frac{1}{(2\pi)^d} \int_{[0, 2\pi)^d} |C(e^{j\boldsymbol{\omega}})|^2 A_\varphi(\boldsymbol{\omega}) d\boldsymbol{\omega} \leq B_{R_\varphi, 2}^2 \|c\|_{\ell_2}^2,$$

and retrace our steps from (2.15) to obtain (2.12) which completes the proof.

A.3 Approximation Error Formula in d -dimensions

Our proof of Theorem 2.1 is essentially similar to that given in [54]. By hypothesis $|\hat{h}(\boldsymbol{\omega})| < \infty$, $|Q(e^{j\boldsymbol{\omega}})| < \infty$ and φ generates a L_2 -stable Riesz basis. Moreover, by substituting $|L(\boldsymbol{\omega})|^2 = (1 + \|\boldsymbol{\omega}\|^2)^\rho - 1$ in Theorem 3.1 and Corollary 3.1 in Chapter 3, it can be verified that

$$\left\{ \left(f \star \frac{1}{T^d} h \left(\frac{\bullet}{T} \right) \right) (\mathbf{k}T) \right\} \in \ell_2(\mathbb{Z}^d).$$

Therefore, the Fourier transform of $\mathcal{Q}_{T\mathbf{I}}\{f\}(\mathbf{x})$ is given by

$$\mathcal{F}\{\mathcal{Q}_{T\mathbf{I}}\{f\}\}(\boldsymbol{\omega}) = Q(e^{j\boldsymbol{\omega}T}) \hat{\varphi}(\boldsymbol{\omega}T) \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{f} \left(\boldsymbol{\omega} + \frac{2\pi}{T} \mathbf{k} \right) \hat{h}(\boldsymbol{\omega}T + 2\pi\mathbf{k})$$

is well-defined in the L_2 sense. Therefore, by Parseval's identity,

$$\begin{aligned} \|\mathcal{Q}_{T\mathbf{I}}\{f\}\|_{L_2}^2 &= \frac{1}{(2\pi)^d} \|\mathcal{F}\{\mathcal{Q}_{T\mathbf{I}}\{f\}\}\|_{L_2}^2 \\ &= \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \left| Q(e^{j\boldsymbol{\omega}T}) \hat{\varphi}(\boldsymbol{\omega}T) \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{f}\left(\boldsymbol{\omega} + \frac{2\pi}{T}\mathbf{k}\right) \hat{h}(\boldsymbol{\omega}T + 2\pi\mathbf{k}) \right|^2 d\boldsymbol{\omega} \\ &= \frac{1}{(2\pi)^d} \int_{[0, \frac{2\pi}{T})^d} \left(\begin{array}{c} |Q(e^{j\boldsymbol{\omega}T})|^2 A_\varphi(\boldsymbol{\omega}T) \\ \times \left| \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{f}\left(\boldsymbol{\omega} + \frac{2\pi}{T}\mathbf{k}\right) \hat{h}(\boldsymbol{\omega}T + 2\pi\mathbf{k}) \right|^2 \end{array} \right) d\boldsymbol{\omega}. \end{aligned}$$

Since the integrand is positive, we can exchange summation and integration to obtain

$$\|\mathcal{Q}_{T\mathbf{I}}\{f\}\|_{L_2}^2 = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \left(\begin{array}{c} \hat{f}^*(\boldsymbol{\omega}) \hat{h}^*(\boldsymbol{\omega}T) |Q(e^{j\boldsymbol{\omega}T})|^2 A_\varphi(\boldsymbol{\omega}T) \\ \times \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{f}\left(\boldsymbol{\omega} + \frac{2\pi}{T}\mathbf{k}\right) \hat{h}(\boldsymbol{\omega}T + 2\pi\mathbf{k}) \end{array} \right) d\boldsymbol{\omega}.$$

Since, $|\langle f, \mathcal{Q}_{T\mathbf{I}}\{f\} \rangle_{L_2}| \leq \|f\|_{L_2} \|\mathcal{Q}_{T\mathbf{I}}\{f\}\|_{L_2} < +\infty$ by Cauchy-Schwarz inequality, we can similarly manipulate $\langle f, \mathcal{Q}_{T\mathbf{I}}\{f\} \rangle_{L_2}$ to get

$$\langle f, \mathcal{Q}_{T\mathbf{I}}\{f\} \rangle_{L_2} = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \left(\begin{array}{c} \hat{f}^*(\boldsymbol{\omega}) Q(e^{j\boldsymbol{\omega}T}) \hat{\varphi}(\boldsymbol{\omega}T) \\ \times \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{f}\left(\boldsymbol{\omega} + \frac{2\pi}{T}\mathbf{k}\right) \hat{h}(\boldsymbol{\omega}T + 2\pi\mathbf{k}) \end{array} \right) d\boldsymbol{\omega}.$$

Therefore, by Parseval's identity,

$$\begin{aligned} \varepsilon_f^2(T) &= \|f\|_{L_2}^2 + \|\mathcal{Q}_{T\mathbf{I}}\{f\}\|_{L_2}^2 - 2\langle f, \mathcal{Q}_{T\mathbf{I}}\{f\} \rangle_{L_2} \\ &= \frac{1}{(2\pi)^d} \left(\|\hat{f}\|_{L_2}^2 + \|\mathcal{F}\{\mathcal{Q}_{T\mathbf{I}}\{f\}\}\|_{L_2}^2 - 2\operatorname{Re} \left\{ \langle \hat{f}, \mathcal{F}\{\mathcal{Q}_{T\mathbf{I}}\{f\}\} \rangle_{L_2} \right\} \right) \\ &= \varepsilon_1^2 + \varepsilon_2^2, \end{aligned}$$

where

$$\begin{aligned} \varepsilon_1^2 &= \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} |\hat{f}(\boldsymbol{\omega})|^2 \left(\begin{array}{c} 1 \\ + |\hat{h}(\boldsymbol{\omega})|^2 |Q(e^{j\boldsymbol{\omega}T})|^2 A_\varphi(\boldsymbol{\omega}T) \\ - 2\operatorname{Re} \left\{ \hat{h}(\boldsymbol{\omega}) Q(e^{j\boldsymbol{\omega}T}) \hat{\varphi}(\boldsymbol{\omega}T) \right\} \end{array} \right) d\boldsymbol{\omega} \\ &= \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} |\hat{f}(\boldsymbol{\omega})|^2 \mathcal{E}_{h,\varphi}(\boldsymbol{\omega}T) d\boldsymbol{\omega}, \end{aligned}$$

and

$$\varepsilon_2^2(f, T) = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \hat{f}^*(\boldsymbol{\omega}) Q(e^{j\boldsymbol{\omega}T}) \left[\begin{array}{l} \left(\begin{array}{c} \hat{h}^*(\boldsymbol{\omega}T) Q^*(e^{j\boldsymbol{\omega}T}) A_\varphi(\boldsymbol{\omega}T) \\ - 2\hat{\varphi}(\boldsymbol{\omega}T) \end{array} \right) \\ \times \sum_{\mathbf{k} \neq \mathbf{0}} \hat{f}\left(\boldsymbol{\omega} + \frac{2\pi}{T}\mathbf{k}\right) \hat{h}(\boldsymbol{\omega}T + 2\pi\mathbf{k}) \end{array} \right] d\boldsymbol{\omega}.$$

Writing $\varepsilon_{\text{corr}}(f, T) = \varepsilon_f - \varepsilon_1$, we see that

$$\varepsilon_{\text{corr}}(f, T) = \frac{\varepsilon_2^2}{\varepsilon_f + \varepsilon_1} = 0, \tag{A.2}$$

whenever $\varepsilon_2 = 0$, which happens when f is bandlimited to

$$\begin{aligned} & \left(-\frac{\pi}{T}(\mathbf{k} + \mathbf{1}), -\frac{\pi}{T}\mathbf{k}\right] \cup \left[\frac{\pi}{T}\mathbf{k}, \frac{\pi}{T}(\mathbf{k} + \mathbf{1})\right) \\ \implies & \hat{f}^*(\boldsymbol{\omega}) \hat{f}\left(\boldsymbol{\omega} + \frac{2\pi}{T}\mathbf{k}\right) = 0, \quad \forall \mathbf{k} \neq \mathbf{0}, \end{aligned} \tag{A.3}$$

or h is bandlimited to $[-\pi, \pi]^d$:

$$\hat{h}(\boldsymbol{\omega}T + 2\pi\mathbf{k}) = 0, \quad \forall \mathbf{k} \neq \mathbf{0}. \tag{A.4}$$

Next, in order to bound $|\varepsilon_{\text{corr}}(f, T)|$, we partition the Fourier transform $\hat{f}(\boldsymbol{\omega})$ as follows:

$$\hat{f}_{\mathbf{k}}(\boldsymbol{\omega}) = \begin{cases} \hat{f}(\boldsymbol{\omega}), & \text{if } \frac{\pi}{T}k_i \leq |\omega_i| < \frac{\pi}{T}(k_i + 1), \quad i = 1, 2, \dots, d, \\ 0, & \text{otherwise,} \end{cases} \tag{A.5}$$

so that $\hat{f}(\boldsymbol{\omega}) = \sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{\mathbf{0}\}} \hat{f}_{\mathbf{k}}(\boldsymbol{\omega})$, where $\mathbb{Z}_+^d = \{\mathbf{k} \in \mathbb{Z}^d : k_i \geq 1, i = 1, 2, \dots, d\}$.

The underlying idea is to exploit properties (A.2) to (A.4) of $\varepsilon_{\text{corr}}$: Since the signal $f_{\mathbf{k}}(\mathbf{x})$ (corresponding to the \mathbf{k} -th bandpass segment $\hat{f}_{\mathbf{k}}(\boldsymbol{\omega})$) is bandlimited to the frequency range

$$\left(-\frac{\pi}{T}(\mathbf{k} + \mathbf{1}), -\frac{\pi}{T}\mathbf{k}\right] \cup \left[\frac{\pi}{T}\mathbf{k}, \frac{\pi}{T}(\mathbf{k} + \mathbf{1})\right),$$

we have $\varepsilon_{\text{corr}}(f_{\mathbf{k}}, T) = 0$, which leads to

$$\varepsilon_{f_{\mathbf{k}}}^2(T) = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} |\hat{f}_{\mathbf{k}}(\boldsymbol{\omega})|^2 \mathcal{E}_{h, \varphi}(\boldsymbol{\omega}T) d\boldsymbol{\omega}.$$

Then, since $\left| \sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{0\}} \hat{f}_{\mathbf{k}}(\boldsymbol{\omega}) \right|^2 = \sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{0\}} |\hat{f}_{\mathbf{k}}(\boldsymbol{\omega})|^2$ from (A.5), using Minkowski's inequality, we obtain

$$\begin{aligned} \varepsilon_1 &= \left(\frac{1}{(2\pi)^d} \sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{0\}} \int_{\mathbb{R}^d} |\hat{f}_{\mathbf{k}}(\boldsymbol{\omega})|^2 \mathcal{E}_{h,\varphi}(\boldsymbol{\omega}T) d\boldsymbol{\omega} \right)^{\frac{1}{2}} \\ &\leq \sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{0\}} \left(\frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} |\hat{f}_{\mathbf{k}}(\boldsymbol{\omega})|^2 \mathcal{E}_{h,\varphi}(\boldsymbol{\omega}T) d\boldsymbol{\omega} \right)^{\frac{1}{2}} \\ &= \sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{0\}} \varepsilon_{f_{\mathbf{k}}}. \end{aligned} \quad (\text{A.6})$$

Also,

$$\varepsilon_1 \geq \left(\frac{1}{(2\pi)^d} \int_{[-\frac{\pi}{T}, \frac{\pi}{T}]^d} |\hat{f}(\boldsymbol{\omega})|^2 \mathcal{E}_{h,\varphi}(\boldsymbol{\omega}T) d\boldsymbol{\omega} \right)^{\frac{1}{2}} = \varepsilon_{f_0}. \quad (\text{A.7})$$

Thus, (A.6) and (A.7) together lead to

$$2\varepsilon_{f_0} - \varepsilon_1 \leq \varepsilon_1 \leq \sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{0\}} \varepsilon_{f_{\mathbf{k}}}$$

or, equivalently,

$$\varepsilon_{f_0} - \sum_{\mathbf{k} \in \mathbb{Z}_+^d} \varepsilon_{f_{\mathbf{k}}} \leq \varepsilon_1 \leq \varepsilon_{f_0} + \sum_{\mathbf{k} \in \mathbb{Z}_+^d} \varepsilon_{f_{\mathbf{k}}}, \quad (\text{A.8})$$

which implies that

$$|\varepsilon_1 - \varepsilon_{f_0}| \leq \sum_{\mathbf{k} \in \mathbb{Z}_+^d} \varepsilon_{f_{\mathbf{k}}}. \quad (\text{A.9})$$

We now bound $\sum_{\mathbf{k} \in \mathbb{Z}_+^d} \varepsilon_{f_{\mathbf{k}}}$ as follows: Since ω_i is bounded below by $\frac{\pi}{T}k_i$ in the

\mathbf{k} -th segment of (A.5), we find that, for $\rho > \frac{d}{2}$,

$$1 \leq \frac{\|\boldsymbol{\omega}\|^{2\rho}}{\|\mathbf{k}\|^{2\rho}} \frac{T^{2\rho}}{\pi^{2\rho} d^\rho}.$$

Therefore,

$$\sum_{\mathbf{k} \in \mathbb{Z}_+^d} \varepsilon_{f_{\mathbf{k}}} \leq \frac{T^\rho}{\pi^\rho \sqrt{d^\rho}} \sum_{\mathbf{k} \in \mathbb{Z}_+^d} \frac{1}{\|\mathbf{k}\|^\rho} \left(\frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \|\boldsymbol{\omega}\|^{2\rho} |\hat{f}_{\mathbf{k}}(\boldsymbol{\omega})|^2 \mathcal{E}_{h,\varphi}(\boldsymbol{\omega}) \, d\boldsymbol{\omega} \right)^{\frac{1}{2}}.$$

Then, using Cauchy-Schwarz inequality for discrete sequences, we obtain

$$\begin{aligned} \sum_{\mathbf{k} \in \mathbb{Z}_+^d} \varepsilon_{f_{\mathbf{k}}} &\leq \frac{T^\rho \sqrt{\|\mathcal{E}_{h,\varphi}\|_{L_\infty}}}{\pi^\rho \sqrt{d^\rho}} \sqrt{\sum_{\mathbf{k} \in \mathbb{Z}_+^d} \frac{1}{\|\mathbf{k}\|^{2\rho}}}, \\ &\quad \times \left(\frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \|\boldsymbol{\omega}\|^{2\rho} \sum_{\mathbf{k} \in \mathbb{Z}_+^d} |\hat{f}_{\mathbf{k}}(\boldsymbol{\omega})|^2 \, d\boldsymbol{\omega} \right)^{\frac{1}{2}}. \end{aligned}$$

Again, $\sum_{\mathbf{k} \in \mathbb{Z}_+^d} |\hat{f}_{\mathbf{k}}(\boldsymbol{\omega})|^2 = \left| \sum_{\mathbf{k} \in \mathbb{Z}_+^d} \hat{f}_{\mathbf{k}}(\boldsymbol{\omega}) \right|^2 = |\hat{f}(\boldsymbol{\omega}) - \hat{f}_0(\boldsymbol{\omega})|^2$ from (A.5). Therefore,

$$\sum_{\mathbf{k} \in \mathbb{Z}_+^d} \varepsilon_{f_{\mathbf{k}}} \leq \frac{T^\rho \sqrt{\|\mathcal{E}_{h,\varphi}\|_{L_\infty}}}{\pi^\rho \sqrt{d^\rho}} \|f^{(\rho)} - f_0^{(\rho)}\|_{L_2} \sqrt{\sum_{\mathbf{k} \in \mathbb{Z}_+^d} \frac{1}{\|\mathbf{k}\|^{2\rho}}, \quad (\text{A.10})$$

where $\|f^{(\rho)}\|_{L_2} = \left(\frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \|\boldsymbol{\omega}\|^{2\rho} |f(\boldsymbol{\omega})|^2 \, d\boldsymbol{\omega} \right)^{\frac{1}{2}} < \infty$, when $f \in W_2^\rho$. In order to bound $\sum_{\mathbf{k} \in \mathbb{Z}_+^d} \|\mathbf{k}\|^{-2\rho}$, we set $\|\mathbf{k}\|^2 = m$ and write

$$\sum_{\mathbf{k} \in \mathbb{Z}_+^d} \frac{1}{\|\mathbf{k}\|^{2\rho}} = \sum_{m \geq 1} \frac{\mathcal{N}(m)}{m^\rho},$$

where $\mathcal{N}(m)$ is the number of points on a sphere of radius \sqrt{m} in d -dimensions. Obviously, $\mathcal{N}(m)$ is upper-bounded by the volume of the hyper-annulus of unit

thickness and radius \sqrt{m} :

$$\begin{aligned} \mathcal{N}(m) &\leq \frac{\pi^{\frac{d}{2}}}{\Gamma\left(\frac{d}{2}+1\right)} \left((m+1)^{\frac{d}{2}} - m^{\frac{d}{2}} \right) \\ &= \frac{\pi^{\frac{d}{2}}}{\Gamma\left(\frac{d}{2}+1\right)} m^{\frac{d}{2}} \left[\left(1 + \frac{1}{m}\right)^{\frac{d}{2}} - 1 \right] \\ &\leq \frac{\pi^{\frac{d}{2}}}{\Gamma\left(\frac{d}{2}+1\right)} m^{\frac{d}{2}-1} (2^{\frac{d}{2}} - 1), \end{aligned}$$

where we have used the fact that $\left(1 + \frac{1}{m}\right)^{\frac{d}{2}} \leq 2^{\frac{d}{2}}$ for $m \geq 1$. Therefore, we have

$$\begin{aligned} \sum_{\mathbf{k} \in \mathbb{Z}_+^d} \frac{1}{\|\mathbf{k}\|^{2\rho}} &\leq \frac{\pi^{\frac{d}{2}} (2^{\frac{d}{2}} - 1)}{\Gamma\left(\frac{d}{2}+1\right)} \sum_{m \geq 1} \frac{1}{m^{\rho - \frac{d}{2} + 1}} \\ &= \frac{\pi^{\frac{d}{2}} (2^{\frac{d}{2}} - 1)}{\Gamma\left(\frac{d}{2}+1\right)} \zeta\left(\rho - \frac{d}{2} + 1\right) < \infty, \end{aligned} \quad (\text{A.11})$$

where ζ is the Riemann-zeta function which, in this case, is finite whenever $\rho > \frac{d}{2}$. Thus, putting together all the results (A.9)-(A.11), we get

$$\begin{aligned} |\varepsilon_1 - \varepsilon_{f_0}| &\leq \frac{T^\rho \sqrt{\|\mathcal{E}_{h,\varphi}\|_{L_\infty}}}{\pi^\rho \sqrt{d^\rho}} \frac{\pi^{\frac{d}{4}} \sqrt{2^{\frac{d}{2}} - 1}}{\sqrt{\Gamma\left(\frac{d}{2}+1\right)}} \sqrt{\zeta\left(\rho - \frac{d}{2} + 1\right)} \|f^{(\rho)} - f_0^{(\rho)}\|_{L_2} \\ &\leq \mathcal{K} T^\rho \|f^{(\rho)}\|_{L_2} \Lambda(f, T), \end{aligned} \quad (\text{A.12})$$

where

$$\mathcal{K} = \frac{\pi^{\frac{d}{4} - \rho}}{\sqrt{d^\rho \Gamma\left(\frac{d}{2}+1\right)}} \sqrt{2^{\frac{d}{2}} - 1} \sqrt{\zeta\left(\rho - \frac{d}{2} + 1\right)} \|\mathcal{E}_{h,\varphi}\|_{L_\infty},$$

and

$$\begin{aligned}
 \Lambda(f, T) &= \frac{\|f^{(\rho)} - f_{\mathbf{0}}^{(\rho)}\|_{L_2}}{\|f^{(\rho)}\|_{L_2}} \\
 &= \frac{1}{\|f^{(\rho)}\|_{L_2}} \left(\frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \|\boldsymbol{\omega}\|^{2\rho} \sum_{\mathbf{k} \in \mathbb{Z}_+^d} |\hat{f}_{\mathbf{k}}(\boldsymbol{\omega})|^2 d\boldsymbol{\omega} \right)^{\frac{1}{2}} \\
 &= \frac{1}{\|f^{(\rho)}\|_{L_2}} \left(\frac{1}{(2\pi)^d} \int_{\mathbb{R}^d \setminus [-\frac{\pi}{T}, \frac{\pi}{T}]^d} \|\boldsymbol{\omega}\|^{2\rho} |\hat{f}(\boldsymbol{\omega})|^2 d\boldsymbol{\omega} \right)^{\frac{1}{2}} \\
 &\leq 1.
 \end{aligned}$$

Following a similar line of reasoning, we can establish an inequality such as (A.8) for ε_f so that

$$|\varepsilon_f - \varepsilon_{f_0}| \leq \mathcal{K} T^\rho \left\| f^{(\rho)} \right\|_{L_2} \Lambda(f, T). \tag{A.13}$$

Thus, combining (A.12) and (A.13), we obtain

$$\begin{aligned}
 |\varepsilon_{\text{corr}}| \leq |\varepsilon_f - \varepsilon_1| &\leq |\varepsilon_f - \varepsilon_0| + |\varepsilon_1 - \varepsilon_0| \\
 &\leq 2\mathcal{K} T^\rho \left\| f^{(\rho)} \right\|_{L_2} \Lambda(f, T) \\
 &\leq 2\mathcal{K} T^\rho \left\| f^{(\rho)} \right\|_{L_2} \quad (\text{since } \Lambda(f, T) \leq 1),
 \end{aligned}$$

so that $\varepsilon_{\text{corr}} = \mathcal{O}(T^\rho)$, which completes the proof.

Appendix B

Proofs of Some Results of Chapter 3

B.1 Sobolev Spaces—Proof of Theorem 3.1

Consider $\mathcal{A}_{\mathcal{I}}(\boldsymbol{\omega}) = \sum_{\mathbf{k} \in \mathcal{I}} \hat{f}(\boldsymbol{\omega} + 2\pi\mathbf{k})$ for some $f(\mathbf{x}) \in W_2^L$ and \mathcal{I} , a bounded subset of \mathbb{Z}^d . Using the Cauchy-Schwarz inequality, we have,

$$\begin{aligned} |\mathcal{A}_{\mathcal{I}}(\boldsymbol{\omega})|^2 &= \left| \sum_{\mathbf{k} \in \mathcal{I}} \hat{f}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \frac{\sqrt{1 + |\hat{L}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2}}{\sqrt{1 + |\hat{L}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2}} \right|^2 \\ &\leq \left(\sum_{\mathbf{k} \in \mathcal{I}} \frac{1}{1 + |\hat{L}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2} \right) \left(\sum_{\mathbf{k} \in \mathcal{I}} (1 + |\hat{L}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2) |\hat{f}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 \right) \\ &\leq C_0 \sum_{\mathbf{k} \in \mathbb{Z}^d} (1 + |\hat{L}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2) |\hat{f}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 = C_0 \mathcal{B}(\boldsymbol{\omega}). \end{aligned} \quad (\text{B.1})$$

But, $\mathcal{B}(\boldsymbol{\omega}) \in L_2([0, 2\pi)^d)$, as seen below: Since $f(\mathbf{x}) \in W_2^L$ and using Fubini-Tonelli theorem we write,

$$\begin{aligned} \int_{[0, 2\pi)^d} \mathcal{B}(\boldsymbol{\omega}) \, d\boldsymbol{\omega} &= \sum_{\mathbf{k} \in \mathbb{Z}^d} \int_{[0, 2\pi)^d} (1 + |\hat{L}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2) |\hat{f}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 \, d\boldsymbol{\omega} \\ &= \int_{\mathbb{R}^d} (1 + |\hat{L}(\boldsymbol{\omega})|^2) |\hat{f}(\boldsymbol{\omega})|^2 \, d\boldsymbol{\omega} < +\infty. \end{aligned} \quad (\text{B.2})$$

Thus, $\mathcal{A}_{\mathcal{I}}(\boldsymbol{\omega}) \in L_2([0, 2\pi)^d)$ because of (B.1) and (B.2). Moreover,

$$|\mathcal{A}_{\mathcal{I}}(\boldsymbol{\omega})| \leq \sqrt{C_0} \sqrt{\mathcal{B}(\boldsymbol{\omega})} \leq \sqrt{C_0} (1 + \mathcal{B}(\boldsymbol{\omega})),$$

and

$$\int_{[0, 2\pi)^d} \sqrt{\mathcal{B}(\boldsymbol{\omega})} \, d\boldsymbol{\omega} \leq (2\pi)^d + \int_{[0, 2\pi)^d} \mathcal{B}(\boldsymbol{\omega}) \, d\boldsymbol{\omega} < +\infty. \quad (\text{B.3})$$

Now, consider the following limit:

$$\lim_{\mathcal{I} \rightarrow \mathbb{Z}^d} \mathcal{A}_{\mathcal{I}}(\boldsymbol{\omega}) = \mathcal{A}(\boldsymbol{\omega}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{f}(\boldsymbol{\omega} + 2\pi\mathbf{k}).$$

Since $|\mathcal{A}_{\mathcal{I}}(\boldsymbol{\omega})|^2$ is bounded by an integrable function (*cf.* (B.1)), we use Lebesgue's dominated convergence theorem and write,

$$\begin{aligned} \int_{[0, 2\pi)^d} |\mathcal{A}(\boldsymbol{\omega})|^2 \, d\boldsymbol{\omega} &= \int_{[0, 2\pi)^d} \lim_{\mathcal{I} \rightarrow \mathbb{Z}^d} |\mathcal{A}_{\mathcal{I}}(\boldsymbol{\omega})|^2 \, d\boldsymbol{\omega} \\ &= \lim_{\mathcal{I} \rightarrow \mathbb{Z}^d} \int_{[0, 2\pi)^d} |\mathcal{A}_{\mathcal{I}}(\boldsymbol{\omega})|^2 \, d\boldsymbol{\omega} < +\infty. \end{aligned}$$

Thus, $\mathcal{A}(\boldsymbol{\omega}) \in L_2([0, 2\pi)^d)$ and therefore has a Fourier series expansion:

$$\mathcal{A}(\boldsymbol{\omega}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} c[\mathbf{k}] e^{-j\boldsymbol{\omega}^T \mathbf{k}} \quad \text{a.e.}$$

Again using Lebesgue's dominated convergence theorem ($|\mathcal{A}_{\mathcal{I}}(\boldsymbol{\omega})|$ is bounded by an integrable function), the coefficient $c[\mathbf{k}]$ can be written as

$$\begin{aligned} c[\mathbf{k}] &= \frac{1}{(2\pi)^d} \int_{[0, 2\pi)^d} \mathcal{A}(\boldsymbol{\omega}) e^{j\boldsymbol{\omega}^T \mathbf{k}} \, d\boldsymbol{\omega} = \frac{1}{(2\pi)^d} \int_{[0, 2\pi)^d} \lim_{\mathcal{I} \rightarrow \mathbb{Z}^d} \mathcal{A}_{\mathcal{I}}(\boldsymbol{\omega}) e^{j\boldsymbol{\omega}^T \mathbf{k}} \, d\boldsymbol{\omega} \\ &= \frac{1}{(2\pi)^d} \lim_{\mathcal{I} \rightarrow \mathbb{Z}^d} \int_{[0, 2\pi)^d} \mathcal{A}_{\mathcal{I}}(\boldsymbol{\omega}) e^{j\boldsymbol{\omega}^T \mathbf{k}} \, d\boldsymbol{\omega} \\ &= \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \hat{f}(\boldsymbol{\omega}) e^{j\boldsymbol{\omega}^T \mathbf{k}} \, d\boldsymbol{\omega} \\ &= f(\mathbf{k}). \end{aligned}$$

Thus, the coefficients $\{c[\mathbf{k}]\}$ are nothing but the samples of the function $f(\mathbf{x})$ evaluated at $\mathbf{x} = \mathbf{k} \in \mathbb{Z}^d$. Finally, because $\mathcal{A}(\boldsymbol{\omega}) \in L_2([0, 2\pi)^d)$, by Parseval's identity we have $\{f(\mathbf{k})\} \in \ell_2(\mathbb{Z}^d)$.

B.2 Orthogonality Principle—Proof of Theorem 3.2

For any $u \in W_2^L$, we have $(h \star u)(\mathbf{x}) \in W_2^L$ by Corollary 3.1, as well as $\{(h \star u)(\mathbf{k})\} \in \ell_2(\mathbb{Z}^d)$, by Theorem 3.1. Correspondingly, Proposition 3.3 ensures that $f_{\text{con}} \in W_2^L$.

Now, consider the function defined as $s(\mathbf{x}) = u(\mathbf{x}) - f_{\text{con}}(\mathbf{x})$. Since W_2^L is closed, $s(\mathbf{x}) \in W_2^L$. Consequently, $(h \star s)(\mathbf{x}) \in W_2^L$ and $(h \star s)(\mathbf{k}) \in \ell_2(\mathbb{Z}^d)$. But, $(h \star s)(\mathbf{k}) = 0, \forall \mathbf{k} \in \mathbb{Z}^d$, because of the consistency of $f_{\text{con}}(\mathbf{x})$. Therefore, the following Poisson's summation formula holds:

$$\sum_{\mathbf{k} \in \mathbb{Z}^d} (h \star s)(\mathbf{k}) e^{-j\boldsymbol{\omega}^T \mathbf{k}} = \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{s}(\boldsymbol{\omega} + 2\pi\mathbf{k}) = 0, \quad \text{a.e.,}$$

and hence,

$$\left\| \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{s}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \right\|_{L_2([0, 2\pi)^d)} = 0. \quad (\text{B.4})$$

Writing,

$$\begin{aligned} \|Lu\|_{L_2}^2 &= \|L\{u - f_{\text{con}} + f_{\text{con}}\}\|_{L_2}^2 \\ &= \|L\{s\}\|_{L_2}^2 + \|L\{f_{\text{con}}\}\|_{L_2}^2 + \langle L\{s\}, L\{f_{\text{con}}\} \rangle_{L_2}, \end{aligned}$$

it only remains to show that $\langle L\{s\}, L\{f_{\text{con}}\} \rangle_{L_2} = 0$. Letting

$$D(e^{j\boldsymbol{\omega}}) = R_{0,\text{opt}}(e^{j\boldsymbol{\omega}}) |Q(e^{j\boldsymbol{\omega}})|^2 \sum_{\mathbf{k} \in \mathbb{Z}^d} \hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{u}(\boldsymbol{\omega} + 2\pi\mathbf{k}),$$

and using Parseval's relation,

$$\begin{aligned} \langle L\{f_{\text{con}}\}, L\{s\} \rangle_{L_2} &= \frac{1}{(2\pi)^d} \langle \hat{L} \hat{f}_{\text{con}}, \hat{L} \hat{s} \rangle_{L_2} \\ &= \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} D(e^{j\boldsymbol{\omega}}) \hat{h}(\boldsymbol{\omega}) \hat{s}(\boldsymbol{\omega}) \, d\boldsymbol{\omega} \\ &= \lim_{\mathcal{I} \rightarrow \mathbb{Z}^d} \frac{1}{(2\pi)^d} \int_{[0, 2\pi)^d} D(e^{j\boldsymbol{\omega}}) \left(\sum_{\mathbf{k} \in \mathcal{I}} \hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{s}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \right) \, d\boldsymbol{\omega}. \end{aligned}$$

Then, using Cauchy-Schwarz inequality, we see that

$$\begin{aligned} \langle L\{f_{\text{con}}\}, L\{s\} \rangle_{L_2} &\leq \left(\frac{1}{(2\pi)^d} \|D(e^{j\boldsymbol{\omega}})\|_{L_2([0, 2\pi]^d)} \right) \\ &\quad \times \left(\lim_{\mathcal{I} \rightarrow \mathbb{Z}^d} \left\| \sum_{\mathbf{k} \in \mathcal{I}} \hat{h}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \hat{s}(\boldsymbol{\omega} + 2\pi\mathbf{k}) \right\|_{L_2([0, 2\pi]^d)} \right) \\ &= 0. \end{aligned}$$

Since $R_{0,\text{opt}}(e^{j\boldsymbol{\omega}})$, $|Q(e^{j\boldsymbol{\omega}})|$ are bounded and $(h \star u)(\mathbf{k}) \in \ell_2$, the first norm is finite and by (B.4), the limit of the second norm tends to zero, thus proving (3.17) in Chapter 3.

B.3 Inverse Fourier Transform of d -dimensional Matérn B-spline

Using the definition of the Gamma function $\Gamma(\bullet)$:

$$\Gamma(\gamma_m)(\alpha_m + \|\boldsymbol{\omega}\|^2)^{-\gamma_m} = \int_0^\infty t^{\gamma_m-1} e^{-t(\alpha_m + \|\boldsymbol{\omega}\|^2)} dt,$$

we write the inverse Fourier transform of $\hat{\beta}_{\alpha_m, \gamma_m}(\boldsymbol{\omega})$ as

$$\begin{aligned} \beta_{\alpha_m, \gamma_m}(\mathbf{x}) &= \frac{(2\pi)^{-d}}{\Gamma(\gamma_m)} \int_0^\infty t^{\gamma_m-1} \exp\left(-\alpha_m t + \frac{\|\mathbf{x}\|^2}{4t}\right) \\ &\quad \times \underbrace{\left(\int_{\boldsymbol{\omega} \in \mathbb{R}^d} \exp\left(-t \left\| \boldsymbol{\omega} - j \frac{\mathbf{x}}{2t} \right\|^2\right) d\boldsymbol{\omega} \right)}_{= \left(\sqrt{\frac{\pi}{t}}\right)^d} dt. \end{aligned}$$

Substituting $t = \frac{\|\mathbf{x}\|}{2\sqrt{\alpha_m}} e^{-u}$ and manipulating, we obtain

$$\begin{aligned} \beta_{\alpha_m, \gamma_m}(\mathbf{x}) &= \mathcal{C} \left(\frac{\|\mathbf{x}\|}{\sqrt{\alpha_m}} \right)^{\nu_m} \int_0^\infty \cosh(\nu_m u) e^{-\|\mathbf{x}\| \sqrt{\alpha_m} \cosh(u)} du \\ &= \mathcal{C} \left(\frac{\|\mathbf{x}\|}{\sqrt{\alpha_m}} \right)^{\nu_m} \mathcal{K}_{\nu_m}(\sqrt{\alpha_m} \|\mathbf{x}\|), \end{aligned}$$

where $\nu_m = \gamma_m - \frac{d}{2}$, $C = (\sqrt{(2\pi)^d} 2^{\gamma_m-1} \Gamma(\gamma_m))^{-1}$ and the integral in the last step is the modified Bessel function of the second kind [87].

B.4 Approximation Error for the Noisy Scenario

To prove Proposition 3.10, we first expand $E \left\{ |f(\mathbf{x}) - \mathcal{Q}_{TI}\{f\}(\mathbf{x})|^2 \right\}$ to get

$$E \left\{ |f(\mathbf{x}) - \mathcal{Q}_{TI}\{f\}(\mathbf{x})|^2 \right\} = E \left\{ |f(\mathbf{x})|^2 \right\} + E \left\{ |\mathcal{Q}_{TI}\{f\}(\mathbf{x})|^2 \right\} - 2 E \left\{ f(\mathbf{x}) \mathcal{Q}_{TI}\{f\}(\mathbf{x}) \right\}.$$

We concentrate on the second term

$$E \left\{ |\mathcal{Q}_{TI}\{f\}(\mathbf{x})|^2 \right\} = E \left\{ \sum_{\substack{\mathbf{k}, \mathbf{l} \in \mathbb{Z}^d \\ \mathbf{m}, \mathbf{n} \in \mathbb{Z}^d}} g[\mathbf{l}] g[\mathbf{n}] r[\mathbf{k} - \mathbf{l}] r[\mathbf{m} - \mathbf{n}] \varphi \left(\frac{\mathbf{x}}{T} - \mathbf{k} \right) \varphi \left(\frac{\mathbf{x}}{T} - \mathbf{m} \right) \right\},$$

where we can exchange the summation and expectation provided

$$E \left\{ \sum_{\substack{\mathbf{k}, \mathbf{l} \in \mathbb{Z}^d \\ \mathbf{m}, \mathbf{n} \in \mathbb{Z}^d}} |g[\mathbf{l}] g[\mathbf{n}]| |r[\mathbf{k} - \mathbf{l}] r[\mathbf{m} - \mathbf{n}]| \left| \varphi \left(\frac{\mathbf{x}}{T} - \mathbf{k} \right) \varphi \left(\frac{\mathbf{x}}{T} - \mathbf{m} \right) \right| \right\} < \infty,$$

which we verify now. Since the argument inside $E\{\bullet\}$ is positive, we move the expectation inside and use the fact that $E\{|\bullet|\} \leq 1 + E\{|\bullet|^2\}$ to get

$$\begin{aligned} & \sum_{\substack{\mathbf{k}, \mathbf{l} \in \mathbb{Z}^d \\ \mathbf{m}, \mathbf{n} \in \mathbb{Z}^d}} E\{|g[\mathbf{l}] g[\mathbf{n}]|\} |r[\mathbf{k} - \mathbf{l}] r[\mathbf{m} - \mathbf{n}]| \left| \varphi \left(\frac{\mathbf{x}}{T} - \mathbf{k} \right) \varphi \left(\frac{\mathbf{x}}{T} - \mathbf{m} \right) \right| \tag{B.5} \\ & \leq \sum_{\substack{\mathbf{k}, \mathbf{l} \in \mathbb{Z}^d \\ \mathbf{m}, \mathbf{n} \in \mathbb{Z}^d}} (1 + E\{|g[\mathbf{l}]|^2\}) (1 + E\{|g[\mathbf{n}]|^2\}) |r[\mathbf{k} - \mathbf{l}] r[\mathbf{m} - \mathbf{n}]| \\ & \quad \times \left| \varphi \left(\frac{\mathbf{x}}{T} - \mathbf{k} \right) \varphi \left(\frac{\mathbf{x}}{T} - \mathbf{m} \right) \right|. \end{aligned}$$

Now, since noise is zero-mean and independent of f , we have $\forall \mathbf{k} \in \mathbb{Z}^d$,

$$E\{|g[\mathbf{k}]|^2\} = E \left\{ \left[\frac{1}{T^d} \left(f \star h \left(\frac{\bullet}{T} \right) \right) (\mathbf{k}T) \right]^2 \right\} + c_n[\mathbf{0}],$$

where \star represents continuous convolution operation and $c_n[\mathbf{k}] = E\{n[\mathbf{l}] n[\mathbf{k}+\mathbf{l}]\}$ is the auto-correlation sequence of the noise. The first term on the r.h.s. of the above equation can be bounded as

$$\begin{aligned}
 & E \left\{ \left[\frac{1}{T^d} \left(f \star h \left(\frac{\bullet}{T} \right) \right) (\mathbf{k}T) \right]^2 \right\} \\
 & \leq \frac{1}{T^{2d}} \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} E \{ |f(\boldsymbol{\xi}_1) f(\boldsymbol{\xi}_2)| \} \left| h \left(\mathbf{k} - \frac{\boldsymbol{\xi}_1}{T} \right) h \left(\mathbf{k} - \frac{\boldsymbol{\xi}_2}{T} \right) \right| d\boldsymbol{\xi}_1 d\boldsymbol{\xi}_2 \\
 & \leq \frac{1}{T^{2d}} \int_{\mathbb{R}^d} \int_{\mathbb{R}^d} (1 + E \{ |f(\boldsymbol{\xi}_1)|^2 \}) (1 + E \{ |f(\boldsymbol{\xi}_2)|^2 \}) \\
 & \quad \times \left| h \left(\mathbf{k} - \frac{\boldsymbol{\xi}_1}{T} \right) h \left(\mathbf{k} - \frac{\boldsymbol{\xi}_2}{T} \right) \right| d\boldsymbol{\xi}_1 d\boldsymbol{\xi}_2 \\
 & \leq (\|h\|_{L_1})^2 (1 + c_f(\mathbf{0}))^2 \stackrel{\text{def}}{=} M_{c_f, h} < \infty, \tag{B.6}
 \end{aligned}$$

since $h(\mathbf{x}) \in L_1(\mathbb{R}^d)$ by hypothesis. Therefore, (B.5) becomes

$$\begin{aligned}
 & \sum_{\substack{\mathbf{k}, \mathbf{l} \in \mathbb{Z}^d \\ \mathbf{m}, \mathbf{n} \in \mathbb{Z}^d}} E \{ |g[\mathbf{l}] g[\mathbf{n}]| \} |r[\mathbf{k} - \mathbf{l}] r[\mathbf{m} - \mathbf{n}]| \left| \varphi \left(\frac{\mathbf{x}}{T} - \mathbf{k} \right) \varphi \left(\frac{\mathbf{x}}{T} - \mathbf{m} \right) \right| \\
 & \leq (1 + M_{c_f, h} + c_n[\mathbf{0}])^2 (\|r\|_{\ell_1})^2 \left(\sup_{\mathbf{x} \in [0, T]^d} \sum_{\mathbf{k} \in \mathbb{Z}^d} \left| \varphi \left(\frac{\mathbf{x}}{T} - \mathbf{k} \right) \right| \right)^2 < \infty,
 \end{aligned}$$

since $\{r[\mathbf{k}]\} \in \ell_1(\mathbb{Z}^d)$ and φ satisfies the **SR** conditions. Therefore, we can exchange the expectation and the summation in $E \{ |\mathcal{Q}_{TI}\{f\}(\mathbf{x})|^2 \}$ to get

$$\begin{aligned}
 E \{ |\mathcal{Q}_{TI}\{f\}(\mathbf{x})|^2 \} & = \sum_{\substack{\mathbf{k}, \mathbf{l} \in \mathbb{Z}^d \\ \mathbf{m}, \mathbf{n} \in \mathbb{Z}^d}} E \{ |g[\mathbf{l}] g[\mathbf{n}]| r[\mathbf{k} - \mathbf{l}] r[\mathbf{m} - \mathbf{n}] \varphi \left(\frac{\mathbf{x}}{T} - \mathbf{k} \right) \varphi \left(\frac{\mathbf{x}}{T} - \mathbf{m} \right) \} \\
 & = \sum_{\substack{\mathbf{k}, \mathbf{l} \in \mathbb{Z}^d \\ \mathbf{m}, \mathbf{n} \in \mathbb{Z}^d}} (g + c_n)[\mathbf{n} - \mathbf{l}] r[\mathbf{k} - \mathbf{l}] r[\mathbf{m} - \mathbf{n}] \varphi \left(\frac{\mathbf{x}}{T} - \mathbf{k} \right) \varphi \left(\frac{\mathbf{x}}{T} - \mathbf{m} \right),
 \end{aligned}$$

where (expectation and integration operations can be interchanged due to (B.6))

$$g[\mathbf{k}] = \frac{1}{T^{2d}} \left(h \left(-\frac{\bullet}{T} \right) \star c_f \star h \left(\frac{\bullet}{T} \right) \right) (\mathbf{k}T).$$

Since $c_f \in W_2^L$, we have $\{\mathbf{g}[\mathbf{k}]\} \in \ell_2(\mathbb{Z}^d)$ from Theorem 3.1 and Corollary 3.1. Therefore, $\{(r * \mathbf{g} * \bar{r})[\mathbf{k}]\} \in \ell_2(\mathbb{Z}^d)$ by Young's inequality (see Section 3.2 in Chapter 3), where $*$ represents discrete convolution operation. Similarly, $\{(r * c_n * \bar{r})[\mathbf{k}]\} \in \ell_2(\mathbb{Z}^d)$ because $\{c_n[\mathbf{k}]\} \in \ell_2(\mathbb{Z}^d)$.

Thus, we have

$$E \left\{ |\mathcal{Q}_{T\mathbf{I}}\{f\}(\mathbf{x})|^2 \right\} = \sum_{\mathbf{k} \in \mathbb{Z}^d} (r * (\mathbf{g} + c_n) * \bar{r})[\mathbf{k}] \times \left(\sum_{\mathbf{m} \in \mathbb{Z}^d} \varphi\left(\frac{\mathbf{x}}{T} - \mathbf{m}\right) \varphi\left(\frac{\mathbf{x}}{T} - \mathbf{k} - \mathbf{m}\right) \right),$$

which is T -period because of the T -periodicity of the bracketed term on the r.h.s. Then,

$$\int_{[0, T]^d} E \left\{ |\mathcal{Q}_{T\mathbf{I}}\{f\}(\mathbf{x})|^2 \right\} \frac{d\mathbf{x}}{T^d} = \int_{[0, T]^d} \sum_{\mathbf{k} \in \mathbb{Z}^d} (r * (\mathbf{g} + c_n) * \bar{r})[\mathbf{k}] \times \left(\sum_{\mathbf{m} \in \mathbb{Z}^d} \varphi\left(\frac{\mathbf{x}}{T} - \mathbf{m}\right) \varphi\left(\frac{\mathbf{x}}{T} - \mathbf{k} - \mathbf{m}\right) \right) \frac{d\mathbf{x}}{T^d}. \quad (\text{B.7})$$

As φ satisfies **SR** conditions and because $\{(r * \mathbf{g} * \bar{r})[\mathbf{k}]\}, \{(r * c_n * \bar{r})[\mathbf{k}]\} \in \ell_2(\mathbb{Z}^d)$, we have (using $|xy| \leq \frac{1}{2}(x + y)^2$)

$$\begin{aligned} & \int_{[0, T]^d} \sum_{\mathbf{k} \in \mathbb{Z}^d} |(r * (\mathbf{g} + c) * \bar{r})[\mathbf{k}]| \left(\sum_{\mathbf{m} \in \mathbb{Z}^d} \frac{|\varphi\left(\frac{\mathbf{x}}{T} - \mathbf{m}\right)|}{|\varphi\left(\frac{\mathbf{x}}{T} - \mathbf{k} - \mathbf{m}\right)|} \right) \frac{d\mathbf{x}}{T^d} \\ & \leq \int_{[0, 1]^d} \sum_{\mathbf{k} \in \mathbb{Z}^d} (|(r * \mathbf{g} * \bar{r})[\mathbf{k}]| + |(r * c_n * \bar{r})[\mathbf{k}]|) \\ & \quad \times \left(\sum_{\mathbf{m} \in \mathbb{Z}^d} \frac{|\varphi(\mathbf{x} - \mathbf{m})|}{|\varphi(\mathbf{x} - \mathbf{k} - \mathbf{m})|} \right) d\mathbf{x} \\ & \leq \frac{1}{2} \sum_{\mathbf{k} \in \mathbb{Z}^d} (|(r * \mathbf{g} * \bar{r})[\mathbf{k}]| + |(r * c_n * \bar{r})[\mathbf{k}]|)^2 \\ & \quad + \frac{1}{2} \int_{[0, 1]^d} \sum_{\mathbf{k} \in \mathbb{Z}^d} \left(\sum_{\mathbf{m} \in \mathbb{Z}^d} \frac{|\varphi(\mathbf{x} - \mathbf{m})|}{|\varphi(\mathbf{x} - \mathbf{k} - \mathbf{m})|} \right)^2 d\mathbf{x} \end{aligned}$$

$$\begin{aligned}
&\leq \frac{1}{2} \|r * \mathbf{g} * \bar{r}\|_{\ell_2}^2 + \frac{1}{2} \|r * c_n * \bar{r}\|_{\ell_2}^2 + \sum_{\mathbf{k} \in \mathbb{Z}^d} |(r * \mathbf{g} * \bar{r})[\mathbf{k}]| |(r * c_n * \bar{r})[\mathbf{k}]| \\
&\quad + \frac{1}{2} \int_{[0,1]^d} \sum_{\substack{\mathbf{m} \in \mathbb{Z}^d \\ \mathbf{n} \in \mathbb{Z}^d}} |\varphi(\mathbf{x} - \mathbf{m})| \varphi(\mathbf{x} - \mathbf{n}) \left(\sum_{\mathbf{k} \in \mathbb{Z}^d} |\varphi(\mathbf{x} - \mathbf{k} - \mathbf{m})| \right. \\
&\quad \left. \times |\varphi(\mathbf{x} - \mathbf{k} - \mathbf{n})| \right) d\mathbf{x} \\
&\leq \|r * \mathbf{g} * \bar{r}\|_{\ell_2}^2 + \|r * c_n * \bar{r}\|_{\ell_2}^2 \\
&\quad + \frac{1}{2} \int_{[0,1]^d} \left(\sum_{\mathbf{m} \in \mathbb{Z}^d} |\varphi(\mathbf{x} - \mathbf{m})| \right)^2 \left(\sum_{\mathbf{k} \in \mathbb{Z}^d} |\varphi(\mathbf{x} - \mathbf{k})|^2 \right) d\mathbf{x} \\
&\leq \|r * \mathbf{g} * \bar{r}\|_{\ell_2}^2 + \|r * c_n * \bar{r}\|_{\ell_2}^2 + \frac{1}{2} \left(\sup_{\mathbf{x} \in [0,1]^d} \sum_{\mathbf{m} \in \mathbb{Z}^d} |\varphi(\mathbf{x} - \mathbf{m})| \right)^4 < \infty,
\end{aligned}$$

where we have used the following inequality in the last step:

$$\sum_{\mathbf{k} \in \mathbb{Z}^d} |\varphi(\mathbf{x} - \mathbf{k})|^2 \leq \left(\sum_{\mathbf{k} \in \mathbb{Z}^d} |\varphi(\mathbf{x} - \mathbf{k})| \right)^2.$$

Therefore, the summation and integration can be exchanged in the r.h.s. of (B.7) to yield

$$\begin{aligned}
\int_{[0,T]^d} E \left\{ |\mathcal{Q}_{\mathbf{T}\mathbf{I}}\{f\}(\mathbf{x})|^2 \right\} \frac{d\mathbf{x}}{T^d} &= \sum_{\mathbf{k} \in \mathbb{Z}^d} (r * (\mathbf{g} + c_n) * \bar{r})[\mathbf{k}] \mathbf{a}_\varphi[\mathbf{k}] \\
&= \langle (r * \mathbf{g} * \bar{r}), \mathbf{a} \rangle_{\ell_2} + \langle (r * c_n * \bar{r}), \mathbf{a} \rangle_{\ell_2} \quad (\text{B.8})
\end{aligned}$$

where

$$\begin{aligned}
\mathbf{a}_\varphi[\mathbf{k}] &= \sum_{\mathbf{m} \in \mathbb{Z}^d} \int_{[0,T]^d} \varphi\left(\frac{\mathbf{x}}{T} - \mathbf{m}\right) \varphi\left(\frac{\mathbf{x}}{T} - \mathbf{k} - \mathbf{m}\right) \frac{d\mathbf{x}}{T^d} \\
&= \int_{\mathbb{R}^d} \varphi(\mathbf{x}) \varphi(\mathbf{x} - \mathbf{k}) d\mathbf{x}.
\end{aligned}$$

Using Parseval's identity, we write¹

$$\begin{aligned} \langle r * \mathbf{g} * \bar{r}, \mathbf{a} \rangle_{\ell_2} &= \frac{T^d}{(2\pi)^d} \int_{[0, \frac{2\pi}{T}]^d} \left(\frac{1}{T^d} \sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{h}(\boldsymbol{\omega}T + 2\pi\mathbf{k})|^2 \hat{c}_f \left(\boldsymbol{\omega} + \frac{2\pi}{T} \mathbf{k} \right) \right) \\ &\quad \times |R(e^{j\boldsymbol{\omega}T})|^2 \left(\sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{\varphi}(\boldsymbol{\omega}T + 2\pi\mathbf{k})|^2 \right) d\boldsymbol{\omega}. \end{aligned}$$

Simplifying the r.h.s. of the above equation, we obtain

$$\langle r * \mathbf{g} * \bar{r}, \mathbf{a} \rangle_{\ell_2} = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} |\hat{h}(\boldsymbol{\omega}T)|^2 \hat{c}_f(\boldsymbol{\omega}) |R(e^{j\boldsymbol{\omega}T})|^2 A_\varphi(\boldsymbol{\omega}T) d\boldsymbol{\omega}, \quad (\text{B.9})$$

where $A_\varphi(\boldsymbol{\omega}) = \sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2$. Using the same technique,¹

$$\begin{aligned} \langle (r * c_n * \bar{r}), \mathbf{a} \rangle_{\ell_2} &= \frac{1}{(2\pi)^d} \int_{[0, 2\pi]^d} C_n(e^{j\boldsymbol{\omega}}) |R(e^{j\boldsymbol{\omega}})|^2 \sum_{\mathbf{k} \in \mathbb{Z}^d} |\hat{\varphi}(\boldsymbol{\omega} + 2\pi\mathbf{k})|^2 d\boldsymbol{\omega} \\ &= \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} C_n(e^{j\boldsymbol{\omega}}) |R(e^{j\boldsymbol{\omega}})|^2 |\hat{\varphi}(\boldsymbol{\omega})|^2 d\boldsymbol{\omega}. \end{aligned} \quad (\text{B.10})$$

We develop $E \{f(\mathbf{x}) \mathcal{Q}_{T\mathbf{I}}\{f\}(\mathbf{x})\}$ in a similar fashion to obtain

$$E \{f(\mathbf{x}) \mathcal{Q}_{T\mathbf{I}}\{f\}(\mathbf{x})\} = \sum_{\mathbf{k} \in \mathbb{Z}^d} r[\mathbf{k}] \left(\sum_{\mathbf{m} \in \mathbb{Z}^d} \begin{pmatrix} \frac{1}{T^d} h\left(-\frac{\bullet}{T}\right) * c_f(\mathbf{x} - \mathbf{m}T) \\ \times \varphi\left(\frac{\mathbf{x}}{T} - \mathbf{k} - \mathbf{m}\right) \end{pmatrix} \right),$$

which is also T -periodic. Therefore,

$$\int_{[0, T]^d} E \{f(\mathbf{x}) \mathcal{Q}_{T\mathbf{I}}\{f\}(\mathbf{x})\} \frac{d\mathbf{x}}{T^d} = \sum_{\mathbf{k} \in \mathbb{Z}^d} r[\mathbf{k}] \left(\frac{1}{T^d} h\left(-\frac{\bullet}{T}\right) * c_f * \frac{1}{T^d} \varphi\left(-\frac{\bullet}{T}\right) \right) (\mathbf{k}T),$$

¹We exchange the summation and the integral since the integrand is positive.

which can be written in the Fourier domain using Parseval's identity as²

$$\begin{aligned}
 & \int_{[0, T]^d} E \{f(\mathbf{x}) \mathcal{Q}_{T\mathbf{I}}\{f\}(\mathbf{x})\} \frac{d\mathbf{x}}{T^d} \\
 &= \frac{T^d}{(2\pi)^d} \int_{[0, \frac{2\pi}{T}]^d} R^*(e^{j\boldsymbol{\omega}T}) \left(\sum_{\mathbf{k} \in \mathbb{Z}^d} \frac{1}{T^d} \hat{h}^*(\boldsymbol{\omega}T + 2\pi\mathbf{k}) \hat{c}_f(\boldsymbol{\omega} + \frac{2\pi}{T}\mathbf{k}) \right) d\boldsymbol{\omega}, \\
 &= \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} R^*(e^{j\boldsymbol{\omega}T}) \hat{h}^*(\boldsymbol{\omega}T) \hat{c}_f(\boldsymbol{\omega}) \hat{\varphi}^*(\boldsymbol{\omega}T) d\boldsymbol{\omega}. \tag{B.11}
 \end{aligned}$$

The first term is nothing but

$$E \left\{ |f(\mathbf{x})|^2 \right\} = c_f(\mathbf{0}) = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \hat{c}_f(\boldsymbol{\omega}) d\boldsymbol{\omega}. \tag{B.12}$$

Therefore, combining (B.8) to (B.12), we obtain the desired result:

$$\begin{aligned}
 \varepsilon_{f, \text{noisy}}(T) &= \frac{1}{T^d} \int_{[0, T]^d} E \left\{ |f(\mathbf{x}) - \mathcal{Q}_{T\mathbf{I}}\{f\}(\mathbf{x})|^2 \right\} d\mathbf{x} \\
 &= \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} [\mathcal{E}_{h, \varphi}(\boldsymbol{\omega}T) \hat{c}_f(\boldsymbol{\omega}) + C_n(e^{j\boldsymbol{\omega}}) |R(e^{j\boldsymbol{\omega}})|^2 |\hat{\varphi}(\boldsymbol{\omega})|^2] d\boldsymbol{\omega}.
 \end{aligned}$$

²The summation and the integration operations can be exchanged since $\int_{\mathbb{R}^d} |R(e^{j\boldsymbol{\omega}T})| |\hat{h}^*(\boldsymbol{\omega}T)| |\hat{c}_f(\boldsymbol{\omega})| |\hat{\varphi}^*(\boldsymbol{\omega}T)| d\boldsymbol{\omega} \leq \|r\|_{\ell_1} \|h\|_{L_1} \frac{1}{2} (\|c_f\|_{L_2}^2 + \|\varphi\|_{L_2}^2) < \infty$.

Appendix C

Poisson-Likelihood-Based Regularized Interpolation

Here, we present an algorithm for regularization interpolation when the data fidelity is governed by the Poisson likelihood $\mathcal{L}_{\text{Poisson}}$. This always leads to a non-linear scheme irrespective of whether the regularization is quadratic or not (see Table 4.1 in Chapter 4). In the case of non-quadratic regularization, the cost to be minimized is

$$J_{\text{PLNQ}}\{g, c\} = - \sum_{\mathbf{k} \in \mathbb{Z}^d} g[\mathbf{k}] \log((c * b)[\mathbf{k}]) + \sum_{\mathbf{k} \in \mathbb{Z}^d} (c * b)[\mathbf{k}] + \lambda \Psi_{\text{NQ}}\{c\}. \quad (\text{C.1})$$

Since J_{PLNQ} contains a term of the form $\log((c * b)[\mathbf{k}])$, it is imperative that $(c * b)[\mathbf{k}] > 0, \forall \mathbf{k} \in \mathbb{Z}^d$. Therefore, whenever we use $\mathcal{L}_{\text{Poisson}}$, we restrict ourselves to basis functions that satisfy $b[\mathbf{k}] = \varphi(\mathbf{k}) \geq 0 \forall \mathbf{k} \in \mathbb{Z}^d$ (e.g., integer-degree polynomial B-splines). Then, the above constraint simply boils down to ensuring the positivity of $\{c[\mathbf{k}]\}$.

Typically, such estimations based on Poisson likelihood are performed in an expectation-maximization (EM) framework [97, 98, 101]. However, it was demonstrated in [96] that every EM algorithm is a particular case of the MM approach. Therefore, we propose to use the MM principle described by (4.32) for minimizing J_{PLNQ} .

To construct the auxiliary cost at the current estimate $\{c_t\}$, we use the

convexity of $\log(\frac{1}{\bullet})$ [96, Equation (10)]) to write

$$-\sum_{\mathbf{k} \in \mathbb{Z}^d} g[\mathbf{k}] \log((c * b)[\mathbf{k}]) \leq -\sum_{\mathbf{k}, \mathbf{m} \in \mathbb{Z}^d} g[\mathbf{k}] \frac{b[\mathbf{k} - \mathbf{m}] c_t[\mathbf{m}]}{(c_t * b)[\mathbf{k}]} \log \left((c_t * b)[\mathbf{k}] \frac{c[\mathbf{m}]}{c_t[\mathbf{m}]} \right).$$

Similarly, Equation (4.34) and the convexity of $(\bullet)^2$ [96, Equation (9)]) result in

$$\begin{aligned} \lambda \Psi_{\text{NQ}}\{c\} &\leq \lambda \frac{p}{2} \sum_{m=1}^{d_L} \sum_{\mathbf{k}, \mathbf{m} \in \mathbb{Z}^d} |\chi_t[\mathbf{k}]|^{p-2} \alpha[\mathbf{k} - \mathbf{m}] \\ &\quad \times \left(\frac{\eta_m[\mathbf{k} - \mathbf{m}]}{\alpha[\mathbf{k} - \mathbf{m}]} (c[\mathbf{m}] - c_t[\mathbf{m}]) + (c_t * \eta_m)[\mathbf{k}] \right)^2 \\ &\quad + \mathcal{A}_{\chi_t}, \end{aligned}$$

where $\{\alpha[\mathbf{k}]\}_{\mathbf{k} \in \mathbb{Z}^d}$ are constants such that

$$\begin{aligned} \alpha[\mathbf{k}] &= 0, \text{ if } \eta_m[\mathbf{k}] = 0, \\ \alpha[\mathbf{k}] &> 0, \text{ if } \eta_m[\mathbf{k}] \neq 0, \end{aligned}$$

and

$$\sum_{\mathbf{k} \in \mathbb{Z}^d} \alpha[\mathbf{k}] = 1.$$

After some algebraic manipulations, we obtain

$$J_{\text{PLNQ}}\{g, c\} \leq \sum_{\mathbf{m} \in \mathbb{Z}^d} J_{\text{AUX}}\{g, c|c_t, \mathbf{m}\}, \quad (\text{C.2})$$

where $J_{\text{AUX}}\{g, c|c_t, \mathbf{m}\}$ is the decoupled auxiliary cost that depends only on $c[\mathbf{m}]$:

$$\begin{aligned} J_{\text{AUX}}\{g, c|c_t, \mathbf{m}\} &= -\log(c[\mathbf{m}]) (c_t[\mathbf{m}] \mathcal{A}_{1\mathbf{m}}) \\ &\quad + p\lambda c[\mathbf{m}] \left(\frac{c[\mathbf{m}] \mathcal{A}_{2\mathbf{m}}}{2} - c_t[\mathbf{m}] \mathcal{A}_{2\mathbf{m}} + \mathcal{A}_{3\mathbf{m}} + \frac{\mathcal{A}_b}{p\lambda} \right) \\ &\quad + \mathcal{A}_{4\mathbf{m}}. \end{aligned}$$

The quantity $\mathcal{A}_b = \sum_{\mathbf{k} \in \mathbb{Z}^d} b[\mathbf{k}] > 0$ is a constant independent of $\{c_t[\mathbf{k}]\}$ while the

$\{\mathcal{A}_{\bullet\mathbf{m}}\}_{\mathbf{m}\in\mathbb{Z}^d}$ depend on $\{c_t[\mathbf{k}]\}$ and are given by

$$\mathcal{A}_{1\mathbf{m}} = (\bar{b} * w_t)[\mathbf{m}], \quad (\text{C.3})$$

$$\mathcal{A}_{2\mathbf{m}} = \sum_{m=1}^{d_{\mathbf{L}}} (\bar{\mu}_m * |\chi_t|^{p-2})[\mathbf{m}], \quad (\text{C.4})$$

$$\mathcal{A}_{3\mathbf{m}} = \sum_{m=1}^{d_{\mathbf{L}}} \sum_{\mathbf{k}\in\mathbb{Z}^d} |\chi_t|^{p-2}[\mathbf{k}] (c_t * \eta_m)[\mathbf{k}] \eta_m[\mathbf{m} - \mathbf{k}], \quad (\text{C.5})$$

$$\begin{aligned} \mathcal{A}_{4\mathbf{m}} &= -c_t[\mathbf{m}] \sum_{\mathbf{k}\in\mathbb{Z}^d} b[\mathbf{k} - \mathbf{m}] w_t[\mathbf{k}] \log \left(\frac{(c_t * b)[\mathbf{k}]}{c_t[\mathbf{m}]} \right) + \mathcal{A}_{\chi_t} \\ &\quad + c_t[\mathbf{m}] (c_t[\mathbf{m}] \mathcal{A}_{2\mathbf{m}} - \mathcal{A}_{3\mathbf{m}}) + \mathcal{A}_{3\mathbf{m}}, \end{aligned} \quad (\text{C.6})$$

where $w_t[\mathbf{k}] = \frac{g[\mathbf{k}]}{(c_t * b)[\mathbf{k}]}$ and $\mu_m[\mathbf{k}] = \frac{\eta_m^2[\mathbf{k}]}{\alpha[\mathbf{k}]}$.

Then, minimizing $J_{\text{PLNQ}}\{g, c\}$ simply amounts to minimizing each of the individual $J_{\text{AUX}}\{g, c|c_t, \mathbf{m}\}$ which is accomplished by setting their derivative with respect to $c[\mathbf{m}]$ to zero. The solution is given by

$$c[\mathbf{m}] = \frac{1}{2\lambda p \mathcal{A}_{2\mathbf{m}}} \left(-\gamma + \sqrt{\gamma^2 + 4\lambda p c_t[\mathbf{m}] \mathcal{A}_{1\mathbf{m}} \mathcal{A}_{2\mathbf{m}}} \right), \quad (\text{C.7})$$

where $\gamma = \lambda p (\mathcal{A}_{3\mathbf{m}} - c_t[\mathbf{m}] \mathcal{A}_{2\mathbf{m}}) + \mathcal{A}_b$. Since, the sequences $\{b[\mathbf{k}]\}$ and $\{\chi[\mathbf{k}]\}$ are non-negative and because $\{g[\mathbf{k}]\}$ represents the data counts under the Poisson noise model, we have that $\mathcal{A}_{1\mathbf{m}} > 0$ and $\mathcal{A}_{2\mathbf{m}} > 0$ whenever $c_t[\mathbf{m}]|_{\mathbf{m}\in\mathbb{Z}^d} > 0$. Thus, the solution given by (C.7) is always positive. Based on (C.2)-(C.7), we present Algorithm C.1 for the minimization of J_{PLNQ} .

In the case of quadratic regularization, we have that

$$\Psi_Q\{c\} = \langle c, q * c \rangle_{\ell_2} = \sum_{m=1}^{d_{\mathbf{L}}} \left\| c * q_{m\frac{1}{2}} \right\|_{\ell_2}^2, \quad (\text{C.8})$$

where $q_{m\frac{1}{2}}$ is specified via its Fourier transform $Q_{m\frac{1}{2}}(e^{j\boldsymbol{\omega}}) = \sqrt{Q_m(e^{j\boldsymbol{\omega}})}$, wherein $Q_m(e^{j\boldsymbol{\omega}})$ is the Fourier transform of q_m . Comparing (4.26), (C.1) and (C.8), we see that Algorithm C.1 can be applied to this case too by setting $p = 2$ and replacing $\{\eta_m[\mathbf{k}]\}$ by $\{q_{m\frac{1}{2}}[\mathbf{k}]\}$ in (C.4) and (C.5).

Algorithm C.1. *MM algorithm for minimizing J_{PLNQ} for a given λ and p*

(1) *Precompute \mathcal{A}_b and $\{\mu_m[\mathbf{k}]\}$*

(2) *Initial estimate = $\{c_0[\mathbf{k}]\}_{\mathbf{k} \in \mathbb{Z}^d}$; $t = 0$*

Repeat Steps 3 to 5 until Stop Criterion is met

(3) *Update constants $\{\mathcal{A}_{1\mathbf{m}}\}_{\mathbf{m} \in \mathbb{Z}^d}$, $\{\mathcal{A}_{2\mathbf{m}}\}_{\mathbf{m} \in \mathbb{Z}^d}$, $\{\mathcal{A}_{3\mathbf{m}}\}_{\mathbf{m} \in \mathbb{Z}^d}$ using (C.3)-(C.5)*

(4) *Compute $\{c_{t+1}[\mathbf{k}]\}_{\mathbf{k} \in \mathbb{Z}^d}$ using (C.7)*

(5) *Set $t = t + 1$*

Appendix D

Risk Estimation

D.1 Proof of Unbiasedness of SURE

Here, we present a proof of Theorem 5.1 in Chapter 5. We start from the l.h.s. of (5.6) and expand it to get

$$\begin{aligned} E_{\mathbf{b}} \left\{ \frac{1}{N} \|\mathbf{x} - \mathbf{f}_{\lambda}(\mathbf{y})\|^2 \right\} &= E_{\mathbf{b}} \left\{ \frac{1}{N} \|\mathbf{y} - \mathbf{b} - \mathbf{f}_{\lambda}(\mathbf{y})\|^2 \right\} \\ &= E_{\mathbf{b}} \left\{ \frac{1}{N} \|\mathbf{y} - \mathbf{f}_{\lambda}(\mathbf{y})\|^2 \right\} - 2 E_{\mathbf{b}} \left\{ \frac{1}{N} \mathbf{b}^T (\mathbf{y} - \mathbf{f}_{\lambda}(\mathbf{y})) \right\} \\ &\quad + E_{\mathbf{b}} \left\{ \frac{1}{N} \|\mathbf{b}\|^2 \right\} \\ &= E_{\mathbf{b}} \left\{ \frac{1}{N} \|\mathbf{y} - \mathbf{f}_{\lambda}(\mathbf{y})\|^2 \right\} + 2 E_{\mathbf{b}} \left\{ \frac{1}{N} \mathbf{b}^T \mathbf{f}_{\lambda}(\mathbf{y}) \right\} - \sigma^2. \end{aligned}$$

We now write the second term on the r.h.s. of the above equation explicitly as

$$E_{\mathbf{b}} \{ \mathbf{b}^T \mathbf{f}_{\lambda}(\mathbf{y}) \} = \int_{\mathbb{R}^N} \mathbf{b}^T \mathbf{f}_{\lambda}(\mathbf{y}) \varphi_{\mathbf{G}}(\mathbf{b}) \, d\mathbf{b},$$

where

$$\varphi_{\mathbf{G}}(\mathbf{b}) = \left(\frac{1}{2\pi\sigma^2} \right)^{\frac{N}{2}} \exp \left(-\frac{1}{2\sigma^2} \|\mathbf{b}\|^2 \right)$$

is the probability density function of the Gaussian random variable \mathbf{b} . An interesting property¹ of $\varphi_G(\mathbf{b})$ is that

$$\mathbf{b} \varphi_G(\mathbf{b}) = -\sigma^2 \nabla_{\mathbf{b}} \varphi_G(\mathbf{b}), \quad (\text{D.1})$$

where $\nabla_{\mathbf{b}} \varphi_G(\mathbf{b})$ is the gradient of $\varphi_G(\mathbf{b})$ with respect to \mathbf{b} . Therefore, we have

$$E_{\mathbf{b}} \{ \mathbf{b}^T \mathbf{f}_{\lambda}(\mathbf{y}) \} = -\sigma^2 \int_{\mathbb{R}^N} (\nabla_{\mathbf{b}} \varphi_G(\mathbf{b}))^T \mathbf{f}_{\lambda}(\mathbf{y}) \, d\mathbf{b}.$$

Performing integration by parts for each component of the gradient, we get

$$-\sigma^2 \int_{\mathbb{R}^N} (\nabla_{\mathbf{b}} \varphi_G(\mathbf{b}))^T \mathbf{f}_{\lambda}(\mathbf{y}) \, d\mathbf{b} = \sigma^2 \int_{\mathbb{R}^N} \varphi_G(\mathbf{b}) \operatorname{div}_{\mathbf{b}} \{ \mathbf{f}_{\lambda}(\mathbf{y}) \} \, d\mathbf{b},$$

where we have made use of the fact that $\lim_{|b_k| \rightarrow \infty} \varphi_G(\mathbf{b}) \mathbf{f}_{\lambda}(\mathbf{y}) = 0$ whenever

$$E_{\mathbf{b}} \left\{ \left| \frac{\partial \mathbf{f}_{\lambda k}(\mathbf{y})}{\partial y_k} \right| \right\} < \infty, \quad k = 1, 2, \dots, N \quad [131, 139].$$

Keeping in mind that $\operatorname{div}_{\mathbf{b}} \{ \bullet \} = \operatorname{div}_{\mathbf{y}} \{ \bullet \}$ for the data model in (5.1), we have

$$\sigma^2 \int_{\mathbb{R}^N} \varphi_G(\mathbf{b}) \operatorname{div}_{\mathbf{b}} \{ \mathbf{f}_{\lambda}(\mathbf{y}) \} \, d\mathbf{b} = \sigma^2 E_{\mathbf{b}} \{ \operatorname{div}_{\mathbf{b}} \{ \mathbf{f}_{\lambda}(\mathbf{y}) \} \} = \sigma^2 E_{\mathbf{b}} \{ \operatorname{div}_{\mathbf{y}} \{ \mathbf{f}_{\lambda}(\mathbf{y}) \} \}.$$

Thus, we obtain

$$E_{\mathbf{b}} \left\{ \frac{1}{N} \mathbf{b}^T \mathbf{f}_{\lambda}(\mathbf{y}) \right\} = \frac{\sigma^2}{N} E_{\mathbf{b}} \{ \operatorname{div}_{\mathbf{y}} \{ \mathbf{f}_{\lambda}(\mathbf{y}) \} \},$$

which completes the result.

D.2 Proof of Unbiasedness of the New URE

Here, we present a proof of Theorem 5.2. Again, we start from the MSE and use the fact that $\mathbf{f}_{\lambda}(\mathbf{y}) = \mathbf{H}^T \bar{\mathbf{f}}_{\lambda}(\mathbf{y})$ to get

$$E_{\mathbf{y}} \left\{ \frac{1}{N} \|\mathbf{x} - \mathbf{f}_{\lambda}(\mathbf{y})\|^2 \right\} = \frac{1}{N} \|\mathbf{x}\|^2 + E_{\mathbf{y}} \left\{ \frac{1}{N} \|\mathbf{f}_{\lambda}(\mathbf{y})\|^2 \right\} - \frac{2}{N} E_{\mathbf{y}} \{ \mathbf{x}^T \mathbf{H}^T \bar{\mathbf{f}}_{\lambda}(\mathbf{y}) \}.$$

¹This is exactly what we meant by exploiting the statistics of noise for risk estimation at the beginning of Chapter 5. This property is the multidimensional version of the well-known result in 1-D: For $\varphi_G(b) = (\sqrt{2\pi\sigma^2})^{-1} \exp(-(2\sigma^2)^{-1} b^2)$, we have $b \varphi_G(b) = -\sigma^2 \varphi'_G(b)$.

We concentrate on the last term in the r.h.s. of the above equation:

$$\begin{aligned} E_{\mathbf{y}} \{ \mathbf{x}^T \mathbf{H}^T \bar{\mathbf{f}}_{\lambda}(\mathbf{y}) \} &= \int_{\mathbb{R}^N} \mathbf{x}^T \mathbf{H}^T \bar{\mathbf{f}}_{\lambda}(\mathbf{y}) \varphi_{\text{PG}}(\mathbf{y}) \, d\mathbf{y} \\ &= \sum_{l=1}^N \int_{\mathbb{R}^N} (\mathbf{H}\mathbf{x})_l \bar{f}_{\lambda l}(\mathbf{y}) \varphi_{\text{PG}}(\mathbf{y}) \, d\mathbf{y}, \end{aligned} \quad (\text{D.2})$$

where $\varphi_{\text{PG}}(\mathbf{y})$ is the (Poisson-Gaussian mixture²) probability density function of \mathbf{y} given by

$$\varphi_{\text{PG}}(\mathbf{y}) = (\varphi_{\text{P}} \star \varphi_{\text{G}})(\mathbf{y}),$$

where φ_{P} and φ_{G} are the probability densities corresponding to pure-Poisson and pure-Gaussian components (in accordance with the data model in Equation (5.14)), respectively:

$$\begin{aligned} \varphi_{\text{P}}(\bullet) &= \sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{\mathbf{0}\}} \underbrace{\exp(-\alpha(\mathbf{H}\mathbf{x})) \frac{(\alpha\mathbf{H}\mathbf{x})^{\mathbf{k}}}{\mathbf{k}!}}_{b_{\text{P}}[\mathbf{k}]}} \delta(\bullet - \alpha\mathbf{k}), \\ \varphi_{\text{G}}(\bullet) &= \left(\frac{1}{2\pi\sigma^2} \right)^{\frac{N}{2}} \exp\left(-\frac{1}{2\sigma^2} \|\bullet - \mu\mathbf{1}\|^2 \right), \end{aligned}$$

where $\mathbb{Z}_+^d = \{\mathbf{k} \in \mathbb{Z}^d : k_i \geq 1, i = 1, 2, \dots, d\}$, $\delta(\bullet)$ is the Dirac-distribution, $\mathbf{1}$ is a column vector of 1's and

$$\begin{aligned} \exp(-\alpha(\mathbf{H}\mathbf{x})) &= \prod_{i=1}^N \exp(-\alpha(\mathbf{H}\mathbf{x})_i), \\ (\alpha\mathbf{H}\mathbf{x})^{\mathbf{k}} &= \prod_{i=1}^N (\alpha)^{k_i} (\mathbf{H}\mathbf{x})_i^{k_i}, \\ \mathbf{k}! &= \prod_{i=1}^N k_i!, \end{aligned}$$

²In order to get $\varphi_{\text{PG}}(\mathbf{y})$ we use the property that the probability density function of a sum of two independent random variables is equal to the convolution of the individual probability densities of the two random variables, respectively.

where $(\mathbf{H}\mathbf{x})_i$ represents the i -th component of the vector $\mathbf{H}\mathbf{x}$. Substituting for $\varphi_{\text{PG}}(\mathbf{y})$ in (D.2), and manipulating the integrand, we have

$$\begin{aligned}
 & (\mathbf{H}\mathbf{x})_l \bar{f}_{\lambda l}(\mathbf{y}) \varphi_{\text{PG}}(\mathbf{y}) \\
 &= \frac{\exp(-(\alpha\mathbf{H}\mathbf{x}))}{\alpha} \bar{f}_{\lambda l}(\mathbf{y}) \\
 & \quad \times \left[\left(\sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{\mathbf{0}\}} \prod_{\substack{i=1 \\ i \neq l}}^N \frac{(\alpha\mathbf{H}\mathbf{x})_i^{k_i}}{k_i!} \frac{(\alpha\mathbf{H}\mathbf{x})_l^{k_l+1}}{k_l!} \delta(\bullet - \alpha\mathbf{k}) \right) \star \varphi_{\text{G}} \right] (\mathbf{y}) \\
 &= \frac{1}{\alpha} \bar{f}_{\lambda l}(\mathbf{y}) \left[\left(\sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{\mathbf{0}\}} k_l \mathbf{b}_{\text{P}}[\mathbf{k}] \delta(\bullet - \alpha\mathbf{k} + \alpha\mathbf{e}_l) \right) \star \varphi_{\text{G}} \right] (\mathbf{y}) \\
 &= \frac{1}{\alpha} \bar{f}_{\lambda l}(\mathbf{y}) \left[\left(\sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{\mathbf{0}\}} k_l \mathbf{b}_{\text{P}}[\mathbf{k}] \delta(\bullet - \alpha\mathbf{k}) \right) \star \varphi_{\text{G}} \right] (\mathbf{y} + \alpha\mathbf{e}_l),
 \end{aligned}$$

where \mathbf{e}_l is a column vector whose components are zero except for the l -th one which is unity. Therefore, (D.2) becomes

$$\begin{aligned}
 & E_{\mathbf{y}} \{ \mathbf{x}^T \mathbf{H}^T \bar{f}_{\lambda}(\mathbf{y}) \} \\
 &= \frac{1}{\alpha} \sum_{l=1}^N \int_{\mathbb{R}^N} \bar{f}_{\lambda l}(\mathbf{y}) \left[\left(\sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{\mathbf{0}\}} k_l \mathbf{b}_{\text{P}}[\mathbf{k}] \delta(\bullet - \alpha\mathbf{k}) \right) \star \varphi_{\text{G}} \right] (\mathbf{y} + \alpha\mathbf{e}_l) \, \text{d}\mathbf{y} \\
 &= \frac{1}{\alpha} \sum_{l=1}^N \int_{\mathbb{R}^N} \bar{g}_{\lambda l}(\mathbf{y}) \left[\left(\sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{\mathbf{0}\}} k_l \mathbf{b}_{\text{P}}[\mathbf{k}] \delta(\bullet - \alpha\mathbf{k}) \right) \star \varphi_{\text{G}} \right] (\mathbf{y}) \, \text{d}\mathbf{y},
 \end{aligned}$$

where $\bar{g}_{\lambda,l}(\mathbf{y}) = \bar{f}_{\lambda,l}(\mathbf{y} - \alpha\mathbf{e}_l)$ corresponds to the Poisson contribution to the risk estimate since we have primarily manipulated the Poisson part of the density φ_{PG} .

Adding and subtracting $(y_l - \mu)$ inside the innermost brackets on the r.h.s.

of the above equation, we get

$$\begin{aligned}
 & E_{\mathbf{y}} \{ \mathbf{x}^T \mathbf{H}^T \bar{\mathbf{f}}_{\lambda}(\mathbf{y}) \} \\
 &= -\frac{1}{\alpha} \sum_{l=1}^N \int_{\mathbb{R}^N} \bar{g}_{\lambda l}(\mathbf{y}) \left[\left(\sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{\mathbf{0}\}} (y_l - k_l - \mu) \mathbf{b}_P[\mathbf{k}] \delta(\bullet - \alpha \mathbf{k}) \right) \star \varphi_G \right] (\mathbf{y}) \, d\mathbf{y} \\
 &\quad + \frac{1}{\alpha} \sum_{l=1}^N \int_{\mathbb{R}^N} \bar{g}_{\lambda l}(\mathbf{y}) (y_l - \mu) \underbrace{\left[\left(\sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{\mathbf{0}\}} \mathbf{b}_P[\mathbf{k}] \delta(\bullet - \alpha \mathbf{k}) \right) \star \varphi_G \right]}_{\varphi_P} (\mathbf{y}) \, d\mathbf{y} \\
 &= +\frac{\sigma^2}{\alpha} \sum_{l=1}^N \int_{\mathbb{R}^N} \bar{g}_{\lambda l}(\mathbf{y}) \frac{\partial(\varphi_P \star \varphi_G)(\mathbf{y})}{\partial y_l} \, d\mathbf{y} + \frac{1}{\alpha} E_{\mathbf{y}} \{ (\mathbf{y} - \mu \mathbf{1})^T \bar{\mathbf{g}}_{\lambda}(\mathbf{y}) \}, \quad (\text{D.3})
 \end{aligned}$$

where we have made use of the following equality that is based on differentiability property of φ_G :

$$\begin{aligned}
 & \left[\left(\sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{\mathbf{0}\}} (y_l - k_l - \mu) \mathbf{b}_P[\mathbf{k}] \delta(\bullet - \alpha \mathbf{k}) \right) \star \varphi_G \right] (\mathbf{y}) \\
 &= \sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{\mathbf{0}\}} (y_l - k_l - \mu) \mathbf{b}_P[\mathbf{k}] \varphi_G(\mathbf{y} - \alpha \mathbf{k}) \\
 &= -\sigma^2 \sum_{\mathbf{k} \in \mathbb{Z}_+^d \cup \{\mathbf{0}\}} \mathbf{b}_P[\mathbf{k}] \frac{\partial \varphi_G(\mathbf{y} - \alpha \mathbf{k})}{\partial y_l} \quad (\text{similar to (D.1)}) \\
 &= -\sigma^2 \frac{\partial(\varphi_P \star \varphi_G)(\mathbf{y})}{\partial y_l}.
 \end{aligned}$$

Thus, in (D.3), the first term on the r.h.s. is the contribution of the Gaussian part of φ_{PG} since we make use of the differentiability property of φ_G to obtain it.

Performing, integration by parts and using the hypothesis

$$E_{\mathbf{y}} \left\{ \left| \frac{\partial \bar{g}_{\lambda k}(\mathbf{y})}{\partial y_k} \right| \right\} = E_{\mathbf{y}} \left\{ \left| \frac{\partial \bar{f}_{\lambda k}(\mathbf{y})}{\partial y_k} \right| \right\} < \infty,$$

for $k = 1, 2, \dots, N$ [131, 139], the first term on the r.h.s. of (D.3) becomes

$$\begin{aligned} \frac{\sigma^2}{\alpha} \sum_{l=1}^N \int_{\mathbb{R}^N} \bar{\mathbf{g}}_{\boldsymbol{\lambda}l}(\mathbf{y}) \frac{\partial(\varrho_{\mathbf{P}} \star \varrho_{\mathbf{G}})(\mathbf{y})}{\partial y_l} d\mathbf{y} &= -\frac{\sigma^2}{\alpha} \sum_{l=1}^N \int_{\mathbb{R}^N} \frac{\partial \bar{\mathbf{g}}_{\boldsymbol{\lambda}l}(\mathbf{y})}{\partial y_l} (\varrho_{\mathbf{P}} \star \varrho_{\mathbf{G}})(\mathbf{y}) d\mathbf{y} \\ &= -\frac{\sigma^2}{\alpha} E_{\mathbf{y}}\{\operatorname{div}_{\mathbf{y}}\{\bar{\mathbf{g}}_{\boldsymbol{\lambda}}(\mathbf{y})\}\}, \end{aligned}$$

so that

$$E_{\mathbf{y}}\{\mathbf{x}^T \mathbf{H}^T \bar{\mathbf{f}}_{\boldsymbol{\lambda}}(\mathbf{y})\} = -\frac{\sigma^2}{\alpha} E_{\mathbf{y}}\{\operatorname{div}_{\mathbf{y}}\{\bar{\mathbf{g}}_{\boldsymbol{\lambda}}(\mathbf{y})\}\} + \frac{1}{\alpha} E_{\mathbf{y}}\{(\mathbf{y} - \boldsymbol{\mu}\mathbf{1})^T \bar{\mathbf{g}}_{\boldsymbol{\lambda}}(\mathbf{y})\},$$

which completes the proof.

Appendix E

Differentiability and Monte-Carlo Divergence Estimation

Here, we elaborate on the solution to the differentiability issue associated with the Monte-Carlo divergence estimation proposed (in Theorem 6.1) in Chapter 6. Firstly, we verify the validity of the Taylor expansion-based argumentation of Theorem 6.1 for algorithms like total-variation denoising (TVD). Following that, we give a proof of the second part of Theorem 6.1 which deals with a weaker hypothesis (using tempered distributions) of the problem.

E.1 Verification of Taylor Expansion-Based Hypothesis

In Chapter 6, we considered the discrete-domain formulation for total-variation denoising (TVD) where we minimize the cost

$$\mathcal{J}_{\text{TV}}(\mathbf{u}) = \|\mathbf{y} - \mathbf{u}\|^2 + \lambda \text{TV}\{\mathbf{u}\}, \quad (\text{E.1})$$

where $\text{TV}\{\mathbf{u}\} = \sum_k \sqrt{(\mathbf{D}_h \mathbf{u})[k]^2 + (\mathbf{D}_v \mathbf{u})[k]^2}$ is the discrete 2D total variation norm and \mathbf{D}_h and \mathbf{D}_v are matrices corresponding to the first order finite difference in the horizontal and vertical directions, respectively. We will concentrate

only on the bounded-optimization (half-quadratic) algorithm developed in [158]. However, since a typical implementation of TVD will always involve finite differences in place of continuous domain derivatives the analysis can be easily extended to other algorithms including the variant in [158] and those based on Euler-Lagrange equations. We show that the bounded-optimization algorithm for TVD admits first and second order derivatives with respect to the data \mathbf{y} and therefore satisfies the stronger hypothesis (Taylor expansion-based) of Theorem 6.1.

The TVD algorithm is described by the following recursive equation [158]: The signal estimate at iteration $k + 1$ denoted by the $N \times 1$ vector \mathbf{f}_λ^{k+1} is obtained by solving the linear system

$$\mathbf{M}^k \mathbf{f}_\lambda^{k+1} = \mathbf{y}, \quad (\text{E.2})$$

where \mathbf{M}^k is the $N \times N$ system matrix at iteration k given by

$$\mathbf{M}^k = \mathbf{I} + \mathbf{D}_h^T \Lambda^k \mathbf{D}_h + \mathbf{D}_v^T \Lambda^k \mathbf{D}_v, \quad (\text{E.3})$$

and $\Lambda^k = \text{diag}\{w_i^k; i = 1, 2, \dots, N\}$, with

$$w_i^k = \frac{\lambda}{2} [(\mathbf{D}_h \mathbf{f}_\lambda^k)_i^2 + (\mathbf{D}_v \mathbf{f}_\lambda^k)_i^2 + \kappa]^{-\frac{1}{2}}, \quad (\text{E.4})$$

where $(\mathbf{D}_* \mathbf{f}_\lambda^k)_i$ is the i -th element of the vector $\mathbf{D}_* \mathbf{f}_\lambda^k$ and $\kappa > 0$ is a small constant that prevents the denominator of w_i^k going to zero (or else the algorithm \mathbf{f}_λ itself may become numerically unstable).

Differentiating (E.2) with respect to \mathbf{y} , we obtain

$$\frac{\partial \mathbf{M}^k}{\partial \mathbf{y}} \mathbf{f}_\lambda^{k+1} + \mathbf{M}^k \underbrace{\frac{\partial \mathbf{f}_\lambda^{k+1}}{\partial \mathbf{y}}}_{\mathbf{J}_{\mathbf{f}_\lambda^{k+1}}} = \mathbf{I}, \quad (\text{E.5})$$

where $\mathbf{J}_{\mathbf{f}_\lambda^{k+1}}$ is the Jacobian matrix of \mathbf{f}_λ^{k+1} at iteration $k + 1$. If M_{mn}^k represents the mn -th element of \mathbf{M}^k and y_m and f_m^k represent the m -th elements of \mathbf{y} and \mathbf{f}_λ^k , respectively, then using Einstein's summation notation (repeated indices will be summed over unless they appear on both sides of an equation) (E.5) can be written as

$$\frac{\partial M_{mn}^k}{\partial y_p} f_n^{k+1} + M_{mn}^k \frac{\partial f_n^{k+1}}{\partial y_p} = \delta_{mp}, \quad (\text{E.6})$$

where, for example, the index n is summer over in both the terms on the l.h.s. of the above equation.

Differentiating (E.6) a second time, we obtain

$$\frac{\partial^2 M_{mn}^k}{\partial y_l \partial y_p} \mathbf{f}_n^{k+1} + \frac{\partial M_{mn}^k}{\partial y_p} \frac{\partial \mathbf{f}_n^{k+1}}{\partial y_l} + \frac{\partial M_{mn}^k}{\partial y_l} \frac{\partial \mathbf{f}_n^{k+1}}{\partial y_p} + M_{mn}^k \frac{\partial \mathbf{f}_n^{k+1}}{\partial y_l \partial y_p} = 0. \quad (\text{E.7})$$

It is clear that the $N \times N \times N$ tensor $\mathbf{r} = \left\{ \frac{\partial \mathbf{f}_n^{k+1}}{\partial y_l \partial y_p} \right\}_{n,l,p=1}^N$ is the desired second derivative in the Taylor expansion in Theorem 6.1. We will show that for a given (l, p) , all the terms in (E.7) are well-defined, so that the $N \times 1$ vector $\mathbf{r}_{lp} = \frac{\partial \mathbf{f}_n^{k+1}}{\partial y_l \partial y_p}$ can be obtained by solving a linear system.

Firstly, we analyze $\frac{\partial M_{mn}^k}{\partial y_p}$ which is given by

$$\frac{\partial M_{mn}^k}{\partial y_p} = (D_{h_{qm}} D_{h_{qn}} + D_{v_{qm}} D_{v_{qn}}) \frac{\partial w_q^k}{\partial y_p}, \quad (\text{E.8})$$

with

$$\frac{\partial w_q^k}{\partial y_p} = - \left(\frac{2}{\lambda} \right)^2 (w_q^k)^3 \left((\mathbf{D}_h \mathbf{f}_\lambda^k)_q D_{h_{qi}} + (\mathbf{D}_v \mathbf{f}_\lambda^k)_q D_{v_{qi}} \right) \frac{\partial \mathbf{f}_i^k}{\partial y_p}, \quad (\text{E.9})$$

where the index q is not summed over on the r.h.s. of the above equation.

Similarly,

$$\frac{\partial^2 M_{mn}^k}{\partial y_l \partial y_p} = (D_{h_{qm}} D_{h_{qn}} + D_{v_{qm}} D_{v_{qn}}) \frac{\partial^2 w_q^k}{\partial y_l \partial y_p}, \quad (\text{E.10})$$

where

$$\begin{aligned} \frac{\partial^2 w_q^k}{\partial y_l \partial y_p} &= - \left(\frac{2}{\lambda} \right)^2 (w_q^k)^3 \\ &\times \left[\begin{aligned} &\frac{1}{3} \left(\frac{\lambda}{2} \right)^2 (w_q^k)^{-4} \frac{\partial w_q^k}{\partial y_l} \frac{\partial w_q^k}{\partial y_p} \\ &- \left((\mathbf{D}_h \mathbf{f}_\lambda^k)_q D_{h_{qi}} D_{h_{qj}} + (\mathbf{D}_v \mathbf{f}_\lambda^k)_q D_{v_{qi}} D_{v_{qj}} \right) \frac{\partial \mathbf{f}_i^k}{\partial y_l} \frac{\partial \mathbf{f}_j^k}{\partial y_p} \\ &- \left((\mathbf{D}_h \mathbf{f}_\lambda^k)_q D_{h_{qi}} + (\mathbf{D}_v \mathbf{f}_\lambda^k)_q D_{v_{qi}} \right) \frac{\partial^2 \mathbf{f}_i^k}{\partial y_l \partial y_p} \end{aligned} \right]. \end{aligned} \quad (\text{E.11})$$

The analysis is then simply as follows:

1. In principle, if we start with a well-defined initial estimate \mathbf{f}_λ^0 , the algorithm described by equations (E.2)-(E.4) is designed so that \mathbf{M}^k and \mathbf{f}_λ^{k+1} are well-defined for all $k \geq 1$. Moreover, \mathbf{M}^k is a full-rank matrix and therefore has a stable inverse $(\mathbf{M}^k)^{-1}$.
2. It should be noted all the elements of \mathbf{M}^k are differentiable because (E.4) is a true function of \mathbf{y} . Thus, all the derivatives involved in this analysis are in the true sense of differentiation and not in the weak sense of distributions.
3. Then, (E.8) and (E.9) ensure that $\frac{\partial M_{mn}^k}{\partial y_p}$ is well-defined $\forall m, n, p = 1, 2, \dots, N$, provided Λ^k is well-conditioned which is the case as long as $w_q^k < +\infty$, $\forall q = 1, 2, \dots, N$, and $k \geq 1$. This can be ensured numerically. Therefore, for a fixed p , a well-defined $N \times 1$ vector $\frac{\partial \mathbf{f}_\lambda^{k+1}}{\partial y_p}$ is obtained from (E.6) as

$$\frac{\partial \mathbf{f}_\lambda^{k+1}}{\partial y_p} = (\mathbf{M}^k)^{-1} (\mathbf{e}_p - \mathbf{S}_p^k \mathbf{f}_\lambda^{k+1}), \quad (\text{E.12})$$

where \mathbf{S}_p^k is a $N \times N$ matrix such that $(\mathbf{S}_p^k)_{mn} = \frac{\partial M_{mn}^k}{\partial y_p}$, \mathbf{e}_p is a $N \times 1$ column vector whose elements are all zeros except the p -th one which is unity.

4. While (E.10) and (E.11) ensure that $\frac{\partial^2 M_{mn}^k}{\partial y_l \partial y_p}$ is well-defined $\forall m, n, l, p = 1, 2, \dots, N$, equations (E.8) and (E.12) ensure that the second and third terms in the l.h.s. of (E.7) are well-defined. Thus, we see that for a given (l, p) a well-defined $N \times 1$ vector $\frac{\partial \mathbf{f}_\lambda^{k+1}}{\partial y_l \partial y_p}$ is obtained from (E.7) as

$$\frac{\partial \mathbf{f}_\lambda^{k+1}}{\partial y_l \partial y_p} = -(\mathbf{M}^k)^{-1} \left(\mathbf{P}_{lp}^k \mathbf{f}_\lambda^{k+1} + \mathbf{S}_p^k \frac{\partial \mathbf{f}_\lambda^{k+1}}{\partial y_p} + \mathbf{S}_l^k \frac{\partial \mathbf{f}_\lambda^{k+1}}{\partial y_p} \right), \quad (\text{E.13})$$

where \mathbf{P}_{lp}^k is a $N \times N$ matrix such that $(\mathbf{P}_{lp}^k)_{mn} = \frac{\partial^2 M_{mn}^k}{\partial y_l \partial y_p}$.

E.2 Divergence Estimation Under a Weaker Hypothesis

We restate the second part of Theorem 6.1 in Chapter 6 which deals with the Monte-Carlo divergence estimation under the weaker hypothesis of tempered distributions and then give a formal proof of this result.

Theorem E.1. *Let \mathbf{b}' be a zero-mean unit variance i.i.d random vector. Assume that $\exists n_0 > 1$ and $C_0 > 0$ such that*

$$\|\mathbf{f}_\lambda(\mathbf{y})\| \leq C_0(1 + \|\mathbf{y}\|^{n_0}), \tag{E.14}$$

that is, \mathbf{f}_λ is tempered. Then,

$$\operatorname{div}_{\mathbf{y}}\{\mathbf{f}_\lambda(\mathbf{y})\} = \lim_{\varepsilon \rightarrow 0} E_{\mathbf{b}'} \left\{ \mathbf{b}'^T \left(\frac{\mathbf{f}_\lambda(\mathbf{y} + \varepsilon \mathbf{b}') - \mathbf{f}_\lambda(\mathbf{y})}{\varepsilon} \right) \right\}$$

in the weak-sense of tempered distributions.

Proof. Let $\psi \in \mathcal{S}$ be a rapidly decaying (test) function that is infinitely differentiable. We have to show that

$$\langle \operatorname{div}\{\mathbf{f}_\lambda(\mathbf{y})\}, \psi(\mathbf{y}) \rangle = \lim_{\varepsilon \rightarrow 0} \left\langle E_{\mathbf{b}'} \left\{ \mathbf{b}'^T \left(\frac{\mathbf{f}_\lambda(\mathbf{y} + \varepsilon \mathbf{b}') - \mathbf{f}_\lambda(\mathbf{y})}{\varepsilon} \right) \right\}, \psi(\mathbf{y}) \right\rangle. \tag{E.15}$$

We note that the l.h.s. of (E.15) can be expressed as (from theory of distributions)

$$\langle \operatorname{div}\{\mathbf{f}_\lambda(\mathbf{y})\}, \psi(\mathbf{y}) \rangle = -\langle \mathbf{f}_\lambda(\mathbf{y}), \nabla \psi(\mathbf{y}) \rangle. \tag{E.16}$$

The r.h.s. of (E.15) involves the double integration

$$I_1(\varepsilon) = \int_{\mathbf{y}} d\mathbf{y} \psi(\mathbf{y}) \int_{\mathbf{b}'} \mathbf{b}'^T \left(\frac{\mathbf{f}_\lambda(\mathbf{y} + \varepsilon \mathbf{b}') - \mathbf{f}_\lambda(\mathbf{y})}{\varepsilon} \right) \wp(\mathbf{b}') d\mathbf{b}',$$

where $\wp(\mathbf{b}')$ represents the probability density function of \mathbf{b}' . The order of the integration can be swapped as soon as (Fubini's theorem)

$$I_2(\varepsilon) = \int_{\mathbf{y}} \int_{\mathbf{b}'} |\psi(\mathbf{y})| \left| \mathbf{b}'^T \left(\frac{\mathbf{f}_\lambda(\mathbf{y} + \varepsilon \mathbf{b}') - \mathbf{f}_\lambda(\mathbf{y})}{\varepsilon} \right) \right| \wp(\mathbf{b}') d\mathbf{b}' d\mathbf{y} < +\infty.$$

Using Triangle inequality, we bound $I_2(\varepsilon)$ as

$$I_2(\varepsilon) \leq \frac{J(\varepsilon) + J(0)}{\varepsilon},$$

where

$$J(\varepsilon) = \int_{\mathbf{y}} \int_{\mathbf{b}'} |\psi(\mathbf{y})| |\mathbf{b}'^T \mathbf{f}_\lambda(\mathbf{y} + \varepsilon \mathbf{b}')| \wp(\mathbf{b}') d\mathbf{b}' d\mathbf{y}.$$

Using Cauchy-Schwarz inequality and the fact that \mathbf{f}_λ is tempered (cf. Equation (E.14)), we get

$$J(\varepsilon) \leq \int_{\mathbf{y}} \int_{\mathbf{b}'} |\psi(\mathbf{y})| \|\mathbf{b}'\| C_0 (1 + \|\mathbf{y} + \varepsilon \mathbf{b}'\|^{n_0}) \varphi(\mathbf{b}') \, d\mathbf{b}' \, d\mathbf{y}.$$

Using the convexity property of the function $(\bullet)^{n_0}$ for $n_0 > 1$, we get

$$\begin{aligned} J(\varepsilon) &\leq \int_{\mathbf{y}} \int_{\mathbf{b}'} |\psi(\mathbf{y})| \|\mathbf{b}'\| C_0 (1 + 2^{n_0-1} \|\mathbf{y}\|^{n_0} + 2^{n_0-1} \|\varepsilon \mathbf{b}'\|^{n_0}) \varphi(\mathbf{b}') \, d\mathbf{b}' \, d\mathbf{y} \\ &= C_0 \left(\int_{\mathbf{y}} |\psi(\mathbf{y})| \, d\mathbf{y} \right) \left(\int_{\mathbf{b}'} \|\mathbf{b}'\| \varphi(\mathbf{b}') \, d\mathbf{b}' \right) \\ &\quad + C_0 2^{n_0-1} \left(\int_{\mathbf{y}} |\psi(\mathbf{y})| \|\mathbf{y}\|^{n_0} \, d\mathbf{y} \right) \left(\int_{\mathbf{b}'} \|\mathbf{b}'\| \varphi(\mathbf{b}') \, d\mathbf{b}' \right) \\ &\quad + C_0 2^{n_0-1} \varepsilon^{n_0} \left(\int_{\mathbf{y}} |\psi(\mathbf{y})| \, d\mathbf{y} \right) \left(\int_{\mathbf{b}'} \|\mathbf{b}'\|^{n_0+1} \varphi(\mathbf{b}') \, d\mathbf{b}' \right) \\ &< +\infty, \end{aligned}$$

under the hypothesis that $E_{\mathbf{b}'}\{\|\mathbf{b}'\|^{n_0}\} < +\infty, \forall n_0 \geq 1$. The ones involving ψ are also finite because ψ is a rapidly decaying function with finite support. Thus, $J(\varepsilon) < \infty, \forall \varepsilon \geq 0$. Hence, we interchange the integrals (with appropriate change of variables) to get

$$\begin{aligned} I_1(\varepsilon) &= \int_{\mathbf{b}'} \varphi(\mathbf{b}') \, d\mathbf{b}' \int_{\mathbf{y}} \psi(\mathbf{y}) \mathbf{b}'^T \left(\frac{\mathbf{f}_\lambda(\mathbf{y} + \varepsilon \mathbf{b}') - \mathbf{f}_\lambda(\mathbf{y})}{\varepsilon} \right) \, d\mathbf{y} \\ &= \int_{\mathbf{b}'} \varphi(\mathbf{b}') \, d\mathbf{b}' \int_{\mathbf{y}} \mathbf{b}'^T \mathbf{f}_\lambda(\mathbf{y}) \left(\frac{\psi(\mathbf{y} - \varepsilon \mathbf{b}') - \psi(\mathbf{y})}{\varepsilon} \right) \, d\mathbf{y}. \end{aligned}$$

Since ψ is infinitely differentiable, we apply Taylor's theorem [166] to $\psi(\mathbf{y} - \varepsilon \mathbf{b}')$ and obtain

$$\frac{\psi(\mathbf{y} - \varepsilon \mathbf{b}') - \psi(\mathbf{y})}{\varepsilon} = - \int_0^1 \mathbf{b}'^T \nabla \psi(\mathbf{y} - t \varepsilon \mathbf{b}') \, dt.$$

Therefore,

$$I_1(\varepsilon) = - \int_{\mathbf{b}'} \varphi(\mathbf{b}') \, d\mathbf{b}' \int_{\mathbf{y}} \mathbf{b}'^T \mathbf{f}_\lambda(\mathbf{y}) \int_0^1 \mathbf{b}'^T \nabla \psi(\mathbf{y} - t \varepsilon \mathbf{b}') \, dt \, d\mathbf{y}.$$

We want to let ε tend to 0 in the above expression. This is accomplished by the application of Lebesgue's dominated convergence theorem. But firstly, we must bound the integrand

$$z(\mathbf{y}, \mathbf{b}', t, \varepsilon) = -\wp(\mathbf{b}') \mathbf{b}'^T \mathbf{f}_\lambda(\mathbf{y}) \mathbf{b}'^T \nabla \psi(\mathbf{y} - t\varepsilon \mathbf{b}') \beta^0(t),$$

by an integrable function $g(\mathbf{y}, \mathbf{b}', t)$, where

$$\beta^0(t) = \begin{cases} 1, & \text{if } t \in (0, 1) \\ 0, & \text{otherwise} \end{cases}.$$

To do that, we start with $|z(\mathbf{y}, \mathbf{b}', t, \varepsilon)|$ and apply Cauchy-Schwarz inequality to obtain

$$0 \leq |z(\mathbf{y}, \mathbf{b}', t, \varepsilon)| \leq \underbrace{\wp(\mathbf{b}') \|\mathbf{b}'\|^2 \|\nabla \psi(\mathbf{y} - t\varepsilon \mathbf{b}')\| \beta^0(t) \|\mathbf{f}_\lambda(\mathbf{y})\|}_{\stackrel{\text{def}}{=} g_0(\mathbf{y}, \mathbf{b}', t, \varepsilon)}.$$

Now, by using convexity of $(\bullet)^{n_1}$ for $n_1 \geq 1$, we have

$$\begin{aligned} 1 + \|\mathbf{y}\|^{n_1} &= 1 + \|\mathbf{y} - \varepsilon t \mathbf{b}' + \varepsilon t \mathbf{b}'\|^{n_1} \\ &\leq 1 + 2^{n_1-1} \|\mathbf{y} - \varepsilon t \mathbf{b}'\|^{n_1} + 2^{n_1-1} \|\varepsilon t \mathbf{b}'\|^{n_1}. \end{aligned}$$

Then, for $\varepsilon \leq 1$,

$$\begin{aligned} (1 + \|\mathbf{y}\|^{n_1}) g_0(\mathbf{y}, \mathbf{b}', t, \varepsilon) &\leq (1 + 2^{n_1-1} \|\mathbf{y} - \varepsilon t \mathbf{b}'\|^{n_1}) \\ &\quad \times \|\nabla \psi(\mathbf{y} - t\varepsilon \mathbf{b}')\| \wp(\mathbf{b}') \|\mathbf{b}'\|^2 \beta^0(t) \\ &\quad + 2^{n_1-1} \|\nabla \psi(\mathbf{y} - t\varepsilon \mathbf{b}')\| \wp(\mathbf{b}') \|\mathbf{b}'\|^{n_1+2} t^{n_1} \beta^0(t) \\ &\leq \underbrace{C_{\psi,1} \wp(\mathbf{b}') \|\mathbf{b}'\|^2 \beta^0(t) + C_{\psi,2} \wp(\mathbf{b}') \|\mathbf{b}'\|^{n_1+2} t^{n_1} \beta^0(t)}_{\stackrel{\text{def}}{=} g_1(\mathbf{b}', t)}, \end{aligned}$$

where $C_{\psi,1} = \sup_{\mathbf{y}} \{(1 + 2^{n_1-1} \|\mathbf{y}\|^{n_1}) \|\nabla \psi(\mathbf{y})\|\}$ and $C_{\psi,2} = 2^{n_1-1} \sup_{\mathbf{y}} \{\|\nabla \psi(\mathbf{y})\|\}$. Since $E_{\mathbf{b}'}\{\|\mathbf{b}'\|^{n_1+2}\} < +\infty$, it is clear that

$$\int_{\mathbf{b}'} \int_t g_1(\mathbf{b}', t) d\mathbf{b}' dt < +\infty. \quad (\text{E.17})$$

Therefore, choosing $n_1 > n_0 + N$, where N is the dimension of \mathbf{y} , and $\forall \varepsilon \leq 1$ we see that

$$\begin{aligned}
 |z(\mathbf{y}, \mathbf{b}', t, \varepsilon)| &\leq g_0(\mathbf{y}, \mathbf{b}', t, \varepsilon) \|\mathbf{f}_\lambda(\mathbf{y})\| \\
 &\leq g_1(\mathbf{b}', t) \frac{\|\mathbf{f}_\lambda(\mathbf{y})\|}{1 + \|\mathbf{y}\|^{n_1}} \\
 &\leq \underbrace{C_0 g_1(\mathbf{b}', t) \frac{1 + \|\mathbf{y}\|^{n_0}}{1 + \|\mathbf{y}\|^{n_1}}}_{\stackrel{\text{def}}{=} g(\mathbf{y}, \mathbf{b}', t)}. \tag{E.18}
 \end{aligned}$$

Then, we notice that

$$\begin{aligned}
 \int_{\mathbf{y}} \frac{1 + \|\mathbf{y}\|^{n_0}}{1 + \|\mathbf{y}\|^{n_1}} d\mathbf{y} &= \sum_{\mathbf{k} \in \mathbb{Z}^N} \int_{[0, 1)^N} \frac{1 + \|\mathbf{y} + \mathbf{k}\|^{n_0}}{1 + \|\mathbf{y} + \mathbf{k}\|^{n_1}} d\mathbf{y} \\
 &= \int_{[0, 1)^N} \left(\sum_{\mathbf{k} \in \mathbb{Z}^N} \frac{1 + \|\mathbf{y} + \mathbf{k}\|^{n_0}}{1 + \|\mathbf{y} + \mathbf{k}\|^{n_1}} \right) d\mathbf{y} \quad (\text{Fubini}) \\
 &\leq \sum_{\mathbf{k} \in \mathbb{Z}^N} \frac{1 + \|\mathbf{1} + \mathbf{k}\|^{n_0}}{1 + \|\mathbf{k}\|^{n_1}} \quad (\mathbf{1} \text{ is a column vector of 1s}) \\
 &\leq \sum_{\mathbf{k} \in \mathbb{Z}^N} \frac{1 + 2^{n_0-1} N^{\frac{n_0}{2}} + 2^{n_0-1} \|\mathbf{k}\|^{n_0}}{1 + \|\mathbf{k}\|^{n_1}} \\
 &\quad (\text{using convexity of } (\bullet)^{n_0}) \\
 &\leq (1 + 2^{n_0-1} N^{\frac{n_0}{2}}) \left(1 + \underbrace{\sum_{\mathbf{k} \in \mathbb{Z}^N \setminus \{0\}} \frac{1}{\|\mathbf{k}\|^{n_1}}}_{< +\infty} \right) \\
 &\quad + 2^{n_0-1} \underbrace{\sum_{\mathbf{k} \in \mathbb{Z}^N \setminus \{0\}} \frac{1}{\|\mathbf{k}\|^{n_1-n_0}}}_{< +\infty} \\
 &< +\infty, \tag{E.19}
 \end{aligned}$$

whenever $n_1 > n_0 + N$. Finally, because of (E.17) and (E.19), we find

$$\begin{aligned} \int_{\mathbf{y}} \int_{\mathbf{b}'} \int_t g(\mathbf{y}, \mathbf{b}', t) \, d\mathbf{y} \, d\mathbf{b}' \, dt &= \left(\int_{\mathbf{b}'} \int_t g(\mathbf{y}, \mathbf{b}', t) \, d\mathbf{b}' \, dt \right) \\ &\times \left(\int_{\mathbf{y}} \frac{1 + \|\mathbf{y}\|^{n_0}}{1 + \|\mathbf{y}\|^{n_1}} \, d\mathbf{y} \right) \\ &< +\infty. \end{aligned} \tag{E.20}$$

Therefore, z qualifies for both Fubini's and Lebesgue's Dominant Convergence theorems (*cf.* Equations (E.18) and (E.20)). Hence, applying the limit with appropriate change of integrals, we get the desired result:

$$\begin{aligned} \lim_{\varepsilon \rightarrow 0} I_1(\varepsilon) &= - \lim_{\varepsilon \rightarrow 0} \int_{\mathbf{b}'} \wp(\mathbf{b}') \, d\mathbf{b}' \int_{\mathbf{y}} \mathbf{b}'^T \mathbf{f}_{\lambda}(\mathbf{y}) \int_0^1 \mathbf{b}'^T \nabla \psi(\mathbf{y} - t\varepsilon \mathbf{b}') \, dt \, d\mathbf{y} \\ &= - \int_{\mathbf{b}'} \wp(\mathbf{b}') \, d\mathbf{b}' \int_{\mathbf{y}} \mathbf{b}'^T \mathbf{f}_{\lambda}(\mathbf{y}) \int_0^1 \lim_{\varepsilon \rightarrow 0} \mathbf{b}'^T \nabla \psi(\mathbf{y} - t\varepsilon \mathbf{b}') \, dt \, d\mathbf{y} \\ &= - \int_{\mathbf{y}} \mathbf{f}_{\lambda}^T(\mathbf{y}) \underbrace{\left(\int_{\mathbf{b}'} \wp(\mathbf{b}') \mathbf{b}' \mathbf{b}'^T \, d\mathbf{b}' \right)}_{=I} \nabla \psi(\mathbf{y}) \, d\mathbf{y} \\ &= - \langle \mathbf{f}_{\lambda}(\mathbf{y}), \nabla \psi(\mathbf{y}) \rangle = \langle \operatorname{div}\{\mathbf{f}_{\lambda}(\mathbf{y})\}, \psi(\mathbf{y}) \rangle \quad (\text{from (E.16)}) . \end{aligned}$$

■

Bibliography

- [1] C. E. Shannon, “A Mathematical Theory of Communication,” *Bell System Technical Journal*, vol. 27, no. 3 and 4, pp. 379–423 and 623–656, 1948.
- [2] C. E. Shannon, “Communication in the Presence of Noise,” *Proceedings of the Institute of Radio Engineers*, vol. 37, no. 1, pp. 10–21, 1949.
- [3] H. Nyquist, “Certain Topics in Telegraph Transmission Theory,” *Transactions of the American Institute of Electrical Engineers*, vol. 47, pp. 617–644, 1928.
- [4] A. J. Jerri, “The Shannon Sampling Theorem—Its Various Extensions and Applications—Tutorial Review,” *Proceedings of the IEEE*, vol. 65, no. 11, pp. 1565–1596, 1977.
- [5] M. Unser, “Sampling—50 Years After Shannon,” *Proceedings of the IEEE*, vol. 88, no. 4, pp. 569–587, 2000.
- [6] E. Meijering, “A Chronology of Interpolation: From Ancient Astronomy to Modern Signal and Image Processing,” *Proceedings of the IEEE*, vol. 90, no. 3, pp. 319–342, 2002.
- [7] E. T. Whittaker, “On the Functions which are Represented by the Expansion of Interpolating Theory,” *Proceedings of the Royal Society of Edinburgh*, vol. 35, pp. 181–194, 1915.
- [8] J. M. Whittaker, “On the Cardinal Function of Interpolation Theory,” *Proceedings of Edinburgh Mathematical Society*, vol. 1, pp. 41–46, 1927–1929.

-
- [9] J. M. Whittaker, “The ‘Fourier’ Theory of the Cardinal Function,” *Proceedings of Edinburgh Mathematical Society*, vol. 1, no. 169–176, 1927–1929.
- [10] J. M. Whittaker, “Interpolation Function Theory,” in *Cambridge Tracts in Mathematics and Mathematical Physics*. Cambridge University Press, Cambridge, U.K, 1935.
- [11] E. Borel, “Sur l’ Interpolation,” *Comptes Rendus des Séances de l’Académie des Sciences*, vol. 124, no. 13, pp. 673–676, 1897.
- [12] K. Ogura, “On a Certain Transcendental Integral Function in the Theory of Interpolation,” *Tohoku Mathematical Journal*, vol. 17, pp. 64–72, 1928.
- [13] K. Ogura, “On Some Central Difference Formulas of Interpolation,” *Tohoku Mathematical Journal*, vol. 17, pp. 232–241, 1920.
- [14] M. Unser, “Splines: A Perfect Fit for Signal and Image Processing,” *IEEE Signal Processing Magazine*, vol. 16, no. 6, pp. 22–38, November 1999.
- [15] P. Thévenaz, T. Blu, and M. Unser, “Image Interpolation and Resampling,” in *Handbook of Medical Imaging, Processing and Analysis*, I. N. Bankman, Ed., pp. 393–420. Academic Press, San Diego, CA, USA, 2000.
- [16] T. M. Lehmann, C. Gönnner, and K. Spitzer, “Survey: Interpolation Methods in Medical Image Processing,” *IEEE Transactions on Medical Imaging*, vol. 18, no. 11, pp. 1049–1075, November 1999.
- [17] T. M. Lehmann, C. Gönnner, and K. Spitzer, “Addendum: B-Spline Interpolation in Medical Image Processing,” *IEEE Transactions on Medical Imaging*, vol. 20, no. 7, pp. 660–665, July 2001.
- [18] R. W. Parrot, M. R. Stytz, P. Amburn, and D. Robinson, “Towards Statistically Optimal Interpolation for 3-D Medical Imaging,” *IEEE Engineering in Medicine and Biology Magazine*, vol. 12, no. 5, pp. 49–59, September 1993.
- [19] E. Meijering, W. J. Niessen, and M. A. Viergever, “Quantitative Evaluation of Convolution-Based Methods for Medical Image Interpolation,” *Medical Image Analysis*, vol. 5, no. 2, pp. 111–126, June 2001.

-
- [20] E. Meijering, K. J. Zuiderveld, and M. A. Viergever, “Image Registration for Digital Subtraction Angiography,” *Internatioal Journal of Computer Vision*, vol. 32, no. 2-3, pp. 227–246, 1999.
- [21] V. Walimbe and R. Shekhar, “Automatic Elastic Image Registration by Interpolation of 3-D Rotations and Translations from Discrete Rigid-Body Transformations,” *Medical Image Analysis*, vol. 10, no. 6, pp. 899–914, 2006.
- [22] M. Holden, “A Review of Geometric Transformations for Nonrigid Body Registration,” *IEEE Transactions on Medical Imaging*, vol. 27, no. 1, pp. 111–128, January 2008.
- [23] J. A. Fessler, “On NUFFT-based Gridding for Non-Cartesian MRI,” *Journal of Magnetic Resonance*, vol. 188, no. 2, pp. 191–195, 2007.
- [24] M. R. Smith and S. T. Nichols, “Efficient Algorithms for Generating Interpolated (Zoomed) MR Images,” *Magnetic Resonance in Medicine*, vol. 7, no. 2, pp. 156–171, June 1988.
- [25] E. B. Ameer, D. Sbibih, A. Almhdie, and C. Léger, “New Spline Quasi-Interpolant for Fitting 3-D Data on the Sphere: Applications to Medical Imaging,” *IEEE Signal Processing Letters*, vol. 14, no. 5, pp. 333–336, May 2007.
- [26] A. Kadosh, D. Cohen-Or, and R. Yagel, “Tricubic Interpolation of Discrete Surfaces for Binary Volumes,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 9, no. 4, pp. 580–586, October-December 2003.
- [27] O. Sommer, A. Dietz, R. Westermann, and T. Ertl, “An Interactive Visualization and Navigation Tool for Medical Volume Data,” *Computers and Graphics*, vol. 23, no. 2, pp. 233–244, 1999.
- [28] B. Csébfalvi, “An Evaluation of Prefiltered Reconstruction Schemes for Volume Rendering,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 2, pp. 289–301, March-April 2008.
- [29] M. Unser, A. Aldroubi, and M. Eden, “B-Spline Signal Processing: Part I—Theory,” *IEEE Transactions on Signal Processing*, vol. 41, no. 2, pp. 821–832, February 1993.

-
- [30] T. Möller, R. Machiraju, K. Mueller, and R. Yagel, “Evaluation and Design of Filters Using a Taylor Series Expansion,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 3, no. 2, pp. 184–199, April–June 1997.
- [31] F. M. Weinhaus and V. Devarajan, “Texture Mapping 3-D Models of Real-World Scenes,” *ACM Computing Surveys*, vol. 29, no. 4, pp. 325–365, December 1997.
- [32] Y. Tang, J. Wang, H. Bao, and Q. Peng, “RBF-Based Constrained Texture Mapping,” *Computers and Graphics*, vol. 27, no. 3, pp. 415–422, 2003.
- [33] P. Brigger, J. Hoeg, and M. Unser, “B-Spline Snakes: A Flexible Tool for Parametric Contour Detection,” *IEEE Transactions on Image Processing*, vol. 9, no. 9, pp. 1484–1496, 2000.
- [34] M. R. Banham and A. K. Katsaggelos, “Digital Image Restoration,” *IEEE Signal Processing Magazine*, vol. 14, no. 2, pp. 24–41, 1997.
- [35] A. Tikhonov and V. Arsenin, *Solution of Ill-Posed Problems*, Wiley, New York NY, USA, 1977.
- [36] I. J. Schoenberg, “Spline Functions and the Problem of Graduation,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 52, no. 4, pp. 947–950, 1964.
- [37] C. M. Reinsch, “Smoothing by Spline Functions,” *Numerische Mathematik*, vol. 10, no. 3, pp. 177–183, 1967.
- [38] G. Wahba, “Smoothing Noisy Data with Spline Functions,” *Numerische Mathematik*, vol. 24, no. 5, pp. 383–393, 1975.
- [39] P. Craven and G. Wahba, “Smoothing Noisy Data with Spline Functions—Estimating the Correct Degree of Smoothing by the Method of Generalized Cross-Validation,” *Numerische Mathematik*, vol. 31, no. 4, pp. 377–403, 1979.
- [40] M. L. Stein, *Interpolation of Spatial Data—Some Theory for Kriging*, Springer Series in Statistics. Springer, New York NY, USA, 1999.

-
- [41] P. L. Butzer and R. L. Stens, "Sampling Theory for Not Necessarily Band-Limited Functions: A Historical Overview," *SIAM Review*, vol. 34, no. 1, pp. 40–53, March 1992.
- [42] C. de Boor, "Quasi-Interpolation and Approximation Power of Multivariate Splines," *Nato Advanced Study INST on Computation of Curves and Surfaces, July 10-21, 1989, Puerto La Cruz Spain*, vol. 307, pp. 313–345, 1990.
- [43] C. de Boor and A. Ron, "Fourier-Analysis of the Approximation Power of Principal Shift-Invariant Sapces," *Constructive Approximation*, vol. 8, no. 4, pp. 427–462, 1992.
- [44] C. de Boor, R. A. DeVore, and A. Ron, "Approximation From Shift-Invariant Subspaces of $L_2(\mathbb{R}^d)$," *Transactions of the American Mathematical Society*, vol. 341, no. 2, pp. 787–806, February 1994.
- [45] G. Strang and G. Fix, "A Fourier Analysis of the Finite Element Variational Method," in *Constructive Aspects of Functional Analysis-Part II*, Cremonese, Ed., pp. 796–830. Centro Internazionale Matematico Estivo, Rome, Italy, 1971.
- [46] C. de Boor and R. Q. Jia, "Controlled Approximation and a Characterization of the Local Approximation Order," *Proceedings of the American Mathematical Society*, vol. 95, no. 4, pp. 547–553, 1985.
- [47] C. de Boor and G. Fix, "Spline Approximation by Quasi-Interpolants," *Journal of Approximation Theory*, vol. 8, no. 1, pp. 19–45, 1973.
- [48] I. Schoenberg, "Contribution to the Problem of Approximation of Equidistant Data by Analytic Functions," *Quarterly of Applied Mathematics*, vol. 4, pp. 45–99 and 112–141, 1946.
- [49] C. de Boor, "On Calculating with B-Splines," *Journal of Approximation Theory*, vol. 6, no. 1, pp. 50–62, 1972.
- [50] M. Unser, A. Aldroubi, and M. Eden, "Polynomial Spline Signal Approximations: Filter Design and Asymptotic Equivalence with Shannon's Sampling Theorem," *IEEE Transactions on Information Theory*, vol. 38, no. 1, pp. 95–103, January 1992.

-
- [51] A. Aldroubi and M. Unser, “Sampling Procedures in Function Spaces and Asymptotic Equivalence with Shannon’s Sampling Theory,” *Numerical Functional Analysis and Optimization*, vol. 15, no. 1 & 2, pp. 1–21, 1994.
- [52] M. Unser and A. Aldroubi, “A General Sampling Theory for Nonideal Acquisition Devices,” *IEEE Transactions on Signal Processing*, vol. 42, no. 11, pp. 2915–2925, 1994.
- [53] T. Blu and M. Unser, “Quantitative Fourier Analysis of Approximation Techniques: Part I—Interpolators and Projectors,” *IEEE Transactions on Signal Processing*, vol. 47, no. 10, pp. 2783–2795, October 1999.
- [54] T. Blu and M. Unser, “Approximation Error for Quasi-Interpolators and (Multi-) Wavelet Expansions,” *Applied and Computational Harmonic Analysis*, vol. 6, no. 2, pp. 219–251, March 1999.
- [55] E. Kreyszig, *Introduction to Functional Analysis with Applications*, Wiley, New York NY, USA, 1978.
- [56] S. Saitoh, V. K. Tuan, and M. Yamamoto, “Convolution Inequalities and Applications,” *Journal of Inequalities in Pure and Applied Mathematics*, vol. 42, no. 3, pp. Article 50, 2003.
- [57] S. H. Izen, “Generalized Sampling Expansion on Lattices,” *IEEE Transactions on Signal Processing*, vol. 53, no. 6, pp. 1949–1963, June 2005.
- [58] H. R. Künsch, E. Agrell, and F. A. Hamprecht, “Optimal Lattices for Sampling,” *IEEE Transactions on Information Theory*, vol. 51, no. 2, pp. 634–647, February 2005.
- [59] C. de Boor, R. A. DeVore, and A. Ron, “The Structure of Finitely Generated Shift-Invariant Spaces,” *Journal of Functional Analysis*, vol. 119, no. 1, pp. 37–78, January 1994.
- [60] M. Unser, “Approximation Power of Biorthogonal Wavelet Expansions,” *IEEE Transactions on Signal Processing*, vol. 44, no. 3, pp. 519–527, March 1996.
- [61] G. H. Hardy, “Notes on Special Systems of Orthogonal Functions (IV): The Orthogonal Functions Series of Whittaker’s Cardinal Series,” *Proceedings of the Cambridge Philosophical Society*, vol. 37, no. 4, pp. 331–348, October 1941.

-
- [62] I. Schoenberg, “Cardinal Interpolation and Spline Functions,” *Journal of Approximation Theory*, vol. 2, pp. 167–206, 1969.
- [63] P. Thévenaz, T. Blu, and M. Unser, “Interpolation Revisited,” *IEEE Transactions on Medical Imaging*, vol. 19, no. 7, pp. 739–758, July 2000.
- [64] A. Ron, “A Characterization of the Approximation Order of Multivariate Spline Spaces,” *Studia Mathematica*, vol. 98, no. 1, pp. 73–90, 1991.
- [65] N. Dyn, I. R. H. Jackson, D. Levin, and A. Ron, “On Multivariate Approximation by the Integer Translates of a Basis Function,” *Israel Journal of Mathematics*, vol. 78, no. 1, pp. 95–130, 1992.
- [66] Y. C. Eldar and M. Unser, “Nonideal Sampling and Interpolation from Noisy Observations in Shift-Invariant Spaces,” *IEEE Transactions on Signal Processing*, vol. 54, no. 7, pp. 2636–2651, 2006.
- [67] D. P. Petersen and D. Middleton, “Sampling and Reconstruction of Wave-Number-Limited Functions in N -Dimensional Euclidean Spaces,” *Information and Control*, vol. 5, no. 4, pp. 279–323, 1962.
- [68] A. Papoulis, “Generalized Sampling Expansion,” *IEEE Transactions on Circuits and Systems*, vol. 24, no. 11, pp. 652–654, 1977.
- [69] B. Pesquet-Popescu and J. L. Véhel, “Stochastic Fractal Models for Image Processing,” *IEEE Signal Processing Magazine*, vol. 19, no. 5, pp. 48–62, 2002.
- [70] C. Rabut, “Elementary m -Harmonic Cardinal B-splines,” *Numerical Algorithms*, vol. 2, pp. 39–62, 1992.
- [71] A. Aldroubi and K. Gröchenig, “Nonuniform Sampling and Reconstruction in Shift-Invariant Spaces,” *SIAM Review*, vol. 43, pp. 585–620, 2001.
- [72] H. H. Fu, M. Ng, M. Nikolova, and J. Barlow, “Efficient Minimization Methods of Mixed $\ell_2 - \ell_1$ and $\ell_1 - \ell_1$ Norms for Image Restoration,” *SIAM Journal on Scientific Computing*, vol. 27, no. 6, pp. 1881–1902, 2006.
- [73] M. Nikolova, “Minimizers of Cost-Functions Involving Nonsmooth Data-Fidelity Terms: Application to the Processing of Outliers,” *SIAM Journal on Numerical Analysis*, vol. 40, no. 3, pp. 965–994, 2002.

-
- [74] S. Osher, L. Rudin, and E. Fatemi, “Nonlinear Total Variation Based Noise Removal Algorithms,” *Physica D.*, vol. 60, pp. 259–268, 1992.
- [75] T. F. Chan and S. Esedoğlu, “Aspects of Total Variation Regularized L^1 Function Approximation,” *SIAM Journal on Applied Mathematics*, vol. 65, no. 5, pp. 1817–1837, 2005.
- [76] Y. Katznelson, *An Introduction to Harmonic Analysis*, Dover Publications, New York, 1976.
- [77] L. Schwartz, *Théorie des Distributions*, Hermann, Paris, 1966.
- [78] S. Tirosh, D. Van De Ville, and M. Unser, “Polyharmonic Smoothing Splines and the Multi-Dimensional Wiener Filtering of Fractal-like Signals,” *IEEE Transactions on Image Processing*, vol. 15, no. 9, pp. 2616–2630, 2006.
- [79] J. Duchon, “Splines Minimizing Rotation-Invariant Semi-Norms in Sobolev Spaces,” in *Multivariate Approximation Theory*, W. Schempp and K. Zeller, Eds., pp. 85–100. Birkhäuser-Verlag, Basel, 1979.
- [80] D. Pierre, *Optimization Theory with Applications*, Dover Publications, New York, 1986.
- [81] 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, 2005.
- [82] M. Unser and T. Blu, “Fractional Splines and Wavelets,” *SIAM Review*, vol. 42, no. 1, pp. 43–67, 2000.
- [83] W. R. Madych and S. A. Nelson, “Polyharmonic Cardinal Spline,” *Journal of Approximation Theory*, vol. 60, pp. 141–156, 1990.
- [84] P. Thévenaz and M. Unser, “Precision Isosurface Rendering of 3-D Image Data,” *IEEE Transactions on Image Processing*, vol. 12, no. 7, pp. 764–775, July 2003.
- [85] D. Van De Ville, T. Blu, and M. Unser, “Isotropic Polyharmonic B-splines: Scaling Functions and Wavelets,” *IEEE Transactions on Image Processing*, vol. 14, no. 11, pp. 1798–1813, 2005.

-
- [86] B. Whitcher, “Wavelet-Based Bootstrapping of Spatial Patterns on a Finite Lattice,” *Computational Statistics and Data Analysis*, vol. 50, no. 9, pp. 2399–2421, 2006.
- [87] M. Abramowitz and I. Stegun, *Handbook of Mathematical Functions*, Dover Publications, New York, 9th edition, 1965.
- [88] S. M. Kay, *Statistical Signal Processing Vol. I: Estimation Theory*, Prentice-Hall Inc., Englewood Cliffs, New Jersey, USA, 1995.
- [89] C. L. Fales, F. O. Huck, J. A. McCormick, and S. K. Park, “Wiener Restoration of Sampled Image Data: End-to-End Analysis,” *Journal of the Optical Society of America A*, vol. 5, no. 3, pp. 300–314, 1988.
- [90] A. J. Gray, J. W. Kay, and D. M. Titterington, “An Empirical Study of the Simulation of Various Models Used for Images,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 16, no. 5, pp. 507–513, May 1994.
- [91] S. Ramani, D. Van De Ville, T. Blu, and M. Unser, “Nonideal Sampling and Regularization Theory,” *IEEE Transactions on Signal Processing*, vol. 56, no. 3, pp. 1055–1070, March 2008.
- [92] R. R. Schultz and R. L. Stevenson, “A Bayesian Approach to Image Expansion for Improved Definition,” *IEEE Transactions on Image Processing*, vol. 3, no. 3, pp. 233–242, May 1994.
- [93] A. Raj and K. Thakur, “Fast and Stable Bayesian Image Expansion Using Sparse Edge Priors,” *IEEE Transactions on Image Processing*, vol. 16, no. 4, pp. 1073–1084, April 2007.
- [94] L. Chen and K. H. Yap, “Regularized Interpolation Using Kronecker Product for Still Images,” *Proceedings of the IEEE International Conference on Image Processing (ICIP), Genova, Italy*, vol. 1-5, pp. 2205–2208, 2005.
- [95] J. H. Shin, J. H. Jung, and J. K. Paik, “Regularized Iterative Image Interpolation and Its Application to Spatially Scalable Coding,” *IEEE Transactions on Consumer Electronics*, vol. 44, no. 3, pp. 1042–1047, August 1998.
- [96] D. R. Hunter and K. Lange, “A Tutorial on MM Algorithms,” *The American Statistician*, vol. 59, no. 1, pp. 30–37, February 2004.

-
- [97] P. J. Green, “On Use of the EM for Penalized Likelihood Estimation,” *Journal of the Royal Statistical Society Series B*, vol. 52, no. 3, pp. 443–452, 1990.
- [98] A. R. De Pierro, “A Modified Expectation Maximization Algorithm for Penalized Likelihood Estimation in Emission Tomography,” *IEEE Transactions on Medical Imaging*, vol. 14, no. 1, pp. 132–137, March 1995.
- [99] M. Yavuz and J. A. Fessler, “Penalized-Likelihood Estimators and Noise Analysis for Randoms-Precorrected PET Transmission Scans,” *IEEE Transactions on Medical Imaging*, vol. 18, no. 8, pp. 665–674, August 1999.
- [100] C. Vonesch, F. Aguet, J. -L. Vonesch, and M. Unser, “The Colored Revolution of Bioimaging,” *IEEE Signal Processing Magazine*, vol. 23, no. 3, pp. 20–31, May 2006.
- [101] L. A. Shepp and Y. Vardi, “Maximum Likelihood Reconstruction for Emission Tomography,” *IEEE Transactions on Medical Imaging*, vol. 1, no. 2, pp. 113–122, October 1982.
- [102] I. M. Gelfand and G. E. Shilov, *Generalized Functions*, vol. 1, Academic Press, New York, 1964.
- [103] T. Chan, S. Esedoğlu, F. Park, and A. Yip, “Recent Developments in Total Variation Image Restoration,” in *Mathematical Models in Computer Vision*, N. Paragios, Y. Chen, and O. Faugeras, Eds. 2005, pp. 17–30, Springer-Verlag.
- [104] T. Le, R. Chartrand, and T. J. Asaki, “A Variational Approach to Reconstructing Images Corrupted by Poisson Noise,” *Journal of Mathematical Vision and Imaging*, vol. 27, no. 3, pp. 257–263, 2007.
- [105] 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, August 2002.
- [106] J. Duchon, “Splines Minimizing Rotation-invariant Semi-norms in Sobolev Spaces,” in *Constructive Theory of Functions of Several Variables*, W. Schempp and K. Zeller, Eds., pp. 85–100. Springer-Verlag, Berlin, Germany, 1977.

-
- [107] D. Geman and G. Reynolds, “Constrained Restoration and Recovery of Discontinuities,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 3, pp. 367–383, 1992.
- [108] P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud, “Deterministic Edge-Preserving Regularization in Computed Imaging,” *IEEE Transactions on Image Processing*, vol. 2, no. 6, pp. 298–311, February 1997.
- [109] M. Nikolova and M. K. Ng, “Analysis of Half-Quadratic Minimization Methods for Signal and Image Recovery,” *SIAM Journal on Scientific Computing*, vol. 27, no. 3, pp. 937–966, 2005.
- [110] M. Nikolova and R. H. Chan, “The Equivalence of Half-Quadratic Minimization and the Gradient Linearization Iteration,” *IEEE Transactions on Image Processing*, vol. 16, no. 6, pp. 1623–1627, June 2007.
- [111] C. T. Kelley, *Iterative Methods for Optimization*, SIAM, Philadelphia, USA., 1999.
- [112] “Brainweb: Simulated MRI Volumes for Normal Brain,” McConnell Brain Imaging Center, www.bic.mni.mcgill.ca/brainweb/selection_normal.html.
- [113] E. Meijering, “RandomJ: A Java Package for Image Randomization,” <http://www.imagescience.org/meijering/software/randomj/>.
- [114] W. C. Karl, “Regularization in Image Restoration and Reconstruction,” in *Handbook of Image & Video Processing*, A. Bovik, Ed., pp. 183–202. ELSEVIER, 2nd edition, 2005.
- [115] P. C. Hansen, “Analysis of Discrete Ill-Posed Problems by Means of the L-Curve,” *SIAM Review*, vol. 34, no. 4, pp. 561–580, 1992.
- [116] P. C. Hansen and D. P. O’Leary, “The Use of the L-Curve in the Regularization of Discrete Ill-Posed Problems,” *SIAM Journal on Scientific Computing*, vol. 14, no. 6, pp. 1487–1503, 1993.
- [117] T. Regińska, “A Regularization Parameter in Discrete Ill-Posed Problems,” *SIAM Journal on Scientific Computing*, vol. 17, no. 3, pp. 740–749, 1996.
- [118] G. Wahba, *Spline Models for Observational Data*, SIAM, Philadelphia, 1990.

-
- [119] G. H. Golub, M. Heath, and G. Wahba, "Generalized Cross-Validation as a Method for Choosing a Good Ridge Parameter," *Technometrics*, vol. 21, no. 2, pp. 215–223, May 1979.
- [120] D. Nychka, "Bayesian Confidence Intervals for Smoothing Splines," *Journal of the American Statistical Association*, vol. 83, no. 404, pp. 1134–1143, 1988.
- [121] A. M. Thompson, J. C. Brown, J. W. Kay, and D. M. Titterton, "A Study of Methods of Choosing the Smoothing Parameter in Image Restoration by Regularization," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 4, pp. 326–339, 1991.
- [122] N. P. Galatsanos and A. K. Katsaggelos, "Methods for Choosing the Regularization Parameter and Estimating the Noise Variance in Image Restoration and Their Relation," *IEEE Transactions on Image Processing*, vol. 1, no. 3, pp. 322–336, 1992.
- [123] K. M. Perry and S. J. Reeves, "A Practical Stopping Rule for Iterative Signal Restoration," *IEEE Transactions on Signal Processing*, vol. 42, no. 7, pp. 1829–1833, July 1994.
- [124] S. J. Reeves, "Generalized Cross-Validation as a Stopping Rule for the Richardson-Lucy Algorithm," *International Journal of Imaging Systems and Technology*, vol. 6, no. 4, pp. 387–391, Winter 1995.
- [125] S. J. Reeves, "Optimal Space-Varying Regularization in Iterative Image Restoration," *IEEE Transactions on Image Processing*, vol. 3, no. 3, pp. 319–324, May 1994.
- [126] S. J. Reeves, "Optimal Estimation of the Regularization Parameter and Stabilizing Functional for Regularized Image Restoration," *Optical Engineering*, vol. 29, no. 5, pp. 446–454, May 1990.
- [127] R. Molina, A. K. Katsaggelos, and J. Mateos, "Bayesian and Regularization Methods for Hyperparameter Estimation in Image Restoration," *IEEE Transactions on Image Processing*, vol. 8, no. 2, pp. 231–246, 1999.
- [128] N. P. Galatsanos, V. Mesarovic, R. Molina, J. Mateos, and A. K. Katsaggelos, "Hyper-Parameter Estimation Using Gamma Hyper-Priors in Image Restoration from Partially-Known Blurs," *Optical Engineering*, vol. 41, pp. 1845–1854, 2002.

-
- [129] G. Archer and D. M. Titterton, "On some Bayesian / Regularization Methods for Image Restoration," *IEEE Transactions on Image Processing*, vol. 4, no. 7, pp. 989–995, 1995.
- [130] A. Mohammad-Djafari, "A Full Bayesian Approach for Inverse Problems," in *Maximum Entropy and Bayesian Methods*, K. Hanson and R. Silver, Eds. Kluwer, 1996.
- [131] C. Stein, "Estimation of the Mean of a Multivariate Normal Distribution," *Annals of Statistics*, vol. 9, pp. 1135–1151, 1981.
- [132] D. L. Donoho and I. M. Johnstone, "Adapting to Unknown Smoothness via Wavelet Shrinkage," *Journal of the American Statistical Association*, vol. 90, no. 432, pp. 1200–1224, 1995.
- [133] X. -P. Zhang and M. D. Desai, "Adaptive Denoising Based on SURE Risk," *IEEE Signal Processing Letters*, vol. 5, no. 10, pp. 265–267, 1998.
- [134] A. Benazza-Benyahia and J. -C. Pesquet, "Building Robust Wavelet Estimators for Multicomponent Images Using Stein's Principle," *IEEE Transactions on Image Processing*, vol. 14, no. 11, pp. 1814–1830, 2005.
- [135] F. Luisier, T. Blu, and M. Unser, "A New SURE Approach to Image Denoising: Interscale Orthonormal Wavelet Thresholding," *IEEE Transactions on Image Processing*, vol. 16, no. 3, pp. 593–606, 2007.
- [136] T. Blu and F. Luisier, "The SURE-LET Approach to Image Denoising," *IEEE Transactions on Image Processing*, vol. 16, no. 11, pp. 2778–2786, 2007.
- [137] D. L. Snyder and A. M. Hammoud, "Image Recovery from Data Acquired with a Charge-Coupled-Device Camera," *Journal of the Optical Society of America A*, vol. 10, no. 5, pp. 1014–1023, 1993.
- [138] H. Lanteri and C. Theys, "Inverse Problems and Image Deconvolution," *EAS Publications*, vol. 22, no. 3–33, 2006.
- [139] H. M. Hudson, "A Natural Identity for Exponential Families with Applications in Multiparameter Estimation," *Annals of Statistics*, vol. 6, no. 3, pp. 473–484, 1978.

-
- [140] W. R. Madych and S. A. Nelson, "Polyharmonic Cardinal Splines: A Minimization Property," *Journal of Approximation Theory*, vol. 60, pp. 141–156, 1990.
- [141] K. C. Li, "Asymptotic Optimality of \mathcal{C}_L and Generalized Cross-Validation in Ridge Regression with Application to Spline Smoothing," *Annals of Statistics*, vol. 14, no. 3, pp. 1101–1112, 1986.
- [142] P. Sarder and A. Nehorai, "Deconvolution Methods for 3-D Fluorescence Microscopy Images," *IEEE Signal Processing Magazine*, vol. 23, no. 3, pp. 32–45, 2006.
- [143] Y. C. Eldar, "Generalized SURE for Exponential Families: Applications to Regularization," *IEEE Transactions on Signal Processing*, in press.
- [144] D. A. Girard, "A Fast 'Monte-Carlo Cross-Validation' Procedure for Large Least Squares Problems with Noisy Data," *Numerische Mathematik*, vol. 56, pp. 1–23, 1989.
- [145] M. F. Hutchinson, "A Stochastic Estimator of the Trace of the Influence Matrix for Laplacian Smoothing Splines," *Communications in Statistics-Simulation and Computation*, vol. 18, 19, no. 3, 2, pp. 1059–1076, 433–450, 1989.
- [146] W. L. Briggs, V. E. Henson, and S. F. McCormick, *A Multigrid Tutorial*, SIAM, Philadelphia, USA., 2000.
- [147] M. Arigovindan, P. Hunziker M. Sühling, and M. Unser, "Variational Image Reconstruction from Arbitrarily Spaced Samples: A Fast Multiresolution Spline Solution," *IEEE Transactions on Image Processing*, vol. 14, no. 4, pp. 460–460, April 2005.
- [148] Z. Bai, M. Fahey, and G. Golub, "Some Large-Scale Matrix Computation Problems," *Journal of Computational and Applied Mathematics*, vol. 74, no. 1-2, pp. 71–89, November 1996.
- [149] S. Dong and K. Liu, "Stochastic Estimation with Z_2 Noise," *Physics Letters B*, vol. 328, no. 1-2, pp. 130–136, May 1994.
- [150] D. A. Girard, "The Fast Monte-Carlo Cross-Validation and \mathcal{C}_L procedures: Comments, New Results and Application to Image Recovery Problems," *Computational Statistics*, vol. 10, no. 3, pp. 205–231, 251–258, 1995.

-
- [151] G. Wahba, D. Johnson, F. Gao, and J. Gong, “Adaptive Tuning of Numerical Weather Prediction Models: Part I: Randomized GCV and Related Methods in Three and Four Dimensional Data Assimilation,” Technical Reprpt 920, Department of Statistics, University of Wisconsin, Madison, WI, USA, 1994.
- [152] G. Wahba, “The Fast Monte-Carlo Cross-Validation and $\mathcal{C}_{\mathcal{L}}$ Procedures: Comments, New Results and Application to Image Recovery Problems - Comments,” *Computational Statistics*, vol. 10, no. 3, pp. 249–250, 1995.
- [153] D. A. Girard, “The Fast Monte-Carlo Cross-Validation and $\mathcal{C}_{\mathcal{L}}$ Procedures: Comments, New Results and Application to Image Recovery Problems - Rejoinder,” *Computational Statistics*, vol. 10, no. 3, pp. 251–258, 1995.
- [154] A. Chambolle, R. A. DeVore, N. -Y. Lee, and B. J. Lucier, “Nonlinear Wavelet Image Processing: Variational Problems, Compression, and Noise Removal Through Wavelet Shrinkage,” *IEEE Transactions on Image Processing*, vol. 7, no. 3, pp. 319–335, 1998.
- [155] M. Lang, H. Guo, J. E. Odegard, C. S. Burrus, and R. O. Wells Jr, “Noise Reduction Using an Undecimated Discrete Wavelet Transform,” *IEEE Signal Processing Letters*, vol. 3, no. 1, pp. 10–12, 1996.
- [156] P. Dutilleul, “An Implementation of the Algorithm à Trous to Compute the Wavelet Transform,” in *Wavelets: Time-Frequency Methods and Phase Space*, J. -M. Combes, A. Grossman, and P. Tchamitchian, Eds., pp. 298–304. Springer-Verlag, Berlin, Germany, 1989.
- [157] M. Holschneider, R. Kronland-Martinet, J. Morlet, and P. Tchamitchian, “A Real-Time Algorithm for Signal Analysis with the Help of the Wavelet Transform,” in *Wavelets: Time-Frequency Methods and Phase Space*, J. -M. Combes, A. Grossman, and P. Tchamitchian, Eds., pp. 286–297. Springer-Verlag, Berlin, Germany, 1989.
- [158] M. A. T. Figueiredo, J. B. Dias, J. P. Oliveira, and R. D. Nowak, “On Total Variation Denoising: A New Majorization-Minimization Algorithm and an Experimental Comparison with Wavelet Denoising,” *Proceedings of IEEE International Conference on Image Processing (ICIP’06), Atlanta, GA, USA*, pp. 2633–2636, October 2006.

-
- [159] D. Geman and G. Yang, “Nonlinear Image Recovery with Half-Quadratic Regularization,” *IEEE Transactions on Image Processing*, vol. 4, pp. 932–946, 1995.
- [160] C. M. Hurvich and C. L. Tsai, “A Crossvalidatory AIC for Hard Wavelet Thresholding in Spatially Adaptive Function Estimation,” *Biometrika*, vol. 85, no. 3, pp. 701–710, 1998.
- [161] E. Zarahn, G. Aguirre, and M. D’Esposito, “Empirical Analyses of BOLD fMRI Statistics. I. Spatially Unsmoothed Data Collected Under Null-Hypothesis Conditions,” *NeuroImage*, vol. 5, pp. 179–197, 1997.
- [162] L. Sendur and I. W. Selesnick, “Bivariate Shrinkage with Local Variance Estimation,” *IEEE Signal Processing Letters*, vol. 9, no. 12, pp. 438–441, 2002.
- [163] J. Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, “Image Denoising Using Scale Mixtures of Gaussians in the Wavelet Domain,” *IEEE Transactions on Image Processing*, vol. 12, no. 11, pp. 1338–1351, 2003.
- [164] A. Pižurica and W. Philips, “Estimating the Probability of the Presence of a Signal of Interest in Multiresolution Single and Multiband Image Denoising,” *IEEE Transactions on Image Processing*, vol. 15, no. 3, pp. 654–665, 2006.
- [165] M. Yuan, “Automatic Smoothing for Poisson Regression,” *Communications in Statistics-Theory and Methods*, vol. 34, pp. 603–617, 2005.
- [166] J. T. Day, “On the Convergence of Taylor Series for Functions of n Variables,” *Mathematics Magazine*, vol. 40, no. 5, pp. 258–260, 1967.

Curriculum Vitæ

BM 4.139, STI-LIB, Station-17
École Polytechnique Fédérale de Lausanne
Ecublens, Lausanne
CH-1015, Switzerland.

Phone: +41 21 693 51 37
Fax: +41 21 693 37 01
Email: mesathish@gmail.ch
sathish.ramani@epfl.ch



PERSONAL DATA

Date of Birth	3 rd August, 1978
Place of Birth	Chennai, India
Nationality	Indian

EDUCATION

2004 – 2009	Ph. D.	Electrical Engineering	École Polytechnique Fédérale de Lausanne Switzerland
2001 – 2004	M. Sc.	Electrical Engineering	Indian Institute of Science Bangalore, India
1998 – 2000	M. Sc.	Physics	Sri Sathya Sai University, India
1995 – 1998	B. Sc.	Physics	Sri Sathya Sai University, India

AWARDS

- **Postdoctoral Fellowship**, Swiss National Science Foundation (awarded for the period 2009-2010).
- **Best Student-Paper Award**, IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Toulouse, France, May 2006.
- All-India **Junior Research Fellow** (JRF) scholarship award in physics, Council of Scientific and Industrial Research (CSIR), 2000 and 2001.
- **Gold Medal Award** for graduating with distinction, Sri Sathya Sai University, 2000.

PROFESSIONAL ACTIVITIES

Member	IEEE Signal Processing Society (since January 2008)
Reviewer	IEEE Transactions on Signal Processing IEEE Transactions on Image Processing Image and Vision Computing Signal Processing

TEACHING

- Teaching Assistant at the Biomedical Imaging Group, École Polytechnique Fédérale de Lausanne
 - 2005-2008: Image Processing Course I & II (Master's Level)
 - 2007-2008: Advanced Image Processing and Analysis Course (Doctoral Level)
- Supervised various student projects

RESEARCH INTERESTS

- Signal and image processing: Inverse problems (Deconvolution of natural and fluorescence microscopy images, Medical Imaging: Tomography, MRI, etc), Interpolation and Denoising
- Wavelets and Variational (partial differential equation-based) methods for inverse problems

- Statistics: Risk estimation for inverse problems, Monte-Carlo methods, CCD noise modeling

TECHNICAL SKILLS

- Programming Languages: C, C++, Java.
- Packages: MATLAB, Mathematica.
- Others: Microsoft products, Mac OS X.

PUBLICATIONS

Journal Articles

1. S. Ramani, P. Thévenaz, M. Unser, “Regularized Interpolation of Noisy Images”, *IEEE Transactions on Medical Imaging*, in peer-review.
2. S. Ramani, T. Blu, M. Unser, “Monte-Carlo SURE: A Black-Box Optimization of Regularization Parameters for General Denoising Algorithms”, *IEEE Transactions on Image Processing*, Vol. 17, No. 9, pp. 1540-1554, September 2008.
3. S. Ramani, D. V. De Ville, T. Blu, M. Unser, “Nonideal Sampling and Regularization Theory”, *IEEE Transactions on Signal Processing*, Vol. 56, No. 3, pp. 1055-1070, March 2008.
4. S. Ramani and M. Unser, “Matérn B-splines and the Optimal Reconstruction of Signals”, *IEEE Signal Processing Letters*, Vol. 13, No. 7, pp. 437-440, July 2006.
5. R. Sathish and G. V. Anand, “Improved Direction-Of-Arrival Estimation using Wavelet based Denoising Techniques”, *Journal of VLSI Signal Processing*, Vol. 45, pp. 29-48, September 2006.

Conference Proceedings

1. S. Ramani, C. Vonesch, M. Unser, “Deconvolution of 3D Fluorescence Micrographs with Automatic Risk Minimization”, *Proceedings of the Fifth IEEE International Symposium on Biomedical Imaging (ISBI 08)*, Paris, France, May 14-17, 2008, pp. 732-735.

2. C. Vonesch, S. Ramani, M. Unser, “Recursive Risk Estimation for Non-linear Image Deconvolution with a Wavelet-Domain Sparsity Constraint”, to appear in the *Proceedings of the Fifteenth IEEE International Conference on Image Processing (ICIP08)*, San Diego, CA, USA, October 12-15, 2008.
3. S. Delpretti, F. Luisier, S. Ramani, T. Blu, M. Unser, “Multiframe SURE-Let Denoising of Timelapse Fluorescence Microscopy Images”, *Proceedings of the Fifth IEEE International Symposium on Biomedical Imaging (ISBI’08)*, Paris, France, May 14-17, 2008, pp. 149-152.
4. S. Ramani, T. Blu, M. Unser, “Blind Optimization of Algorithm Parameters for Signal Denoising by Monte-Carlo SURE”, *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing*, Las Vegas, Nevada, USA, March 30-April 4, 2008.
5. S. Ramani, P. Thévenaz, M. Unser, “Regularized Interpolation for Noisy Data”, *Proceedings of the Fourth IEEE International Symposium on Biomedical Imaging (ISBI’07)*, Arlington VA, USA, April 12-15, 2007, pp. 612-615.
6. S. Ramani, D. V. De Ville, M. Unser, “Non-Ideal Sampling and Adapted Reconstruction Using the Stochastic Matérn Model”, **Best Student-paper Award**, *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*, Toulouse, France, Vol. 2, pp. II-73 - II-76, 2006.
7. S. Ramani, D. V. De Ville, Michael Unser, “Sampling in Practice: is the Best Reconstruction Space Bandlimited ?”, *Proceedings of IEEE International Conference on Image Processing*, Genoa, Italy, Vol. 2, pp. II-153-156, 2005.
8. R. Sathish and G. V. Anand, “Spatial Wavelet Packet Denoising for Improved DOA Estimation”, *Proceedings of Fourteenth IEEE Signal Processing Society Workshop on Machine Learning for Signal Processing*, pp. 745 - 754, 2004.
9. R. Sathish and G.V. Anand, “Improved Localisation of Underwater Acoustic Sources by Wavelet Array Denoising”, *Proceedings of International Conference on Underwater Defense Technology (UDT)*, Nice, France, June 2004.

-
10. R. Sathish and G. V. Anand, "Improved Bearing Estimation in Ocean using Wavelet Denoising", *Wavelets-X, Proceedings of SPIE*, San Diego, CA, USA, Vol. 5207, No. 1, pp. 47-58, 2003.
 11. R. Sathish and G. V. Anand, "Wavelet Denoising for Plane Wave DOA Estimation by MUSIC", *Proceedings of IEEE TENCON 2003 Conference on Convergent Technologies for Asia-Pacific Region*, Bangalore, India, Vol. 1, pp. 104-108, 2003.