

Function-valued Mappings and SSIM-based Optimization in Imaging

by

Daniel Otero

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Applied Mathematics

Waterloo, Ontario, Canada, 2015

© Daniel Otero 2015

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

In a few words, this thesis is concerned with two alternative approaches to imaging, namely, Function-valued Mappings (FVMs) and Structural Similarity Index Measure (SSIM)-based Optimization. Briefly, a FVM is a mathematical object that assigns to each element in its domain a function that belongs to a given function space. The advantage of this representation is that the infinite dimensionality of the range of FVMs allows us to give a more accurate description of complex datasets such as hyperspectral images and diffusion magnetic resonance images, something that can not be done with the classical representation of such data sets as vector-valued functions. For instance, a hyperspectral image can be described as a FVM that assigns to each point in a spatial domain a spectral function that belongs to the function space $L^2(\mathbb{R})$; that is, the space of functions whose energy is finite. Moreover, we present a Fourier transform and a new class of fractal transforms for FVMs to analyze and process hyperspectral images.

Regarding SSIM-based optimization, we introduce a general framework for solving optimization problems that involve the SSIM as a fidelity measure. This framework offers the option of carrying out SSIM-based imaging tasks which are usually addressed using the classical Euclidean-based methods. In the literature, SSIM-based approaches have been proposed to address the limitations of Euclidean-based metrics as measures of visual quality. These methods show better performance when compared to their Euclidean counterparts since the SSIM is a better model of the human visual system; however, these approaches tend to be developed for particular applications. With the general framework that it is presented in this thesis, rather than focusing on particular imaging tasks, we introduce a set of novel algorithms capable of carrying out a wide range of SSIM-based imaging applications. Moreover, such a framework allows us to include the SSIM as a fidelity term in optimization problems in which it had not been included before.

Acknowledgements

This thesis would not have been completed without the support of all the persons that have been along the path that I have traversed since I came to Waterloo back in January of 2011. Among them, I want to first thank my supervisor, Professor Edward Vrscay, for his constant guidance and support throughout all these years. His expertise, kindness, and sense of humour will be always remembered. I feel fortunate for having had him as my academic advisor

Also, I want to thank my cosupervisor, Professor Oleg Michailovich, who introduced me to the “wonders” of the field of optimization and its algorithms. His vast knowledge and energetic disposition towards research are examples to follow. Special thanks also go to Professor Davide La Torre, who is the embodiment of a true mathematician. His expertise and valuable comments are key ingredients of this thesis.

I would also like to thank the other Ph.D. committee members, Professor David Siegel and Profesor Justin Wan, as well as the external examiner, Professor Abdol-Reza Mansouri. Their feedback and comments highly contributed to improve the final version of the thesis.

Further, I want to gratefully acknowledge the financial support from the Faculty of Mathematics, University of Waterloo, the Department of Applied Mathematics and the Natural Sciences and Engineering Research Council of Canada Discovery Grant (ERV).

My deepest gratitude goes to the friends who have been around all these years: Iván Camilo Salgado, Andree Susanto, Vivek Kant, Shiv Vyas, Satya Kumar and Rahul Sharma. Your company and all the moments that have been shared have made this experience more significant.

To the members of the family Frasica Perdomo, Alejito, Mari, Sofi and Sarita, thank you so much for opening the doors of your home. You have not only being friends but also like a family to me.

My colombian friends, the sapoiguanacaimanes Cami, Caro, Dianis, Pachito, Betty, Pedro, Gloria, Glenda, Andy and Jose, your company is priceless. Thank you so much for bringing more happiness to this world. My other colombian friends, Mati, Ferchito, Angelú, Angie, Paola and all the mistapeos, melguistas and satsanguis, it is a blessing that our friendship endures distance and time.

Finally, special thanks to my entire family, specially to my grandparents, Fa, Andrés and my beloved nephew, Tomi. To my beautiful family, Padre querido, Mamita hermosa, Manita, Crespis and Cuchis, all of you are and endless source of love, laughter and smiles. Sometimes we are thousands of kilometers apart, but our hearts are always close. Thank you for always being there. I love you very much.

Dedication

To Sant Sadhu Ram Ji, Who is the manifestation of True Love on this Earth.

*“God is love, our soul is a drop
from the Ocean of love, and the
Way back to God is also through
love.”*

Sant Kirpal Singh

Table of Contents

List of Tables	x
List of Figures	xiv
1 Introduction	1
I Function-valued Mappings	5
2 Function-valued Mappings	6
2.1 Image Models	6
2.1.1 Stochastic Modelling	6
2.1.2 Mumford-Shah Model	7
2.1.3 Images as Distributions	7
2.1.4 L^p Images	8
2.1.5 Sobolev Images	8
2.1.6 Images of Bounded Variation	9
2.2 The Function-valued Mapping Approach	9
2.3 Applications	11
2.3.1 Hyperspectral Imaging	11
2.3.2 Diffusion MRI	14
2.3.3 Sensor networks	16
2.4 Definition	18

3	Mathematical Foundation	20
3.1	Calculus	20
3.1.1	Limit	21
3.1.2	Continuity	23
3.1.3	Differentiation	24
3.1.4	Best Linear Approximation	27
3.1.5	Integration	28
3.1.6	Fundamental Theorem of Calculus	33
3.2	L^p Spaces	35
3.2.1	The Dual Space of $L^p(\mathbf{X}; \mathcal{F}(\mathbf{Y}))$	37
3.3	Euler-Lagrange Equation	39
4	Fourier Transform	41
4.1	Some Current Definitions	41
4.2	Fourier Transform for FVMs	43
4.2.1	Properties	45
4.3	Applications	51
5	A Simple Class of Fractal Transforms for Hyperspectral Images	54
5.1	A Complete Metric Space $(\mathbf{Z}, \mathbf{d}_Z)$ of Hyperspectral Images	55
5.2	A Class of Fractal Transforms on $(\mathbf{Z}, \mathbf{d}_Z)$	56
5.3	Inverse problem for fractal transforms on $(\mathbf{Z}, \mathbf{d}_Z)$	58
5.4	Block fractal transforms on digital hyperspectral images	59
5.5	Some Numerical Results	61
5.6	Final Remarks	66

6	Denoising of Hyperspectral Images	69
6.1	Denoising Approaches	71
6.1.1	Denoising of hyperspectral images as a collection of spectra	71
6.1.2	Denoising of hyperspectral images as a collection of bands	72
6.1.3	Denoising of hyperspectral images as a whole	73
6.1.4	Denoising of hyperspectral images as a collection of both bands and spectra	74
6.1.5	Denoising of hyperspectral images using FVMs	75
6.2	Experiments	78
II	SSIM-based Optimization	82
7	Preliminaries of SSIM-based Optimization	83
7.1	Current SSIM-based Optimization Approaches	84
7.2	General Framework for SSIM-based Optimization	86
7.3	SSIM	87
7.3.1	Definition	87
7.3.2	The SSIM as a normalized metric	90
7.3.3	Quasiconvexity and Quasiconcavity	92
8	Constrained SSIM-based optimization	95
8.1	Optimizing $\mathbf{T}(\Phi(\mathbf{x}), \mathbf{y})$	96
8.2	Applications	98
8.2.1	SSIM with Tikhonov constraint	98
8.2.2	\mathbf{L}^1 -constrained SSIM-based optimization	99
8.2.3	SSIM and Total Variation	99
8.2.4	Deblurring	100
8.2.5	Zooming	100
8.3	Experiments	101

9	Unconstrained SSIM-based optimization	107
9.1	The special case $\lambda = \mathbf{0}$	108
9.2	The general case $\lambda > \mathbf{0}$	111
9.2.1	Differentiable $\mathbf{h}(\mathbf{x})$	111
9.2.2	Non-differentiable $\mathbf{h}(\mathbf{x})$	114
9.3	Experiments	117
10	SSIM-based Optimization Using Mollifiers	121
10.1	Smoothing via Mollifiers	122
10.2	SSIM-based Optimization with Sparsity	128
10.3	Experiments	129
11	Conclusions and Future Work	133
	APPENDICES	136
A	Proofs of Chapter 3	137
B	Proofs of Chapter 8	149
C	MATLAB Codes	153
C.1	Algorithm I	153
C.2	Algorithm II	155
C.3	Algorithm III	157
C.4	Algorithm IV	159
	References	171

List of Tables

- 6.1 Numerical results for the different approaches. Numbers in bold identify the best results with respect to each of the four measures of performance considered. In all cases, the PSNR prior to denoising was 30.103 dB. . . . 80
- 8.1 Numerical results for the different approaches and applications. Numbers in bold identify the best results with respect to each measure of performance.103

List of Figures

2.1	Hyperspectral images are sometimes called <i>image cubes</i> . This cube was taken by the AVIRIS hyperspectral sensor made by NASA. The image corresponds to the mining district in Colorado, where the front of the cube shows the areas from acid mine drainage highlighted in orange, yellow and red. Taken from [107].	12
2.2	These spectra are from a hyperspectral image captured by the AVIRIS sensor. The red spectrum is from a pixel that corresponds to the location of a place filled with the mineral Alunite; the black spectrum is from a pixel associated with a location filled with the mineral Kaolinite. Taken from [107].	13
2.3	Diffusion MRI in a patient with a low grade astrocytoma, which is a kind of tumor found in the brain (top). The tumor shows decreased diffusion (arrows). Image taken from http://en.wikipedia.org/wiki/File:AstrocytomaDiffMRI.jpg	15
2.4	The color coded tracks show major nerve cell axons in the white matter (cortex) of a resting brain. An MRI cross sectional image of the brain bisects the diffusion image. Image taken from http://newscenter.berkeley.edu/2011/01/05/functionalmri/	16
2.5	The most common network topologies. Image taken from http://en.wikipedia.org/wiki/File:NetworkTopologies.svg	17
4.1	The Fourier transform of the rectangular function given in Eq. (4.51). . . .	50
4.2	The Fourier transform of the Gaussian function of Example 4.2.2. In this case, $\sigma = \sqrt{3}$	50

4.3	Matrices of correlation coefficients between bands of the hyperspectral image Indian Pines. Image on the left corresponds to the correlations in the spatial domain, whereas the image on the right shows how this correlation between bands tends to be greater in the frequency domain.	52
4.4	Matrices of correlation coefficients between bands of the hyperspectral image Salinas. As in the previous figure, the image on the left correspond to the correlations in the spatial domain, whereas the image on the right shows the correlation between bands in the frequency domain.	53
4.5	The image on the left is the band number 23 of the hyperspectral image Indian Pines. The image on the right corresponds to band number 57 of the hyperspectral image Salinas.	53
5.1	Distributions of errors Δ_{ij} over the interval $[0, 0.02]$ obtained by approximating non-overlapping 8×8 -pixel range blocks $u(R_i)$ with all possible decimated 16×16 -pixel domain blocks of AVIRIS hyperspectral image. Top: Using affine mapping, Eq. (5.13). Bottom: No affine mapping, Eq. (5.21).	63
5.2	α_i parameters over the interval $[-1, 1]$ employed in the “best fit” of AVIRIS HS image range subblocks $u(R_i)$	64
5.3	Some channels of the attractor \bar{u} of the fractal transform T obtained by fractally coding the AVIRIS hyperspectral image using 8×8 -pixel range blocks and 16×16 -domain blocks.	65
5.4	8×8 -pixel block approximations to Channel 20 of the AVIRIS hyperspectral image.	66
5.5	Correlations C_{kl} between channels of AVIRIS hyperspectral image, demonstrating the existence of several subgroups of highly correlated channels.	68
5.6	Correlations C_{kl} between channels of hyperspectral fern image, which demonstrates a lesser degree of grouping of correlated channels.	68
6.1	Visual results for Band No. 23 of the Indian Pines HS image. Beside the original (noiseless) image in the lower row are shown the various reconstruction results. Beside the noisy image in the upper row are shown the corresponding SSIM maps between the reconstructed (denoised) images and the original image.	80

6.2	Denoising results for a particular spectral function of the Indian Pines HS image. In the top row, for visual comparison, the reconstructions (blue plots) and noisy spectra (red plots) are shown. The original spectra (green plots) along with the corresponding reconstructions (blue plots) can be observed in the bottom row.	81
7.1	This is the plot of $T(x, 3)$ for the one-dimensional case with $C = 0$. Notice that when $x \geq 0$ the dissimilarity measure is quasiconvex. Indeed, the sub-level sets are convex as long as $T(x, 3) \leq \alpha$, for any $\alpha \in [0, 1]$. In a similar manner, the super-level sets are convex if $1 \leq T(x, 3) \leq 2$, which implies that the function is quasiconcave over that region of the real line.	94
7.2	The plot of $T(x, y)$ in the two-dimensional case. Once again, $C = 0$. Here, $x = [x_1, x_2]^T$ and $y = [3, 2]^T$. As it was shown, the line $3x_1 + 2x_2 = 0$ delimits the regions of quasiconvexity and quasiconcavity of $T(x, y)$	94
8.1	Visual results for the sparse reconstructions. In this case, for each pixel block, the maximum allowed value for the L^1 norm of the coefficients that are to be recovered is 1. In the top row, SSIM maps are shown. Original and recovered images can be seen in the bottom row.	104
8.2	Visual results for the deblurring application. In the top row, the blurred and noisy image along with the SSIM maps are presented. As above, the recovered and original images are seen in the bottom row.	105
8.3	Visual results of the zooming experiments. The low resolution image and the SSIM maps can be seen in the top row. Original image along with the SSIM and L^2 reconstructions are shown in the bottom row.	106
9.1	The first two plots from left to right correspond to the average SSIM versus the L^0 norm of the recovered coefficients for the test images <i>Lena</i> and <i>Mandrill</i> . In the last two plots, a visual comparison between the original and recovered coefficients from a particular block of the <i>Lena</i> image can be appreciated. Regularization is carried out so that the two methods being compared induce the same sparseness in their recoveries. In the two shown examples, the same block was processed but subjected to different amounts of regularization. In particular, the L^0 norm of the set of DCT coefficients that were recovered by both the proposed method and ST is 3 for the first example (third plot), and 15 for the second (fourth plot).	119

9.2	Visual results for a patch from the test image <i>Lena</i> . In all cases, the L^0 norm of the recovered DCT coefficients for each non-overlapping block is 13. In the upper row, the SSIM maps between the reconstructions and the original patch are shown. Reconstructed and original patches can be seen in the lower row. The average $T(Dx, y)$ of all non overlapping blocks for the proposed method is 0.8864, whereas for ST is 0.8609.	120
10.1	The plot on the left shows an example of the different solutions that were obtained by the three methods that were compared. The plot on the right shows a visual example of how a sequence of minimizers x_ϵ^* of the mollified SSIM-based optimization problem (10.5) converges to a minimizer x^* of the non-smooth problem (10.1).	131
10.2	Visual results for a sub-image from the test image <i>Lena</i> . In all cases, regularization is carried out to induce the same degree of sparsity for all methods at each non-overlapping pixel block. In the bottom row the original image and its reconstructions are shown. The corresponding SSIM maps can be seen in the top row.	132

Chapter 1

Introduction

There are two important paradigms in image processing upon which imaging tools are developed: (i) representation of images as valued-vector functions and (ii) using the Euclidean distance as either a metric or a norm for images. The former offers a good framework for applications since some type of images, e.g., colour and hyperspectral images, can be well represented by vector-valued functions [60, 81]. As for the latter, this is a convex and mathematically tractable metric that is frequently used in a variety of imaging tasks in which an optimization problem is to be solved (e.g., deblurring, denoising, sparse reconstruction, etc.) [13, 12].

Although many efficient imaging tools have been developed which are based on these paradigms [24, 15, 4, 8], these approaches have some limitations. For instance, consider the case of a video sequence, which can be represented as a vector-valued function $u : \Omega \rightarrow \mathbb{R}^n$. Here, Ω is the spatial domain, and each component u_i of the range of u , $1 \leq i \leq n$, is a time frame. Due to the finite dimensionality of the range, it is not possible to propose a suitable space of functions as a model for the class of time functions defined at each pixel location of the video. A similar situation is encountered when complex data sets such as hyperspectral images are to be modelled.

As for the Euclidean distance, measures of visual quality such as the Mean Square Error (MSE) and Peak to Signal Noise Ratio (PSNR) are Euclidean-based metrics, however, it has been shown that these metrics are not the best choices when it comes to measure the visual quality of images. This drawback of Euclidean-based metrics has been overcome up to some point by the Structural Similarity Index Measure (SSIM), which has proved to be a good model of the Human Visual System (HVS) [122]. For this reason, the SSIM has been employed as a fidelity term in several types of optimization problems for carrying

out a variety of imaging tasks—e.g., filtering, best-SSIM approximation, video coding, etc. [18, 29, 118]. Nevertheless, these methods tend to be developed for particular applications and are not general.

In this thesis we propose two alternative approaches to overcome the difficulties mentioned above, namely, *Function-valued Mappings* (FVMs) and *SSIM-based Optimization*. With FVMs we address the problem of modelling complex data sets properly. In this case, we define a FVM as follows:

$$u : X \rightarrow \mathcal{F}(Y), \tag{1.1}$$

where X is the support of the FVM (for digital images, the “pixel space”) and $\mathcal{F}(Y)$ is a Banach space of either real- or complex-valued images supported on the set Y . Observe that the range of u is infinite dimensional, therefore, depending on the application, an appropriate space $\mathcal{F}(Y)$ can be chosen to describe the functions that are assigned to each $x \in X$. For instance, a hyperspectral image can be represented as a FVM of the form $u : X \subset \mathbb{R}^2 \rightarrow L^2(\mathbb{R})$, where $L^2(\mathbb{R})$ is the space of square integrable functions supported on the real line—that is, we are interested in the spectral functions that have finite energy. Moreover, spaces of FVMs can be defined, such as the *Lebesgue-Bochner L^p spaces* [40], therefore, some of the classical theorems of functional analysis can be employed for the development of imaging tools based on this approach. In particular, we introduce a Fourier transform and a new class of fractal transforms for FVMs, which we employ to analyze and process hyperspectral images.

It must be mentioned that the FVM approach is not a novelty in other fields such as partial differential equations [126], harmonic analysis [97, 85], statistics [3], and others [40]. In fact, FVMs are known in the mathematical community as *Banach-valued functions*, the latter being studied mainly by analysts who have been interested in seeing if the classical results of real-valued functions still hold in the Banach-valued setting [40, 23, 39]. Nevertheless, in imaging, this methodology has been barely explored. Despite this, some contributions can be found which employ the concept of a function taking values in an infinite dimensional Banach space. For instance, in an effort of closing the gap between the mathematical formalism of Banach-valued functions and practical applications in imaging, in [84], the authors use the FVM approach to provide a solid mathematical platform to describe and treat diffusion magnetic resonance images. Also, in [75], an analogue of FVMs is introduced, namely, measure-valued images, which are well suited for non-local image processing. Indeed, non-local means denoising [21] and fractal image coding [72] are the two applications that are addressed in [75] using this measure-valued methodology. Given the latter, rather than introducing state-of-the-art algorithms for imaging applications, we propose the FVM approach as a mathematical framework that may offer interesting possibilities for the image processing community.

Regarding SSIM-based optimization, we present a general framework for optimization problems that involve the SSIM as a fidelity term. Such a framework is divided into two approaches, namely, *unconstrained and constrained SSIM-based optimization*. In the unconstrained approach, an optimization problem is defined as follows:

$$\min_x \{T(\Phi(x), y) + \lambda h(x)\}, \quad (1.2)$$

where $T(\cdot, \cdot) := 1 - \text{SSIM}(\cdot, \cdot)$ is a SSIM-based dissimilarity measure, Φ is a linear transformation, y a given observation, and $h(x)$ a convex regularizing term along with its regularization parameter λ . As for the constrained counterpart, the minimization problem is given by

$$\begin{aligned} \min_x \quad & T(\Phi(x), y) \\ \text{subject to} \quad & h_i(x) \leq 0, \quad i = 1, \dots, m \\ & Ax = b, \end{aligned} \quad (1.3)$$

where the $h_i(x)$ define a set of convex constraints, and $Ax = b$ is an equality constraint. We also introduce a set of algorithms to solve both (1.2) and (1.3). The advantage of this framework is that it allows us to involve the SSIM in optimization problems in which it had not been employed before, in particular, SSIM optimization with L^1 -norm regularization. Moreover, we complement this general framework with the usage of *mollifiers* [50] for solving non-smooth SSIM-based optimization problems with classical gradient-based methods. Also, we provide experimental results to contrast the performance of these novel algorithms with the classical Euclidean-based approaches.

Regarding the structure of the thesis, this is divided into two parts. The first part is about the FVM methodology, which is composed by five chapters. In the first of these chapters, Chapter 2, some of the current approaches to imaging are reviewed and the FVM approach is presented in more detail. Further, the mathematical definition of a FVM is also presented. To continue with the mathematical formalism, in Chapter 3, we provide a brief review of part of the mathematical theory of Banach-valued functions within the context of FVMs. Further, the mathematical concepts that are used throughout this first part of the thesis are also included in this chapter. Chapters 4 and 5 correspond to the Fourier and the fractal transforms for FVMs respectively. In Chapter 4, definition and properties of the Fourier transform for FVMs are provided, which is a special case of the Fourier transform for Banach-valued functions. In Chapter 5, a fractal transform for hyperspectral images is introduced as well as a complete space of FVMs in which such a transform is a contractive mapping. Finally, in Chapter 6, a review of several approaches for denoising

hyperspectral images is provided. In particular, we contrast some classical methods with a FVM-based denoising approach.

The second part of the thesis is composed of four chapters and deals with the subject of SSIM-based optimization. In the first chapter of this second part, Chapter 7, the classical Euclidean-based methodology is contrasted with some of the current SSIM-based approaches that are found in the literature. Also, we present formally our general framework for SSIM-based optimization and provide a definition of the SSIM along with some of its properties. Following this chapter, the subject of constrained SSIM-based optimization is presented in Chapter 8. An algorithm that solves problem (1.3) is introduced along with several SSIM-based imaging applications. Also, in Chapter 9, unconstrained SSIM-based optimization is addressed. In particular, two algorithms that solve both the differentiable and non-differentiable counterparts of problem (1.2) are formulated. Ultimately, in Chapter 10, we address the problem of solving non-smooth SSIM-based optimization problems by means of a special class of smooth functions known as mollifiers. Experimental results are provided in the last sections of Chapters 8, 9 and 10.

Part I

Function-valued Mappings

Chapter 2

Function-valued Mappings

In the first section of this chapter, we give a brief review of some of the classical image models that are employed in image processing, which we contrast with the usage of FVMs as representations of images. Moreover, some potential applications in which such mappings can be employed are also presented. We conclude the chapter by introducing a formal definition of the concept of FVM.

2.1 Image Models

A key task in image processing is to find proper mathematical models of images. Over time, distinct approaches have been proposed, which usually differ in how these translate into mathematical terms what is understood as an image. What model is the best choice depends on the type of visual data at hand and the image processing task that is to be carried out [25]. In the following sections some of the most known image models are reviewed.

2.1.1 Stochastic Modelling

In some cases, images are interpreted as the result of one or more underlying random processes. This is commonly seen in natural scenery, where the different elements in the landscape haven been shaped by distinct random factors: mountains “carved” by the erosion produced by both wind and water; the formation of clouds, which is affected by the changes in temperature in the atmosphere, the evaporation of water and winds; the

patterns of a distant forest, which are usually created by the way the trees have been distributed over the soil; and so on. As such, it makes sense to model the probability distributions that govern the random processes associated to a particular image—this is precisely the main goal of stochastic modelling. Methods such as Gibbs’ ensembles, Markov random fields, visual filters and others aim to this goal [25, 96]. These are well suited for representing not only images of natural scenes, but also textures and tissues [25].

2.1.2 Mumford-Shah Model

Although the stochastic approach gives special attention to details and textures, it is usually seen that images can be well approximated by either piecewise constant or smooth functions—this fact is the main idea behind the well known Mumford-Shah model. In this approach, an image is modelled as a collection of elements with almost uniform textures and well defined boundaries. Given this, an approximation of a given image f supported on Ω is obtained by minimizing the Mumford-Shah functional [92, 25]:

$$E[u, \Gamma] = \alpha \int_{\Omega \setminus \Gamma} (K(u) - f)^2 dx + \beta \int_{\Omega \setminus \Gamma} \|\nabla u\|^2 dx + \gamma \mathcal{H}^1(\Gamma), \quad (2.1)$$

where K is a blur operator, Γ is an edge set containing the boundaries of the objects present in the image, $\mathcal{H}^1(\cdot)$ is the one-dimensional Hausdorff measure, and α , β and γ are regularization parameters. It can be noted in Eq. (2.1) that non-smooth piecewise functions and excessive long boundaries are penalized, as well as solutions that are not close to the observed image f . Notice that this model is well suited for smoothing and segmentation.

2.1.3 Images as Distributions

Along the lines of deterministic models we have the interpretation of images as distributions. In this instance, images are linear functionals acting on an space of test functions. That is, given an image u supported on Ω , where Ω is an open and bounded subset of \mathbb{R}^2 , and a set of test functions defined as

$$D(\Omega) = \{\phi : \phi \in C_c^\infty(\Omega), \text{supp } \phi \subseteq \Omega\}, \quad (2.2)$$

we have that u is a linear functional on $D(\Omega)$:

$$u : \phi \rightarrow \langle u, \phi \rangle, \quad (2.3)$$

where $\langle \cdot, \cdot \rangle$ is an inner product. This implies that the space of images is $D^*(\Omega)$, i.e., the dual space of $D(\Omega)$. This interpretation is quite useful since it allows to handle in formal mathematical terms images that are not functions in the classical sense, e.g., a bright point at the origin ($u(x, y) = \delta(x, y)$) or a uniform bright line ($u(x, y) = \delta(y)$) [89, 25].

2.1.4 L^p Images

Although distributions can represent basically any kind of image, this model tends to be too general, thus it is more difficult to extract key features of images under this interpretation. It is, therefore, desirable to introduce models that have more structure—one example of this is the case of L^p images. Mathematically speaking, an image u with support Ω belongs to the space $L^p(\Omega)$ if its L^p norm is finite. In other words, the space that contains all L^p images is the set given by

$$L^p(\Omega) = \{u : \|u\|_p^p < \infty\}, \quad (2.4)$$

where

$$\|u\|_p = \left(\int_{\Omega} |u(x)|^p dx \right)^{\frac{1}{p}}. \quad (2.5)$$

One key feature of this model is that any element in $L^p(\Omega)$ has to be bounded with respect to the L^p norm. Indeed, this characteristic leads to the fact that

$$L^q(\Omega) \subseteq L^p(\Omega), \quad \forall q \geq p \geq 1, \quad (2.6)$$

as long as Ω is bounded; therefore, $L^1(\Omega)$ is the most general class of L^p images [25]. Furthermore, when $p = 2$, we have the special case $L^2(\Omega)$, which is a Hilbert space endowed with an inner product. This important characteristic allows to, for instance, represent any element in $L^2(\Omega)$ as a linear combination of orthonormal basis functions.

2.1.5 Sobolev Images

There are other models that are even more specific than the L^p class. Among these we have Sobolev images, which are denoted as $W_p^n(\Omega)$; that is, the images whose n -th order distributional derivatives have finite L^p norm. In the special case when $n = 1$ and $p = 2$, we have that an image u belongs to $W_2^1(\Omega)$ as long as its first order derivatives belong to $L^2(\Omega)$ in the distributional sense. The norm used for $W_2^1(\Omega)$ is given by

$$\|u\|_{W_2^1} = \left(\|u\|_{L^2}^2 + \|\nabla u\|_{L^2 \times L^2}^2 \right)^{\frac{1}{2}} < \infty. \quad (2.7)$$

In general, the norm of a Sobolev image that belongs to $W_p^n(\Omega)$ is defined in a similar fashion, with the difference that it has higher order distributional derivatives [25].

2.1.6 Images of Bounded Variation

Another important model for images with a nice structure is the space of images with bounded variation $BV(\Omega)$ [105, 24, 25]. In the distributional sense, an L^1 image is in $BV(\Omega)$ if and only if its distributional gradient Du satisfies

$$\|Du\| = \sup_{\phi \in \Phi} \langle \nabla u, \phi \rangle < \infty, \quad (2.8)$$

where Φ is defined as

$$\Phi = \{\phi : \phi \in C_c^\infty(\Omega, \mathbb{R}^2), \|\phi(\mathbf{x})\|_2 \leq 1 \ \forall \mathbf{x} \in \Omega\}. \quad (2.9)$$

If such supremum exists, then this is called the Total Variation (TV) of u , which is denoted as $\|Du\|$. The norm of $BV(\Omega)$ is defined as

$$\|u\|_{BV} = \|u\|_1 + \|Du\|. \quad (2.10)$$

One of the relevant attributes of this model is that images in $BV(\Omega)$ may have edges, which are an important visual feature in images, whereas Sobolev images do not. Also, irregularities, which are usually caused by noise, are penalized by the norm of this space. Notice that L^p images might have edges as well, however, the L^p norm is not a good measure of oscillatory irregularities¹[25].

2.2 The Function-valued Mapping Approach

Other important models are the level set representation of images, introduced by Osher and Sethian in [94], and wavelets and multi scale representation [79, 117, 25]. It is worth to point out that these models, and most of the models mentioned above, rely on the interpretation of an image as a function, which is the most common representation found in the literature [79, 62, 24, 25]. For instance, a grey scale image is usually described as a real-valued function u that assigns to each point in its domain an intensity value that belongs to its range. In mathematical terms, this is usually written as follows:

$$u : \Omega \rightarrow Y, \quad (2.11)$$

¹For functions with support $\Omega = (0, \pi)$, $\|\sin(nx)\|_2$ tends to $\sqrt{\pi/2}$ as $n \rightarrow \infty$, whereas $\|\sin(nx)\|_{BV}$ diverges as n goes to infinity.

where Ω and Y are subsets of \mathbb{R}^2 and \mathbb{R} , respectively. To represent colour images this definition is extended to vector-valued functions. In this case, the set Y becomes a vector space well suited for the colour space that is being employed. In the case of RGB, Y may be defined as a subset of \mathbb{R}^3 .

Although functions—along with their generalized counterpart: distributions—are the foundation of a large class of image models and image processing tasks, these may not be sufficient to describe in all their richness more complex visual objects such as magnetic resonance images or hyperspectral images. For instance, hyperspectral images are usually represented as vector-valued functions [120, 83, 81], nevertheless, the finite dimensionality of the range of this representation does not allow us to model properly the spectral functions that are found at each spatial location. This difficulty can be overcome by introducing the concept of function-valued mapping (FVM), which is a mathematical relation that assigns to every element in its domain a function that belongs to a certain space. In more formal terms, a FVM u is defined as follows:

$$u : X \rightarrow \mathcal{F}(Y), \quad (2.12)$$

where both X and $\mathcal{F}(Y)$ are Banach spaces. In particular, $\mathcal{F}(Y)$ is a space of functions whose support is the set Y .

The FVM approach offers several advantages over the classical representations of images. First, it can be employed in a wide range of applications such as diffusion MRI (Magnetic Resonance Imaging) [84], sensor networks, video [118] and hyperspectral imaging [115, 116]. Also, a proper function space $\mathcal{F}(Y)$ can be defined to model more faithfully the class of functions that are embedded in these data sets (e.g., the spectral functions of hyperspectral images). Moreover, functions (e.g., grayscale images) and vector-valued functions (e.g., hyperspectral images) can be seen as samples of a FVM, thus (2.12) can be considered as a generalization of these classical representations: in other words, given a FVM $u : X \rightarrow \mathcal{F}(Y)$, the i -th entry of a vector-valued image $v : X \rightarrow \mathbb{R}^n$ can be defined as

$$v_i(x) = u(x)(z_i) = \int_Y u(x)\delta(y - z_i)dy, \quad (2.13)$$

where each $z_i \in Y$, $1 \leq i \leq n$, is a given element of Y , and $\delta(\cdot)$ is the Dirac delta distribution. By setting $n = 1$, one obtains the representation of images as real-valued functions.

Indeed, in [84], this approach has already been used to provide a solid mathematical foundation for treating and modelling diffusion MRI images. Also, in [115, 116], FVMs

are used as representations of hyperspectral images for studying the self-similarity of hyperspectral data [115] and defining a simple class of fractal transforms [6, 7] in a complete metric space [116].

Given the latter, it is worthwhile to discuss in more detail how the FVM approach can be employed to represent the different kinds of signals that are found in some applications. This is the subject of the next section.

2.3 Applications

Although in this study we focus our attention on hyperspectral imaging, in this section we review some of the applications in which the FVM approach can be employed. The definitions that may be used to represent the distinct data sets that are mentioned are also provided.

2.3.1 Hyperspectral Imaging

In a few words, hyperspectral imaging may be defined as image processing across the electromagnetic spectrum. In this field, the hyperspectral sensors that collect the information divide the spectrum into many narrow bands, and all the data is usually organized in a three-dimensional matrix [26, 107]. A processed image of this matrix is shown in figure 2.1.

Hyperspectral images are produced by sensors called *imaging spectrometers*. Their development comes from two related but distinct technologies: *spectroscopy* and *remote imaging* of the Earth. Spectroscopy is the study of light that is reflected or emitted by objects and its variation in energy with wavelength. It deals with the spectrum of sunlight that is reflected by an object in the surface of the Earth. Remote imaging has to do with the measure of light reflected from many areas on the Earth surface. Its name comes from the fact that the images are taken from an object that is at considerable distance from the Earth, usually a satellite [99].

In spectroscopy, the fundamental property that is measured is *spectral reflectance*, which is the ratio of reflected to incident energy as a function of wavelength. In general, reflectance varies with wavelength because energy at certain frequencies is absorbed in different degrees. For instance, visible light is reflected by many objects, but X-rays are usually absorbed by solid objects that reflect visible light [99, 107].




Figure 2.1: Hyperspectral images are sometimes called *image cubes*. This cube was taken by the AVIRIS hyperspectral sensor made by NASA. The image corresponds to the mining district in Colorado, where the front of the cube shows the areas from acid mine drainage highlighted in orange, yellow and red. Taken from [107].

Different materials have different spectral reflectances. If a graph of the reflectance versus wavelength is plotted, it is seen that each material has a different graph. This graph is sometimes called the *spectral signature* of the material. In general, there are regions in the spectrum where a material selectively absorbs the incident energy. These regions are commonly called *absorption bands*. The shape of the signature of a material, together with the position and “depth” of the absorption bands, are useful features to identify different materials. In figure 2.2, the spectra of two minerals may be seen [99, 76].

Applications

Originally hyperspectral imaging was developed for mining and geology since it is ideal to identify different types of minerals or fossilized organic materials such as ore and oil. However, given that it is easy to identify materials using their spectra, hyperspectral imaging is now used in many fields such as ecology, mineralogy, hydrology, surveillance, archeology, among others.

In hyperspectral applications, the following objectives are usually pursued: Target detection, material mapping, material identification and mapping details of surface properties [107]. Even though hyperspectral imaging may be used for the inspection of harvested products, pharmaceutical studies [63], or for MRI [26], we will focus on the description of some applications that are classified as hyperspectral remote sensing [107]. These are grouped according to their objectives.




Figure 2.2: These spectra are from a hyperspectral image captured by the AVIRIS sensor. The red spectrum is from a pixel that corresponds to the location of a place filled with the mineral Alunite; the black spectrum is from a pixel associated with a location filled with the mineral Kaolinite. Taken from [107].

Target detection: The aim in this application is to locate known target materials. For example, hyperspectral imaging has been used to detect military vehicles under partial vegetation, and also to detect small military objects within relatively larger pixels. Another application is the identification of different vegetation species, and also the detection of vegetation stress and disease.

Material mapping: This application is also performed when the materials present in the scene are known beforehand. For example, hyperspectral images have been used by geologists for mapping economically interesting minerals. They have also been used to map heavy metals and other toxic wastes within mine tailings in active and historic mining districts including superfund sites.

Material identification: The aim of this application is to determine the unknown materials that are present in the scene. This analysis may also be accompanied by the geographical location of the materials throughout the image. For instance, a project in Australia used hyperspectral imaging to identify roofs susceptible to hail damage [107]. The overall shape of the spectral curve was detected and also the position and strength of distinguishing absorption features in these roofing materials. These spectral characteristics provided information that was useful to identify the locations more susceptible to hail damage.

Mapping details of surface properties: In this case, the details of surface properties that are undetectable using other types of imaging are studied. For instance, hyperspectral images have been used to detect soil properties including moisture, organic content, and salinity.

Hyperspectral image as a function-valued mapping

As expected, a hyperspectral image may be understood as a discrete version of a function valued mapping. From this perspective, the discrete data that is organized in a three-dimensional matrix may be seen as a continuous mapping that assigns to every element in its domain a spectrum or a function of wavelength. Moreover, the FVM approach allows us to choose a function space as a model for these spectral functions. For instance, a natural choice is $\mathcal{F}(Y) = L^2(\mathbb{R})$; that is, the space of functions with finite energy. Mathematically this may be written as follows:

Definition 2.3.1. *A hyperspectral image can be represented by the function-valued mapping $u : X \subset \mathbb{R}^2 \rightarrow L^2(\mathbb{R})$, where $L^2(\mathbb{R})$ is the space of square integrable functions supported on the real line.*

2.3.2 Diffusion MRI

Magnetic resonance imaging (MRI) is a medical imaging technique usually used to produce detailed images of the internal biological structures of the human body. MRI scanners produce a strong magnetic field that changes the magnetic configuration of some of the atoms in the body. Very briefly, this causes the atoms in the body to rotate in this field. When the field is momentarily turned off, these rotations decay to zero, but at different rates according to the different environments in which the nuclei are found. From these differing rates, an image of the biological tissue may be obtained. One advantage of MRI is that provides good contrast for the images of soft tissues such as brain, muscles, nerves, among others—i.e., the tissues that are not considered bone [68].

In contrast with MRI, diffusion MRI (dMRI) is a method based on the MRI technology which measures the diffusion of water molecules in biological tissues. In a glass of water for example, water molecules move randomly (without taking into account temperature gradients or turbulence), but in a biological tissue, the diffusion of the water molecules depends strongly on the structure of the tissue. For instance, under normal conditions, a water molecule in the axon² of a neuron tends to move inside the neural fiber since

²The axon is a long projection of a nerve cell or a neuron that conduces electrical impulses.

it is not very likely that it will cross the myelin membrane³. For this reason, dMRI is a key technology determining the connections of the nervous system [65][68]. Its main applications are *medical diagnosis* and *connectomics*.

Applications

Medical diagnosis: Its relevance in medical diagnosis is due to the fact that various brain pathologies may be best detected by looking at particular measures of anisotropy and diffusivity [65, 68]. It can be seen in figure 2.3 how a tumor can be detected by means of a MRI scanning.




Figure 2.3: Diffusion MRI in a patient with a low grade astrocytoma, which is a kind of tumor found in the brain (top). The tumor shows decreased diffusion (arrows). Image taken from <http://en.wikipedia.org/wiki/File:AstrocytomaDiffMRI.jpg>.

Connectomics: This one is the science that assembles and analyzes comprehensive maps of the neural connections of the brain, usually know as *connectomes*. The principal object of study is the brain, but any neural connection could be mapped. It is believed that as *genomic*⁴ brought many advances to genetics, *connectomic* may benefit greatly different fields, mainly *neuroscience*⁵. An example of an image showing some connections of a resting brain can be seen in figure 2.4.

³It is a layer with dielectric properties that covers the axon.

⁴Discipline on genetics concerned with the study of the genomes of organisms.

⁵The scientific study of the nervous system.




Figure 2.4: The color coded tracks show major nerve cell axons in the white matter (cortex) of a resting brain. An MRI cross sectional image of the brain bisects the diffusion image. Image taken from <http://newscenter.berkeley.edu/2011/01/05/functionalmri/>.

dMRI image as a function-valued mapping:

Intuitively, it may be seen that the way water diffuses in a biological tissue depends on the location of the water molecules. Given this, one also would think that at each point a function that describes the diffusion of water exists. Then, a dMRI image can be understood as a function-valued mapping that assigns to every element in its three-dimensional domain a function that measures the diffusivity of water in every direction. This can be defined as follows:

Definition 2.3.2. *A dMRI image may be represented by a mapping $u : X \subset \mathbb{R}^3 \rightarrow \mathbb{L}^2(\mathbb{S}^2)$, where $\mathcal{F}(Y) = \mathbb{L}^2(\mathbb{S}^2)$ is the space of functions supported on the three-dimensional unit sphere \mathbb{S}^2 . Such a space is a good model for functions that contain the information of the diffusivity of water in every direction in \mathbb{R}^3 .*

2.3.3 Sensor networks

A sensor network (SN) is a collection of autonomous sensors spatially distributed that detect and collect data from a variety of environments. The sensors usually monitor or measure physical quantities in a geographical area and can be connected to one or several

sensors in the network. The fashion in which the sensors are connected is called *the topology of the network*. There are different network topologies, each one with its own characteristics and advantages, but all of them represent how the information is flowing through the network. In some cases a network may have a hybrid topology, which is a combination of two or more topologies that create a new one. In figure 2.5 different network topologies are shown.




Figure 2.5: The most common network topologies. Image taken from <http://en.wikipedia.org/wiki/File:NetworkTopologies.svg>.

Nowadays, thanks to advances in hardware and wireless network technologies, low-cost, low-power, multi-functional miniature sensor devices are available. For this reason, many SNs are connected via a wireless connection. This special type of SNs are known as wireless sensor networks (WSN) and are of a major importance in the field. Also, these advances have made SNs more reliable, accurate, cheap and flexible, which in turn have allowed SNs to be used in a larger number of applications [112].

Applications

Since a SN collects data from a geographical area or space, this may monitor the space itself, the elements that are present in the area that is being monitored, or the interactions between these elements. Given this, the wide range of applications of a SN can be divided roughly into three groups [36]:

Monitoring space: environmental and habitat monitoring, precision agriculture, indoor climate control, surveillance, among others.

Monitoring elements in a space: structural monitoring, condition-based equipment maintenance, urban terrain mapping, etc.

Monitoring interactions among the elements of a space: wildlife habitats, disaster management, emergency response and more.

SN as a function-valued mapping

The information collected by a sensor may be understood as a function of time, therefore, a SN can be seen as a graph where at each node there is a function that represents the data gathered by each sensor. Then, a SN may be represented as a function-valued mapping that assigns to each node in the graph a function. This can be defined more formally as follows:

Definition 2.3.3. *A SN may be represented a function-valued mapping $u : X \rightarrow \mathcal{F}(Y)$, where X is a graph and $\mathcal{F}(Y)$ may be the space of square integrable functions that are supported on a subset of \mathbb{R}_+ .*

2.4 Definition

We define a FVM u as a mathematical relation that assigns to every element x in its domain X a function f that belongs to a function space $\mathcal{F}(Y)$, where X and Y , unless otherwise stated, are compact sets in \mathbb{R}^n . That is,

$$u : X \rightarrow \mathcal{F}(Y). \quad (2.14)$$

Here, the pairs $(X, \|\cdot\|_X)$ and $(\mathcal{F}(Y), \|\cdot\|_{\mathcal{F}(Y)})$ are Banach spaces and the elements of $\mathcal{F}(Y)$ are either real or complex-valued functions supported on Y . These spaces are defined as follows:

$$X = \{x \in \mathbb{R}^n : \|x\|_X < \infty\}, \quad (2.15)$$

and

$$\mathcal{F}(Y) = \{f : Y \rightarrow \mathbb{R}, \mathbb{C} : \|f\|_{\mathcal{F}(Y)} < \infty\}. \quad (2.16)$$

Since the domain of u will be usually defined as a compact set of \mathbb{R}^n , we have that X has an *interior* and a *closure*. The interior of X is defined as

$$\text{int}(X) = \{x \in X : \exists B_r(x) \subseteq X, r > 0\}, \quad (2.17)$$

where $B_r(x) = \{\hat{x} \in X : \|\hat{x} - x\|_X < r\}$. The closure of X is denoted as \bar{X} and is given by

$$\bar{X} = \left\{ \hat{x} \in \mathbb{R}^n : \hat{x} = \lim_{n \rightarrow \infty} x_n, x_n \in X, n \in \mathbb{N} \right\}, \quad (2.18)$$

that is, the set that has all the limits of the sequences that are in X . Both the interior and the closure of X are useful for defining its boundary, which we denote as $\partial(X)$. This is defined as

$$\partial(X) = \bar{X} \setminus \text{int}(X); \quad (2.19)$$

or in other words, $\partial(X)$ is the relative complement of $\text{int}(X)$ respect to \bar{X} . If X is open, we simply have that $\partial(X) = \bar{X} \setminus X$. Throughout this first part of the thesis, (2.19) is the definition of boundary that we will use for sets.

Chapter 3

Mathematical Foundation

The mathematical foundation that is presented in this chapter consists of basically three parts, namely, calculus, L^p spaces of FVMs and the Euler-Lagrange equation. In each one of these, some of the mathematical concepts that are already found in the literature are presented within the context of FVMs. The literature on functions that assume values in a Banach space is vast, so it is not our objective to give a comprehensive review of this field; nevertheless, it is intended to introduce the concepts that are considered necessary for having a mathematical cornerstone solid enough upon which some imaging tools can be developed.

3.1 Calculus

Although the origins of calculus go back in time at least more than two thousand years to the ancient Greeks, who considered the problem of the area of certain figures, it is in the XVII century when finally it was discovered the relation that exists between the problem of finding the tangent of a curve, which was first analyzed by Pierre Fermat, and the problem of area [108]. This realization lead to the birth of infinitesimal calculus, which is evidently one of the greatest accomplishments of the XVII century. From it new branches of mathematics sprouted, such as differential equations, differential geometry, calculus of variations and many others [71]. Given this, it is clear why calculus plays a central role in mathematics and why it is usually the first building block of other mathematical fields. For these reasons, it is the first subject that is presented in the first section of this chapter¹.

¹Some of the ideas presented in this section follow the same approach that can be found in [80] and [30].

The reader who is familiar with the calculus of functions that take values in a Banach space may skip this section.

3.1.1 Limit

The concept of limit is fundamental for the study of functions and an important tool in Calculus and other branches of mathematics. It is not only used to define derivatives, but also for finding asymptotes, calculating improper integrals, and for studying other important features of functions [80]. The basic idea of limit can be traced back to Isaac Newton, but it was not until the XIX century when the german mathematician Karl Weierstrass introduced the modern definition of limit as we know it today [108]. This classical definition is valid for real-valued functions and is as follows:

Definition 3.1.1. *Let $u : \Omega \rightarrow \mathbb{R}$ be a function defined on the open set $\Omega \subset \mathbb{R}$. Let x_0 be a point that belongs to either Ω or its boundary $\partial(\Omega)$. We say that the limit of u is L as x tends to x_0 , which we write as*

$$\lim_{x \rightarrow x_0} u(x) = L,$$

if and only if for all $\varepsilon > 0$ there exists $\delta > 0$ such that for $x \in \Omega$ that satisfies $0 < |x - x_0| < \delta$, we have that $|u(x) - L| < \varepsilon$.

Given the importance of this concept, it becomes evident that a definition of limit for FVMs is necessary for defining the basic mathematical tools of the FVM approach. For this reason, the concept of limit for FVMs is introduced in the following section.

Definition of limit

It is tempting at this point to use 3.1.1 as a “stepping stone” and continue directly to define what it means when one says that the limit of a FVM u is the function f when one is approaching a limit point x_0 in the domain of u ; however, before going any further, we introduce first the following definition and prove that there exists an equivalent $\varepsilon - \delta$ version of it for FVMs.

Definition 3.1.2 (Definition of limit for a FVM). *Let $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$. Let x_0 be a point that belongs to either X or its boundary $\partial(X)$, and let N_f be an open neighbourhood of $f \in \mathcal{F}(Y)$. We say that f is eventually in N_f as $x \in X$ approaches x_0 in the $\|\cdot\|_X$ sense*

if there exists an open neighbourhood N_{x_0} of x_0 such that $x \neq x_0$, $x \in (X \cap N_{x_0})$, implies that $f \in N_f$. We say that **u tends to f when x tends to x₀**, which we write as

$$\lim_{x \rightarrow x_0} u(x) = f,$$

when, for any neighbourhood N_f of f , $u(x)$ is eventually in N_f as x approaches x_0 . If, as x tends to x_0 , $u(x)$ does not approach to any particular function $f \in \mathcal{F}(Y)$ in the $\|\cdot\|_{\mathcal{F}(Y)}$ sense, then we say that **the limit of u as $x \rightarrow x_0$ does not exist**.

Notice that in definition 3.1.2 we have stated that $x \neq x_0$, which is the standard convention when the limit of either a real or a vector-valued function is being defined [108, 80]. Now we prove in the following theorem that definition 3.1.2 has an equivalent $\varepsilon - \delta$ counterpart.

Theorem 3.1.1. *Let $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ and let x_0 be a point in either X or its boundary $\partial(X)$. Then $\lim_{x \rightarrow x_0} u(x) = f$ if and only if for all $\varepsilon > 0$ there exists a $\delta > 0$ such that for $x \in X$ that satisfies $0 < \|x - x_0\|_X < \delta$, we have that $\|u(x) - f\|_{\mathcal{F}(Y)} < \varepsilon$.*

Proof. See Appendix A, Theorem A.0.1. □

Properties of Limits

For the sake of completeness, two important properties of limits are included in this section, namely, the limit of a FVM times a constant and the limit of the sum of two FVMs. These properties are presented in the following theorem.

Theorem 3.1.2. *Let $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$, $v : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$, $f, g \in \mathcal{F}(Y)$, $\alpha \in \mathbb{R}$ and $x_0 \in X$ a point in either X or its boundary $\partial(X)$. Then, the following affirmations hold:*

1. *If $\lim_{x \rightarrow x_0} u(x) = f$, then $\lim_{x \rightarrow x_0} \alpha u(x) = \alpha f$, where $\alpha u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ is defined as $x \mapsto \alpha u(x)$.*
2. *If $\lim_{x \rightarrow x_0} u(x) = f$ and $\lim_{x \rightarrow x_0} v(x) = g$, then $\lim_{x \rightarrow x_0} (u + v)(x) = f + g$, where $(u + v)(x) : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ is defined as $x \mapsto u(x) + v(x)$.*

Proof. See Appendix A, Theorem A.0.2 □

3.1.2 Continuity

As with real and vector-valued functions, FVMs may also be continuous. In this case, intuitively speaking, we consider that the FVM u is continuous at x_0 if the limit of u is its function-value at x_0 . If this limit exists for any x_0 in a domain X of u , we say simply that u is continuous on X . In the following definition, we state formally the concept of continuity for FVMs.

Definition 3.1.3 (Continuity of a FVM). *A FVM $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ is continuous at $x_0 \in X$ if and only if*

$$\lim_{x \rightarrow x_0} u(x) = u(x_0).$$

If this limit exists for any $x_0 \in X$, we simply say that u is continuous on X .

Now, we introduce an $\varepsilon - \delta$ criterion for continuity of a FVM u using Theorem 3.1.1. This result is given in the following theorem.

Theorem 3.1.3. *A FVM $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ is continuous at $x_0 \in X$ if and only if for all $\varepsilon > 0$ there exists $\delta > 0$ such that for $x \in X$ that satisfies $0 < \|x - x_0\|_X < \delta$, we have that $\|u(x) - u(x_0)\|_{\mathcal{F}(Y)} < \varepsilon$.*

Proof. See Appendix A, Theorem A.0.3. □

Example 3.1.1. *We may have different types of continuity depending on the norm of the function space $\mathcal{F}(Y)$. Let $u : X \subset \mathbb{R} \rightarrow \mathcal{F}(Y)$ be defined as*

$$u(x)(y) = \begin{cases} e^{-(x^2+y^2)}, & x \neq 0 \\ f(y), & x = 0, \end{cases} \quad (3.1)$$

where

$$f(y) = \begin{cases} e^{-y^2}, & y \neq 0 \\ 0, & y = 0. \end{cases} \quad (3.2)$$

If $\mathcal{F}(Y) = L^1(\mathbb{R})$, then u is continuous; however, if $\mathcal{F}(Y) = L^\infty(\mathbb{R})$, where $L^\infty(Y) = \{f : Y \rightarrow \mathbb{R} : \|f\|_\infty < \infty\}$, u is not continuous at $x = 0$.

Example 3.1.2. *Although continuity of FVMs is in general “weaker” as compared to continuity of functions, there exist FVMs that are definitely not continuous in any sense. For instance, let $u : X \subset \mathbb{R} \rightarrow L^2(\mathbb{R})$ be defined as*

$$u(x)(y) = \begin{cases} \frac{\sin(y)}{y}, & x = 0 \\ 0, & x \neq 0. \end{cases} \quad (3.3)$$

In this case, $\lim_{x \rightarrow 0} u(x) = 0$, nevertheless, $u(0)(y) = \frac{\sin(y)}{y}$; hence u is not continuous.

Absolute Continuity

It is a known result that if the derivative of a continuous real-valued function G of one variable is an integrable function f , such a function does not necessarily differ from the indefinite integral of f by a constant. However, if the function G is *absolutely continuous*, then the latter statement is true [30]. Interestingly enough, FVMs defined over \mathbb{R} have also the property of being absolutely continuous [126]. This type of continuity is introduced in the following definition.

Definition 3.1.4. [126] *Let $u : X \subset \mathbb{R} \rightarrow \mathcal{F}(Y)$. The FVM \mathbf{u} is **absolutely continuous** if and only if for all $\varepsilon > 0$ there exists a $\delta > 0$ so that*

$$\sum_n \|u(x_n) - u(z_n)\|_{\mathcal{F}(Y)} < \varepsilon \quad (3.4)$$

holds for all finite systems of pairwise disjoint intervals (x_n, z_n) such that $|x_n - z_n| < \delta$.

An interesting property of absolutely continuous FVMs is that these are also Lipschitz continuous; that is, there exists $K > 0$ such that for any $x, z \in X$, we have that $\|u(x) - u(z)\|_{\mathcal{F}(Y)} \leq K|x - z|$. This follows from the definition of absolute continuity [126].

3.1.3 Differentiation

As mentioned previously, a FVM can be considered as a mapping between two Banach spaces. Given this, the derivative of a FVM u at a given point $x_0 \in X \subset \mathbb{R}^n$ can be defined by means of the Fréchet derivative:

Definition 3.1.5 (Fréchet Derivative). *Let $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$. We say that $\mathbf{Du}(x_0)$ is the derivative of \mathbf{u} at $x_0 \in X$ if*

$$\lim_{h \rightarrow 0} \frac{\|u(x_0 + h) - u(x_0) - \mathbf{Du}(x_0)h\|_{\mathcal{F}(Y)}}{\|h\|_X} = 0. \quad (3.5)$$

*Moreover, $\mathbf{Du} : X \rightarrow \mathcal{F}(Y)$ is bounded and linear. Also, if (3.5) exists for any $x_0 \in X$, we simply say that \mathbf{u} is **differentiable in the $\|\cdot\|_{\mathcal{F}(Y)}$ sense**.*

As for properties of the derivative, we have that it is a linear operator. Moreover, it also possesses a product rule of differentiation as its classical counterpart. In the following theorem such results are presented.

Theorem 3.1.4. Let $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ and $v : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ be differentiable, and suppose that $(uv)(x) \in \mathcal{F}(Y)$ for all $x \in X$. Also, let $f \in \mathcal{F}(Y)$, $\alpha, \beta \in \mathbb{R}$ and $x \in X$. Then, the following assertions are true:

1. **Derivative of a Constant:** If $u(x) = f$, then $Du(x) = 0$ for all $x \in X$.
2. **Sum Rule:** If $w(x) = \alpha u(x) + \beta v(x)$, then $Dw(x) = \alpha Du(x) + \beta Dv(x)$.
3. **Product Rule:** If $w(x) = (uv)(x)$, then $Dw(x) = Du(x)v(x) + u(x)Dv(x)$.

Proof. See Appendix A, Theorem A.0.4. □

As with real-valued functions, FVMs have also a directional derivative. Such a derivative can be defined by means of the well known *Gâteaux derivative*.

Definition 3.1.6 (Gâteaux Derivative). Let $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$. We say that $Du(x_0; \mathbf{h})$ is the directional derivative of u at $x_0 \in X$ in the direction of $\mathbf{h} \in X$ if the limit

$$Du(x_0; h) = \lim_{\varepsilon \downarrow 0} \frac{u(x_0 + \varepsilon h) - u(x_0)}{\varepsilon} \quad (3.6)$$

exists. If (3.6) holds for all $x \in X$, we say that u is **Gâteaux differentiable in X** .

As usual, by definition and properties of the Gâteaux derivative, we also have that the directional derivative of a FVM u may not always be linear. This implies that if a FVM u is Gâteaux differentiable, it may not be Fréchet differentiable; however, Fréchet differentiability of u —or simply differentiability—always implies Gâteaux differentiability thereof.

Example 3.1.3. Let $u(x) = |x|$. If $x \neq 0$, we have that

$$\lim_{\varepsilon \rightarrow 0} \frac{|x + \varepsilon h| - |x|}{\varepsilon} = \lim_{\varepsilon \rightarrow 0} \frac{(x + \varepsilon h)^2 - x^2}{\varepsilon(|x + \varepsilon h| + |x|)} \quad (3.7)$$

$$= \lim_{\varepsilon \rightarrow 0} \frac{2\varepsilon x h + (\varepsilon h)^2}{\varepsilon(|x + \varepsilon h| + |x|)} \quad (3.8)$$

$$= h \frac{x}{|x|}. \quad (3.9)$$

When $x = 0$, by the definition of directional derivative, it can be easily seen that $Du(0; h) = |h|$. Then,

$$Du(x; h) = \begin{cases} h \frac{x}{|x|}, & x \neq 0 \\ |h|, & x = 0, \end{cases} \quad (3.10)$$

Notice that the Gâteaux derivative of u exists at $x = 0$, however, u is not differentiable at this point. Moreover, at $x = 0$, $Du(0; h)$ is a non-linear function of h .

We now review some of the most important properties of the directional derivative. These are presented in the following theorem.

Theorem 3.1.5. *Let $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ and $v : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ be Gâteaux differentiable, and suppose that $(uv)(x) \in \mathcal{F}(Y)$ for all $x \in X$. Also, let $f \in \mathcal{F}(Y)$, $\alpha, \beta \in \mathbb{R}$ and $x \in X$. Then, the following assertions are true:*

1. **Derivative of a Constant:** *Let $u(x) = f$, where, $f \in \mathcal{F}(Y)$. Then, $Du(x; h) = 0$.*
2. **Sum Rule:** *If $w(x) = (u + v)(x)$, then $D(u + v)(x; h) = Du(x; h) + Dv(x; h)$.*
3. **Product Rule:** *If $w(x) = (uv)(x)$, then $D(w)(x; h) = Du(x; h)v(x) + u(x)Dv(x; h)$.*

Proof. See Appendix A, Theorem A.0.5. □

Let $X \subset \mathbb{R}^n$. As usual, the partial derivative of u with respect to the component x_i of $x \in X$ can be defined as a directional derivative. We denote such a derivative as $\frac{\partial u}{\partial x_i}$, which is given by

$$\frac{\partial u}{\partial x_i} = \lim_{h \rightarrow 0} \frac{u(x_1, \dots, x_i + h, \dots, x_n) - u(x_1, \dots, x_i, \dots, x_n)}{h}. \quad (3.11)$$

Moreover, if all the partial derivatives of u exist and are continuous at $x_0 \in X$, then u is differentiable at this point.

Theorem 3.1.6. *If the partial derivatives $\frac{\partial u}{\partial x_i}$ exist at every $x \in X$, and if the mappings $\frac{\partial u}{\partial x_i} : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ are continuous at a point $x_0 \in X$, then u is differentiable at x_0 .*

Proof. See [23], Proposition 3.7.2, p. 46. □

This theorem implies that $Du(x)$ is a vector in \mathbb{R}^n such that its i -th entry is the partial derivative of u with respect to the variable x_i . As expected, $Du(x)$ is linear and bounded. We denote this operator as $\nabla_x u(x)$. In the particular case in which $X \subset \mathbb{R}$, we have that

$$Du(x) = \lim_{h \rightarrow 0} \frac{u(x + h) - u(x)}{h}. \quad (3.12)$$

Furthermore, if u is differentiable at $x_0 \in X$ with respect to $\|\cdot\|_{\mathcal{F}(Y)}$, $Du(x_0)$ is also the derivative of u with respect to any other norm $\|\cdot\|$, provided that $\|\cdot\|_{\mathcal{F}(Y)}$ and $\|\cdot\|$ are equivalent.

Theorem 3.1.7. *If u is differentiable at $x_0 \in X$ for a given norm, then u is also differentiable at this point for any other equivalent norm and its derivative remains the same.*

Proof. See [23], Proposition 2.1.1, p. 27. □

The previous results motivate the following definition.

Definition 3.1.7. *Let $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$. If u is differentiable at $x_0 \in X$, we say that $\nabla_{\mathbf{x}}\mathbf{u}(\mathbf{x}_0)$ is the derivative of \mathbf{u} at $\mathbf{x}_0 \in \mathbf{X}$. Also, if u is differentiable for any $x_0 \in X$, we simply say that $\nabla_{\mathbf{x}}\mathbf{u}(\mathbf{x})$ is the derivative of \mathbf{u} and that \mathbf{u} is differentiable.*

3.1.4 Best Linear Approximation

Thanks to definition 3.1.5, it is now known that u is differentiable at x_0 if

$$\lim_{h \rightarrow 0} \frac{\|u(x_0 + h) - u(x_0) - Du(x_0)h\|_{\mathcal{F}(Y)}}{\|h\|_X} = 0. \quad (3.13)$$

By making the substitution $h = x - x_0$, the following equivalent condition for differentiability is obtained:

$$\lim_{x \rightarrow x_0} \frac{\|u(x) - u(x_0) - Du(x_0)(x - x_0)\|_{\mathcal{F}(Y)}}{\|x - x_0\|_X} = 0. \quad (3.14)$$

The latter equation implies that if u is differentiable at x_0 , the quantity $R_{x_0}(x) = u(x) - u(x_0) - Du(x_0)(x - x_0)$ approaches zero faster than x gets closer to x_0 . As expected, this implies that the linear form $L_{x_0}(x) = u(x_0) + Du(x_0)(x - x_0)$ is the *best linear approximation* of u at the point x_0 with respect to the norm of the function space $\mathcal{F}(Y)$.

Definition 3.1.8 (Best Linear Approximation). *Let $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$. If u is differentiable at x_0 , then $\mathbf{L}_{\mathbf{x}_0}\mathbf{u}(\mathbf{x}) = \mathbf{u}(\mathbf{x}_0) + \nabla_{\mathbf{x}}\mathbf{u}(\mathbf{x}_0)(\mathbf{x} - \mathbf{x}_0)$ is the *best linear approximation of \mathbf{u} at $\mathbf{x}_0 \in \mathbf{X}$.**

Example 3.1.4. *Let $u : X \subset \mathbb{R} \rightarrow L^1(\mathbb{R})$ be defined as*

$$u(x)(y) = e^{-(x^2+y^2)}. \quad (3.15)$$

Say one wishes to obtain $L_1(x)$. First, one has that

$$\frac{du}{dx}(y) = -2xe^{-(x^2+y^2)}, \quad (3.16)$$

then $u'(1)(y) = -2e^{-(1+y^2)}$. Also, $u(1)(y) = e^{-(1+y^2)}$. Thus, the linear form

$$L_1(x)(y) = e^{-(1+y^2)}(3 - 2x) \quad (3.17)$$

is the best linear approximation of u at $x = 1$.

3.1.5 Integration

Integration of FVMs can be defined by means of the *Bochner integral* [11, 3]. Such an integral is a natural generalization of the Lebesgue integral for mappings taking values in a Banach space. As expected, this integral inherits some of the properties of the Lebesgue integral, such as linearity.

Definition 3.1.9. [3] Let $v : \Omega \rightarrow Z$ be a mapping and let (Ω, Σ, μ) be a measure space. We say that v is **strongly μ -measurable** if there exists a sequence $\{\varphi_n\}$ of Ω -simple mappings such that $\lim_{n \rightarrow \infty} \|v(\omega) - \varphi_n(\omega)\|_Z = 0$ for almost all $\omega \in \Omega$ in the sense of the measure μ .

Definition 3.1.10 (Bochner Integral). [3] Let $v : \Omega \rightarrow Z$ be a mapping between the Banach spaces Ω and Z . Also, let Σ be a σ -algebra of Ω and (Ω, Σ, μ) be a measure space. Then, a strongly μ -measurable v is **Bochner integrable** if there exists a sequence $\{\varphi_n\}$ of Ω -step mappings defined below such that the real-valued function $\|v(\omega) - \varphi_n(\omega)\|_Z$ is Lebesgue integrable for each n and

$$\lim_{n \rightarrow \infty} \int \|v(\omega) - \varphi_n(\omega)\|_Z d\mu = 0. \quad (3.18)$$

Moreover, for each $D \in \Sigma$, the **Bochner integral** of v over D is defined as

$$\int_D v d\mu = \lim_{n \rightarrow \infty} \int_D \varphi_n d\mu. \quad (3.19)$$

In this case, an Ω -step mapping is a measurable mapping $\varphi : \Omega \rightarrow Z$ such that its range is finite; i.e.,

$$R(\varphi) = \{z_i : z_i \in Z, i \in \mathbb{N}, 1 \leq i \leq n\}, \quad (3.20)$$

and

$$A_i = \varphi^{-1}(z_i) \in \Sigma, \quad (3.21)$$

for all i , with $\mu(A_i) < \infty$ for every non-zero z_i . If the condition $\mu(A_i) < \infty$ is not considered, then φ is called an Ω -simple mapping. Both Ω -simple and Ω -step mappings can be represented with the following formula:

$$\varphi(\omega) = \sum_{i=1}^n z_i \chi_{A_i}(\omega), \quad (3.22)$$

where χ is the characteristic function². In particular, the integral of an Ω -step mapping is given by

$$\int \varphi \, d\mu = \sum_{i=1}^n \mu(A_i) z_i. \quad (3.23)$$

Furthermore, the integral of φ over any $D \in \Sigma$ is defined as

$$\int_D \varphi \, d\mu = \int \varphi \chi_D \, d\mu. \quad (3.24)$$

At this point it is clear that definitions 3.1.9 and 3.1.10 can be employed easily to define measurability and integrability of FVMs. In few words, we say that a FVM u is *measurable* if it is strongly measurable. Similarly, if u is measurable, we say that u is *integrable* if it is Bochner integrable. These concepts are formally introduced in the following definitions:

Definition 3.1.11 (Measurable FVM). *Let $u : X \subseteq \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ and let (X, Σ, μ) be a measure space, where μ is the canonical Lebesgue measure and Σ is a σ -algebra of X . We say that \mathbf{u} is **measurable** if there exists a sequence $\{\varphi_n\}$ of X -simple mappings, where $\varphi_n : X \subseteq \mathbb{R}^n \rightarrow \mathcal{F}(Y)$, such that $\lim_{n \rightarrow \infty} \|u(x) - \varphi_n(x)\|_{\mathcal{F}(Y)} = 0$ for almost all $x \in X$ in the sense of the Lebesgue measure.*

Definition 3.1.12 (Integral of a FVM). *Let $u : X \subseteq \mathbb{R}^n \rightarrow \mathcal{F}(Y)$. Also, let Σ be a σ -algebra of X and (X, Σ, μ) be a measure space endowed with the Lebesgue measure. If u is measurable, we say that \mathbf{u} is **integrable** if there exists a sequence $\{\varphi_n\}$ of X -simple mappings such that the real-valued function $\|u(x) - \varphi_n(x)\|_{\mathcal{F}(Y)}$ is Lebesgue integrable for each n and*

$$\lim_{n \rightarrow \infty} \int \|u(x) - \varphi_n(x)\|_{\mathcal{F}(Y)} \, dx = 0. \quad (3.25)$$

Moreover, for each $D \in \Sigma$, the **integral of \mathbf{u} over D** is defined as

$$\int_D u(x) \, dx := \lim_{n \rightarrow \infty} \int_D \varphi_n(x) \, dx. \quad (3.26)$$

The latter definition of integrability might be difficult to apply; nevertheless, for finite measure spaces, integrability of a measurable FVM u can be checked more easily by employing the following criterion:

Theorem 3.1.8. *Let $u : X \subseteq \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ be a measurable FVM and let (X, Σ, μ) be a finite measure space. Then, u is integrable if and only if $\|u(x)\|_{\mathcal{F}(Y)}$ is Lebesgue integrable, that is, $\int_X \|u(x)\|_{\mathcal{F}(Y)} \, dx < \infty$.*

²As usual, $\chi_D(x)$ is equal to one if $x \in D$ and zero elsewhere.

Proof. See [40], Theorem II.2.2, p. 45. □

The collection of all Bochner integrable functions is a vector subspace of the set of all the strongly measurable mappings of the form $v : \Omega \rightarrow Z$; moreover, the Bochner integral acts as a linear operator from this space into Z . Therefore, we also have that the set of all the integrable FVMs of the form $u : X \subseteq \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ is a vector subspace of the set of all the measurable FVMs, which we denote as $\mathcal{M}_{\mathcal{F}}$:

$$\mathcal{M}_{\mathcal{F}} = \{u \in (\mathcal{F}(Y))^X : u \text{ is measurable}\}. \quad (3.27)$$

As expected, the integral of FVMs is also a linear operator from the set $\mathcal{M}_{\mathcal{F}}$ into the function space $\mathcal{F}(Y)$.

Theorem 3.1.9. *If u and v are integrable and $\alpha, \beta \in \mathbb{R}$, then $\alpha u + \beta v$ is also integrable and*

$$\int_D (\alpha u + \beta v) dx = \alpha \int_D u dx + \beta \int_D v dx \quad (3.28)$$

for all $D \in \Sigma$.

Proof. See [3], Theorem 11.43, p. 426. □

Moreover, the integral of FVMs also has the property of *additivity of integrals on sets*. This is proved in the following theorem.

Theorem 3.1.10. *Let $u : X \subseteq \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ be an integrable FVM. Also, let $D, E \in \Sigma$ and $D \cap E = \emptyset$. Then, we have that*

$$\int_{D \cup E} u dx = \int_D u dx + \int_E u dx. \quad (3.29)$$

Proof. See Appendix A, Theorem A.0.6. □

As with the Lebesgue integral, there is also a *dominated convergence theorem* for the Bochner integral. We present the FVM-version of this result in the following theorem.

Theorem 3.1.11 (Dominated Convergence Theorem). *Let (X, Σ, μ) be a measure space. Also, let $u : X \subseteq \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ be a measurable FVM and $\{u_n\}$ be a sequence of integrable FVMs such that $\|u_n(x) - u(x)\|_{\mathcal{F}(Y)} \rightarrow 0$ for μ -almost all $x \in X$. If there exists*

a positive real-valued Lebesgue integrable function such that $\|u_n(x)\|_{\mathcal{F}(Y)} \leq g(x)$ μ -almost everywhere, then u is integrable and

$$\lim_{n \rightarrow \infty} \int_D u_n \, dx = \int_D u \, dx \quad (3.30)$$

for all $D \in \Sigma$.

Proof. See [3], Theorem 11.46, p. 427. □

Another important property of the Bochner integral is that it commutes with bounded operators. As expected, this implies that integrals and bounded operators that act on FVMs also commute.

Theorem 3.1.12. *Let $u : X \subseteq \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ be integrable and let Z be a Banach space. If $T : \mathcal{F}(Y) \rightarrow Z$ is a bounded operator, then the mapping $Tu : \mathcal{F}(Y) \rightarrow Z$ is integrable and*

$$\int_X Tu \, dx = T \left(\int_X u \, dx \right). \quad (3.31)$$

Proof. See [109], Theorem 2.12, p. 17. □

Furthermore, for FVMs of the form $u : X \subseteq \mathbb{R}^n \rightarrow L^p(Y)$, we have the following important result:

Theorem 3.1.13. *Let (X, Σ, μ) and (Y, \mathcal{T}, ν) be finite measurable spaces. Then, for a FVM of the form $u : X \subseteq \mathbb{R}^n \rightarrow L^p(Y)$ the following results hold:*

1. *If u is integrable and $1 \leq p \leq \infty$, then there exists a $\mu \times \nu$ measurable real-valued function $v : X \times Y \rightarrow \mathbb{R}$ —which is uniquely determined up to a $\mu \times \nu$ null set—such that for μ -almost all $x \in X$ we have that $u(x) = v(x, \cdot)$. Moreover, for ν -almost all $y \in Y$, the real-valued function $v(\cdot, y)$ is μ -integrable and*

$$\left(\int_X u(x) \, d\mu \right) (y) = \int_X v(x, y) \, d\mu(x). \quad (3.32)$$

2. *If $u : X \subseteq \mathbb{R}^n \rightarrow L^1(Y)$ is integrable, then there exists a $\mu \times \nu$ integrable real-valued function $v : X \times Y \rightarrow \mathbb{R}$ such that for μ -almost all $x \in X$ we have that $u(x) = v(x, \cdot)$. Furthermore, $\int_X v \, d\mu$ exists for ν -almost all $y \in Y$ and*

$$\left(\int_X u(x) \, d\mu \right) (y) = \int_X v(x, y) \, d\mu(x). \quad (3.33)$$

3. If v is $\mu \times \nu$ measurable and $u(x) = v(x, \cdot)$ is in $L^p(Y)$ for μ -almost all $x \in X$, where $1 \leq p < \infty$, then u is μ -measurable.

Proof. See [44], Theorem 17, p. 198. □

It is also worthwhile to mention some results regarding integration of FVMs over sets of measure zero and integrals of FVMs that are equal almost everywhere. These results are presented in the following theorem.

Theorem 3.1.14. *Let $u, v \in \mathcal{M}_{\mathcal{F}}$ such that $u, v \geq 0$ for all $x \in X$ and all $y \in Y$. Then the following affirmations hold:*

1. If $\mu(Z) = 0$, then $\int_Z u \, dx = 0$.
2. If $u = v$ almost everywhere on X , then $\int_X u \, dx = \int_X v \, dx$.

Proof. See Appendix A, Theorem A.0.7. □

Example 3.1.5. *Let $u : [-a, a] \subset \mathbb{R} \rightarrow L^1([-a, a])$ be given by*

$$u(x)(y) = x^2 + y^2. \tag{3.34}$$

The integral of u respect to x over the interval $[-a, a]$ is given by

$$\int_{-a}^a (x^2 + y^2) dx = \frac{2}{3}a^3 + 2ay^2. \tag{3.35}$$

Notice that the result of carrying out this integration is a function that depends only on y .

Example 3.1.6. *Another interesting example is the integral respect to x of the following FVM $u : \mathbb{R} \rightarrow L^1(\mathbb{R})$:*

$$u(x)(y) = \begin{cases} e^{-y^2}, & x = 0 \\ 0, & x \neq 0. \end{cases} \tag{3.36}$$

In this case one has that

$$\int_{-\infty}^{\infty} u(x) dx = 0. \tag{3.37}$$

The reason for this is that u is equal to the FVM $v = 0$ almost everywhere. Since $\int_X v \, dx = 0$, by Theorem 3.1.14 we have that $\int_X u \, dx = \int_X v \, dx = 0$.

It is worthwhile to mention that all the results that are presented in this section also hold if integration is carried out over the field of complex numbers \mathbb{C} [44].

Iterated Integral of u

Integration of FVMs can be carried out with respect to both variables x and y . As expected, the result of such an integration can be interpreted as the “total volume” of a given FVM u . Moreover, under some circumstances, the order of integration does not affect the final result. The conditions under which this result holds are shown in the following theorem

Theorem 3.1.15. *Let $u : X \subseteq \mathbb{R}^n \rightarrow L^1(Y)$ be an integrable FVM. Then, we have that*

$$\int_Y \left(\int_X u \, dx \right) dy = \int_X \left(\int_Y u \, dy \right) dx. \quad (3.38)$$

Proof. See Appendix A, Theorem A.0.8. □

3.1.6 Fundamental Theorem of Calculus

As with real-valued functions whose support is a subset of the real line, there is also a fundamental theorem of calculus for FVMs that depend only on one variable. This result is presented in the following theorem.

Theorem 3.1.16. *Let (X, Σ, μ) and (Y, T, ν) be finite measure spaces, where $X = [a, b] \subset \mathbb{R}$. Also, let $u : X \subset \mathbb{R} \rightarrow \mathcal{F}(Y)$ be a continuous FVM. Then, the FVM U defined as*

$$U(x) = \int_a^x u(z) dz \quad (3.39)$$

is differentiable on (a, b) and $U'(x) = u(x)$ for ν -almost every $y \in Y$.

Proof. See Appendix A, Theorem A.0.9. □

Corollary 3.1.16.1. *If $u : X \subset \mathbb{R} \rightarrow \mathcal{F}(Y)$ is a continuous FVM, then*

$$\lim_{h \rightarrow 0} \frac{1}{h} \int_x^{x+h} u(z) \, dz = u(x) \quad (3.40)$$

for ν -almost all $y \in Y$ and all $x \in (a, b) \subset X$.

Proof. See Appendix A, Theorem A.0.9. □

Observe that the latter corollary is analogous to the Lebesgue differentiation theorem for real-valued functions [30]. Moreover, the definite integral of the derivative of a FVM u is also equal to the difference of the values of u evaluated at the limit points of its integral.

Theorem 3.1.17. *Let $u : [a, b] \subset \mathbb{R} \rightarrow \mathcal{F}(Y)$ be an absolutely continuous FVM and $\mathcal{F}(Y)$ a reflexive Banach space. Then, the classical derivative $u'(x)$ exists for μ -almost X . Moreover, we have that*

$$\int_a^x u'(z) dz = u(x) - u(a) \quad (3.41)$$

for all $x \in [a, b]$.

Proof. See [126], Volume II-A, Problem 23.5c, p. 443. □

Integration by Parts

An immediate consequence of theorems 3.1.16 and 3.1.17 is the possibility of defining the integration by parts of the product of two differentiable FVMs. Let $u, v : X \subset \mathbb{R} \rightarrow \mathcal{F}(Y)$ be differentiable. By properties of the derivative of FVMs, we have that

$$(uv)' = u'v + uv', \quad (3.42)$$

provided that $(uv)(x) \in \mathcal{F}(Y)$ for all $x \in X$. Integrating both sides of the previous equation over $[a, b]$, and using the results from the previous section, we obtain that

$$\int_a^b (uv)' dx = \int_a^b u'v dx + \int_a^b uv' dx, \quad (3.43)$$

which is equivalent to

$$\int_a^b u'v dx = uv|_a^b - \int_a^b uv' dx. \quad (3.44)$$

Example 3.1.7. *Let $u : [-1, 1] \subset \mathbb{R} \rightarrow L^1([-1, 1])$ be defined as*

$$u(x)(y) = (xy)e^{-(x+y)}. \quad (3.45)$$

Its integral over $[-1, 1]$ can be computed using integration by parts:

$$\int_{-1}^{+1} (xy)e^{-(x+y)} dx = -(xy)e^{-(x+y)}|_{-1}^{+1} + \int_{-1}^{+1} ye^{-(x+y)} dx \quad (3.46)$$

$$= -(xy)e^{-(x+y)}|_{-1}^{+1} - ye^{-(x+y)}|_{-1}^{+1} \quad (3.47)$$

$$= -(x+1)ye^{-(x+y)}|_{-1}^{+1} \quad (3.48)$$

$$= -2ye^{-(1+y)}. \quad (3.49)$$

Notice that

$$\frac{d}{dx}(-(x+1)ye^{-(x+y)}) = (xy)e^{-(x+y)}, \quad (3.50)$$

as expected.

3.2 L^p Spaces

In the literature, the L^p spaces of functions that take values on a Banach space are known as *Lebesgue-Bochner L^p spaces* [126, 40]; however, in this thesis, we refer to them simply as L^p spaces. These are defined in the following way.

Definition 3.2.1 (L^p Space). *Let (X, Σ, μ) be a measure space and $1 \leq p < \infty$. We define the $\mathbf{L}^p(\mathbf{X}; \mathcal{F}(\mathbf{Y}))$ **space** as the set of all equivalence classes of μ -measurable FVMs of the form $u : X \subseteq \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ such that*

$$\|u\|_p := \left(\int_X \|u(x)\|_{\mathcal{F}(Y)}^p dx \right)^{\frac{1}{p}} < \infty, \quad (3.51)$$

If $p = \infty$, then we have **the space $\mathbf{L}^\infty(\mathbf{X}; \mathcal{F}(\mathbf{Y}))$** , whose elements are the set of measurable FVMs such that

$$\|u\|_\infty := \operatorname{ess\,sup}_{x \in X} \{\|u(x)\|_{\mathcal{F}(Y)}\} < \infty, \quad (3.52)$$

where $\operatorname{ess\,sup}_{x \in X} \{\cdot\}$ is the essential supremum with respect to the measure μ .

As expected, both $\|\cdot\|_\infty$ and $\|\cdot\|_p$ fulfill the properties of a norm, therefore, L^p spaces are normed spaces, which implies that $\|u\|_p$ is the norm of $u \in L^p(X; \mathcal{F}(Y))$.

Theorem 3.2.1. *The functionals $\|u\|_p : L^p(X; \mathcal{F}(Y)) \rightarrow \mathbb{R}$ and $\|u\|_\infty : L^\infty(X; \mathcal{F}(Y)) \rightarrow \mathbb{R}$ are norms.*

Proof. See Appendix A, Theorem A.0.10. □

Moreover, if $\mathcal{F}(Y)$ is a Banach space, $L^p(X; \mathcal{F}(Y))$ is also a Banach space.

Theorem 3.2.2. *Let (X, Σ, μ) and (Y, T, ν) be finite measure spaces and let $1 \leq p \leq \infty$. If $\mathcal{F}(Y)$ is a Banach space, then $L^p(X; \mathcal{F}(Y))$ is also a Banach space.*

Proof. See Appendix A, Theorem A.0.11. □

Given the previous theorems, it can be seen that L^p spaces of FVMs are Banach spaces; in other words, L^p spaces “inherit” the completeness of their function space $\mathcal{F}(Y)$. This also happens if the function space is separable. If this is so, then $L^p(X; \mathcal{F}(Y))$ is also a separable space.

Theorem 3.2.3. *If $\mathcal{F}(Y)$ is separable and $X \subset \mathbb{R}$, then $L^p(X; \mathcal{F}(Y))$ is separable as well.*

Proof. See [126], Volume II-A, Problem 23.2, p. 440. □

Moreover, if $\mathcal{F}(Y)$ is a Hilbert space, we have that $L^2(X; \mathcal{F}(Y))$ is also a Hilbert space.

Theorem 3.2.4. *If $\mathcal{F}(Y)$ is a Hilbert space with scalar product $\langle \cdot, \cdot \rangle_{\mathcal{F}(Y)}$, then $L^2(X; \mathcal{F}(Y))$ is also a Hilbert space with scalar product defined as*

$$\langle u, v \rangle = \int_X \langle u(x), v(x) \rangle_{\mathcal{F}(Y)} dx, \quad (3.53)$$

for all $u, v \in L^2(X; \mathcal{F}(Y))$.

Proof. See Appendix A, Theorem A.0.12. □

Another property of L^p spaces is that the set of X -step mappings is dense in $L^p(X; \mathcal{F}(Y))$. This follows from the fact that any element of L^p is a μ -measurable FVM, therefore, there is a sequence $\{\varphi_n\}$ of X -step mappings such that for any $u \in L^p(X; \mathcal{F}(Y))$ we have that $\lim_{n \rightarrow \infty} \|u(x) - \varphi_n(x)\|_{\mathcal{F}(Y)}^p = 0$ for μ -almost $x \in X$. Moreover, if $X \subset \mathbb{R}$, the space of continuous FVMs $C(X; \mathcal{F}(Y))$ with norm defined as

$$\|u\| := \sup_{x \in X} \|u(x)\|_{\mathcal{F}(Y)}, \quad (3.54)$$

for all $u \in C(X; \mathcal{F}(Y))$, is dense in $L^p(X; \mathcal{F}(Y))$ as well.

Theorem 3.2.5. [126] *The space $C(X; \mathcal{F}(Y))$ is dense in $L^p(X; \mathcal{F}(Y))$. Even more, the embedding $C(X; \mathcal{F}(Y)) \subseteq L^p(X; \mathcal{F}(Y))$ is continuous.*

Proof. See Appendix A, Theorem A.0.13. □

The previous embedding can be extended not only to the space of continuous FVMs but also to $L^\infty(X; \mathcal{F}(Y))$. In this case, we can use the following inequality

$$\|u\|_p \leq \mu(X)^{\frac{1}{p}} \operatorname{ess\,sup}_{x \in X} \{\|u(x)\|_{\mathcal{F}(Y)}\} \quad (3.55)$$

$$= \mu(X)^{\frac{1}{p}} \|u\|_\infty, \quad (3.56)$$

which implies that the embedding $L^\infty(X; \mathcal{F}(Y)) \subseteq L^p(X; \mathcal{F}(Y))$ is continuous if $\mu(X) < \infty$. An immediate consequence of this is that any bounded FVM is in L^p . Along these lines, if $1 \leq q \leq p \leq \infty$, we also have that the embedding $L^p(X; \mathcal{F}(Y)) \subseteq L^q(X; \mathcal{G}(Y))$ is continuous, provided that the embedding $\mathcal{F}(Y) \subseteq \mathcal{G}(Y)$ is continuous as well. This can be proved using the classical Hölder inequality.

Theorem 3.2.6. [126] *Let (X, Σ, ν) and (Y, T, ν) be finite measure spaces. If the embedding $\mathcal{F}(Y) \subseteq \mathcal{G}(Y)$ is continuous and $1 \leq q \leq p \leq \infty$, then the embedding $L^p(X; \mathcal{F}(Y)) \subseteq L^q(X; \mathcal{G}(Y))$, is also continuous.*

Proof. See Appendix A, Theorem A.0.14. □

3.2.1 The Dual Space of $L^p(X; \mathcal{F}(Y))$

If $1 \leq p < \infty$ and $\frac{1}{p} + \frac{1}{q} = 1$, we have that the dual of $L^p(X; \mathcal{F}(Y))$ is $L^q(X; \mathcal{F}(Y)^*)$; that is, $L^p(X; \mathcal{F}(Y))^* = L^q(X; \mathcal{F}(Y)^*)$, where $\mathcal{F}(Y)^*$ is the dual of $\mathcal{F}(Y)$. This assertion holds if $\mathcal{F}(Y)$ has the *Radon-Nikodým property*; that is, given a σ -algebra Σ on X , $\mathcal{F}(Y)$ has the Radon-Nikodým property with respect to any finite measure μ on X .

Definition 3.2.2. [40] *Let Σ be a σ -algebra of subsets of X , $\rho : \Sigma \rightarrow \mathcal{F}(Y)$ be a vector measure, and μ be a finite non-negative real-valued measure on Σ . If $\lim_{\mu(E) \rightarrow 0} \rho(E) = 0$, then ρ is μ -continuous.*

Definition 3.2.3 (Variation of a vector measure). [40] *Let Σ be a σ -algebra of subsets of X , $\rho : \Sigma \rightarrow \mathcal{F}(Y)$ be a vector measure. The variation of ρ is defined as the extended nonnegative function $|\rho|$ whose value on a set $E \in \Sigma$ is given by*

$$|\rho|(E) := \sup_{\Pi} \sum_{F \in \Pi} \|\rho(F)\|_{\mathcal{F}(Y)}, \quad (3.57)$$

where the supremum is taken over all partitions Π of E into a finite number of pairwise disjoint elements of Σ . If $|\rho|(E)$ is finite, then it is said that ρ has bounded variation

Definition 3.2.4 (Radon-Nikodým property). [5] Let (X, Σ, μ) be a measure space. $\mathcal{F}(Y)$ has the Radon-Nikodým property with respect to μ if for every countably additive μ -continuous vector measure of bounded variation $\rho : \Sigma \rightarrow \mathcal{F}(Y)$ there exists an integrable FVM $v : X \rightarrow \mathcal{F}(Y)$ such that

$$\rho(E) = \int_E v \, d\mu \quad (3.58)$$

for any measurable set $E \in \Sigma$.

The fact that $L^q(X; \mathcal{F}(Y)^*)$ is the dual of $L^p(X; \mathcal{F}(Y))$ relies on two important results, namely, the Hölder inequality and the Radon-Nikodým theorem for Banach-valued functions. These are presented in the following theorems within the context of FVMs.

Theorem 3.2.7 (Hölder Inequality). [126] Let $u \in L^p(X; \mathcal{F}(Y))$ and $v \in L^q(X; \mathcal{F}(Y)^*)$, where $\mathcal{F}(Y)$ is the dual space of $\mathcal{F}(Y)$. Also, let $1 < p < \infty$ and $\frac{1}{p} + \frac{1}{q} = 1$. Then, the following Hölder inequality holds for all the elements of both L^p and L^q :

$$\int_X |\langle u(x), v(x) \rangle_{\mathcal{F}(Y)}| dx \leq \left(\int_X \|u(x)\|_{\mathcal{F}(Y)}^p dx \right)^{\frac{1}{p}} \left(\int_X \|v(x)\|_{\mathcal{F}(Y)^*}^q dx \right)^{\frac{1}{q}}. \quad (3.59)$$

In particular, all the integrals of this inequality exist.

Proof. See Appendix A, Theorem A.0.15. □

Theorem 3.2.8 (Radon-Nikodým theorem). If $\nu : \Sigma \rightarrow \mathcal{F}(Y)$ is a μ -continuous vector-measure of bounded variation, then there exists an integrable FVM $u \in L^1(X; \mathcal{F}(Y))$ such that $\nu(D) = \int_D u \, d\mu$ for all $D \in \Sigma$.

Proof. See [40], Chapter III. □

Theorem 3.2.9 (Dual of $L^p(X; \mathcal{F}(Y))$). Let (X, Σ, μ) be a finite measure space and let $1 \leq p < \infty$. If $\frac{1}{p} + \frac{1}{q} = 1$, then $L^p(X; \mathcal{F}(Y))^* = L^q(X; \mathcal{F}(Y)^*)$ if and only if $\mathcal{F}(Y)^*$ has the Radon-Nikodým property with respect to μ .

Proof. See [40], Theorem I, p. 98. □

Thanks to Theorem 3.2.9, we have that $L^p(X; \mathcal{F}(Y))^*$ is the space of all continuous linear functionals of the form $\Phi_v : L^p(X; \mathcal{F}(Y)) \rightarrow \mathbb{R}$. Such functionals are defined as

$$\Phi_v(u) := \int_X \langle u(x), v(x) \rangle_{\mathcal{F}(Y)} dx \quad (3.60)$$

for all $u \in L^p(X; \mathcal{F}(Y))$. As expected, each $v \in L^q(X; \mathcal{F}(Y)^*)$ defines a unique continuous linear functional on $L^p(X; \mathcal{F}(Y))$ and $\|\Phi\| = \|v\|_q$.

For the sake of completeness, we close this section by presenting the following important corollary that follows from the previous theorem. Its proof can also be found in [40].

Corollary 3.2.9.1. *Let (X, Σ, μ) be a finite measure space and let $1 \leq p \leq \infty$. Then, the space $L^p(X; \mathcal{F}(Y))$ is reflexive if and only if both $L^p(X)$ and $\mathcal{F}(Y)$ are reflexive.*

Proof. See [40], Corollary 2, p. 100. □

3.3 Euler-Lagrange Equation

The origins of calculus of variations go back to the XVII century, when Johann Bernoulli posed the *brachistochrone curve problem*³ in a prize competition. Several important mathematicians from that time replied with a solution, among them Jacob Bernoulli and Gottfried Wilhelm Leibniz; however, it was Leonhard Euler who provided a general method to solve optimization problems in which an unknown curve is to be found. Euler’s approach proved to be useful, but it was Lagrange’s simpler and more general method which became the standard approach and it is still being used even nowadays. As opposed to Euler, Lagrange was able to address multi-dimensional variational problems, which lead to partial differential equations [126].

As is well known, the field of calculus of variations was “nourished” by the contributions of other influential mathematicians, namely, Weierstrass, Legendre and Jacobi [126]. Moreover, developments have not stopped. Consequently, the calculus of variations has become an important branch of mathematics with important applications in different fields such as physics, economics and image processing. Given this, it is not our goal to provide a comprehensive extension of this field to include FVMs, but to simply present the *Euler-Lagrange equation* of a given functional whose argument u is a FVM of the form $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$. In particular, we focus our attention on expressions of the following form:

$$I(u) = \int_X f(x, u, \nabla_x u) dx, \tag{3.61}$$

³Johann Bernoulli published this problem in the June edition of *Leipzig Acta Euditorium* in 1696. It read as follows: “Two points, at different distances from the ground and not in a vertical line, should be connected by such a curve so that a body under the influence of gravitational forces passes in the shortest possible time from the upper to the lower point.”

where $f : X \times \mathcal{F}(Y) \times \mathcal{G}^n(Y) \rightarrow \mathbb{R}$ is a mapping that is Fréchet differentiable with respect to all of its arguments, and $\mathcal{G}^n(Y)$ is the Cartesian product of the range of $\nabla_x u$; that is, $\mathcal{G}^n(Y) = \mathcal{G}(Y) \times \cdots \times \mathcal{G}(Y)$, where $\frac{\partial u}{\partial x_i} : X \subset \mathbb{R}^n \rightarrow \mathcal{G}(Y)$. As expected, the solution of the Euler-Lagrange equation is a FVM that either maximizes or minimizes $I(u)$.

Theorem 3.3.1. *Let (X, Σ, μ) and (Y, T, ν) be finite measure spaces. Also, let $u : X \subseteq \mathbb{R}^n \rightarrow \mathcal{F}(Y)$, $\frac{\partial u}{\partial x_i} : X \subset \mathbb{R}^n \rightarrow \mathcal{G}(Y)$, and assume that the function*

$$\Phi(x) := f(x, u(x), \nabla_x u(x)) \quad (3.62)$$

is integrable over X . In addition, suppose that the Fréchet derivatives of $f : X \times \mathcal{F}(Y) \times \mathcal{G}^n(Y) \rightarrow \mathbb{R}$ with respect to all of its arguments are continuous. Define the functional $I(u) : Z(\mathcal{F}(Y), \mathcal{G}(Y)) \rightarrow \mathbb{R}$ as follows:

$$I(u) := \int_X f(x, u, \nabla_x u) dx, \quad (3.63)$$

where $Z(\mathcal{F}(Y), \mathcal{G}(Y))$ is a Banach space of FVMs that depends on the function spaces $\mathcal{F}(Y)$ and $\mathcal{G}(Y)$. If $u_0 : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ is a stationary point of $I(u)$, u_0 is the solution of the equation

$$\frac{\partial f}{\partial u}(u_0) - \nabla \cdot \frac{\partial f}{\partial \nabla_x u}(\nabla_x u_0) = 0. \quad (3.64)$$

where $\frac{\partial f}{\partial u} \in \mathcal{F}(Y)^$ and $\frac{\partial f}{\partial \nabla_x u} \in \mathcal{G}^n(Y)^*$ are the Fréchet derivatives of f with respect to u and $\nabla_x u$ respectively, $\nabla \cdot$ is the classical divergence operator, and $\mathcal{F}(Y)^*$ and $\mathcal{G}(Y)^*$ are the dual spaces of $\mathcal{F}(Y)$ and $\mathcal{G}(Y)$ respectively.*

Proof. See Appendix A, Theorem A.0.16. □

We consider Eq. (3.64) as the Euler-Lagrange equation of the functional $I(u)$ defined in (3.63). As its classical counterpart, it is also a necessary condition for the solutions of the variational problem stated in the previous theorem, however, it is not a sufficient condition for the existence of such solutions. To determine if such solutions exist, the standard *sufficient conditions* from calculus of variations can be employed to such an end [126].

Chapter 4

Fourier Transform

Motivated by the study of the diffusion of heat, Joseph Fourier proposed in the beginning of the XIX century that any periodic function can be decomposed as a series of harmonically related sine functions. With this idea alone, Fourier revolutionized both mathematics and physics; however, it took one and a half centuries to understand the convergence of Fourier series and complete the theory of Fourier integrals [79].

Although Fourier used his breakthrough idea to solve the partial differential equation that governs the diffusion of heat over a surface, the extensions of his work have been employed in many fields such as signal and image processing, quantum physics (e.g., Heisenberg's uncertainty principle) [79], number theory [110] and others. Moreover, the well known Fourier transform diagonalizes all linear-time invariant operators, which are the building blocks of signal processing [79]. Given the clear importance of Fourier analysis, in this chapter, we present a Fourier transform for FVMs, as well as its inverse counterpart. Properties of this transform, examples and some applications are also provided.

4.1 Some Current Definitions

Over time, Fourier's work has been extended much further by many mathematicians, among them, Peter Dirichlet, one of Fourier's doctoral students, and Bernhard Riemann. A classical example of these generalizations is the Fourier transform, which allows to analyze non-periodic functions. For instance, let $u : \mathbb{R} \rightarrow \mathbb{R}$. Its Fourier transform is usually defined as

$$U(\omega) := \int_{-\infty}^{\infty} u(x)e^{-i\omega x} dx, \quad \omega \in \mathbb{R}, \quad (4.1)$$

provided that $u \in L^1(\mathbb{R})$. As is known, this definition can be extended for functions that belong to $L^2(\mathbb{R})$ by using the fact that the space $L^1(\mathbb{R}) \cap L^2(\mathbb{R})$ is dense in $L^2(\mathbb{R})$ [79]. On the other hand, if both u and U belong to $L^1(\mathbb{R})$, the inverse Fourier transform is given by

$$u(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} U(\omega) e^{i\omega x} d\omega. \quad (4.2)$$

A natural extension of the previous result is the Fourier transform of real-valued functions of several variables. Let $u : \mathbb{R}^n \rightarrow \mathbb{R}$, $u \in L^1(\mathbb{R}^n)$. Its Fourier transform is defined as

$$U(\omega) := \int_{\mathbb{R}^n} u(x) e^{-i\omega \cdot x} dx, \quad (4.3)$$

where $\omega \cdot x$ is the dot product of the vectors $\omega, x \in \mathbb{R}^n$. The inverse counterpart is equal to

$$u(x) = \frac{1}{(2\pi)^n} \int_{\mathbb{R}^n} U(\omega) e^{i\omega \cdot x} d\omega, \quad (4.4)$$

provided that both u and U belong to $L^1(\mathbb{R}^n)$.

Another interesting generalization is the *fractional Fourier transform*. If the classical Fourier transform is understood as an operator of order one, its fractional counterpart is an operator of order n , where n may not be an integer. Let $u : \mathbb{R} \rightarrow \mathbb{R}$. The fractional Fourier transform of u is defined as

$$\mathfrak{F}_\alpha \{u(x)\} := \sqrt{\frac{1 - i \cot(\alpha)}{2\pi}} e^{i \cot(\alpha) \frac{\omega^2}{2}} \int_{-\infty}^{\infty} u(x) e^{-i \csc(\alpha) \omega \cdot x + i \cot(\alpha) \frac{x^2}{2}} dx, \quad (4.5)$$

where $\alpha \in \mathbb{R}$. Notice that when $\alpha = \pi/2$, we obtain the classical symmetric Fourier transform. Moreover, $\mathfrak{F}_{-\pi/2} \{U(\omega)\}$ is equal to the usual inverse Fourier transform. This transform is also employed for solving certain classes of ordinary and partial differential equations that arise in quantum mechanics [90].

Furthermore, in [97], Peetre provides perhaps one of the first generalizations of the Fourier transform for Banach-valued functions. In fact, for $p \in (1, 2]$, Peetre proves that the Fourier transform is a bounded operator from $L^p(\mathbb{R}; Z)$ to $L^q(\mathbb{R}; Z)$, where q is the Hölder conjugate of p and Z is a Banach space. This result was extended further by Milman in [85]. Here, it is proved that the Fourier transform is a well defined operator from $L^p(G; Z)$ to $L^q(\hat{G}; Z)$, where G is a locally compact abelian group and \hat{G} its *Pontryagin dual*¹.

¹In few words, the Pontryagin dual is the set of all frequencies ω . The “nature” of this dual is determined by how the set G is defined: e.g., the real line, a finite cyclic group, etc..

Along these lines, in [109], as an application of the Bochner integral [3], the following definition of the Fourier transform is provided:

$$U(\omega) := \int_{\mathbb{R}^n} e^{i\omega \cdot x} u(x) dx, \quad (4.6)$$

where $u : \mathbb{R}^n \rightarrow H$ is an element of $L^1(\mathbb{R}^n; H)$. Here, H is a separable Hilbert space. Note that in [109], the exponent of the complex exponential has a positive sign. Moreover, it is shown that $U(\omega)$ is well defined and that it is a bounded operator from $L^1(\mathbb{R}^n; H)$ to $L^\infty(\mathbb{R}^n; H)$. However, no definition of the inverse Fourier transform is presented.

4.2 Fourier Transform for FVMs

The already existing definitions of the Fourier transforms for Banach-valued functions provide the foundation for defining the Fourier transform of FVMs. In particular, we focus our attention on the elements of the space $L^1(\mathbb{R}^n; \mathcal{F}(Y))$; that is, the space of integrable FVMs.

Definition 4.2.1. *Let $u \in L^1(\mathbb{R}^n; \mathcal{F}(Y))$, where $\mathcal{F}(Y)$ is a complex-valued space. We define **the Fourier transform of u** as the integral*

$$\mathfrak{F}\{u\} := \int_{\mathbb{R}^n} e^{-i\omega \cdot x} u(x) dx, \quad (4.7)$$

where $\omega \in \mathbb{R}^n$. In some cases, we will denote $\mathfrak{F}\{u\}$ as $U(\omega)$ as well.

In technical terms, observe that ω must belong to the Pontryagin dual of \mathbb{R}^n , however, the Euclidean space \mathbb{R}^n is a locally compact abelian group that it is self-dual, therefore, ω is indeed an element of \mathbb{R}^n . On the other hand, in a general setting, it is customary to employ the Haar measure when integration is carried out over a locally compact topological group, nevertheless, in the case of \mathbb{R}^n , the Haar measure is the Lebesgue measure.

Theorem 4.2.1. [109] \mathfrak{F} is a bounded operator of the form

$$\mathfrak{F} : L^1(\mathbb{R}^n; \mathcal{F}(Y)) \rightarrow L^\infty(\mathbb{R}^n; \mathcal{F}(Y)). \quad (4.8)$$

Proof. We have that

$$\|\mathfrak{F}\{u\}(\omega)\|_{\mathcal{F}(Y)} \leq \int_{\mathbb{R}^n} \|e^{-i\omega \cdot x} u(x)\|_{\mathcal{F}(Y)} dx \quad (4.9)$$

$$= \int_{\mathbb{R}^n} \|u(x)\|_{\mathcal{F}(Y)} dx \quad (4.10)$$

$$= \|u\|_1. \quad (4.11)$$

By taking the supremum at both sides of the inequality over all $\omega \in \mathbb{R}^n$ we obtain that

$$\sup_{\omega \in \mathbb{R}^n} \|\mathfrak{F}\{u\}(\omega)\|_{\mathcal{F}(Y)} = \|\mathfrak{F}\{u\}\|_{\infty} \leq \|u\|_1. \quad (4.12)$$

□

Regarding the inverse transform, this operator is not well defined for all the elements that belong to $L^\infty(\mathbb{R}^n; \mathcal{F}(Y))$ since not all the FVMs of this space are in $L^1(\mathbb{R}^n; \mathcal{F}(Y))$. Given this, as is customary in harmonic analysis [53], we define this transform under the assumption that both u and U belong to $L^1(\mathbb{R}^n; \mathcal{F}(Y))$.

Definition 4.2.2. *If both u and U are elements of $L^1(\mathbb{R}^n; \mathcal{F}(Y))$, we define the **inverse Fourier transform of U** as*

$$\mathfrak{F}^{-1}\{U\} := \frac{1}{(2\pi)^n} \int_{\mathbb{R}^n} e^{i\omega \cdot x} U(\omega) d\omega. \quad (4.13)$$

As for the measure of the above integral, this is defined as the *dual measure of μ* , which we denote as $\hat{\mu}$. In our particular case, we have that $\hat{\mu} = \frac{1}{(2\pi)^n} \mu$ [53].

Theorem 4.2.2. *If both u and U belong to $L^1(\mathbb{R}^n; \mathcal{F}(Y))$, then*

$$u(x) = \frac{1}{(2\pi)^n} \int_{\mathbb{R}^n} e^{i\omega \cdot x} U(\omega) d\omega. \quad (4.14)$$

Proof. The proof is completely analogous to the real-valued case, which can be found in [79], Theorem 2.1, p. 23. First, let us define the following FVM $I_\varepsilon : \mathbb{R}^n \rightarrow \mathcal{F}(Y)$:

$$I_\varepsilon(x) := \frac{1}{(2\pi)^n} \int_{\mathbb{R}^n} \int_{\mathbb{R}^n} e^{-\frac{\varepsilon^2 \|\omega\|_2^2}{2} + i\omega(x-z)} u(z) dz d\omega. \quad (4.15)$$

Integration with respect to z yields

$$I_\varepsilon(x) = \frac{1}{(2\pi)^n} \int_{\mathbb{R}^n} e^{-\frac{\varepsilon^2 \|\omega\|_2^2}{2} + i\omega x} U(\omega) d\omega. \quad (4.16)$$

Observe that the integrand is bounded by $\|U(\omega)\|_{\mathcal{F}(Y)}$, therefore, by the dominated convergence Theorem 3.1.11, we have that

$$\lim_{\varepsilon \rightarrow 0} I_\varepsilon(x) = \frac{1}{(2\pi)^n} \int_{\mathbb{R}^n} e^{i\omega \cdot x} U(\omega) d\omega. \quad (4.17)$$

Now, let us apply Fubini's theorem to Eq. (4.15):

$$I_\varepsilon(x) = \frac{1}{(2\pi)^n} \int_{\mathbb{R}^n} \int_{\mathbb{R}^n} e^{-\frac{\varepsilon^2 \|\omega\|_2^2}{2} + i\omega(x-z)} d\omega u(z) dz. \quad (4.18)$$

Notice that the inner integral is equal to a multivariate Gaussian of variance ε^2 and mean equal to x . Therefore,

$$I_\varepsilon(x) = \frac{1}{(2\pi)^n} \int_{\mathbb{R}^n} \frac{1}{(2\pi\varepsilon^2)^{\frac{n}{2}}} e^{-\frac{\|z-x\|_2^2}{2\varepsilon^2}} u(z) dz. \quad (4.19)$$

This Gaussian converges to a multivariate Dirac delta as $\varepsilon \rightarrow 0$, then

$$\lim_{\varepsilon \rightarrow 0} I_\varepsilon(x) = u(x), \quad (4.20)$$

which implies that

$$\lim_{\varepsilon \rightarrow 0} \int_{\mathbb{R}^n} \|I_\varepsilon(x) - u(x)\|_{\mathcal{F}(Y)} dx = 0 \quad (4.21)$$

for μ -almost all $x \in X$. This completes the proof. \square

4.2.1 Properties

It is clear that the operator \mathfrak{F} is linear, which is, of course, a consequence of the linearity of the integral of the Fourier transform. Let $\alpha, \beta \in \mathbb{C}$ and $u, v \in L^1(\mathbb{R}^n; \mathcal{F}(Y))$. We claim, without proof, that the following equality holds:

$$\mathfrak{F}\{\alpha u + \beta v\} = \alpha \mathfrak{F}\{u\} + \beta \mathfrak{F}\{v\}. \quad (4.22)$$

As expected, linearity also holds for the inverse operator \mathfrak{F}^{-1} . Let $\alpha, \beta \in \mathbb{C}$ and $U, V \in L^1(\mathbb{R}^n; \mathcal{F}(Y))$. Then,

$$\mathfrak{F}^{-1}\{\alpha U + \beta V\} = \alpha \mathfrak{F}^{-1}\{U\} + \beta \mathfrak{F}^{-1}\{V\}. \quad (4.23)$$

As with the classical Fourier transform, in the FVM setting, we also have properties such as translation, scaling, modulation, differentiation with respect to x and integration. These are presented in the following theorems.

Theorem 4.2.3. *Let both $u \in L^1(\mathbb{R}^n; \mathcal{F}(Y))$. Also, let $\omega_0, x_0 \in \mathbb{R}^n$ and $a \in \mathbb{R}$, $a \neq 0$. Then, the following assertions hold:*

1. **Translation:** $\mathfrak{F}\{u(x - x_0)\} = e^{-i\omega \cdot x_0} U(\omega).$

2. **Modulation:** $\mathfrak{F}\{e^{-i\omega_0 \cdot x} u(x)\} = U(\omega - \omega_0).$

3. **Scaling:** $\mathfrak{F}\{u(ax)\} = \frac{1}{|a|^n} U\left(\frac{\omega}{a}\right).$

4. **Integration:** $\int_{\mathbb{R}^n} u(x) dx = U(0).$

Proof. Assertions one and two follow easily by making a simple substitution. Regarding assertion three, substitution is also required, but two cases need to be considered separately: $a > 0$ and $a < 0$. If a is positive, we have that

$$\mathfrak{F}\{u(ax)\} = \int_{\mathbb{R}^n} e^{-i\omega \cdot x} u(ax) dx \quad (4.24)$$

$$= \frac{1}{a^n} \int_{\mathbb{R}^n} e^{-i\left(\frac{\omega}{a}\right) \cdot z} u(z) dz \quad (4.25)$$

$$= \frac{1}{a^n} U(\omega). \quad (4.26)$$

For negative a , $\mathfrak{F}\{u(ax)\}$ is given by

$$\mathfrak{F}\{u(ax)\} = \int_{\mathbb{R}^n} e^{-i\omega \cdot x} u(-|a|x) dx \quad (4.27)$$

$$= \frac{(-1)^n}{|a|^n} \int_{-\mathbb{R}^n} e^{i\left(\frac{\omega}{|a|}\right) \cdot z} u(z) dz. \quad (4.28)$$

Notice that $\int_{-\mathbb{R}^n} (\cdot) dx = (-1)^n \int_{\mathbb{R}^n} (\cdot) dx$. Then,

$$\mathfrak{F}\{u(ax)\} = \frac{1}{|a|^n} \int_{\mathbb{R}^n} e^{i\left(\frac{\omega}{|a|}\right) \cdot z} u(z) dz \quad (4.29)$$

$$= \frac{1}{|a|^n} U\left(-\frac{\omega}{|a|}\right). \quad (4.30)$$

Both Eqs. (4.26) and (4.30) can be combined into the next single expression:

$$\mathfrak{F}\{u(ax)\} = \frac{1}{|a|^n} U\left(\frac{\omega}{a}\right). \quad (4.31)$$

As for the property of integration, this is evident. □

Theorem 4.2.4 (Differentiation). *Assume all $\frac{\partial^l u}{\partial x_j^l}$ and $(i\omega_j)^l U$ are elements of $L^1(\mathbb{R}^n; \mathcal{F}(Y))$ whenever $0 \leq l \leq k$. Then,*

$$\mathfrak{F} \left\{ \frac{\partial^k u(x)}{\partial x_j^k} \right\} = (i\omega_j)^k U(\omega). \quad (4.32)$$

Proof. Since $(i\omega_j)^l U$ is in $L^1(\mathbb{R}^n; \mathcal{F}(Y))$ for all $0 \leq l \leq k$, we obtain that

$$\frac{\partial^k u(x)}{\partial x_j^k} = \frac{1}{(2\pi)^n} \int_{\mathbb{R}^n} e^{i\omega \cdot x} (i\omega_j)^k U(\omega) d\omega. \quad (4.33)$$

On the other hand, we have that

$$G(\omega) = \int_{\mathbb{R}^n} e^{-i\omega \cdot x} \frac{\partial^k u(x)}{\partial x_j^k} dx. \quad (4.34)$$

Given that both $\frac{\partial^k u}{\partial x_j^k}$ and $(i\omega_j)^k U$ are in $L^1(\mathbb{R}^n; \mathcal{F}(Y))$, we must have that $G = (i\omega)^k U$. Therefore,

$$\mathfrak{F} \left\{ \frac{\partial^k u(x)}{\partial x_j^k} \right\} = (i\omega_j)^k U(\omega), \quad (4.35)$$

This completes the proof. □

Observe that the last assertion of the previous theorem shows that differential operators are multiplication operators in the frequency domain, therefore, as with the real-valued case, we also have that the Fourier transform of FVMs “diagonalizes” differential operators. Moreover, we have a convolution theorem as well.

Theorem 4.2.5 (Convolution). *Let $(\mathbb{R}^n, \Sigma, \mu)$ be a measure space. Also, let $u, v \in L^1(\mathbb{R}^n; \mathcal{F}(Y))$. Then,*

$$u * v := \int_{\mathbb{R}^n} u(x-z)v(z)dz = \int_{\mathbb{R}^n} u(z)v(x-z)dz \quad (4.36)$$

is in $L^1(\mathbb{R}^n; \mathcal{F}(Y))$ as well and

$$\mathfrak{F}\{u * v\} = \mathfrak{F}\{u\}\mathfrak{F}\{v\}. \quad (4.37)$$

Proof. Once again, the proof is analogous to the real-valued case. First, let us verify that $u * v \in L^1(\mathbb{R}^n; \mathcal{F}(Y))$. Let $w : \mathbb{R}^{2n} \rightarrow \mathcal{F}(Y)$ be equal to $w(x, z) = u(x - z)v(z)$. Since both u and v are integrable, there exist sequences $\{\varphi_n\}$ and $\{\psi_n\}$ such that

$$\lim_{n \rightarrow \infty} \|f(x) - \varphi_n(x)\|_{\mathcal{F}(Y)} = 0 \quad \text{and} \quad \lim_{n \rightarrow \infty} \|v(x) - \psi_n(x)\|_{\mathcal{F}(Y)} = 0 \quad (4.38)$$

for μ -almost all $x \in \mathbb{R}^n$. Define $\zeta_n(x, z) = \varphi_n(x - z)\psi_n(z)$, which is equivalent to

$$\zeta_n(x, z) = \left(\sum_{j=1}^q f_j \chi_{A_j}(x - z) \right) \left(\sum_{k=1}^r g_k \chi_{B_k}(z) \right) \quad (4.39)$$

$$= \sum_{l=1}^s h_l \chi_{C_l}(x, z). \quad (4.40)$$

Thus,

$$\lim_{n \rightarrow \infty} \|w(x, z) - \zeta_n(x, z)\|_{\mathcal{F}(Y)} = 0 \quad (4.41)$$

for $\mu \times \mu$ -almost all $(x, z) \in \mathbb{R}^{2n}$; which implies that $w(x, z)$ is measurable. On the other hand, we have that

$$\int_{\mathbb{R}^n} \int_{\mathbb{R}^n} \|w(x, z)\|_{\mathcal{F}(Y)} dx dz = \int_{\mathbb{R}^n} \int_{\mathbb{R}^n} \|u(x - z)\|_{\mathcal{F}(Y)} \|v(z)\|_{\mathcal{F}(Y)} dx dz \quad (4.42)$$

$$= \int_{\mathbb{R}^n} \|u\|_1 \|v(z)\|_{\mathcal{F}(Y)} dz \quad (4.43)$$

$$= \|u\|_1 \|v\|_1. \quad (4.44)$$

Given that both u and v are in $L^1(\mathbb{R}^n; \mathcal{F}(Y))$, we have that $w \in L^1(\mathbb{R}^{2n}; \mathcal{F}(Y))$. By Fubini's Theorem for Banach-valued functions (see [39]), we have that

$$\int_{\mathbb{R}^n} \int_{\mathbb{R}^n} u(x - z)v(z) dx dz = \int_{\mathbb{R}^n} \int_{\mathbb{R}^n} u(x - z)v(z) dz dx, \quad (4.45)$$

therefore, $u * v$ is integrable. Furthermore,

$$\int_{\mathbb{R}^n} \|(u * v)(x)\|_{\mathcal{F}(Y)} dx \leq \int_{\mathbb{R}^n} \int_{\mathbb{R}^n} \|u(x - z)\|_{\mathcal{F}(Y)} \|v(z)\|_{\mathcal{F}(Y)} dz dx < \infty. \quad (4.46)$$

Thus, $u * v \in L^1(\mathbb{R}^n; \mathcal{F}(Y))$; i.e., $L^1(\mathbb{R}^n; \mathcal{F}(Y))$ is closed under the convolution operation.

Regarding $\mathfrak{F}\{u * v\}$, this is given by

$$\mathfrak{F}\{u * v\} = \int_{\mathbb{R}^n} e^{-i\omega \cdot x} \int_{\mathbb{R}^n} u(x-z)v(z)dzdx \quad (4.47)$$

$$= \int_{\mathbb{R}^n} \int_{\mathbb{R}^n} e^{-i\omega \cdot (t+z)} u(t)v(z)dzdt \quad (4.48)$$

$$= \left(\int_{\mathbb{R}^n} e^{-i\omega \cdot t} u(t)dt \right) \left(\int_{\mathbb{R}^n} e^{-i\omega \cdot z} v(z)dz \right) \quad (4.49)$$

$$= \mathfrak{F}\{u\}\mathfrak{F}\{v\}. \quad (4.50)$$

□

Example 4.2.1 (Rectangular Function). Let $u : \mathbb{R} \rightarrow L^\infty(\mathbb{R})$ be given by

$$u(x)(y) = \begin{cases} 1, & |x| \leq 1 \\ 0, & |x| > 1. \end{cases} \quad (4.51)$$

Its Fourier transform is given by

$$U(\omega)(y) = \int_{-1}^{+1} e^{-i\omega x} dx = \frac{\sin(\omega)}{\omega}. \quad (4.52)$$

In other words,

$$U(\omega)(y) = \begin{cases} \frac{\sin(\omega)}{\omega}, & \omega \neq 0 \\ 1, & \omega = 0. \end{cases} \quad (4.53)$$

Observe that U is an element of $L^\infty(\mathbb{R}, L^\infty(\mathbb{R}))$ but it is not in $L^1(\mathbb{R}, L^\infty(\mathbb{R}))$, thus, the inverse Fourier transform is not defined for this FVM.

Example 4.2.2 (Gaussian Function). Let $u : \mathbb{R} \rightarrow L^1(\mathbb{R})$ be the following FVM

$$u(x)(y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}. \quad (4.54)$$

We have that the Fourier transform is equal to

$$U(\omega)(y) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{y^2}{2\sigma^2}} \int_{-\infty}^{+\infty} \left(\frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{x^2}{2\sigma^2}} \right) e^{-i\omega x} dx = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{\sigma^2\omega^2}{2} - \frac{y^2}{2\sigma^2}}. \quad (4.55)$$

Notice that in this case $U \in L^1(\mathbb{R}, L^1(\mathbb{R}))$, therefore, its inverse Fourier transform can be computed:

$$\mathfrak{F}^{-1}\{U\}(y) = \left(\frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{y^2}{2\sigma^2}} \right) \left(\frac{1}{2\pi} \int_{-\infty}^{+\infty} e^{-\frac{\sigma^2\omega^2}{2}} e^{i\omega x} d\omega \right) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}. \quad (4.56)$$




Figure 4.1: The Fourier transform of the rectangular function given in Eq. (4.51).




Figure 4.2: The Fourier transform of the Gaussian function of Example 4.2.2. In this case, $\sigma = \sqrt{3}$.

4.3 Applications

Under certain assumptions, the Fourier transform can be employed to find the particular solutions of certain ordinary differential equations (ODE). Suppose we wish to find $u : \mathbb{R} \rightarrow \mathcal{F}(Y)$ such that

$$u' + u = v. \quad (4.57)$$

Assuming that all u' , u and v are in $L^1(\mathbb{R}; \mathcal{F}(Y))$, we can apply the Fourier transform to both sides of the ODE, which yields

$$U(\omega) = \frac{V(\omega)}{\omega + 1}, \quad (4.58)$$

where $V(\omega) = \mathfrak{F}\{v\}$. If $U \in L^1(\mathbb{R}; \mathcal{F}(Y))$, we obtain that

$$u(x) = \mathfrak{F}^{-1} \left\{ \frac{V(\omega)}{\omega + 1} \right\}, \quad (4.59)$$

Observe that this approach is equivalent to employ the Fourier transform for solving the classical partial differential equations that arise in practical applications—e.g., the heat equation [64]. In fact, in [69], the authors employ Fourier techniques to study the periodic solutions of second order differential equations that take values in a Banach space.

Regarding practical implementations, to get some insight into the behaviour of the continuous Fourier transform of real images, we applied the *fast Fourier transform* (FFT) to hyperspectral datasets of size $M \times N \times P$; that is, the hyperspectral image has P bands and its spatial domain is an array of size $M \times N$. In the continuous setting, we define a hyperspectral image as a FVM $u : X \subset \mathbb{R}^2 \rightarrow L^2(\mathbb{R})$; i.e., u is a mapping that assigns to each element $x \in X$ a spectral function of finite energy. The continuous Fourier transform of such a mapping is also a FVM of the form $U : \Omega \subset \mathbb{R}^2 \rightarrow L^2(\mathbb{R})$, where Ω is the appropriate set of frequencies. Notice that the elements on $L^2(\mathbb{R})$ may be complex-valued functions. The discrete approximation of U can be obtained by computing the FFT of each band of the discrete hyperspectral image independently.

In this setting, it is interesting to see how the correlation between bands changes before and after computing the discrete approximation of U of a given hyperspectral dataset. In Figure 4.3, the image on the left corresponds to a depiction of the correlation matrix of all the bands of the hyperspectral image Indian Pines. This image has 224 bands, thus the matrix of correlation coefficients is of size 224×224 . Each entry (k, l) of this matrix is the value of the correlation coefficient C_{kl} between bands u_k and u_l , which is given by

$$C_{kl} = \frac{\sum_{i_1=1}^N \sum_{i_2=1}^M (u_k(i_1, i_2) - \bar{u}_k)(u_l(i_1, i_2) - \bar{u}_l)}{\sqrt{\sum_{i_1=1}^N \sum_{i_2=1}^M (u_k(i_1, i_2) - \bar{u}_k)^2} \sqrt{\sum_{i_1=1}^N \sum_{i_2=1}^M (u_l(i_1, i_2) - \bar{u}_l)^2}}, \quad (4.60)$$

where \bar{u}_k and \bar{u}_l are the mean values of u_k and u_l respectively:

$$\bar{u}_k = \frac{1}{MN} \sum_{i_1=1}^N \sum_{i_2=1}^M u_k(i_1, i_2), \quad \text{and} \quad \bar{u}_l = \frac{1}{MN} \sum_{i_1=1}^N \sum_{i_2=1}^M u_l(i_1, i_2). \quad (4.61)$$

As expected, this matrix is symmetric. The image on the right is the matrix of correlation coefficients upon computation of the magnitude of the FFT of each band. It can be observed that upon transformation, the bands tend to be more correlated in the frequency domain as opposed to their spatial counterparts. This suggests that the functions $U(\omega)(y)$ tend to be more regular across the spectral domain. This should not be surprising since the operator \mathfrak{F} may transform discontinuous functions in the spatial domain into continuous functions in the frequency domain (see example 4.2.1). On the other hand, this example also shows that information across bands in the frequency domain has less “variation” than in the spatial domain, which is an example of the compression capabilities of the Fourier transform. Another example is shown in Figure 4.4. In this figure, the correlation matrices correspond to the hyperspectral image Salinas. Both Indian Pines and Salinas can be found in [38].




Figure 4.3: Matrices of correlation coefficients between bands of the hyperspectral image Indian Pines. Image on the left corresponds to the correlations in the spatial domain, whereas the image on the right shows how this correlation between bands tends to be greater in the frequency domain.




Figure 4.4: Matrices of correlation coefficients between bands of the hyperspectral image Salinas. As in the previous figure, the image on the left correspond to the correlations in the spatial domain, whereas the image on the right shows the correlation between bands in the frequency domain.



Figure 4.5: The image on the left is the band number 23 of the hyperspectral image Indian Pines. The image on the right corresponds to band number 57 of the hyperspectral image Salinas.

Chapter 5

A Simple Class of Fractal Transforms for Hyperspectral Images

In [115], we examined some basic self-similarity properties of hyperspectral (HS) images, considering them as FVMs of a base (or pixel) space X to a suitable (spectral) function space. At each location or pixel $x \in X$, the hyperspectral image mapping $u(x)$ is a *function* that is supported on a suitable domain Y . In practical applications, of course, HS images are digitized: Both the base space X and spectral domain Y are discretized so that $u(x)$ is a vector.

Earlier studies of greyscale images [2, 19] have shown that most subblocks of natural images are well approximated (using various forms of affine greyscale mappings) by a number of other subblocks of the image. Such *image self-similarity* is responsible, at least in part, for the effectiveness of various non-local image processing schemes, including nonlocal-means denoising [22], fractal image coding [125, 78] and a variety of other methods devoted to image enhancement, e.g., [37, 45, 48, 51, 55]. The study in [115] shows that HS images are also quite self-similar, in the sense that “data cubes”, namely, M -channel vectors supported over $n \times n$ -pixel subblocks of the HS image are well approximated by a number of other data cubes of the image. Moreover, the spectral functions over individual pixels demonstrate a remarkable degree of correlation with each other, not only locally but over the entire image. This suggests that various nonlocal image processing schemes which rely on self-similarity should be quite effective for HS images.

In Section 5.2 of this chapter, we provide the mathematical formalism for a particular class of *affine fractal transforms* on the space of function-valued HS images and show that under certain conditions, a fractal transform T can be contractive. From Banach’s Fixed

Point Theorem, this implies the existence of a fixed point HS image \bar{u} such that $T\bar{u} = \bar{u}$. This leads to the *inverse problem* of fractal image coding, namely, given an HS image u , find a fractal transform T with fixed point \bar{u} that approximates u to a sufficient degree of accuracy. As in the case of fractal coding of greyscale images, this problem can be solved by means of *collage coding*, i.e., find a fractal transform T that maps the HS image u as close to itself as possible.

One of the original motivations for fractal image coding was *image compression* [7, 125, 78]. As in the case of standard transform coding of images, it was found that much less computer memory was required to store the parameters defining the block-based fractal transform T of an image u . Moreover, the fixed-point approximation \bar{u} to u can be constructed by iteration of the transform T . Fractal image coding has been shown to be effective in performing a number of other image processing tasks, for example, denoising [58] and super-resolution [82].

In Section 5.4, we examine in more detail a block fractal coding scheme briefly introduced in [115], deriving sufficient conditions for contractivity of the associated fractal transform T . We also present the results of some computations on a hyperspectral image. However, it is not our purpose to investigate the compression capabilities of this fractal coding scheme nor to compare it with other compression schemes.

Acknowledging the tremendous amount of work that has been done on hyperspectral images, e.g., [27, 88], we mention that our work is intended to complement the well-established notion that hyperspectral images generally exhibit a high degree of correlation which can be exploited for the purposes of image enhancement.

5.1 A Complete Metric Space $(\mathbf{Z}, d_{\mathbf{Z}})$ of Hyperspectral Images

We consider hyperspectral images as function-valued mappings of a *base space* X to an appropriate space of *spectral functions* $\mathcal{F}(Y)$, along the lines established in [115, 84]. In this chapter, the ingredients of our formalism are as follows:

- **The base space X :** The compact support of the hyperspectral images, with metric d_X . For convenience, $X = [0, 1]^n$, where $n = 1, 2$ or 3 .
- **The range or spectral space $\mathcal{F}(Y)$:** The space $L^2(\mathbb{R}_s)$ of square-integrable functions supported on a compact set $\mathbb{R}_s \subset \mathbb{R}_+$, where $\mathbb{R}_+ = \{y \in \mathbb{R} \mid y \geq 0\}$. $L^2(\mathbb{R}_s)$ is

a Hilbert space with the standard definition of the inner product, i.e.,

$$\langle f, g \rangle = \int_{\mathbb{R}_s} f(y) g(y) dy, \quad \forall f, g \in L^2(\mathbb{R}_s). \quad (5.1)$$

This inner product defines a norm on $L^2(\mathbb{R}_s)$, to be denoted as $\|\cdot\|_{L^2(\mathbb{R}_s)}$.

We now let Z denote the set of all FVMs from X to $L^2(\mathbb{R}_s)$. Given a hyperspectral image $u \in Z$, its value $u(x)$ at a particular location $x \in X$ will be a function—more precisely, an element of the space $L^2(\mathbb{R}_s)$. Following the same prescription as in [84], the norm $\|\cdot\|_{L^2(\mathbb{R}_s)}$ arising from Eq. (5.1) may be used to define a norm $\|\cdot\|_Z$ on Z which, in turn, defines a metric d_Z on Z . The distance between two hyperspectral images $u, v \in Z$ will then be defined as

$$d_Z(u, v) = \sqrt{\int_X \|u(x) - v(x)\|_{L^2(\mathbb{R}_s)}^2 dx}. \quad (5.2)$$

By Theorem 3.2.2, we have that the space $L^2(X; L^2(\mathbb{R}_s))$ is complete, which implies that the metric space (Z, d_Z) of hyperspectral images is complete as well.

5.2 A Class of Fractal Transforms on (Z, d_Z)

We now list the ingredients for a class of fractal transforms on the space of HS images introduced above. For simplicity (especially as far as notation is concerned), we assume that our HS images are “one-dimensional,” i.e., $X = [0, 1]$. The extension to $[0, 1]^n$, in particular, $n = 2$, is straightforward.

1. A set of N one-to-one, affine contraction mappings $w_i : X \rightarrow X$, $w_i(x) = s_i x + a_i$, $x \in X$, with the condition that $\cup_{i=1}^N w_i(X) = X$. In other words, the contracted copies, or “tiles” of X , $w_i(X)$, cover X .
2. Associated with each map w_i are the following:
 - (a) A scalar $\alpha_i \in \mathbb{R}$ and
 - (b) A function $\beta_i : \mathbb{R}_s \rightarrow \mathbb{R}_+$, $\beta_i \in L^2(\mathbb{R}_s)$.

The action of the fractal transform $T : Z \rightarrow Z$ defined by the above is as follows: For a $u \in Z$ and any $x \in X$,

$$v(x) = (Tu)(x) = \sum_{i=1}^N ' [\alpha_i u(w_i^{-1}(x)) + \beta_i]. \quad (5.3)$$

The prime on the summation signifies that we sum over only those $i \in \{1, 2, \dots, N\}$ for which the preimage $w_i^{-1}(x)$ exists, i.e., those i for which $x \in w_i(X)$.

The above formulation represents a generalization of the standard fractal transform for greyscale images. The “value” of the HS image $v(x) = (Tu)(x)$ at a point $x \in X$ is a *spectral function*, i.e., $v(x) \in L^2(\mathbb{R}_s)$. Furthermore, the values of $v(x)$ at $y \in \mathbb{R}_s$ are given by

$$v(x)(y) = (Tu)(x)(y) = \sum_{i=1}^N ' [\alpha_i u(w_i^{-1}(x))(y) + \beta_i(y)]. \quad (5.4)$$

The function $\beta_i(y)$ replaces the traditional constant β_i employed in standard fractal transforms for (single-valued) images [125, 78].

Another way of viewing this procedure is as follows: For each $x \in X$, N copies of the function $u(x)$ are first placed at the points $w_i(x)$, $1 \leq i \leq N$. Each of these copies is then altered in the spectral direction by multiplication by the appropriate α_i factor followed by the addition of the β_i function. If two or more modified copies are situated at a point $x \in X$, then they are added together to produce the function $v(x) = (Tu)(x)$.

Theorem 5.2.1. *Given the fractal transform T defined above, for any $u, v \in Z$,*

$$d_Z(Tu, Tv) \leq K d_Z(u, v), \quad (5.5)$$

where

$$K = \sum_{i=1}^N |s_i|^{1/2} |\alpha_i| \geq 0. \quad (5.6)$$

Proof. This can be proved by simply finding a bound for $d_Z(Tu, Tv)$:

$$d_Z(Tu, Tv) = \left(\int_X \left\| \sum_{i=1}^N \alpha_i (u(w_i^{-1}(x)) - v(w_i^{-1}(x))) \right\|_{L^2(\mathbb{R}_s)}^2 dx \right)^{1/2} \quad (5.7)$$

$$\leq \sum_{i=1}^N \left(\int_{X_i} \left\| \alpha_i (u(w_i^{-1}(x)) - v(w_i^{-1}(x))) \right\|_{L^2(\mathbb{R}_s)}^2 dx \right)^{1/2} \quad (5.8)$$

$$= \sum_{i=1}^N \left(|s_i| \int_X \left\| \alpha_i (u(z) - v(z)) \right\|_{L^2(\mathbb{R}_s)}^2 dz \right)^{1/2} \quad (5.9)$$

$$= \left(\sum_{i=1}^N |s_i|^{1/2} |\alpha_i| \right) \left(\int_X \|u(z) - v(z)\|_{L^2(\mathbb{R}_s)}^2 dz \right)^{1/2} \quad (5.10)$$

$$= K d_Z(u, v). \quad (5.11)$$

□

The following is a consequence of Banach's Fixed Point Theorem.

Corollary 5.2.1.1. *If $K < 1$, i.e., T is contractive on Z , then there exists a unique $\bar{u} \in Z$, the fixed point of T , such that $\bar{u} = T\bar{u}$. Furthermore, let $u_0 \in Z$ be any "seed" for the iteration sequence $u_{n+1} = Tu_n$. Then $u_n \rightarrow \bar{u}$ as $n \rightarrow \infty$, i.e., $d_Z(u_n, \bar{u}) \rightarrow 0$.*

5.3 Inverse problem for fractal transforms on (Z, d_Z)

We now wish to consider the following inverse problem, which includes fractal image coding [125, 78] as a special case: Given a target element $u \in Z$, find a contractive fractal transform $T : Z \rightarrow Z$ such that its fixed point \bar{u} approximates u to a desired accuracy, i.e., $d_Z(\bar{u}, u)$ is sufficiently small. Given the complicated nature of the fractal transform, such direct inverse problems are very difficult. An enormous simplification is yielded by the following consequence of Banach's Fixed Point Theorem, known in the fractal coding literature as the *Collage Theorem* [6].

Theorem 5.3.1 (Collage Theorem). *If $K < 1$, i.e., $T : Z \rightarrow Z$ is contractive with fixed point $\bar{u} \in Y$, then for any $u \in Z$,*

$$d_Z(u, \bar{u}) \leq \frac{1}{1-K} d_Z(u, Tu). \quad (5.12)$$

In *collage coding* [72], one looks for a contractive fractal transform T that maps the target u as close as possible to itself, in an effort to make the so-called *collage error*, $d_Z(u, Tu)$, as small as possible.

As mentioned in the Introduction, one of the original motivations for fractal image coding was *image compression* [7, 125, 78].

5.4 Block fractal transforms on digital hyperspectral images

The remainder of this chapter will be concerned with digital HS images supported on an $N_1 \times N_2$ -pixel array, M channels per pixel. Formally, a digital HS image can be represented by a vector-valued image function, $u : X \rightarrow \mathbb{R}_+^M$, where $X = \{1, 2, \dots, N_1\} \times \{1, 2, \dots, N_2\}$ is the *base* or *pixel space* and \mathbb{R}_+^M , the nonnegative orthant of \mathbb{R}^M , is the *spectral space*. At a pixel location $(i_1, i_2) \in X$, the hyperspectral image function $u(i_1, i_2)$ is a non-negative M -vector with components $u_k(i_1, i_2)$, $1 \leq k \leq M$. We shall refer to this vector as the *spectral function* at pixel (i_1, i_2) .

Most, if not all, fractal image coding methods employ *block-based* transforms, where subblocks of an image are mapped onto small subblocks of the image, following the original method of Jacquin [67]. Here it might be tempting to simply consider an HS as a “cube” of data and simply move 3-D sub-cubes to other sub-cubes. This, however, is contrary to our spirit of function-valued HS images. We wish to see what can be accomplished by keeping the spectral functions intact, or perhaps “partially intact” as we discuss later in this chapter.

It would also be tempting to perform fractal image coding on each channel of an HS image separately. Once again, this is contrary to the spirit of FVMs and the desire to keep spectral functions intact. Our goal is to exploit both the spatial self-similarity of channels as well as the correlation between them.

Here we outline a very simple block-based fractal transform for HS images that keeps spectral functions intact. As done in [2] for greyscale images, we let $\mathcal{R}^{(n)}$ denote a set of non-overlapping $n \times n$ -pixel *range* subblocks R_i , such that $X = \cup_i R_i$, i.e., $\mathcal{R}^{(n)}$ forms a partition of the pixel space X . Furthermore, let $u(R_i)$ denote the portion of the HS image function u that is supported on subblock $R_i \subset X$. In this discrete setting, $u(R_i)$ is an $n \times n \times M$ “data cube” of nonnegative real numbers. We also introduce an associated set $\mathcal{D}^{(m)}$ of $m \times m$ -pixel *domain* subblocks $D_k \subset X$, where $m = 2n$. This set need not be non-overlapping, but the blocks should cover the support X , i.e., $\cup_i D_i = X$.

Given an M -channel digital HS image u , the fractal transform operator T will be constructed as follows: For each image subblock $R_i \in \mathcal{R}^{(n)}$, we choose from $\mathcal{D}^{(m)}$ a domain block $u(D_{j(i)})$ in order to produce an approximation of the form,

$$u(R_i) \approx (Tu)(R_i) = \alpha_i \widehat{u(D_{j(i)})} + \underline{\beta}_i, \quad 1 \leq i \leq N_R. \quad (5.13)$$

(The choice of the *best* domain block will be discussed a little later in this section.) Here, $\underline{\beta}_i = (\beta_{i1}, \beta_{i2}, \dots, \beta_{iM})$ is an M -vector, which plays the role of the $\beta_i(t)$ function in Eq. (5.4). N_R denotes the cardinality of the set $\mathcal{R}^{(n)}$ and the wide hat denotes an appropriate $2n \times 2n \rightarrow n \times n$ pixel decimation operation which produces the geometric contraction in discrete pixel space. Note that only **one** constant α_i is employed for all M -channels supported on the range block R_i . (In other words, the M -channels are **not** coded separately.)

For notational convenience, the approximation problem in Eq. (5.13) may be expressed in the form

$$y_{lm} \approx \alpha x_{lm} + \beta_m, \quad 1 \leq l \leq N, \quad 1 \leq m \leq M, \quad (5.14)$$

where $N = N_1 \times N_2$. (For simplicity, the $N_1 \times N_2$ matrices in pixel space have been converted into N -vectors.) The “stack” of M N -vectors y_{lm} contain the elements of the range block $u(R_i)$ being approximated. The parameters α and $\beta_m, 1 \leq m \leq M$, which minimize the squared L^2 distance,

$$\Delta^2 = \sum_{l=1}^N \sum_{m=1}^M (y_{lm} - \alpha x_{lm} - \beta_m)^2, \quad (5.15)$$

are given by (details in [115])

$$\alpha = \frac{\sum_{m=1}^M \sum_{l=1}^N x_{lm} (y_{lm} - \bar{y}_m)}{\sum_{l=1}^N \sum_{m=1}^M x_{lm}^2 - N \sum_{m=1}^M \bar{x}_m^2} \quad (5.16)$$

and

$$\beta_m = \bar{y}_m - \alpha \bar{x}_m, \quad 1 \leq m \leq M, \quad (5.17)$$

where

$$\bar{x}_m = \frac{1}{N} \sum_{l=1}^N x_{lm}, \quad \bar{y}_m = \frac{1}{N} \sum_{l=1}^N y_{lm}, \quad (5.18)$$

denote the (spatial) means over each channel.

The spectral map parameters, $(\alpha_i, \beta_{i1}, \beta_{i2}, \dots, \beta_{iM})$, for $1 \leq i \leq N_R$, and the set of range-domain assignments $(i, j(i))$, comprise the *fractal code* which defines a fractal transform T . If T is contractive (see below), then its fixed point hyperspectral image \bar{u} may be computed by the iteration procedure $u_{n+1} = Tu_n$, where u_0 is any “seed” image. For convenience, one may employ the zero image $u_0 = 0$. Note that in the special case that $M = 1$, Eq. (5.13) reduces to the usual fractal block transform method for greyscale images.

Technically speaking, Theorem 5.2.1 of Section 5.2 does not apply to block-based fractal transforms since one is not mapping the *entire* image $u(X)$ onto a range block R_i . The determination of L^2 Lipschitz factors such as K in Eq. (5.6) is quite complicated. Fortunately, we may resort to a simplification which is employed in most block coding methods. In the case of digitized images, it is easy to show that the condition $|\alpha_i| < 1$ for all image range blocks $u(R_i)$ is sufficient to guarantee contractivity of the fractal transform T in the L^∞ norm, from which the existence of a unique fixed point \bar{u} of T follows.

We now return to the question of determining the “best” fractal transform T associated with a given image u , i.e., the transform T with fixed point \bar{u} that approximates u as best as possible. Because the range blocks R_i are non-overlapping, the sum of the errors associated with the approximations in Eq. (5.13) defines the total collage error $d_Z(u, Tu)$ on the RHS of Eq. (5.12). Since our goal is to make the approximation error $d_Z(u, \bar{u})$ on the LHS of Eq. (5.12) as small as possible, we choose, for each range block $u(R_i)$, the domain block $u(D_{j(i)})$ which *best approximates* $u(R_i)$. If we let Δ_{ik} denote the error in approximating a range block $u(R_i)$ with a domain block $u(D_k)$, i.e.,

$$\Delta_{ik} = \min_{\alpha, \beta} \left\| u(R_i) - \alpha \widehat{u(D_k)} - \underline{\beta} \right\|_2, \quad (5.19)$$

then the index $j(i)$ of the optimal domain block $u(D_{j(i)})$ associated with $u(R_i)$ is

$$j(i) = \arg \min_k \Delta_{ik}. \quad (5.20)$$

Once again, this is the essence of collage-based fractal image coding.

5.5 Some Numerical Results

We now show some results of our block-based fractal transform as applied to the AVIRIS (Airborne Visible/Infrared Imaging Spectrometer) image, “Yellowstone calibrated scene

0,” a 512 line, 677 samples per line, 224-channel image, available from the Jet Propulsion Laboratory site [91]. The computations reported below were performed on a 512×512 -pixel section of this image. The range blocks R_i used were $N_R = 4096$ nonoverlapping 8×8 -pixel blocks. The domain blocks D_i were $N_D = 1048$ nonoverlapping 16×16 -pixel blocks.

In Figure 5.1 (top) is presented a histogram plot of the distribution of errors Δ_{ij} associated with *all* possible approximations in Eq. (5.19), i.e., for each range block R_i , we consider *all* possible domain blocks D_k . This distribution is very similar to the distribution of “Case 3” same-scale approximation errors in [115], where the domain blocks D_k were the same size as the range blocks R_i . The strong peaking of these distributions near zero error shows that many range blocks $u(R_i)$ of the AVIRIS hyperspectral image are well approximated by affinely modified domain blocks $u(D_k)$.

For reference purposes, Figure 5.1 (bottom) shows a histogram plot of the error distribution that results when no affine mapping is employed in Eq. (5.13), i.e., $\alpha_i = 1$ and $\beta_{i1} = \dots = \beta_{iM} = 0$. Here, each range block $u(R_i)$ is simply approximated by the decimated domain block $u(D_k)$ with error given by

$$\Delta_{ik}^0 = \|u(R_i) - \widehat{u(D_k)}\|_2. \quad (5.21)$$

The distribution of these errors is more diffuse and, in fact, very similar to the “Case 1” same-scale approximation errors (also with no affine mappings) presented in [115]. Clearly, the use of affine maps yields a significant improvement in approximation.

In Figure 5.2 is presented the distribution of α_i coefficients associated with the optimal range blocks $u(R_i)$ employed in the fractal code. The spikes at ± 0.95 are caused by “clamping”. For a relatively small number of range blocks, the optimal value of α lies outside the range $(-1, 1)$. In such cases, the α parameter is “clamped” to ± 0.95 and the β vector is computed from this value. This clamping is performed in an effort to insure the stability and numerical convergence of the iteration procedure $u_{n+1} = Tu_n$ used to construct the fixed point attractor \bar{u} of the fractal transform T . The effect of this sub-optimal fitting of a few blocks is negligible.

The absence of a strong peak at $\alpha = 0$ in Figure 5.2 represents a significant difference between the α -parameter distributions for HS images and those observed for most (single-valued) greyscale images. The near-zero peaking in greyscale images is generally due to the fact that most of their blocks are quite “flat,” i.e., have low variance. As such, they can be well approximated by constant blocks which are produced by α -values close to zero, if not equal to, zero. This does not seem to be the case for HS images, mostly because the image subblocks are “cubes,” i.e., collections of vectors.




Figure 5.1: Distributions of errors Δ_{ij} over the interval $[0, 0.02]$ obtained by approximating non-overlapping 8×8 -pixel range blocks $u(R_i)$ with all possible decimated 16×16 -pixel domain blocks of AVIRIS hyperspectral image. **Top:** Using affine mapping, Eq. (5.13). **Bottom:** No affine mapping, Eq. (5.21).





Figure 5.2: α_i parameters over the interval $[-1, 1]$ employed in the “best fit” of AVIRIS HS image range subblocks $u(R_i)$.


Some of the results of this fractal coding procedure are presented in Figure 5.3. The fixed point attractor function \bar{u} of the fractal transform T defined by the fractal code for the AVIRIS image—the domain-range assignments $j(i)$, the scaling parameters α_i and the shift vectors β_i for $1 \leq i \leq 4096$ —was generated using the iteration procedure $u_{n+1} = Tu_n$, starting with the image $u_0 = \underline{0}$. Reasonable convergence was achieved at u_{10} . Channels 20, 120 and 220 of u_{10} are presented in Figure 5.3 on the right along with corresponding channels of the original (uncoded) AVIRIS image on the left.

A closer inspection of these fractal-based approximations shows that the most noticeable errors are located in regions of high image variation, i.e., detailed textures. Furthermore, because of the block-based nature of the coding method, blurriness is also accompanied by blockiness. Such degradations (which are also observed in the fractal coding of standard greyscale images [125, 78]) are also seen in images that have been highly compressed using the JPEG. (Recall that the JPEG compression method employs 8×8 -pixel blocks.) For purposes of comparison, a JPEG-encoded approximation to Channel 20 is presented in Figure 5.4. The Quality Factor of 25 was chosen in order to yield a compression ratio that was roughly equal to that estimated for the fractal coding method (about 40:1).


As might be expected, the accuracy of the simple fractal coding method employed



Channel 20. Left: Original. Right: Fractal-based approximation.



Channel 120. Left: Original. Right: Fractal-based approximation.



Channel 220. Left: Original. Right: Fractal-based approximation.

Figure 5.3: Some channels of the attractor \bar{u} of the fractal transform T obtained by fractally coding the AVIRIS hyperspectral image using 8×8 -pixel range blocks and 16×16 -domain blocks.



Channel 20. Left: Fractal. Right: JPEG (QF 25).

Figure 5.4: 8×8 -pixel block approximations to Channel 20 of the AVIRIS hyperspectral image.

here may be increased by using smaller range blocks, R_i , for example, 4×4 -pixel blocks. However, this increase in accuracy is accomplished at a price—the need to store more fractal code parameters. Such a trade-off between accuracy and storage is the basis of *rate-distortion analysis*—a fundamental issue of image and data compression—which is beyond the scope of this thesis.

Admittedly, the use of a regular grid of $n \times n$ -pixel blocks is suboptimal. There are more sophisticated methods of partitioning the pixel space X , e.g., quadtree decomposition, rectangular and triangular blocks—for more details, the reader is referred to [125, 78]. We simply mention here that the mathematical framework presented here can easily be adapted to accommodate such schemes.

5.6 Final Remarks

In this chapter, we have presented a complete metric space (Z, d_Z) of FVMs that is suitable for the representation of hyperspectral images. As well, a class of fractal transforms over this space has been introduced. Under appropriate conditions, a fractal transform T can be contractive. This leads to an inverse problem in which an element $u \in Y$ is approximated by the fixed point \bar{u} of a fractal transform.

We then consider a simple class of block-based fractal transforms—a slight modification of the transforms introduced earlier—which are particularly suited for digital hyperspectral images, approximating spectral function vectors at one point $x \in X$ by modified vectors

from other points. Such an approach is consistent with the method of *nonlocal image processing*.

A block-based transform employing 8×8 -pixel range blocks R_i was then employed to fractally code a standard hyperspectral image, the 224-channel AVIRIS “Yellowstone” image. Of course, this transform is in no way optimal. The results could be improved slightly by examining all eight possibilities of mapping square blocks to square blocks, i.e., rotations and inversions. One could also employ smaller range blocks, or a variety of range block sizes, as is done in quadtree-based fractal coding of greyscale images [125, 78]. Of course, the price for any improvement is increased computational time.

A further improvement may also be obtained if significant correlations exist between groups of channels in a hyperspectral image. For example, Figure 5.5 shows a plot of the correlations between all pairs of channels u_k and u_l , $1 \leq k, l \leq M$, of the AVIRIS image employed in this study, where $M = 224$. Recalling that each channel u_k is an $N_1 \times N_2$ array, the correlation C_{kl} between channels u_k and u_l is computed in the standard fashion, i.e.,

$$C_{kl} = \frac{\sigma_{kl}}{\sigma_k \sigma_l}. \quad (5.22)$$

Here

$$\sigma_k = \sqrt{\sum_{i_1=1}^{N_1} \sum_{i_2=1}^{N_2} (u_k(i_1, i_2) - \bar{u}_k)^2} \quad (5.23)$$

and

$$\sigma_{kl} = \sum_{i_1=1}^{N_1} \sum_{i_2=1}^{N_2} (u_k(i_1, i_2) - \bar{u}_k)(u_l(i_1, i_2) - \bar{u}_l), \quad (5.24)$$

where \bar{u}_k denotes the mean of the array u_k . (The usual factor $1/(N_1 N_2 - 1)$ accompanying each double summation can be omitted from the formulas for σ_k and σ_{kl} since it cancels out of the formula for C_{kl} .)

The block nature of the plot in Figure 5.5 shows that the channels can be divided into at least three subgroups, the channels within each subgroup having higher correlations than with those of other subgroups. As such, it may be advantageous to consider separate fractal transforms, each of which operates within a particular group of channels.

The plot in Figure 5.6 of the correlations between channels in the “hyperspectral fern” image [54], which was also examined [115], shows a much lesser amount of internal grouping of channels. As such, a single fractal transform over all channels may suffice.




Figure 5.5: Correlations C_{kl} between channels of AVIRIS hyperspectral image, demonstrating the existence of several subgroups of highly correlated channels.




Figure 5.6: Correlations C_{kl} between channels of hyperspectral fern image, which demonstrates a lesser degree of grouping of correlated channels.

Chapter 6

Denoising of Hyperspectral Images

In this chapter, we consider the problem of denoising digital hyperspectral (HS) images obtained from remote sensing of the Earth’s surface. In this case, the HS image associated with a given surface region R is comprised of a set of reflectance values—ratios of reflected energy vs. incident energy—of electromagnetic radiation at a number of frequencies (or, equivalently, wavelengths) at each pixel location in R . The number of frequencies depends upon the spectral resolution of the sensor of the hyperspectral camera and may range from tens to hundreds. For example, the well-known AVIRIS (Airborne Visible/Infrared Imaging Spectrometer) satellite images [123] typically contain 224 frequencies.

Suppose that a region R of the Earth’s surface is represented by an $M \times N$ pixel array and that associated with each pixel in the array there are P reflectance values. The first (i), and most obvious, way of viewing this HS data set is as a $M \times N \times P$ “data cube.” The correlations between neighbouring entries of this cube give rise to two additional and complementary ways of viewing this 3-D data set: (ii) as a collection of P images of region R at different frequencies—often referred to as *spectral channels* or *frequency/wavelength bands*—and (iii) as a collection of $M \times N$ P -vectors, each of which corresponds to a given pixel location (i, j) of R —frequently referred to as the *spectral function* or, simply, *spectrum* at (i, j) . These three views of the HS image will play a central theme in this chapter.

Let us very briefly recall the importance of spectral functions. Since different materials, e.g., minerals, water, vegetation, exhibit different reflectance spectra, the latter serve as “spectral fingerprints”. The spectrum at a pixel (i, j) makes it possible to determine the composition/nature of material situated at that location. This makes the study of HS images useful in a variety of applications, including agriculture, mineralogy, geography and surveillance, the latter involving hyperspectral imaging tasks such as target detection and

classification [107]. In light of the acquisition process, HS images are, as in the case of many other images, prone to contamination by noise which can compromise the performance of such tasks. As a result, it is desirable to develop reconstruction techniques that recover good approximations of noise-free HS images.

Indeed, many different methods for denoising HS data have been proposed. For example, in [81], diffusion-based filtering is adapted to HS images. The proposed method consists of two diffusion processes, one confined to each band of the HS image, and the another restricted to the spectral domain. The overall anisotropic diffusion is basically a combination of these two processes, which are carried out in a controlled fashion. In [95], a rather novel wavelet-based denoising approach is proposed. This method transforms the HS data set into a spectral-derivative domain, in which the irregularity of noise is more easily detected. The transformed HS image is denoised using wavelet shrinkage (WS) independently in both the spatial and spectral domains. A reconstruction is then obtained by first computing the corresponding inverse wavelet transforms of the denoised data followed by an integration in the spectral direction. Another method that carries out denoising employing WS is presented in [31]. Here, principal component analysis (PCA) is used to decorrelate the most relevant HS data from the noise, most of which is assumed to be contained in the lowest energy components of the transformed data. The noise is removed from these components using WS in both spatial and spectral domains. The denoised HS data set is then retrieved by means of the inverse PCA transform. Variational approaches are proposed in [124, 28]. In [124], a total variation (TV) model that considers the changes of noise intensity present across the bands and pixels of an HS image is presented. In [28], a method that employs a TV model along with sparse representations of each band is also introduced. More approaches can also be found in [103, 98].

In this study, we wish to examine the roles of both spatial (pixel) and spectral domains in the denoising of HS images. For example, is it preferable to focus the denoising in one domain at the expense of the other, or should both domains be considered? In order to shed some light on this and related questions, we compare five different denoising approaches. The main difference between these approaches lies in the way that the HS image is treated, using the three views mentioned earlier, i.e., as a (i) 3-D data “cube”, (ii) a set of frequency bands or (iii) a set of spectral functions.

In our first approach, the denoising process is performed in the spectral domain, corresponding to (iii) above. In particular, we apply L^1 -norm regularization [8, 113, 4] to the spectral functions. In the second approach, the denoising process is performed in the spatial domain, corresponding to (ii) above. As expected, any denoising technique applicable to 2-D signals/images can be employed—here, we focus our attention on the TV approach [105, 24]. Our third approach employs a formulation of vectorial TV to denoise the entire

HS image at once [60, 15], corresponding to (i) above. As for the fourth approach, an HS image is viewed as a collection of both spectra and frequency bands. Our method involves a combination of the first two approaches so that denoising is carried out by regularization in both the spatial and spectral domains. To solve this inverse problem we employ the Alternating Direction Method of Multipliers (ADMM) [12]. Finally, in the fifth approach, we present a simple denoising method based on the FVM approach. In this case, HS data sets are considered as a whole, which corresponds to (i) above; however, it will be seen that approaches (ii) and (iii) arise naturally depending on how the HS images are represented. Experimental results are then presented so that the performance of these methods can be compared.

6.1 Denoising Approaches

In practice, the strengths of the denoising process across spatial and spectral domains of an HS image should be different. Even within the spatial domain, different features such as edges and flat regions should not be denoised with the same intensity. In addition, it is quite common that the power of noise across bands is not constant [10, 124]. Some methods that address these possible scenarios can be found in the literature, e.g., [124, 95]. Nevertheless, in this study, for the sake of simplicity we assume that the power of the noise is constant over the entire HS data set, i.e., it is independent of the location and band of a given voxel. As such, we consider the simple degradation model $f = u + n$, where f is the noisy observation, u is the noiseless HS data we wish to recover, and n is additive white Gaussian noise (AWGN). In this case, f , n and u are considered as $M \times N \times P$ HS data cubes. Moreover, for the remainder of this chapter, this interpretation of HS images as 3-D discrete data sets is the one that we will consider, unless otherwise stated.

Despite that the model mentioned above may not always be well suited for noise in HS images, it will be seen that some of the methods presented below can be easily adapted for different scenarios in which the noise characteristics change over space and wavelength.

6.1.1 Denoising of hyperspectral images as a collection of spectra

Here, we view f as a collection of $M \times N$ spectral functions, each of which is represented by a P -vector. As such, the denoising problem is split into $M \times N$ independent subproblems, each involving the denoising of a particular spectral function.

Each spectral function may be denoised with any of the available denoising techniques that can be applied to 1-D signals, e.g., wavelet shrinkage, linear filtering, 1-D total variation denoising, etc.. Here, however, we investigate the effectiveness of L^1 -norm regularization, primarily because of the good performance of methods that exploit sparse representations of signals [8, 113, 4].

In this approach, we solve the following sparse approximation problem independently at each pixel (i, j) ,

$$\min_{c_{ij}} \left\{ \frac{1}{2} \|Dc_{ij} - s_{ij}\|_2^2 + \lambda \|c_{ij}\|_1 \right\}, \quad (6.1)$$

where s_{ij} denotes the noisy spectrum, D is an appropriate transformation matrix (e.g., frame, random matrix, etc.), and c_{ij} is the set of coefficients that is to be recovered at the pixel location (i, j) .

In the literature, many algorithms for solving (6.1) can be found [8, 113, 4], however, we focus our attention on the special case in which the matrix D is an orthogonal transformation (e.g., DCT, wavelet transform, Fourier matrix, etc.). In this particular case, problem (6.1) can be solved by means of the soft thresholding (ST) operator [113, 12].

It is worth pointing out that this approach allows us to change the strength of the denoising process across the spatial domain, i.e., different regularization parameters can be used at different pixels or in different regions of the HS image.

6.1.2 Denoising of hyperspectral images as a collection of bands

In this approach, the denoising process takes place in the spatial domain. Each frequency/wavelength band is treated independently and the denoising problem is split into P independent subproblems. Here, we consider each k -th band u_k as a scalar function $u_k : \Omega \rightarrow \mathbb{R}$, where $\Omega \subset \mathbb{R}^2$ and $1 \leq k \leq P$.

As expected, any denoising method for 2-D images can be employed here, e.g., linear filtering, non-local means denoising, total variation, non-linear filtering, etc.. Nevertheless, in this study, we employed a TV denoising approach for which a number of fast algorithms exist, e.g., [105, 24, 61]. As well, some TV-based denoising methods for HS images have yielded promising results [28, 124].

Our approach, a channel-by-channel TV method in the spatial domain where each band u_k is treated independently, translates to the following approximation problem,

$$\min_{u_k} \left\{ \frac{1}{2\rho} \|u_k - f_k\|_2^2 + \|u_k\|_{TV} \right\}, \quad (6.2)$$

where $\|\cdot\|_{TV}$ is the total variation norm and f_k is the k -th noisy band or channel. To solve this problem numerically, we employ the method introduced by Chambolle in [24], which has received special attention because of its excellent performance. Here, the following definition of the isotropic TV norm is employed,

$$\|u_k\|_{TV} = \int_{\Omega} \|Du_k\|_2 dx = \sup_{\xi_k \in \Xi_k} \left\{ \int_{\Omega} u_k \nabla \cdot \xi_k \, dx \right\}, \quad (6.3)$$

where $\Xi_k = \{\xi_k : \xi_k \in C_c^1(\Omega, \mathbb{R}^2), \|\xi_k(x)\|_2 \leq 1 \, \forall x \in \Omega\}$, and $\nabla \cdot$ is the divergence operator. If $u_k \in C_c^1(\Omega, \mathbb{R})$, $Du_k = \nabla u_k$ in the distributional sense. This approach is convenient since only the integrability, and not the differentiability, of u_k is required.

By using (6.3), Chambolle shows that the optimal solution u_k^* of (6.2) is given by $u_k^* = f_k - \Pi_{\rho\Gamma_k}(f_k)$, where $\Pi_{\rho\Gamma_k}(f_k)$ is the non-linear projection of f_k onto the convex set $\rho\Gamma_k$, and Γ_k is the closure of the set $\{\nabla \cdot \xi_k : \xi_k \in C_c^1(\Omega, \mathbb{R}^2), \|\xi_k(x)\|_2 \leq 1 \, \forall x \in \Omega\}$. Such projection is obtained by solving the following minimization problem:

$$\min_{\|\xi_k(x)\|_2 \leq 1} \{\rho \nabla \cdot \xi_k - f_k\}. \quad (6.4)$$

Thus, we have that for each band the optimal reconstruction is given by $u_k^* = f_k - \rho \nabla \cdot \xi_k^*$.

This approach may easily be modified to accommodate the case in which the power of the noise is not constant throughout the bands. In this case, one can specify the degree of regularization to be applied to each channel independently by means of the parameter ρ .

6.1.3 Denoising of hyperspectral images as a whole

In this case, we view a HS image as a vector-valued function $u : \Omega \rightarrow \mathbb{R}^P$, where $\Omega \subset \mathbb{R}^2$. To denoise it, we follow a variational approach, employing a definition of the Vectorial TV seminorm (VTV).

Given the effectiveness of TV for denoising images—along with its applicability to other image processing tasks such as inpainting, zooming, etc.—many extensions for vector-valued functions have been proposed [15, 60]. Indeed, a practical application already exists for colour images, which are essentially low-dimensional HS images. This approach can easily be extended to HS images, with no required changes to the definitions presented in the literature. In particular, we use Bresson and Chan’s approach [15], which is a generalization of Chambolle’s algorithm for vector-valued functions. Here, the authors extend the Rudin-Osher-Fatemi model [105] as follows,

$$\min_u \left\{ \frac{1}{2\rho} \|u - f\|_{L_2(\Omega; \mathbb{R}^P)}^2 + \|u\|_{VTV} \right\}, \quad (6.5)$$

where f is the noisy observation. The VTV seminorm is defined as

$$\|u\|_{VTV} = \int_{\Omega} \|Du\| d\mathbf{x} = \sup_{\xi \in \Xi} \left\{ \int_{\Omega} \langle u, \nabla \cdot \xi \rangle d\mathbf{x} \right\}; \quad (6.6)$$

where $\Xi = \{\xi : \xi \in C_c^1(\Omega, \mathbb{R}^{P \times 2}), \|\xi(\mathbf{x})\|_2 \leq 1 \ \forall \mathbf{x} \in \Omega\}$; $\langle \cdot, \cdot \rangle$ is the standard Euclidean scalar product in \mathbb{R}^P ; and $\|Du\|^2 = \sum_{k=1}^P \|\nabla u_k\|_2^2$ if $u \in C_c^1(\Omega, \mathbb{R}^P)$, that is, the L^2 norm of the TV norm of all the bands u_k of the HS image u . Substitution of (6.6) into (6.5) yields the following minimization problem:

$$\min_{\|\xi(\mathbf{x})\|_2 \leq 1} \left\{ \left\| \nabla \cdot \xi - \frac{f}{\rho} \right\|_{L_2(\Omega; \mathbb{R}^P)}^2 \right\}, \quad (6.7)$$

whose solution ξ^* is computed using a semi-implicit gradient descent scheme. The solution u^* of the original problem in (6.5) is obtained using $u^* = f - \rho \nabla \cdot \xi^*$. In our case, f is the noisy HS image.

6.1.4 Denoising of hyperspectral images as a collection of both bands and spectra

In this fourth approach, we perform regularization in both the spectral and spatial domains. This can be done in various ways, but we focus our attention on the methods employed previously, that is, TV and L^1 -norm regularization.

Borrowing from our previous methods, we consider the following hybrid scheme. Firstly, denoising in the spectral domain is carried out by solving an optimization problem in which the L^1 norm of a set of coefficients is used as a regularizing term. Secondly, denoising in the spatial domain is performed using a variational approach. In other words, we consider a good reconstruction of the original HS data u from the noisy observation f to be one with bounded variation across bands and with spectral functions that possess sparse representations in a certain domain. In order to find such a reconstruction, we solve the following optimization problem,

$$\min_c \left\{ \frac{1}{2} \|S(c) - f\|_2^2 + \rho \|S(c)\|_{VTV} + \lambda \|c\|_1 \right\}, \quad (6.8)$$

where $S(\cdot)$ is a synthesis operator that reconstructs the HS image from the set of coefficients c . More specifically, at each pixel (i, j) , the operator $S(\cdot)$ recovers the spectrum located at

that pixel location by computing Dc_{ij} , where c_{ij} is the set of coefficients associated with its spectrum.

For solving problem (6.8) we employ ADMM, which is a method well suited for convex optimization and large scale problems [12]. We first need to express (6.8) in ADMM form:

$$\min_{c,u} \left\{ \frac{1}{2} \|u - f\|_2^2 + \rho \|u\|_{\text{VTV}} + \lambda \|c\|_1 \right\} \quad (6.9)$$

subject to $S(c) - u = 0$.

It is well known [12] that this new problem can be solved by forming the augmented Lagrangian and minimizing with respect to the variables c and u in an alternate fashion. Given this, we propose the following ADMM iterations for solving (6.8):

$$c^{n+1} := \min_c \left\{ \frac{1}{2} \left\| S(c) - \frac{f + \delta(u^n - p^n)}{\delta + 1} \right\|_2^2 + \frac{\lambda}{\delta + 1} \|c\|_1 \right\} \quad (6.10)$$

$$u^{n+1} := \min_u \left\{ \frac{\delta}{2\rho} \|u - (S(c^{n+1}) + p^n)\|_2^2 + \|u\|_{\text{VTV}} \right\} \quad (6.11)$$

$$p^{n+1} := p^n + S(c^{n+1}) - u^{n+1}, \quad (6.12)$$

where p is the dual variable associated to the augmented Lagrangian, (6.12) its update, and δ is a penalty parameter. Problem (6.10) can be solved by any algorithm capable of carrying out sparse reconstruction using the L^1 norm as a regularizing term. Problem (6.11) can be addressed using any method employing the vectorial TV norm. In particular, we have used ST to solve (6.10) at each pixel and Bresson and Chan's algorithm for problem (6.11).

It is important to mention that different regularization terms can be used in problem (6.10) since it is solved at each pixel independently. Moreover, problem (6.11) can be solved using our second approach in Section 6.1.2, that is, denoising each band independently, in which case the regularization can be changed from band to band. In other words, our fourth approach may be adapted for denoising with different intensities across both spatial as well as spectral domains.

6.1.5 Denoising of hyperspectral images using FVMs

If HS images are represented as FVMs there are several ways in which the inverse problem of denoising can be addressed. In this section we review some of the different possibilities that arise when the FVM approach is employed. Moreover, it will be seen that some

methods can be understood as the continuous counterparts of the approaches examined earlier.

Let us begin with the simple assumption that the HS images in which we are interested belong to $C^2(X; L^2(Y))$, where X and Y are subsets of \mathbb{R} and \mathbb{R}^2 respectively. Admittedly, this is a quite strong requirement on the HS image in the spatial direction—piecewise $C^2(X; L^2(Y))$ would be more “realistic”. Here, however, we simply wish to examine the consequences of our simple Euler-Lagrange formalism. Also, let (X, Σ, μ) and (Y, T, ν) be finite measure spaces. In order to recover a denoised reconstruction of the contaminated HS image u we minimize the following functional $I(u) : C^2(X; L^2(Y)) \rightarrow \mathbb{R}$:

$$\min_u \left\{ \frac{1}{2} \int_X (\rho \|u(x) - f(x)\|_2^2 + \|u'(x)\|_2^2) dx \right\}; \quad (6.13)$$

that is, we wish to recover HS images that have a smooth variation across the spectral domain. By Theorem 3.3.1, the Euler-Lagrange equation is given by

$$\rho \int_Y (u(x) - f(x))(y) dy - \frac{d}{dx} \int_Y u'(x)(y) dy = 0. \quad (6.14)$$

Assuming that the second derivative of u is integrable with respect to y , we have that

$$\int_Y (\rho(u(x) - f(x)) - u''(x))(y) dy = 0. \quad (6.15)$$

Observe that any solution of the integrand such that its integral over Y is zero is a minimizer of (6.13). Proceeding heuristically, we choose the trivial case, that is, we seek for a solution u such that

$$(\rho(u(x) - f(x)) - u''(x))(y) = 0 \quad (6.16)$$

for μ -almost all $x \in X$ and ν -almost all $y \in Y$. This choice is motivated by the fact that in the finite dimensional case, e.g., $u : X \rightarrow \mathbb{R}^n$, we have a system of n Euler-Lagrange equations in which each equation is equal to zero. On the other hand, it is worthwhile to mention that the second variation of problem (6.13) is given by $D^2I(u; \varphi \chi_Y) = \nu(Y)(\|\varphi\|_2^2 + \|\varphi'\|_2^2)$, where $\varphi \in C^\infty(X)$ and χ_Y is the characteristic function of the set $Y \subset \mathbb{R}^2$. This expression is clearly greater or equal than zero provided that $\nu(Y) \geq 0$, thus, minimization of (6.13) is feasible.

Given the latter, we are to solve the following differential equation

$$\rho(u - f) - u'' = 0. \quad (6.17)$$

A simple approach for addressing this problem is to use the Fourier transform for finding the non-homogeneous solution of Eq. (6.17). We have, then, that the Fourier transform of u is given by

$$U(\omega) = \rho \frac{F(\omega)}{\omega^2 + \rho}, \quad (6.18)$$

where F is the Fourier transform of f . Therefore,

$$u(x) = \mathfrak{F}^{-1} \left\{ \rho \frac{F(\omega)}{\omega^2 + \rho} \right\}. \quad (6.19)$$

Notice that in this case we are filtering each spectral function $u(x)(y)$ independently, therefore, this approach can be seen as the continuous counterpart of denoising HS images as a collection of spectra.

Furthermore, it is also possible to carry out filtering of each band in an independent fashion by defining HS images as FVMs that assign to each $x \in X$ a spectral function; that is, FVMs of the form $u : X \subset \mathbb{R}^2 \rightarrow L^2(Y)$, where $Y \subset \mathbb{R}$. Once again, we assume that HS images are regular enough so that these belong to $C^2(X; L^2(Y))$. In this case, we minimize the following functional:

$$\min_u \left\{ \frac{1}{2} \int_X (\rho \|u(x) - f(x)\|_2^2 + \|\nabla_x u(x)\|_2^2) dx \right\}. \quad (6.20)$$

The corresponding Euler-Lagrange equation is given by

$$\rho \int_Y (u(x) - f(x))(y) dy - \frac{\partial}{\partial x_1} \int_Y \frac{\partial u(x)}{\partial x_1}(y) dy - \frac{\partial}{\partial x_2} \int_Y \frac{\partial u(x)}{\partial x_2}(y) dy = 0, \quad (6.21)$$

which implies that

$$\int_Y \left(\rho(u(x) - f(x)) - \frac{\partial^2 u(x)}{\partial x_1^2} - \frac{\partial^2 u(x)}{\partial x_2^2} \right) (y) dy = 0, \quad (6.22)$$

We consider a particular solution u such that

$$\left(\rho(u(x) - f(x)) - \frac{\partial^2 u(x)}{\partial x_1^2} - \frac{\partial^2 u(x)}{\partial x_2^2} \right) (y) = 0 \quad (6.23)$$

for μ -almost all $x \in X$ and ν -almost all $y \in Y$. We find such a u by means of the Fourier transform. That is,

$$u(x) = \mathfrak{F}^{-1} \left\{ \rho \frac{F(\omega)}{\|\omega\|_2^2 + \rho} \right\}. \quad (6.24)$$

This approach can be considered as a continuous counterpart of the method of denoising HS images as a collection of bands.

As with the previous methods, the FVM approach also permits to denoise with different intensities across either the spatial or the spectral domain. This can be done by setting the regularization parameter as a function that depends on y . For instance, to control the strength of the denoising process for each spatial function $u(x)(y)$, we can substitute in Eq. (6.24) the regularization parameter ρ by a function $\rho(y)$:

$$u(x) = \mathfrak{F}^{-1} \left\{ \rho(y) \frac{F(\omega)}{\|\omega\|_2^2 + \rho(y)} \right\}. \quad (6.25)$$

6.2 Experiments

In order to compare the performance of the four methods described above, they were applied to noisy versions of the Indian Pines and Salinas-A HS images. The latter is a subset of the Salinas HS image—both of them can be downloaded from [38]. The sizes of the 3-D Indian Pines and Salinas-A data sets are $145 \times 145 \times 220$ and $83 \times 86 \times 224$, respectively. White additive Gaussian noise was added to these HS data (assumed to be noiseless). In all experiments, the Peak Signal-to-Noise Ratio (PSNR) before denoising was 30.103 dB.

In the approaches where a set of optimal coefficients was to be determined, the transformation matrix D employed was the Karhunen-Loève Transform (KLT), which was computed for each HS image. The KLT was chosen since it gives a very sparse representation of the HS data (as compared to DCT, wavelet and other transforms), as well as being optimal in the L^2 sense. When the KLT is used, the mean of the HS data must be subtracted prior to processing.

Regarding the implementation of the FVM approach, we employed the *fast Fourier transform* to get some insight into the continuous behaviour of this approach. In particular, we present the results obtained by the method defined in Eq. (6.24). We refer to this method as “FOURIER” in the results presented below.

As for measures of performance, we employed the Mean Square Error (MSE), PSNR, and the Structural Similarity Index Measure (SSIM) [122]. For the latter, we computed the SSIM between the original and recovered HS images in both the spatial and spectral domains. For the spatial case, the SSIM is computed between bands; whereas in the spectral case, the SSIM is computed between spectra. An overall SSIM is obtained by

simply averaging all the computed SSIMs for both the spatial and the spectral cases. Note that the greater the similarity between two images, the closer their SSIM is to 1. In Table 6.1, a summary of these quantitative results is shown.

According to these results, the fourth approach (ADMM) outperforms all the other methods with respect to any of the metrics of performance that were considered. Only in the denoising results for Indian Pines, the spectral-oriented method described in section 2.1 performs as well as ADMM in the “spectral-SSIM sense”. The latter suggests that methods that carry out regularization in both the spectral and spatial domains may perform better than methods in which the denoising process is not carried out in this fashion. We believe this to be the case because the fourth approach captures best the “nature” of HS data, that is, data that is correlated in both the spatial and spectral domains.

On the other hand, the fifth approach does not show a good performance when compared to the other methods. This may be due mainly for two reasons. Firstly, the filtering carried out by this method attenuates, but does not set to zero, the high frequency coefficients, which results in a lower PSNR. Secondly, the space $C^2(X; L^2(Y))$, as mentioned earlier, is an over-simplified model for images, and therefore the channels which comprise HS images. As is known, the space of functions of bounded variation is a good model for images. However the mathematical foundations for FVMs with bounded variation in the spatial direction is beyond the scope of this thesis. Despite all this, it is worth to point out that these results would most likely be improved by minimizing functionals that are defined over spaces of FVMs that are better models of HS data sets.

For visual comparison, some results are presented in Figures 6.1 and 6.2. Figure 6.1 demonstrates how the methods achieve the denoising in the spatial domain. The SSIM maps, shown in the top row of Figure 6.1, illustrate the similarity between the reconstructions (denoised) and the original (noiseless) HS data for a particular band. The brightness of these maps is an indication of the magnitude of the local SSIM, i.e., the brighter a given location the greater the similarity between the retrieved and the original bands at that point [122]. Figure 6.2 shows the denoising yielded by different methods in the spectral domain. A visual comparison between reconstructions (blue plots) and noisy spectra (red plots) is shown in the top row. In the bottom row, original spectra (green plots) along with the corresponding reconstructions (blue plots) are compared.

SALINAS-A					
	ST	TV	VTV	ADMM	FOURIER
MSE	4860.8284	4944.6999	4853.3592	4218.4193	47798.1633
PSNR (dB)	41.5905	41.5162	41.5972	42.2061	31.6635
SPATIAL SSIM	0.9812	0.9575	0.9658	0.9855	0.6537
SPECTRAL SSIM	0.9977	0.9977	0.9979	0.9980	0.9792
INDIAN PINES					
MSE	13803.4370	15516.8153	18415.5653	13268.2362	60229.8927
PSNR (dB)	38.2492	37.7410	36.9972	38.3046	31.8509
SPATIAL SSIM	0.9533	0.9338	0.9132	0.9556	0.7436
SPECTRAL SSIM	0.9972	0.9970	0.9963	0.9972	0.9881

Table 6.1: Numerical results for the different approaches. Numbers in bold identify the best results with respect to each of the four measures of performance considered. In all cases, the PSNR prior to denoising was 30.103 dB.




Figure 6.1: Visual results for Band No. 23 of the Indian Pines HS image. Beside the original (noiseless) image in the lower row are shown the various reconstruction results. Beside the noisy image in the upper row are shown the corresponding SSIM maps between the reconstructed (denoised) images and the original image.




Figure 6.2: Denoising results for a particular spectral function of the Indian Pines HS image. In the top row, for visual comparison, the reconstructions (blue plots) and noisy spectra (red plots) are shown. The original spectra (green plots) along with the corresponding reconstructions (blue plots) can be observed in the bottom row.

Part II

SSIM-based Optimization

Chapter 7

Preliminaries of SSIM-based Optimization

Many image processing tasks, e.g., denoising, inpainting, deblurring, are usually carried out by solving an appropriate optimization problem. In most cases, the objective function associated with such problems is expressed as the sum of a fidelity term (or terms) $f(x)$ and a regularization term (or terms) $h(x)$. The optimization problem then assumes the form

$$\min_x \{f(x) + \lambda h(x)\}, \quad (7.1)$$

where the constant λ is a regularization parameter.

The role of the fidelity term $f(x)$ is to keep the solution to (7.1) close to the observed data. A typical choice is $f(x) = \frac{1}{2}\|x - y\|_2^2$, where y is the (corrupted) observation, e.g., a noisy image. The regularization term $h(x)$ has a twofold purpose: (i) It prevents overfitting to the observed data and (ii) it imposes constraints on the solution based upon prior information or assumptions. For instance, if the optimal solution is assumed to be sparse, a typical regularization term is $h(x) = \|x\|_1$ [4, 8, 113].

Using the squared Euclidean distance as a measure of closeness is convenient since it is convex, differentiable, and usually mathematically tractable, not to mention easily computed. Furthermore, widely used metrics of visual quality such as Mean Squared Error (MSE) and Peak to Signal Noise Ratio (PSNR) are based on this definition of closeness. Nevertheless, it has been shown that such distortion measures are not the best choice when it comes to quantify visual quality [121, 122]. For this reason, many measures of visual

quality have been proposed in an attempt to model the Human Visual System (HVS). The Structural Similarity (SSIM) image quality measure, originally proposed by Wang *et al.* [122], was based upon the assumption that the HVS evolved to perceive visual errors as changes in structural information. On the basis of subjective quality assessments involving large databases, SSIM has been generally accepted to be one of the best measures of visual quality/closeness.

With these comments in mind, when visual quality is important, it would seem natural to consider the SSIM as a replacement for the widely-used squared Euclidean distance in the fidelity term $f(x)$ of Eq. (7.1), given the limitations of the latter to measure visual closeness. Indeed, from a practical point of view, it is easy to make such a replacement since the mathematical expression for the SSIM between x and the observed data y is rather straightforward. One may then be tempted to simply start computing. There is a problem, however, in that the actual mathematical framework behind such an SSIM-based optimization, which would be important for the establishment of existence and uniqueness of solutions, is more complicated due to the fact that the SSIM is not a convex function. Notwithstanding these obstacles, optimization problems that employ the SSIM as a fitting term have already been addressed.

7.1 Current SSIM-based Optimization Approaches

Let us now review some of the SSIM-based imaging tasks that can be found in the literature. For instance, in [18] the authors find the best approximation coefficients in the SSIM sense when an orthogonal transformation is used (e.g., Discrete Cosine Transform (DCT), Fourier, etc.). Finding the best SSIM approximation coefficients is equivalent to minimizing the function

$$T(\Phi(x), y) = 1 - \text{SSIM}(\Phi(x), y), \quad (7.2)$$

where $\Phi(\cdot)$ is an orthonormal matrix, and y the signal being approximated. The function $T(x, y)$, which will be used in this part of the thesis, may be considered as a measure of the visual dissimilarity between x and y .

Based on this result, Rehman *et al.* [102] address the SSIM version of the image restoration problem proposed by Elad *et al.* in [47]. In this work, a novel method for denoising images using sparse and redundant representations over learned dictionaries is introduced. This is the K-SVD algorithm, which is a generalization of the K-means algorithm that uses singular value decomposition to update the learned dictionary. The denoised image is the

solution to the following problem:

$$\alpha_{ij} = \underset{\alpha}{\operatorname{argmin}} \{ \|D_{ij}\alpha - R_{ij}x\|_2^2 + \mu_{ij}\|\alpha\|_0 \} \quad (7.3)$$

$$\hat{x} = \underset{x}{\operatorname{argmin}} \left\{ \sum_{ij} \|D_{ij}\alpha_{ij} - R_{ij}x\|_2^2 + \lambda\|x - y\|_2^2 \right\}, \quad (7.4)$$

where α is the set of coefficients that it is being sought, D_{ij} is an overcomplete dictionary, R_{ij} is a matrix that extracts the ij -th block of x , α_{ij} and μ_{ij} are regularization parameters, and y is the noisy observation. Elad *et al.* use orthogonal marching pursuit (OMP) to solve (7.4), followed by an update of the dictionaries D_{ij} using SVD. The vector \hat{x} is found by employing the closed form solution of (7.4). Upon this approach, Rehman *et al.* redefine the latter problem by replacing the quadratic terms with the SSIM:

$$\alpha_{ij} = \underset{\alpha}{\operatorname{argmin}} \{ 1 - \operatorname{SSIM}(D_{ij}\alpha, R_{ij}x) + \mu_{ij}\|\alpha\|_0 \} \quad (7.5)$$

$$\hat{x} = \underset{x}{\operatorname{argmax}} \left\{ \sum_{ij} \operatorname{SSIM}(D_{ij}\alpha_{ij}, R_{ij}x) + \lambda \operatorname{SSIM}(MBx, y) \right\}, \quad (7.6)$$

where B is a blurring filter, and M is a downsampling operator. To solve (7.6) the authors use a modified version of OMP based on the SSIM. For solving (7.6) they employ gradient ascent. The dictionaries D_{ij} are updated using K-SVD as well. Moreover, in [102], apart from the proposed OMP algorithm and the approach for denoising images using K-SVD and the SSIM, the authors also introduce a super-resolution algorithm—also based on the SSIM—to recover from a given low resolution image its high resolution version.

Another interesting application for reconstruction and denoising was proposed in [29]. Here, the authors define the statistical SSIM index (statSSIM), an extension of the SSIM for wide-sense stationary random processes. By optimizing the statSSIM, an optimal filter in the SSIM sense is found. The optimization problem they address is

$$\max_g \{ \operatorname{statSSIM}(\hat{x}(g), x) \}. \quad (7.7)$$

Here, x is the source random process; g is the optimal filter that is to be found; and \hat{x} is given by $g*y$, where y is the observed process that is defined as $y = h*x + n$ (h is an invariant linear time filter and n is additive white noise). The non-convex nature of the statSSIM is overcome by reformulating its maximization as a quasiconvex optimization problem, which is solved using the bisection method [13, 29]. Nevertheless, it is not mentioned that the SSIM—under certain conditions—is a quasiconvex function (see [20]). As a result, it can

be minimized using quasiconvex programming techniques, which permits the consideration of a much broader spectrum of SSIM-based optimization problems.

More imaging techniques based on the SSIM can also be found in [118, 101]. In these works, optimization of rate distortion, video coding and image classification are explored using the SSIM as a measure of performance.

7.2 General Framework for SSIM-based Optimization

Note that maximizing $\text{SSIM}(x, y)$ is equivalent to minimizing the function,

$$T(x, y) = 1 - \text{SSIM}(x, y), \quad (7.8)$$

which may be viewed as a kind of distance function or *dissimilarity* between x and y , i.e., $T(x, y) = 0$ if and only if $x = y$ (see Theorem 7.3.3). Furthermore, many SSIM-based imaging tasks, including all of the applications mentioned above, may now be expressed in terms of the following optimization problem,

$$\min_x \{T(\Phi(x), y) + \lambda h(x)\}, \quad (7.9)$$

where Φ is usually a linear transformation. As such, we consider Eq. (7.9) to define a general set of problems involving *unconstrained SSIM-based optimization*. On the other hand, the constrained version of this problem is given by

$$\begin{aligned} \min_x \quad & T(\Phi(x), y) \\ \text{subject to} \quad & h(x) \leq \lambda. \end{aligned} \quad (7.10)$$

We call this optimization problem a *constrained SSIM-based optimization problem*. As it will be seen later, SSIM-based imaging tasks can be cast as either (7.9) or (7.10), therefore, we consider these problems as the general framework of what we call *SSIM-based optimization*.

In this part of the thesis, in the effort of providing a unified framework as opposed to developing specific methods that address particular applications, which has been the tendency of research literature to date, we introduce a set of algorithms to solve the general problems (7.9) and (7.10) in chapters 9 and 8 respectively. In particular, we focus our attention on the case in which $h(x)$ is convex. Moreover, in chapter 8, we show that (7.10) can still be solved if $T(\Phi(x), y)$ is subjected to a set of convex constraints [13]. Applications such as Total Variation and L^1 norm regularization are discussed, as well as comparisons between the L^2 and SSIM approaches.

7.3 SSIM

7.3.1 Definition

The SSIM provides a measure of visual closeness between an image and a distorted or corrupted version of it. Since it is assumed that the distortionless image is always available, the SSIM is considered a *full-reference* measure of image quality assessment (IQA) [122]. Its definition is based on two assumptions: (i) images are highly structured—that is, pixels tend to be correlated, specially if they are spatially close—and (ii), that the HVS is adapted to extract structural information. For these reasons, SSIM measures similarity by quantifying changes in perceived structural information. This measurement is done by comparing luminance, contrast and structure of the two images being compared. Given two images x and y , changes in luminance are measured by quantifying relative changes in the means of the images. This is done by means of the following formula:

$$l(x, y) = \frac{2\mu_x\mu_y + C_1}{\mu_x^2 + \mu_y^2 + C_1}, \quad (7.11)$$

where C_1 is a positive constant added for stability purposes. It is worth to point out that $l(x, y)$ is sensitive to the relative change of luminance, not to its absolute change. This is consistent with Weber's law, a model for light adaptation of the HVS [122]. As for contrast, this comparison is carried out by measuring relative variance:

$$c(x, y) = \frac{2\sigma_x\sigma_y + C_2}{\sigma_x^2 + \sigma_y^2 + C_2}. \quad (7.12)$$

When there is a change in contrast, $c(x, y)$ is more sensitive if the base contrast is low than when this is high; i.e., the HVS perceives objects better when they contrast with the background. As expected, $C_2 > 0$ is added to avoid division by zero. Regarding structure, this is compared simply by calculating the correlation coefficient between the two images x and y :

$$s(x, y) = \frac{\sigma_{xy} + C_3}{\sigma_x\sigma_y + C_3}. \quad (7.13)$$

Once again, the positive constant C_3 is included for stability. Finally, these three components are combined using the function $f(x_1, x_2, x_3) = x_1^\alpha x_2^\beta x_3^\gamma$:

$$\text{SSIM}(x, y) = l(x, y)^\alpha \cdot c(x, y)^\beta \cdot s(x, y)^\gamma, \quad (7.14)$$

where α , β and γ are positive parameters that control the relative influence of the three comparisons. In [122], the authors simplify (7.14) by setting $\alpha = \beta = \gamma = 1$ and $C_3 = C_2/2$. This leads to the well known formula of the SSIM:

$$\text{SSIM}(x, y) = \left(\frac{2\mu_x\mu_y + C_1}{\mu_x^2 + \mu_y^2 + C_1} \right) \left(\frac{2\sigma_{xy} + C_2}{\sigma_x^2 + \sigma_y^2 + C_2} \right). \quad (7.15)$$

This definition of the SSIM will be employed for the remainder of this second part of the thesis.

Although Eq. (7.15) is neither convex nor a metric, it has some appealing properties such as symmetry, boundedness and unique maximum.

Theorem 7.3.1 (Symmetry). [17] *The SSIM(x, y) is symmetric; that is, SSIM(x, y) = SSIM(y, x).*

Proof. Each component of the SSIM, namely, $l(x, y)$, $c(x, y)$ and $s(x, y)$, is symmetric, which immediately implies that $\text{SSIM}(x, y) = \text{SSIM}(y, x)$. \square

Theorem 7.3.2 (Boundedness). *The SSIM(x, y) is bounded: $-1 \leq \text{SSIM}(x, y) \leq 1$.*

Proof. Consider the function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ given by

$$f(w, z) = \frac{2wz + C}{w^2 + z^2 + C}, \quad (7.16)$$

where C is some positive constant. Since for any real numbers w and z we have that $2wz \leq w^2 + z^2$, it follows that

$$\frac{2wz + C}{w^2 + z^2 + C} \leq 1. \quad (7.17)$$

Moreover,

$$-\frac{2|wz| + C}{w^2 + z^2 + C} \leq \frac{2wz + C}{w^2 + z^2 + C}. \quad (7.18)$$

Thus, $-1 \leq f(w, z)$.

Given the latter, we immediately obtain that $-1 \leq l(x, y) \leq 1$. As for the second term of Eq. (7.15), the Cauchy-Schwarz inequality implies that $|\sigma_{xy}| \leq \sigma_x\sigma_y$. Thus, we have that

$$-\frac{2\sigma_x\sigma_y + C}{\sigma_x^2 + \sigma_y^2 + C} \leq \frac{2\sigma_{xy} + C}{\sigma_x^2 + \sigma_y^2 + C} \leq \frac{2\sigma_x\sigma_y + C}{\sigma_x^2 + \sigma_y^2 + C}. \quad (7.19)$$

By using the bounds of the function $f(w, z)$, with $w = \sigma_x$ and $z = \sigma_y$, we obtain that

$$-1 \leq \frac{2\sigma_{xy} + C}{\sigma_x^2 + \sigma_y^2 + C} \leq 1. \quad (7.20)$$

Therefore, we conclude that $-1 \leq \text{SSIM}(x, y) \leq 1$. \square

Theorem 7.3.3 (Uniqueness). *SSIM(x, y) = 1 if and only if x = y.*

Proof. Let $v(x, y)$ be given by

$$v(x, y) = \frac{2\sigma_{xy} + C_2}{\sigma_x^2 + \sigma_y^2 + C_2}. \quad (7.21)$$

If $\text{SSIM}(x, y) = 1$, we have that either $l(x, y) = v(x, y) = 1$ or $l(x, y) = v(x, y) = -1$. However, since both x and y are images, their means μ_x and μ_y are both greater or equal than zero, which implies that $l(x, y) = v(x, y) = 1$. Now, let us consider again the inequality

$$2wz + C_1 \leq w^2 + z^2 + C_1 \quad (7.22)$$

and let $w = \mu_x$ and $z = \mu_y$. Since $l(x, y) = 1$, we have that $\mu_x = \mu_y$. On the other hand, $v(x, y) = 1$ implies that $2\sigma_{xy} = \sigma_x^2 + \sigma_y^2$. Observe that if $x, y \in \mathbb{R}^n$, we obtain that

$$2\sigma_{xy} = 2x^T y - 2n\mu_x\mu_y, \quad (7.23)$$

and

$$\sigma_x^2 + \sigma_y^2 = \|x\|_2^2 - n\mu_x^2 + \|y\|_2^2 - n\mu_y^2. \quad (7.24)$$

Further, let $x, y : \Omega \rightarrow \mathbb{R}$ be functions that belong to a Hilbert space and let (Ω, Σ, ν) be a finite measure space. Then,

$$2\sigma_{xy} = 2\langle x, y \rangle - 2\nu(\Omega)\mu_x\mu_y, \quad (7.25)$$

and

$$\sigma_x^2 + \sigma_y^2 = \|x\|_2^2 - \nu(\Omega)\mu_x^2 + \|y\|_2^2 - \nu(\Omega)\mu_y^2. \quad (7.26)$$

Given that $\mu_x = \mu_y$, we have that either

$$2x^T y = \|x\|_2^2 + \|y\|_2^2, \text{ or } 2\langle x, y \rangle = \|x\|_2^2 + \|y\|_2^2. \quad (7.27)$$

In any case, we obtain that $\|x - y\|_2^2 = 0$, which implies that $x = y$. As for the converse, if $x = y$, this immediately implies that $\text{SSIM}(x, y) = 1$ and the assertion follows. \square

Observe that in the previous theorems we have not specified the nature of both x and y ; these might be either vectors in \mathbb{R}^n or functions. This fact highlights another feature of the SSIM: its generality.

We close this section by pointing out that, in practice, the similarity between two images is computed using the MSSIM, which is basically an average of local SSIM measurements. The reason for this is that the statistics of images vary greatly across their spatial domain. By using a sliding window that moves pixel by pixel the overall SSIM is computed. If M local windows are used in the process, the MSSIM for the images X and Y is simply the average of the individual SSIM measurements:

$$\text{MSSIM}(X, Y) = \sum_{i=1}^M \text{SSIM}(x_i, y_i), \quad (7.28)$$

where x_i and y_i are the sub-images associated to each i -th local window. It is worthwhile to mention that for avoiding “blocking” artifacts, in [122], the authors employ a circular Gaussian function. The effect of this is that the sums, in the discrete case, of the local statistics become weighted sums; analogously, we have “weighted” integrals in the continuous instance.

7.3.2 The SSIM as a normalized metric

Let x and y be vectors in \mathbb{R}^n . In the special case that both x and y have equal means, i.e., $\mu_x = \mu_y$, the luminance component $l(x, y)$ of Eq. (7.15) is equal to one, which implies that the SSIM is now given by

$$\text{SSIM}(x, y) = \frac{2\sigma_{xy} + C_2}{\sigma_x^2 + \sigma_y^2 + C_2}. \quad (7.29)$$

This less cumbersome version of the SSIM can be simplified even further if both x and y have zero mean, i.e., $\mu_x = \mu_y = 0$. If this is so, we have that

$$\sigma_{xy} = \frac{1}{n-1} \sum_{i=1}^n x_i y_i = x^T y \quad \text{and} \quad \sigma_x^2 = \frac{1}{n-1} \sum_{i=1}^n x_i^2 = \|x\|_2^2. \quad (7.30)$$

Substitution of these equations into Eq. (7.29) yields the following simplified formula for the SSIM,

$$\text{SSIM}(x, y) = \frac{2x^T y + C}{\|x\|_2^2 + \|y\|_2^2 + C}, \quad (7.31)$$

where $C = (n - 1)C_2$. The corresponding distance/dissimilarity function $T(x, y)$ in Eq. (7.8) is given by

$$T(x, y) = 1 - \text{SSIM}(x, y) = \frac{\|x - y\|_2^2}{\|x\|_2^2 + \|y\|_2^2 + C}. \quad (7.32)$$

Note that $0 \leq T(x, y) \leq 2$ if $C = 0$. Furthermore, $T(x, y) = 0$ if and only if $x = y$, which can be easily seen from Theorem 7.3.3.

As mentioned before, since $\text{SSIM}(x, y)$ is a measure of similarity, $T(x, y)$ can be considered as a measure of dissimilarity between x and y . In fact, Eq. (7.32) is an example of a (squared) normalized metric, which has been discussed in [17, 20]. Furthermore, $T(x, y)$ gives us a sense of how far x is from a given observation y in the SSIM sense, whereas $\text{SSIM}(x, y)$ tells us about how correlated or similar are x and y . Also, $1 - \text{SSIM}(x, y)$ can be thought as a kind of normalized error. Since in the majority of optimization problems the fidelity term conveys the information about the distance or the error with respect to a given observation, we will state the SSIM based optimization problems using the dissimilarity measure defined in equation (7.32).

On the other hand, for the remainder of this thesis, we shall be working with zero mean vectors, so that Eq. (7.31) and Eq. (7.32) will be employed in all computations of the SSIM. Although this condition might seem restrictive, it is possible to recover non-zero mean vectors from their zero mean counterparts, provided that the non-zero mean vectors are related by a linear degradation model.

Theorem 7.3.4. *Let D be an $m \times n$ matrix and let n be some zero-mean random degradation process. Also, let $x^* \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$ be non-zero mean vectors, where y is given and x^* is to be found. If $y = Dx^* + n$ and the inverse of $D^T D$ exists, then $x^* = x^* + \hat{x}$, $\mathbb{E}(Dx^*) = 0$, and*

$$\hat{x} = \bar{y}(D^T D)^{-1} D^T \mathbf{1} \quad (7.33)$$

Here, $\bar{y} = \mathbb{E}(y)$ and $\mathbf{1}$ is a vector in \mathbb{R}^n whose components are all equal to one.

Proof. We have that

$$\mathbb{E}(y) = \mathbb{E}(Dx^* + n) \quad (7.34)$$

$$\mathbb{E}(y) = \mathbb{E}(Dx^*). \quad (7.35)$$

Moreover, y can be expressed as $y^* + \bar{y}\mathbf{1}$, where $\mathbb{E}(y^*) = 0$. This implies that there exist vectors x^* and \hat{x} such that $\mathbb{E}(Dx^*) = 0$ and $\bar{y}\mathbf{1} = D\hat{x}$. Therefore, we obtain that $x^* = x^* + \hat{x}$, where

$$\hat{x} = \bar{y}(D^T D)^{-1} D^T \mathbf{1} \quad (7.36)$$

□

7.3.3 Quasiconvexity and Quasiconcavity

An interesting property of the dissimilarity measure T is that it is a quasiconvex function. To see under what conditions this property of T holds, we use the fact that a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is quasiconvex if its domain and all its sub-level sets

$$S_\alpha = \{x \in \mathbf{dom} f \mid f(x) \leq \alpha\}, \quad (7.37)$$

for $\alpha \in \mathbb{R}$, are convex [13].

Theorem 7.3.5. *Let $y \in \mathbb{R}^n$ be fixed. Then, $T(x, y)$ is quasiconvex if $x^T y \geq -\frac{C}{2}$, where C is the stability constant of the dissimilarity measure T (see Eq. (7.32)).*

Proof. Since the domain of (7.32) is \mathbb{R}^n , which is convex, we just have to see if its sub-level sets are convex as well. This can be done using (7.37):

$$T(x, y) = \frac{\|x - y\|_2^2}{\|x\|_2^2 + \|y\|_2^2 + C} \leq \alpha \quad (7.38)$$

$$\|x\|^2 - 2x^T y + \|y\|^2 \leq \alpha(\|x\|_2^2 + \|y\|_2^2 + C) \quad (7.39)$$

$$(1 - \alpha)\|x\|_2^2 - 2x^T y + (1 - \alpha)\|y\|_2^2 - \alpha C \leq 0 \quad (7.40)$$

The set of x vectors that make this inequality hold defines a convex set as long as $0 \leq \alpha \leq 1$. This implies that $T(x, y)$ is quasiconvex if it is less or equal than one. Then,

$$\frac{\|x - y\|_2^2}{\|x\|_2^2 + \|y\|_2^2 + C} \leq 1 \quad (7.41)$$

$$\frac{\|x\|_2^2 + \|y\|_2^2 - 2x^T y}{\|x\|_2^2 + \|y\|_2^2 + C} \leq 1 \quad (7.42)$$

$$\frac{\|x\|_2^2 + \|y\|_2^2 + C - C - 2x^T y}{\|x\|_2^2 + \|y\|_2^2 + C} \leq 1. \quad (7.43)$$

Thus, if $2x^T y + C \geq 0$, $T(x, y)$ is quasiconvex. In other words, quasiconvexity of $T(x, y)$ holds if $x^T y \geq -\frac{C}{2}$. \square

Corollary 7.3.5.1. *SSIM(x, y) is quasiconcave if $x^T y \geq -\frac{C}{2}$.*

Proof. We have that $T(x, y) = 1 - \text{SSIM}(x, y)$, which is equivalent to $T(x, y) - 1 = -\text{SSIM}(x, y)$. The function $T(x, y) - 1$ is also quasiconvex since subtracting a constant does not affect quasiconvexity. Given that function a f is quasiconcave if $-f$ is quasiconvex, we immediately obtain that SSIM(x, y) is quasiconcave if $x^T y \geq -\frac{C}{2}$. \square

Moreover, $T(x, y)$ is quasiconcave if $x^T y \leq -\frac{C}{2}$. This can be proved by using the definition of quasiconcavity of a real-valued function: a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is quasiconcave if its domain and all its super-level sets

$$S_\alpha = \{x \in \mathbf{dom} f \mid f(x) \geq \alpha\}, \quad (7.44)$$

for $\alpha \in \mathbb{R}$, are convex [13].

Theorem 7.3.6. *Let $y \in \mathbb{R}^n$ be fixed. Then, $T(x, y)$ is quasiconcave if $x^T y \leq -\frac{C}{2}$, where C is the stability constant of the dissimilarity measure T (see Eq. (7.32)).*

Proof. By the previous theorem, we can see easily that the super-level sets of $T(x, y)$ are given by

$$(1 - \alpha)\|x\|_2^2 - 2x^T y + (1 - \alpha)\|y\|_2^2 - \alpha C \geq 0 \quad (7.45)$$

If $\alpha \geq 1$, the set of x vectors for which this inequality holds is a convex set. Then, $T(x, y)$ is a quasiconcave function over the region in which this inequality is true. Thus, as it was proved in the previous theorem, we obtain that this region is defined by the inequality $x^T y \leq -\frac{C}{2}$. This completes the proof. \square

Corollary 7.3.6.1. *SSIM(x, y) is quasiconvex if $x^T y \leq -\frac{C}{2}$.*

Proof. The proof is completely analogous to the proof of the previous corollary. In this case, we use the fact that a function f is quasiconvex if $-f$ is quasiconcave and the assertion follows. \square




Figure 7.1: This is the plot of $T(x, 3)$ for the one-dimensional case with $C = 0$. Notice that when $x \geq 0$ the dissimilarity measure is quasiconvex. Indeed, the sub-level sets are convex as long as $T(x, 3) \leq \alpha$, for any $\alpha \in [0, 1]$. In a similar manner, the super-level sets are convex if $1 \leq T(x, 3) \leq 2$, which implies that the function is quasiconcave over that region of the real line.




Figure 7.2: The plot of $T(x, y)$ in the two-dimensional case. Once again, $C = 0$. Here, $x = [x_1, x_2]^T$ and $y = [3, 2]^T$. As it was shown, the line $3x_1 + 2x_2 = 0$ delimits the regions of quasiconvexity and quasiconcavity of $T(x, y)$.

Chapter 8

Constrained SSIM-based optimization

A standard quasiconvex problem is defined as follows:

$$\begin{aligned} \min_x \quad & f(x) \\ \text{subject to} \quad & h_i(x) \leq 0, \quad i = 1, \dots, m \\ & Ax = b, \end{aligned} \tag{8.1}$$

where $f(x)$ is quasiconvex, $Ax = b$ is an equality constraint, and the $h_i(x)$ are a set of convex inequality constraints. Optimization problems of this type can be addressed by solving a sequence of feasibility problems. The key ingredients of this technique are a family of convex inequalities, which represent the sub-level sets of $f(x)$; a convex feasibility problem that is to be solved at each step; and the bisection method for finding the optimal value of (8.1) up to a certain accuracy [13]. Given this, thanks to the quasiconvex properties of the dissimilarity measure $T(x, y)$, we can cast constrained SSIM-based optimization problems as quasiconvex optimization problems.

Taking the latter into account, we shall define a constrained SSIM-based optimization problem as follows,

$$\begin{aligned} \min_x \quad & T(\Phi(x), y) \\ \text{subject to} \quad & h_i(x) \leq 0, \quad i = 1, \dots, m \\ & Ax = b, \end{aligned} \tag{8.2}$$

where $\Phi(\cdot)$ is some linear transformation. Also, for the sake of simplicity, we set to zero the stability constant C of the dissimilarity measure $T(\Phi(x), y)$. This implies that $T(\Phi(x), y)$ is now given by

$$T(\Phi(x), y) = \frac{\|\Phi(x) - y\|_2^2}{\|\Phi(x)\|_2^2 + \|y\|_2^2}. \quad (8.3)$$

One might argue that this compromise the stability of the SSIM-based fidelity term, nevertheless, we assume that we always work with a non-zero observation y . This ensures stability and differentiability of $T(\Phi(x), y)$ for all $x \in \mathbb{R}^n$.

In this chapter, we introduce a simple and robust algorithm based on the bisection method to solve the constrained SSIM-based optimization problem (8.2), as well as a review of some of the SSIM-based imaging tasks that can be carried out by solving (8.2). Experimental results are also presented.

8.1 Optimizing $\mathbf{T}(\Phi(\mathbf{x}), \mathbf{y})$

Assuming that the optimal zero-mean solution x^* is in the region where $T(\Phi(x), y)$ is quasiconvex, i.e., $(\Phi(x^*))^T y \geq 0$, the problem in (8.2) can be solved by solving a sequence of feasibility problems. For this, we require a family of convex inequalities that represent the sub-level sets of $T(\Phi(x), y)$ and a convex feasibility problem that is to be solved at each step. The bisection method may be employed to determine the optimal value of (8.2) up to a certain accuracy [13].

The family of convex inequalities is defined by means of functions $\phi_\alpha : \mathbb{R}^n \rightarrow \mathbb{R}$ such that

$$f(x) \leq \alpha \iff \phi_\alpha(x) \leq 0. \quad (8.4)$$

Also, for every x , $\phi_\beta(x) \leq \phi_\alpha(x)$, whenever $\alpha \leq \beta$. The following functions satisfy such conditions:

$$\phi_\alpha(x) = (1 - \alpha)\|\Phi(x) - y\|_2^2 - 2\alpha(\Phi(x))^T y. \quad (8.5)$$

The feasibility problems then assume the form

$$\begin{aligned} & \text{Find} && x \\ & \text{subject to} && \phi_\alpha(x) \leq 0 \\ & && h_i(x) \leq 0, \quad i = 1, \dots, m \\ & && Ax = b. \end{aligned} \quad (8.6)$$

If (8.6) is feasible, then $p^* \leq \alpha$, else, $p^* > \alpha$, where p^* is the optimal value of (8.2).

Using the fact that $0 \leq T(\Phi(x), y) \leq 2$, and defining $\mathbf{1}$ and $\mathbf{0}$ as vectors in \mathbb{R}^n whose entries are all equal to one and zero respectively, we propose the following algorithm for solving (8.2):

Algorithm 1: Bisection method for constrained SSIM-based optimization

```

initialize  $x = \mathbf{0}$ ,  $l = 0$ ,  $u = 2$ ,  $\epsilon > 0$ ;
data preprocessing  $\bar{y} = \frac{1}{n}\mathbf{1}^T y$ ,  $y = y - \bar{y}\mathbf{1}$ ;
while  $u - l > \epsilon$  do
     $\alpha := (l + u)/2$ ;
    Solve (8.6);
    if (8.6) is feasible,  $u := \alpha$ ;
    elseif  $\alpha = 1$ , (8.2) can not be solved, break;
    else  $l := \alpha$ ;
end
return  $x$ ,  $y = y + \bar{y}\mathbf{1}$ .

```

Notice that this method will find a solution x^* such that $l \leq f(x) \leq l + \epsilon$ in exactly $\lceil \log_2((u - l)/\epsilon) \rceil$ iterations [13], provided that such solution lies in the quasiconvex region of $T(\Phi(x), y)$. In other words, if this is so, the algorithm converges to an optimal value p^* . This condition may seem restrictive, however, one is normally interested in solutions that are positively correlated to the given observation y . The MATLAB code for this algorithm is presented in Appendix C, Section C.1.

It is worthwhile to mention that it is not always possible to recover the mean of the non-zero-mean optimal solution x^* . This is because the luminance component of the SSIM has not been taken into account. Nevertheless, in many circumstances (e.g., denoising of a signal corrupted by zero-mean additive white Gaussian noise), the mean of y and $\Phi(x^*)$ coincide. In this case, thanks to theorem 7.3.4, we have that $x^* = x^* + \hat{x}$, where x^* is the zero-mean optimal solution and \hat{x} is given by:

$$\hat{x} = \bar{y}(D^T D)^{-1} D^T \mathbf{1}, \tag{8.7}$$

provided that the inverse of $D^T D$ exists.

8.2 Applications

Clearly, different sets of constraints lead to different SSIM-based optimization problems. For instance, by making the substitution $\Phi(x) = Dx$, where D is an $m \times n$ matrix and $x \in \mathbb{R}^n$, and redefining (8.6) as

$$\begin{aligned} & \text{find} && x \\ & \text{subject to} && (1 - \alpha)\|Dx - y\|_2^2 - \alpha 2x^T D^T y \leq 0 \\ & && \|Ax\|_p^p - C \leq 0, \end{aligned} \tag{8.8}$$

we have this optimization problem:

$$\begin{aligned} & \min_x && T(Dx, y) \\ & \text{subject to} && \|Ax\|_p^p \leq C, \end{aligned} \tag{8.9}$$

where A is some linear operator (e.g., a difference matrix). For different choices of A and p , different interesting SSIM-based imaging tasks emerge. We review some of these in this section.

8.2.1 SSIM with Tikhonov constraint

A common method used for ill-posed problems is *Tikhonov regularization* or *ridge regression*. This is basically a constrained version of least squares and it is found in different fields such as statistics and engineering. It is stated as follows

$$\begin{aligned} & \min_x && \|Dx - y\|_2^2 \\ & \text{subject to} && \|Ax\|_2^2 \leq C, \end{aligned} \tag{8.10}$$

where A is called the *Tikhonov matrix*. A common choice for the matrix A is the identity matrix, however, other choices may be a scaled finite approximation of a differential operator or a scaled orthogonal projection [66, 56, 87]. By making the substitution $p = 2$ in (8.9), we have the SSIM version of Tikhonov regularization:

$$\begin{aligned} & \min_x && T(Dx, y) \\ & \text{subject to} && \|Ax\|_2^2 \leq C. \end{aligned} \tag{8.11}$$

8.2.2 L^1 -constrained SSIM-based optimization

Making the substitution $\Phi(x) = Ax$, where A is a $m \times n$ matrix and $x \in \mathbb{R}^n$ (although x may be complex), and by using the convex constraint $h(x) = \|x\|_1 - \lambda$, we obtain the following SSIM based optimization problem:

$$\begin{aligned} \min_x \quad & T(Ax, y) \\ \text{subject to} \quad & \|x\|_1 \leq \lambda. \end{aligned} \tag{8.12}$$

This particular problem is appealing because it combines the concepts of similarity and sparseness. As with the classical regularized version of the least squares method known as LASSO (Least Absolute Shrinkage and Selection Operator) [46, 4], the solution of (8.12) is also sparse. Observe that by decreasing the L^1 norm of the solution vector (which is accomplished by decreasing λ), a greater number of coefficients are “thresholded” to zero, thereby increasing the “sparsity” of the solution vector. As expected, problem (8.12) can be employed in applications such as compressed sensing, in which sparsity is the assumed underlying model for signals. To the best of our knowledge, this is the first reported optimization problem where the SSIM is optimized having the L^1 norm as a constraint.

8.2.3 SSIM and Total Variation

By employing the constraint $h(x) = \|Dx\|_1 - \lambda$, where D is a difference matrix and $\Phi(x) = x$, we can define informally a SSIM-total-variation-denoising method for one-dimensional discrete signals. Given a noisy signal y , its denoised version is the solution of the problem,

$$\begin{aligned} \min_x \quad & T(x, y) \\ \text{subject to} \quad & \|Dx\|_1 \leq \lambda. \end{aligned} \tag{8.13}$$

Here, we consider $\|Dx\|_1$ as a measure of the total variation (TV) of the vector x . Notice that instead of minimizing the TV norm, we employ it as a constraint. This approach is not new, it can also be found in [32, 52]. As with the classical TV optimization problems, solutions of (8.13) have bounded variation as well.

Moreover, images can also be denoised by minimizing the dissimilarity measure $T(x, y)$ subject to the following convex constraint:

$$h(x) = \|D_1(x)\|_1 + \|D_2(x)\|_1 - \lambda, \tag{8.14}$$

where the linear operators D_1 and D_2 are difference matrices used to compute the spatial derivatives of the digital image x . Notice that the anisotropic TV norm is being used in this case. As far as we are concerned, the work reported in [106] and these applications are the only approaches in the literature that combine TV and the SSIM.

8.2.4 Deblurring

The blurring of an image is usually modelled as the convolution of an undistorted image x and a blur kernel τ . Nevertheless, in practice, the blurred observation y may have been degraded by either additive noise or errors in the acquisition process. For this reason, the following model is used to represent the degradation process [57, 60]:

$$y = \tau * x + \eta, \tag{8.15}$$

where η is usually white Gaussian noise.

The problem of recovering x can be addressed by the proposed approach by using the convex constraint $h(x) = \|D_x(x)\|_1 + \|D_y(x)\|_1 - \lambda$, and by defining $\Phi(x) = Kx$, where K is a linear operator that performs the blurring process. That is, the unblurred image x can be estimated by solving the following SSIM-based optimization problem:

$$\begin{aligned} \min_x \quad & T(Kx, y) \\ \text{subject to} \quad & \|D_x(x)\|_1 + \|D_y(x)\|_1 \leq \lambda. \end{aligned} \tag{8.16}$$

8.2.5 Zooming

In this case, given an image y , assumed to be of “lower resolution”, we desire to find an approximation x to a higher resolution version of y . This inverse problem can be solved in a manner very similar to the one described in the previous section; that is, by defining $\Phi(x) = Sx$, where S is a subsampling matrix, and using the same convex constraint that is employed for the deblurring application. We claim that a good estimate of the high resolution image x is the solution of the SSIM-based optimization problem given by

$$\begin{aligned} \min_x \quad & T(Sx, y) \\ \text{subject to} \quad & \|D_x(x)\|_1 + \|D_y(x)\|_1 \leq \lambda. \end{aligned} \tag{8.17}$$

Observe that, in general, the matrix $S^T S$ is not invertible, therefore, equation (8.7) can not be used to recover the optimal non-zero mean solution x^* . Nevertheless, the mean of the low-resolution observation y can be used as a good estimate of the mean of the high resolution image that is being sought.

The problem of zooming using the SSIM approach has also been addressed in [102], in which sparse representations of non-overlapping blocks of the image are used in the reconstruction process; however, the variational approach is not considered. Methods that employ the TV norm for estimating the high resolution image x can be found in [24, 60], nevertheless, the fitting term is the commonly used square Euclidean distance. Problem (8.17) can be considered as a method that combines the SSIM and the variational approach for addressing this inverse problem.

8.3 Experiments

In a series of numerical experiments, we have compared the performance of optimization methods employing (i) the usual squared Euclidean distance and (ii) Structural Similarity as fitting terms. For simplicity, we refer to these methods as (i) L^2 -based and (ii) SSIM-based methods, respectively. This is done by comparing the structural similarities between an undistorted given image and both L^2 and SSIM reconstructions. The structural similarities are calculated using the definition given by (8.3). By averaging the SSIM values of all non-overlapping 8×8 pixel blocks, the total SSIM for each recovered image is obtained. The reconstructions are obtained by solving either a SSIM-based or an L^2 -based optimization problem over each pixel block. Finally, for each application, the corresponding constraint $h(x)$ being employed is the same for all non-overlapping blocks.

In all the applications that are presented below, the estimated mean from each block is removed prior to processing. Once the zero-mean optimal block x^* is obtained, the optimal non-zero-mean block x^* is recovered by means of Eq. (8.7), except in the case of zooming. In this case, the means of the high resolution blocks are approximated by the means of their corresponding low resolution counterparts. This is necessary since quasiconvexity of $T(\Phi(x), y)$ is guaranteed for zero-mean vectors. This approach is also applied in the classical L^2 -based optimization method even though it is not required, for the sake of fair comparison of the two approaches.

For the L^1 -constrained optimization problems of Section 8.2.2, both (8.12) and its L^2 version are solved over each non-overlapping block. Here, $\Phi(x) = Dx$, where D is a $n \times n$ discrete cosine transform (DCT) matrix and $x \in \mathbb{R}^n$ is the set of DCT coefficients that is

to be recovered. As expected, the fitting term of the L^2 counterpart of (8.12) is substituted by $\|Dx - y\|_2^2$. The result is that for a given λ , a sparse approximation problem is solved at each block. In this experiment, a 96×96 sub-image of the test image *Mandrill* was used.

For the TV-based problems of Section 8.2.3, (8.13) is solved over each non-overlapping block. Since we are working with images, the constraint in (8.13) is replaced by the following convex term:

$$\|D_1(x)\|_1 + \|D_2(x)\|_1 \leq \lambda. \quad (8.18)$$

The fidelity term $\|x - y\|_2^2$ is employed in the L^2 -based version of (8.13). In this experiment, a 96×96 noisy sub-image of the test image *Mandrill*, corrupted with additive zero-mean Gaussian noise (AWGN) ($\sigma = 1/32$) was employed.

For the deblurring problem of Section 8.2.4, a blurred and noisy 104×104 pixel sub-image of the test image *Lena* was processed. The reconstructions were obtained by solving problem (8.16) and its L^2 version over each non-overlapping block. The blurring kernel employed was a Gaussian with unit standard deviation. The blurred image was also contaminated with AGWN with $\sigma = \frac{1}{64}$.

Finally, with regard to the zooming problem of Section 8.2.5, the estimated high resolution images are obtained by solving both problem (8.17) and its L^2 version over each pixel block. The fitting term employed in the L^2 -based method was $\|Sx - y\|_2^2$.

Some results of sparse reconstruction, deblurring and zooming are presented in Figures 8.1, 8.2 and 8.3, respectively. In each figure are shown the original (uncorrupted) image, its corrupted version and SSIM- and L^2 -based reconstructions. For each set of experiments, the SSIM maps between reconstructions and the original image are presented. The brightness of regions in the SSIM maps indicates the degree of similarity between corresponding image blocks—the brighter a given point the greater the SSIM, hence visual similarity, at that location [122].

In Figure 8.1, where results of the sparse reconstruction problem are shown, the SSIM- and L^2 -based reconstructions are very similar. However, we notice that the SSIM reconstruction enhances the contrast at some locations (e.g., the wrinkles below the eye of *Mandrill*).

With regard to Figure 8.2, where the deblurring results are shown, we see that edges in the reconstructions tend to be sharper than those in the blurred and noisy image. However, the reconstruction of textures is not very good in both methods. We expect that improved results may be obtained by tuning the constraints of each pixel block for optimal performance.

With regard to Figure 8.3, where zooming results are shown, we observe that both SSIM- and L^2 -based reconstructions are quite good. Some vertical and horizontal artifacts can be seen in both reconstructions, located mainly near edges. This may be due to the fact that anisotropic total variation was employed as a constraint.

A summary of quantitative results is presented in Table 8.1. The effectiveness of both SSIM- and L^2 -based approaches was quantified using the mean squared error (MSE) and the SSIM defined in Eq. (7.31). The best results with respect to each measure of performance are denoted in bold. We observe that the effectiveness of both methods is almost the same, although the proposed approach performs better with respect to the SSIM measure $\text{SSIM}(x, y) = 1 - T(x, y)$, as expected. We also observe that a low MSE does not necessarily imply a high visual similarity (measured in terms of SSIM) between the reconstructions and the original images, as is well known in the literature [122].

	PROPOSED		L^2	
	SSIM	MSE	SSIM	MSE
SPARSE RECONS.	0.6779	7.6729	0.6656	8.0950
TV DENOISING	0.8903	1.5601	0.8894	1.5549
DEBLURRING	0.6911	2.0389	0.6878	2.0197
ZOOMING	0.8167	2.0753	0.8142	2.1067

Table 8.1: Numerical results for the different approaches and applications. Numbers in bold identify the best results with respect to each measure of performance.




Figure 8.1: Visual results for the sparse reconstructions. In this case, for each pixel block, the maximum allowed value for the L^1 norm of the coefficients that are to be recovered is 1. In the top row, SSIM maps are shown. Original and recovered images can be seen in the bottom row.




Figure 8.2: Visual results for the deblurring application. In the top row, the blurred and noisy image along with the SSIM maps are presented. As above, the recovered and original images are seen in the bottom row.




Figure 8.3: Visual results of the zooming experiments. The low resolution image and the SSIM maps can be seen in the top row. Original image along with the SSIM and L^2 reconstructions are shown in the bottom row.

Chapter 9

Unconstrained SSIM-based optimization

We define an unconstrained SSIM-based optimization problem as follows

$$\min_x \{T(\Phi(x), y) + \lambda h(x)\}, \quad (9.1)$$

where $\Phi(\cdot)$ is a linear transformation, λ is a regularization parameter, and $h(x)$ is a convex regularizing term. As usual, when the regularization parameter λ is different from zero, the regularizing term $h(x)$ plays a role in the solution of (9.1), which is controlled by the magnitude of λ . Notice that, since the dissimilarity measure is not convex—although it is locally quasiconvex—, problem (9.1) is not convex. This implies that the existence of a global optimal point cannot be guaranteed; nevertheless, algorithms that converge to either a local or global minimum can be developed. In particular, in section 9.2 of this chapter, we present a set of algorithms for solving problem (9.1) when λ is greater than zero.

Even though possibly the most interesting applications arise when the regularization parameter is positive, it is worthwhile to take a look at the case when λ is zero. That is, the solution of a problem of the form

$$\min_x \{T(\Phi(x), y)\}. \quad (9.2)$$

This instance of (9.1) was already considered in [18] when Φ is an orthogonal transformation, however, we present an alternative derivation of the same result in the following section. Also, we extend this result by considering the case in which Φ is an $m \times n$ matrix ($m \geq n$).

9.1 The special case $\lambda = 0$

Given that x is a discrete signal or a vector in \mathbb{R}^n , we will focus our attention on the case where $\Phi(x) = Dx$; that is, D is an $m \times n$ matrix. For the sake of simplicity, we assume that y is a non-zero vector in \mathbb{R}^m , thus we set the stability constant C to zero. We have, then, that problem (9.2) is equivalent to

$$\min_x \left\{ \frac{\|Dx - y\|_2^2}{\|Dx\|_2^2 + \|y\|_2^2} \right\}. \quad (9.3)$$

Interestingly enough, the optimal x^* that minimizes (9.3) is the optimal L^2 solution of

$$\min_x \|Dx - y\|_2^2 \quad (9.4)$$

times a constant. This is proved in the following theorem.

Theorem 9.1.1. *If $D^T D$ is invertible, the optimal zero-mean solution x^* of problem (9.3) is given by*

$$x^* = \gamma (D^T D)^{-1} D^T y, \quad (9.5)$$

where the constant γ is equal to

$$\gamma = \frac{\|y\|_2}{\|y^T D (D^T D)^{-1} D^T y\|_2}. \quad (9.6)$$

Proof. To find the x^* that minimizes (9.3) we simply compute the gradient and find the x vectors that make it equal to zero. This can be done easily using logarithmic differentiation and the chain rule. This leads to

$$\nabla(T(Dx, y)) = \left(\frac{2D^T(Dx - y)}{\|Dx - y\|_2^2} - \frac{2D^T Dx}{\|Dx\|_2^2 + \|y\|_2^2} \right) T(Dx, y). \quad (9.7)$$

Equation (9.7) is zero if $Dx = y$. This case is trivial and it does not provide any interesting information about x^* . However, as for term in parentheses, this does give us more insight into the optimal x^* . Equating it to zero leads to

$$\frac{2D^T(Dx^* - y)}{\|Dx^* - y\|_2^2} = \frac{2D^T Dx^*}{\|Dx^*\|_2^2 + \|y\|_2^2} \quad (9.8)$$

$$D^T(Dx^* - y) = (D^T Dx^*) T(Dx^*, y) \quad (9.9)$$

$$x^* = \frac{(D^T D)^{-1} D^T y}{\text{SSIM}(Dx^*, y)}. \quad (9.10)$$

Notice that x^* is simply the x that minimizes $\|Dx - y\|_2^2$ divided by the SSIM of Dx^* and y . In other words, the optimal zero-mean solution x^* is the L^2 optimal solution multiplied by a constant. Indeed, we can find this constant by defining $x^* = \gamma(D^T D)^{-1} D^T y$ and using (7.31) with $C = 0$:

$$\text{SSIM}(Dx^*, y) = \frac{1}{\gamma} = \frac{2x^{*T} D^T y}{\|Dx^*\|_2^2 + \|y\|_2^2} \quad (9.11)$$

$$\frac{1}{\gamma} = \frac{2\gamma((D^T D)^{-1} D^T y)^T D^T y}{\gamma^2 \|D(D^T D)^{-1} D^T y\|_2^2 + \|y\|_2^2} \quad (9.12)$$

$$\|D(D^T D)^{-1} D^T y\|_2^2 \gamma^2 = \|y\|_2^2 \quad (9.13)$$

$$\gamma = \pm \frac{\|y\|_2}{\|D(D^T D)^{-1} D^T y\|_2} \quad (9.14)$$

We have, then, two possible values for the constant γ . Observe that the negative γ corresponds to an optimal x^* that is negatively correlated with the observation y , therefore, this is an optimal zero-mean solution that maximizes the dissimilarity measure $T(Dx, y)$. Thus, we conclude that to minimize problem (9.3) the constant γ must be given by

$$\gamma = \frac{\|y\|_2}{\|D(D^T D)^{-1} D^T y\|_2}. \quad (9.15)$$

□

Corollary 9.1.1.1. *Let D be an orthogonal matrix and let $y \in \mathbb{R}^n$. The optimal zero-mean solution x^* of problem (9.3) subject to $\|x\|_0 = k$, $k \leq n$, is given by $x^* = \gamma \hat{z}$, where \hat{z} is the vector of basis coefficients that best approximates y in the L^2 sense—that is, the non-zero entries of \hat{z} are the k basis coefficients of greatest magnitude of y —and γ is equal to*

$$\gamma = \frac{\|z\|_2}{\|\hat{z}\|_2}. \quad (9.16)$$

Here, $z = D^T y$.

Proof. Let us rewrite this SSIM-based optimization problem as follows:

$$\min_{w, E \in F} \left\{ \frac{\|Ew - y\|_2^2}{\|Ew\|_2^2 + \|y\|_2^2} \right\}, \quad (9.17)$$

where $w \in \mathbb{R}^k$ is the vector of k coefficients that is to be found, and F is the set of submatrices that are obtained from the matrix D by choosing k out of its n columns.

By the previous theorem, we know that for each element of F there exists an optimal zero-mean SSIM solution given by

$$w_E^* = \frac{(E^T E)^{-1} E^T y}{\text{SSIM}(Ew_E^*, y)}. \quad (9.18)$$

for all $E \in F$ —that is, we have minimized (9.17) respect to w . Since D is orthogonal, we have that $E^T E$ is an $k \times k$ identity matrix. Then,

$$w_E^* = \frac{E^T y}{\text{SSIM}(Ew_E^*, y)}. \quad (9.19)$$

Out of the $\frac{n!}{(n-k)!k!}$ possible choices for the submatrix E , we have to find the optimal E in the SSIM sense. Let us consider the following inequality

$$\frac{|\|Ew\|_2 - \|y\|_2|^2}{\|Ew\|_2^2 + \|y\|_2^2} \leq \frac{\|Ew - y\|_2^2}{\|Ew\|_2^2 + \|y\|_2^2}. \quad (9.20)$$

Observe that for $0 \leq k \leq n$, we have that $0 \leq \|Ew\|_2 \leq \|y\|_2$. If we consider the left-hand side of the above inequality as a function of $\|Ew\|_2$, this term is a monotone decreasing function in the interval $0 \leq \|Ew\|_2 \leq \|y\|_2$. Therefore, the lowest bound for the right-hand side of this inequality is determined by the greatest value that $\|Ew\|_2$ can attain. This implies that the best candidate for minimizing problem (9.17) is the submatrix E whose columns correspond to the greatest basis coefficients of y , that is, the optimal L^2 coefficients. Let \hat{E} be such a matrix. Then, we have that the minimizer of problem (9.17) is given by

$$w_{\hat{E}}^* = \frac{\hat{E}^T y}{\text{SSIM}(\hat{E}w_{\hat{E}}^*, y)}. \quad (9.21)$$

which is equivalent to

$$w_{\hat{E}}^* = \gamma \hat{E}^T y, \quad (9.22)$$

where γ is given by

$$\gamma = \frac{\|y\|_2}{\|\hat{E}\hat{E}^T y\|_2}. \quad (9.23)$$

Since D is orthogonal, we have that $\|y\|_2 = \|Dz\|_2 = \|z\|_2$. Moreover, $\|\hat{E}\hat{E}^T y\|_2 = \|\hat{E}^T y\|_2 = \|\hat{z}\|_2$. Hence,

$$\gamma = \frac{\|z\|_2}{\|\hat{z}\|_2}. \quad (9.24)$$

Given the latter, the vector of basis coefficients in \mathbb{R}^n that best approximates y in the SSIM sense is the vector x^* whose non-zero components are determined by $w_{\hat{E}}^*$. Thus, we obtain that

$$x^* = \gamma \hat{D}y, \quad (9.25)$$

where \hat{D} is a matrix whose only non-zero columns are defined by the k columns of D associated to the k basis coefficients of y of the greatest magnitude. Observe that $\hat{D}y = \hat{z}$, thus

$$x^* = \gamma \hat{z}. \quad (9.26)$$

This completes the proof. \square

The previous theorem can be considered a generalization of the result that can be found in [18] when the stability constant is zero. It shows that if the matrix $D^T D$ is invertible, in general, the optimal SSIM solution of problem (9.3) is a scaled version of the optimal L^2 solution. In fact, when D is an orthogonal transformation, this scaling has the property of preserving the norm of the observation y :

$$\|x^*\|_2 = \left\| \frac{\|z\|_2}{\|\hat{z}\|_2} \hat{z} \right\|_2 = \|z\|_2 = \|y\|_2. \quad (9.27)$$

In other words, regardless of how many k basis coefficients are considered for approximating y , the norm of the optimal SSIM basis coefficients is always equal to the norm of y . The visual effect of this is that, given an image block y , the best SSIM approximation has better contrast as compared to its L^2 counterpart [18].

9.2 The general case $\lambda > 0$

The development of algorithms for solving (9.1) depend on whether the regularizing term $h(x)$ is differentiable or not. For this reason, we consider these two cases separately below.

9.2.1 Differentiable $h(x)$

When the regularizing term is differentiable, root-finding algorithms can be employed to find a local zero-mean solution x^* to (7.9). For example, if Tikhonov regularization is used, we have the following SSIM-based optimization problem,

$$\min_x \{T(Dx, y) + \lambda \|Ax\|_2^2\}, \quad (9.28)$$

where D is an $m \times n$ matrix, $\lambda > 0$ is a regularization parameter, and A is the so called Tikhonov matrix. By computing the gradient of (9.28), we find that the solution x^* must satisfy

$$[(\text{SSIM}(Dx^*, y)D^T D + \lambda(\|Dx^*\|_2^2 + \|y\|_2^2 + C)A^T A)x^* = D^T y. \quad (9.29)$$

If we define the following function $f : X \subset \mathbb{R}^n \rightarrow \mathbb{R}$,

$$f(x) = [(\text{SSIM}(Dx, y)D^T D + \lambda(\|Dx\|_2^2 + \|y\|_2^2 + C)A^T A)x - D^T y, \quad (9.30)$$

then x^* is a (zero-mean) vector in \mathbb{R}^n such that $f(x^*) = 0$.

To find x^* we may use the *generalized Newton method* [93]. The convergence of this method is guaranteed, under certain conditions, by the *Newton-Kantorovich theorem*.

Theorem 9.2.1 (Newton-Kantorovich Theorem). *Let X and Y be Banach spaces and $g : X \subset A \rightarrow Y$. Assume g is Fréchet differentiable on an open convex set $D \subset X$ and*

$$\|g'(x) - g'(z)\| \leq K\|x - z\|_2, \text{ for all } x, z \in D, \quad (9.31)$$

Also, for some $x_0 \in D$, suppose that $g'(x_0)^{-1}$ is defined on all Y and that

$$h := L\|g'(x_0)^{-1}\|\|g'(x_0)^{-1}g(x_0)\| \leq \frac{1}{2}, \quad (9.32)$$

where $\|g'(x_0)^{-1}\| \leq \beta$ and $\|g'(x_0)^{-1}g(x_0)\| \leq \eta$. Set

$$K_1 = \frac{1 - \sqrt{1 - 2h}}{K\beta} \text{ and } K_2 = \frac{1 + \sqrt{1 - 2h}}{K\beta}, \quad (9.33)$$

and assume that $S := \{x : \|x - x_0\| \leq K_1\} \subset D$. Then, the Newton iterates

$$x_{k+1} = x_k - g'(x_k)^{-1}g(x_k), \quad k \in \mathbb{N}, \quad (9.34)$$

are well defined, lie in S and converge to a solution x^* of $g(x) = 0$, which is unique in $D \cap \{x : \|x - x_0\| \leq K_2\}$. Moreover, if $h < \frac{1}{2}$, the order of convergence is at least quadratic.

Proof. See [93]. □

In our particular case, the Fréchet derivative of f is its Jacobian, which we denote as J_f and is given by

$$J_f(x) = D^T Dx(\nabla s(x))^T + \lambda A^T Ax(\nabla r(x))^T + s(x)D^T D + \lambda r(x)A^T A, \quad (9.35)$$

where $s(x) = \text{SSIM}(Dx, y)$ and $r(x) = \|Dx\|_2^2 + \|y\|_2^2 + C$. By the previous theorem, we have that convergence of the generalized Newton's method in any open convex subset Ω of X , where $X \subset \mathbb{R}^n$, is guaranteed if the initial guess x_0 satisfies the following condition,

$$L \|J_f(x_0)^{-1}\| \|J_f(x_0)^{-1} f(x_0)\| \leq \frac{1}{2}. \quad (9.36)$$

Here, $J_f(\cdot)^{-1}$ denotes the inverse of $J_f(\cdot)$, and $L > 0$ is a constant less or equal than the Lipschitz constant of $J_f(\cdot)$. In fact, it can be proved that for any open convex subset $\Omega \subset X$, $J_f(\cdot)$ is Lipschitz continuous.

Theorem 9.2.2. *Let $f : X \subset \mathbb{R}^n \rightarrow \mathbb{R}$ be defined as in Eq. (9.30). Then, its Jacobian is Lipschitz continuous on any open convex set $\Omega \subset X$; that is, there exists a constant $L > 0$ such that for any $x, z \in \Omega$,*

$$\|J_f(x) - J_f(z)\|_F \leq L \|x - z\|_2. \quad (9.37)$$

Here, $\|\cdot\|_F$ denotes the Frobenius norm¹ and

$$L = C_1 \|D^T D\|_F + \lambda C_2 \|A^T A\|_F, \quad C_1, C_2 > 0. \quad (9.38)$$

Proof. See appendix B, Theorem B.0.17. □

From this discussion, and the notation $\mathbf{1} = [1, 1, \dots, 1]^T \in \mathbb{R}^n$, we propose the following algorithm for solving the problem in Eq. (9.28).

Algorithm II: Generalized Newton's Method for unconstrained SSIM-based optimization with Tikhonov regularization

initialize Choose $x = x_0$ according to (9.36);
data preprocessing $\bar{y} = \frac{1}{n} \mathbf{1}^T y$, $y = y - \bar{y} \mathbf{1}$;
repeat
 $x = x - J_f(x)^{-1} f(x)$;
until stopping criterion is met (e.g., $\|x^{(new)} - x^{(old)}\|_\infty < \epsilon$);
return x , $y = y + \bar{y} \mathbf{1}$.

Notice that this algorithm can be used for any unconstrained SSIM-based optimization problem by defining f and J_f accordingly. The MATLAB code can be found in Appendix C, Section C.2.

¹The Frobenius norm of an $m \times n$ matrix A is defined as $\|A\|_F = \sqrt{\sum_{i=1}^m \sum_{j=1}^n |a_{ij}|^2}$.

9.2.2 Non-differentiable $h(\mathbf{x})$

In this case, a different approach must be taken. Let us consider the particularly important example $h(x) = \|x\|_1$, i.e., we minimize the following functional

$$\min_x \{T(Dx, y) + \lambda \|x\|_1\}, \quad (9.39)$$

where $y \in \mathbb{R}^m$ is a given observation, λ is a regularization parameter, and D is an $m \times n$ matrix. In this case, the optimal x^* satisfies

$$D^T D x^* \in \frac{D^T y}{\text{SSIM}(D x^*, y)} - \lambda \left(\frac{\|D x^*\|_2^2 + \|y\|_2^2 + C}{2 \text{SSIM}(D x^*, y)} \right) \partial \|x^*\|_1, \quad (9.40)$$

where $\partial \|x^*\|_1$ is the subdifferential of $\|x\|_1$ at x^* [4].

Definition 9.2.1 (Subdifferential). *Let g be a convex function $g : \mathbb{R}^n \rightarrow \mathbb{R}$ and x_0 a fixed vector in \mathbb{R}^n . The subdifferential of g at x_0 is the set*

$$\partial g(x_0) := \{z \in \mathbb{R}^n : g(x_0) + z^T(x - x_0) \leq g(x) \ \forall x \in \mathbb{R}^n\}. \quad (9.41)$$

Example 9.2.1. *Let us consider the subdifferential of the absolute value on \mathbb{R} . According to Definition 9.2.1, the subdifferential of $|x|$ is given by*

$$\partial |x| = \begin{cases} -1, & x < 0, \\ [-1, 1], & x = 0, \\ 1, & x > 0. \end{cases} \quad (9.42)$$

This simple example is important in the discussion to follow.

To find x^* we employ a coordinate descent approach [111], that is, we minimize (9.39) along each component of x while the other components are fixed. From (9.40), for the i -th entry of $x \in \mathbb{R}^n$, the optimal coordinate x_i is given by

$$x_i \in \frac{D_i^T y}{\text{SSIM}(Dx, y) \|D_i^T\|_2^2} - D_i^T D x_{-i} - \lambda \left(\frac{\|Dx\|_2^2 + \|y\|_2^2 + C}{2 \text{SSIM}(Dx, y) \|D_i^T\|_2^2} \right) \partial |x_i|, \quad (9.43)$$

where D_i^T is the i -th row of the transpose of D , x_{-i} is the vector x whose i -th component is set to zero, and $\partial |x_i|$ is the subdifferential of $|x_i|$.

The value of x_i can be found by examining the different cases that arise in (9.43). To begin with, we define

$$\tau_i(x_i) = \lambda \left(\frac{\|Dx\|_2^2 + \|y\|_2^2 + C}{2\text{SSIM}(Dx, y)\|D_i^T\|_2^2} \right) \quad (9.44)$$

and

$$a_i(x_i) = \frac{D_i^T y}{\text{SSIM}(Dx, y)\|D_i^T\|_2^2} - D_i^T Dx_{-i}. \quad (9.45)$$

Then, $x_i = 0$ if

$$a_i(0) \in \tau_i(0)[-1, 1]. \quad (9.46)$$

As expected, $x_i > 0$ if $a_i(0) > \tau_i(0)$, so that

$$x_i = a_i(x_i) - \tau_i(x_i). \quad (9.47)$$

Similarly, we obtain $x_i < 0$ if $a_i(0) < -\tau_i(0)$, in which case x_i is given by

$$x_i = a_i(x_i) + \tau_i(x_i). \quad (9.48)$$

Notice that when $x_i \neq 0$, we have an expression of the form $x_i = g(x_i)$. This non-linear equation may be solved by employing different approaches. For instance, let $g_{\pm}(x_i)$ be defined as

$$g_{\pm}(x_i) := a_i(x_i) \mp \tau_i(x_i). \quad (9.49)$$

Such equation may be solved using either a fixed-point iteration scheme, provided that $g_{\pm}(x_i)$ is a contraction. Further, by defining $f_{\pm}(x)$ as

$$f_{\pm}(x) := x_i - g_{\pm}(x_i), \quad (9.50)$$

we can use a root finding algorithm for finding the value of x_i such that $f_{\pm}(x_i) = 0$. Regardless of which method is chosen for finding x_i , equations (9.46), (9.47) and (9.48) can be combined into the following single operator:

$$\Phi_{\tau_i(0)}(a_i(0)) := \begin{cases} \text{Solve } x_i = a_i(x_i) - \tau_i(x_i), & \text{if } a_i(0) > \tau_i(0), \\ \text{Solve } x_i = a_i(x_i) + \tau_i(x_i), & \text{if } a_i(0) < -\tau_i(0), \\ x_i = 0, & \text{if } |a_i(0)| \leq \tau_i(0). \end{cases} \quad (9.51)$$

Eq. (9.51) is an important result since it may be considered as an extension of the widely used soft-thresholding (ST) operator [41, 113] for the purpose of solving the unconstrained SSIM-based optimization problem (9.39).

It is worthwhile to mention that we follow a fixed-point approach for solving the non-linear equations of the soft-thresholding operator defined above. Indeed, it can be proved that $g_{\pm}(x_i)$ has a fixed-point if the regularization parameter λ is less than a certain value.

Theorem 9.2.3. Let $g_{\pm}(x_i) := a_i(x_i) \mp \tau_i(x_i)$, $y \in \mathbb{R}^m$ be a given observation, $x_0 \in \mathbb{R}^n$ an initial condition, and D an $m \times n$ matrix whose columns have L^2 norm equal to one. Then, $g_{\pm}(x_i)$ has a fixed point provided that the regularization parameter λ fulfills the following condition:

$$\lambda < \min_{x \in \Omega} \left\{ \frac{(SSIM(Dx, y))^2 - \delta \|y\|_2}{2 \|D\|_2 \|x\|_2} \right\}, \quad (9.52)$$

where δ is a small positive constant, and $\Omega \subset \mathbb{R}^n$ is the set of vectors that lie on the path followed by an algorithm that solves problem (9.39) starting at x_0 .

Proof. See appendix B, Theorem B.0.18. □

The latter theorem shows that the higher the SSIM between Dx and y for all $x \in \Omega$, and the smaller the L^2 norms of x , y and D , the greater the regularization parameter λ can be. For instance, if $\|y\|_2 = 1$ and D is orthogonal ($\|D\|_2 = 1$), by definition, an algorithm that solves problem (9.39) will try to maximize the SSIM, therefore, as such an algorithm progresses, the term $SSIM(Dx, y)$ will be close to one. On the other hand, since the similarity between Dx and y tends to be maximized, the norm of all $x \in \Omega$ will also be approximately one. Moreover, experimental results show that typical values of δ are of the order of 10^{-3} and less. Thus, by the previous theorem, a good estimate for the maximum value of λ under these conditions is $\frac{1}{2}$.

From the above discussion, we introduce the following algorithm to determine the optimal x^* for problem (9.39).

Algorithm III: Coordinate Descent algorithm for unconstrained SSIM-based optimization with L^1 norm regularization

```

initialize  $x = x_0$ ;
data preprocessing  $\bar{y} = \frac{1}{n} \mathbf{1}^T y$ ,  $y = y - \bar{y} \mathbf{1}$ ;
repeat
  for  $i = 1$  to  $n$  do
     $x_i = \Phi_{\tau_i(0)}(a_i(0))$ ;
  end
until stopping criteria is met (e.g.,  $\|x^{(new)} - x^{(old)}\|_{\infty} < \epsilon$ );
return  $x$ ,  $y = y + \bar{y} \mathbf{1}$ .

```

As expected, the result proved in Theorem 7.3.4 can be used to recover the non-zero mean optimal solution x^* , provided that the means of y and Dx^* are equal and that $(D^T D)^{-1}$ exists. The MATLAB code for this algorithm is presented in Appendix C, Section C.3.

We close this section by mentioning that if the fixed-point approach is employed it is recommended to work with images whose grayscale values have been normalized. Also, Theorem 9.2.3 suggests that if Dx^* is close to y with respect to the SSIM, Dx_0 should be close to y in the SSIM sense as well. Experimental results show that the optimal L^2 solution of the unconstrained problem $\|Dx - y\|_2^2$ is a good initial guess that satisfies this condition—i.e., $x_0 = (D^T D)^{-1} D^T y$, as long as the inverse of $D^T D$ is well defined.

9.3 Experiments

Algorithms I and II can be used for many different SSIM-based applications. In the results presented below, however, we have focussed our attention on the performance of Algorithm II for solving problem (9.39) when D is an orthogonal transformation. To measure its efficacy, we compare the solutions obtained by the proposed method with the set of solutions of the L^2 version of problem (9.39), namely,

$$\min_x \left\{ \frac{1}{2} \|Dx - y\|_2^2 + \lambda \|x\|_1 \right\}, \quad (9.53)$$

which can be solved by means of the soft thresholding (ST) operator [4, 113] if D is an orthogonal matrix.

The experiments reported below were concerned with the recovery of discrete cosine transform (DCT) coefficients. All images were divided into non-overlapping 8×8 pixel blocks, the means of which were subtracted prior to processing. After a block has been processed, its mean is added. Although this procedure is not required for L^2 approaches, it has been performed for the sake of a fair comparison between the two methods.

In Figure 9.1, the first two plots from left to right corresponds to the average SSIM of all the reconstructions versus the L^0 norm² of the recovered coefficients for the test images *Lena* and *Mandrill*. The average SSIM was computed by combining and averaging all the computed SSIMs from all 4096 non-overlapping blocks for both *Lena* and *Mandrill* (both test images have 512×512 pixels). It can be clearly seen that the proposed algorithm

²The L^0 norm of a vector $x \in \mathbb{R}^n$, which is denoted as $\|x\|_0$, is defined as the number of non-zero components thereof.

outperforms the L^2 -based method (ST). This is because minimization of the dissimilarity measure $T(x, y)$ in Eq. (7.2) is equivalent to maximization of $\text{SSIM}(Dx, y)$, which produces an enhancement in contrast [18]. This effect is demonstrated in the nature of the recovered coefficients. Firstly, the degree of shrinking and thresholding of DCT coefficients by our proposed method is not as strong as ST. Secondly, in some cases, there are DCT coefficients which are thresholded (i.e., set to zero) by the L^2 approach, but kept non-zero by the SSIM-based method for the sake of contrast. These effects are demonstrated in the third and fourth plots in Figure 9.1. In these two plots, the same block from the image *Lena* was processed, but subjected to two different amounts of regularization.

In addition, some visual results are shown in Figure 9.2. In this case, a sub-image from the test image *Lena* was used. The original and recovered images are presented in the bottom row. Regularization was carried out so that the sparsity induced by each method is the same; that is, the L^1 norm of the set of recovered coefficients is 13 in all cases. In the top row of the figure are shown the SSIM maps that illustrate the similarity between the reconstructions and the original image. The brightness of these maps indicates the degree of similarity between corresponding image blocks—the brighter a given point the greater the magnitude of the SSIM between the retrieved and the original image blocks at that location [122]. It can be seen that the performance of the proposed method and the L^2 approach is very similar. However, the proposed algorithm does perform better than ST in terms of SSIM. This can be seen at some locations in the SSIM maps. For instance, note that the upper left corner of the SSIM map of the proposed method is slightly brighter than the corresponding regions of the other two SSIM maps. This is also evident at other locations. Moreover, the enhancement of contrast is clearly seen when the pupils of the left eyes are compared. With regard to numerical results, the average $T(Dx, y)$ for the L^2 approach is 0.8609, whereas for the proposed method is 0.8864, which is moderately better.




Figure 9.1: The first two plots from left to right correspond to the average SSIM versus the L^0 norm of the recovered coefficients for the test images *Lena* and *Mandrill*. In the last two plots, a visual comparison between the original and recovered coefficients from a particular block of the *Lena* image can be appreciated. Regularization is carried out so that the two methods being compared induce the same sparseness in their recoveries. In the two shown examples, the same block was processed but subjected to different amounts of regularization. In particular, the L^0 norm of the set of DCT coefficients that were recovered by both the proposed method and ST is 3 for the first example (third plot), and 15 for the second (fourth plot).




Figure 9.2: Visual results for a patch from the test image *Lena*. In all cases, the L^0 norm of the recovered DCT coefficients for each non-overlapping block is 13. In the upper row, the SSIM maps between the reconstructions and the original patch are shown. Reconstructed and original patches can be seen in the lower row. The average $T(Dx, y)$ of all non overlapping blocks for the proposed method is 0.8864, whereas for ST is 0.8609.

Chapter 10

SSIM-based Optimization Using Mollifiers

In section 9.2.2 an algorithm for solving the following non-smooth problem was presented,

$$\min_x \{T(Dx, y) + \lambda \|x\|_1\}, \quad (10.1)$$

where D is an $m \times n$ matrix, $y \in \mathbb{R}$ a given observation, and λ a regularization parameter. As was shown, it was necessary to introduce an extension of the classical soft-thresholding operator in order to find a zero-mean optimal solution of this SSIM-based optimization problem. Although this approach is efficient for solving problem (10.1), it is not the only option with which non-smooth SSIM-based optimization problems can be addressed.

For instance, the L^2 counterpart of (10.1), which is known as the LASSO (Least Absolute Shrinkage and Selection Operator) problem [4],

$$\min_x \left\{ \frac{1}{2} \|Dx - y\|_2^2 + \lambda \|x\|_1 \right\}, \quad (10.2)$$

can be solved without relying on techniques from subdifferential calculus. One approach is to cast problem (10.2) as a Quadratic Program (QP) and employ generic methods to solve the QP version of (10.2) [4]; another possibility is to approximate the L^1 norm by a family of smooth functions known as *mollifiers*, and use gradient-based methods for solving the “mollified” counterpart of problem (10.2) [50]. In particular, in [114], the Gaussian distribution is used as an approximate mollifier to solve a smooth version of (10.2), which is obtained by convolving each component of the L^1 norm with a standard

one-dimensional Gaussian density function of variance ε^2 . This technique allows the usage of gradient-based methods for approximating the optimal solution x^* of (10.2). In this case, the smooth approximation of the original problem is given by

$$\min_x \left\{ \frac{1}{2} \|Dx - y\|_2^2 + \lambda \|x\|_1^\varepsilon \right\}, \quad (10.3)$$

where $\|x\|_1^\varepsilon$ is equal to

$$\|x\|_1^\varepsilon = \sum_{i=1}^n \int_{\mathbb{R}} |x_i - z_i| \hat{\phi}_\varepsilon \left(\frac{z_i}{\varepsilon} \right) dz. \quad (10.4)$$

Here, $\hat{\phi}(x)$ is the standard normal distribution in one dimension.

In this chapter, inspired by the idea of optimizing non-smooth functions via mollifiers, we introduce a gradient-based method that solves the following smooth version of (10.1):

$$\min_x \{T(Dx, y) + \lambda \|x\|_1^\varepsilon\}, \quad (10.5)$$

where $\|x\|_1^\varepsilon$ is obtained by convolving the L^1 norm with a multivariate Gaussian distribution of variance ε^2 . Moreover, we show in Theorem 10.1.6 that the functional of (10.5) epi-converges to the functional of (10.1), which ensures that the sequence of minimizers x_ε^* of (10.5) converges to an optimal solution x^* of (10.1) when $\varepsilon \rightarrow 0$. Numerical results can be found in section 10.3 of this chapter.

10.1 Smoothing via Mollifiers

In this section, we recall some of the basic notions and properties of mollifiers and introduce the smoothing approach. Let us first introduce the following definition for a family of mollifiers.

Definition 10.1.1. [50] *For each $\varepsilon > 0$, a **family of mollifiers** is a family of functions $\varphi_\varepsilon \in C_0^\infty(\mathbb{R}^n)$ that satisfies the following properties:*

1. $\varphi_\varepsilon(x) \geq 0$, for all $x \in \mathbb{R}^n$,
2. $\text{supp } \varphi_\varepsilon \subseteq \{x \in \mathbb{R}^n : \|x\| \leq \varepsilon\}$,
3. $\int_{\mathbb{R}^n} \varphi_\varepsilon(x) dx = 1$.

A family of mollifiers that is used quite often in applications is the following:

$$\varphi_\varepsilon(x) = \begin{cases} e^{\frac{1}{\|x\|_2 - \varepsilon^2}}, & \|x\|_2 \leq \varepsilon, \\ 0, & \|x\|_2 > \varepsilon. \end{cases} \quad (10.6)$$

One of the advantages of such mollifiers is that any function in $L^1_{loc}(\mathbb{R}^n)$ —that is, any locally integrable function—can be approximated via convolution by these kind of functions. Given a family of mollifiers $\{\varphi_\varepsilon : \mathbb{R}^n \rightarrow \mathbb{R}_+, \varepsilon \in \mathbb{R}_+\}$, we can define a smooth function approximation f^ε of f through the convolution

$$f_\varepsilon(x) := (f * \varphi_\varepsilon)(x) = \int_{\mathbb{R}^n} f(x - z)\varphi_\varepsilon(z)dz = \int_{\mathbb{R}^n} f(z)\varphi_\varepsilon(x - z)dz. \quad (10.7)$$

The sequence $f * \varphi_\varepsilon$ is said to be a sequence of mollified functions.

Moreover, from a computational perspective, this convolution can be computed as an expected value. Let us notice that if $Y_\varepsilon(x, z)$ is a random vector with density defined by $z \rightarrow \varphi_\varepsilon(x - z)$, the above definition can be written as

$$(f * \varphi_\varepsilon)(x) := \mathbb{E}(f(Y_\varepsilon(x, z))), \quad (10.8)$$

where \mathbb{E} is the expected value of the random variable $f(Y_\varepsilon(x, z))$. The advantage of this stochastic interpretation is that allows us to compute the above integral by estimating the expected value of $f(Y_\varepsilon(x, z))$, which is numerically much easier than calculating such a convolution directly.

On the other hand, we also have the following important result regarding the convergence of the smooth approximation of f via convolution.

Theorem 10.1.1. *Let $f \in C(\mathbb{R}^n)$. Then $f * \varphi_\varepsilon$ converges pointwise to f , i.e. $f * \varphi_{\varepsilon_m}(x_m) \rightarrow f(x)$ for all $x_m \rightarrow x$. In fact, $f * \varphi_\varepsilon$ converges uniformly to f on every compact subset of \mathbb{R}^n as $\varepsilon_m \rightarrow 0$.*

Proof. See [16], Proposition 4.21, p. 108. □

Furthermore, the previous convergence property can be generalized.

Definition 10.1.2. [50] *A sequence of functions $\{f_m\} : \mathbb{R}^n \rightarrow \mathbb{R}$ epi-converges to $f : \mathbb{R}^n \rightarrow \mathbb{R}$ at x if:*

1. $\liminf_{m \rightarrow +\infty} f_m(x_m) \geq f(x)$ for all $x_m \rightarrow x$;

2. $\lim_{m \rightarrow +\infty} f_m(x_m) = f(x)$ for some sequence $x_m \rightarrow x$.

The sequence $\{f_m\}$ epi-converges to f if this holds for all $x \in \mathbb{R}^n$. In this case, we say that **f is the epi-limit of $\{f_m\}$** .

It is worthwhile to mention that epi-convergence is a special case of Γ -convergence [14]. In fact, given a sequence $\{f_n\}$ of functions that epi-converges to f , epi-convergence of $\{f_n\}$ amounts to the *Painlevé-Kuratowski convergence* (PK-convergence) of the epigraphs of f_n to the epigraph of f [9]. As is known, the epigraphs of f_n PK-converge to the epigraph of f if and only if the sequence $\{f_n\}$ Γ -converges to f [14].

On the other hand, it can be checked that if the sequence $\{f_m\}$ converges continuously, then it also epi-converges. The importance of the notion of epi-convergence is that it ensures the convergence of minimizers of $\{f_m\}$ to the minimizers of f [104]. Moreover, if f is the epi-limit of some sequence $\{f_m\}$, then f is lower semicontinuous.

Definition 10.1.3. [50] A function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is **strongly lower semicontinuous** (s.l.s.c.) at x if it is lower semicontinuous at x and there exists a sequence $x_m \rightarrow x$ with f continuous at x_m (for all m) such that $f(x_m) \rightarrow f(x)$. The function **f is strongly lower semicontinuous** if this holds at all x .

Theorem 10.1.2. Let $\varepsilon_m \rightarrow 0$ as $m \rightarrow +\infty$. For any s.l.s.c. locally integrable function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, and any associated sequence $\{f_{\varepsilon_m}\}$ of mollified functions, we have that f is the epi-limit of $\{f_{\varepsilon_m}\}$.

Proof. See [50], Theorem 3.7, p. 5. □

Observe that the functional of problem (10.5) is a function of the form $g + f * \varphi_\varepsilon$, where g is differentiable. This implies that the previous results do not necessarily hold in this case. Nevertheless, it can be proved that such a function is the epi-limit of a sequence of smooth approximations of the functional of problem (10.1). To see this, we use the following lemma.

Lemma 10.1.1. Let $\{\varphi_\varepsilon : \mathbb{R}^n \rightarrow \mathbb{R}_+, \varepsilon \in \mathbb{R}_+\}$ be a family of mollifiers. Also, let z be a random variable whose density function is given by $\varphi_\varepsilon(z)$. If $\mathbb{E}(z) = 0$ for all $\varepsilon \in \mathbb{R}_+$, then the mollified norm $\|x\|_1^\varepsilon = \|x\|_1 * \varphi_\varepsilon(x)$ is greater or equal than its non-smooth counterpart $\|x\|_1$ for any $x \in \mathbb{R}^n$ and all $\varepsilon \in \mathbb{R}_+$.

Proof. Let $f(z) = \|x - z\|_1$. Then, by convexity of f and using Jensen's inequality, we have that

$$\|x - \mathbb{E}(z)\|_1 \leq \int_{\mathbb{R}^n} \|x - z\|_1 \varphi_\varepsilon(z) dz. \quad (10.9)$$

Given that $\mathbb{E}(z) = 0$ for all $\varepsilon \in \mathbb{R}_+$, we immediately obtain that $\|x\|_1 \leq \|x\|_1^\varepsilon$ for all $x \in \mathbb{R}^n$ and any $\varepsilon \in \mathbb{R}_+$. \square

Theorem 10.1.3. *Let $g : \mathbb{R}^n \rightarrow \mathbb{R}$ and $\{\varepsilon_m\}$ be a sequence of positive real numbers such that $\varepsilon_m \rightarrow 0$. The function $g(x) + \gamma\|x\|_1$ is the epi-limit of the sequence of functions $h_m : \mathbb{R}^n \rightarrow \mathbb{R}$ defined as*

$$h_m(x) := g(x) + \lambda\|x\|_1^{\varepsilon_m}. \quad (10.10)$$

Proof. Let $\{x_m\}$ be a sequence in \mathbb{R}^n such that $x_m \rightarrow x$. Since $\|x\|_1^{\varepsilon_m}$ converges to $\|x\|_1$ as m tends to infinity, we have that

$$\lim_{m \rightarrow \infty} h_m(x_m) = g(x) + \lambda\|x\|_1. \quad (10.11)$$

Also, by lemma 10.1.1, it follows that for any $x_m \in \mathbb{R}^n$ and any $\varepsilon_m \in \mathbb{R}_+$

$$g(x_m) + \gamma\|x_m\|_1^{\varepsilon_m} \geq g(x_m) + \lambda\|x_m\|_1. \quad (10.12)$$

Taking \liminf of both sides over all sequences $x_m \rightarrow x$ we obtain that

$$\liminf_{m \rightarrow \infty} g(x_m) + \gamma\|x_m\|_1^{\varepsilon_m} \geq g(x) + \lambda\|x\|_1. \quad (10.13)$$

This completes the proof. \square

By means of mollified functions it is possible to define generalized directional derivatives for a non-smooth function f , which, under suitable regularity of f , coincide with *Clarke's subdifferential*. Such a subdifferential is defined as

$$\partial_C f(x) := \{s \in \mathbb{R}^n : \langle s, t \rangle \leq d_C f(x, t), \forall t \in \mathbb{R}^n\}, \quad (10.14)$$

where $d_C f(x, t)$ is the *Clarke subderivative*, which is given by

$$d_C f(x, t) := \limsup_{m \rightarrow \infty, \rho \downarrow 0} \frac{f(x_m + \rho t) - f(x_m)}{\rho}. \quad (10.15)$$

Here, the \limsup is calculated with respect to all sequences $\{x_m\}$ such that $x_m \rightarrow x$ and $\rho \rightarrow 0$. In particular, in [50] (see also [73, 34, 33] for alternative definitions of generalized derivatives through mollified functions), a generalized gradient with respect to the mollifier sequence $\{f_{\varepsilon_m}\}$ has also been defined in the following way:

Definition 10.1.4. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be locally integrable. Also, let $\{f_{\varepsilon_m}\}$ be a sequence of averaged functions that are obtained from the convolution of f with a sequence of mollifiers $\{\varphi_{\varepsilon_m}\}$, where $\varepsilon_m \rightarrow 0$ as $m \rightarrow \infty$. Assume that the mollifiers are such that the averaged functions have continuous first order derivatives. The set of the φ -mollifier subgradients of f at x is by definition

$$\partial_\varphi f(x) := \text{Lim sup}_{m \rightarrow +\infty} \{\nabla f_{\varepsilon_m}(x_m), x_m \rightarrow x\}, \quad (10.16)$$

that is, $\partial_\varphi f(x)$ is the outer superior limit of the sequence of sets $\{\nabla f_{\varepsilon_m}(x_m), x_m \rightarrow x\}$; i.e., the cluster points of all possible sequences $\{\nabla f_{\varepsilon_m}(x_m)\}$ such that $x_m \rightarrow x$. Here, ∇f_ε is equal to

$$\nabla f_\varepsilon(x) = \int_{\mathbb{R}^n} f(z) \nabla \varphi_\varepsilon(x - z) dz. \quad (10.17)$$

The full Φ -mollifier subgradient set is

$$\partial_\Phi f(x) := \bigcup_{\varphi} \partial_\varphi f(x), \quad (10.18)$$

where φ ranges over all possible sequences of mollifiers that generate smooth averaged functions.

Moreover, this generalized gradient is equal to the Clarke's subdifferential if f is locally integrable and locally Lipschitz.

Theorem 10.1.4. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be locally integrable and locally Lipschitz at x . Then, $\partial_\Phi f(x)$ coincides with the Clarke's subdifferential of f at x ; that is, $\partial_\Phi f(x) = \partial_C f(x)$.

Proof. See [50], Theorem 4.10, p. 15. □

The above result can be extended to include functions of the form $g(x) + \lambda \|x\|_1$, provided that the function $g : \mathbb{R}^n \rightarrow \mathbb{R}$ is differentiable, locally integrable and locally Lipschitz. This is shown in the following theorem.

Theorem 10.1.5. Let $g : \mathbb{R}^n \rightarrow \mathbb{R}$ be differentiable, locally integrable and locally Lipschitz at x . Also, let $h(x) := g(x) + \lambda \|x\|_1$. Then, $\partial_\Phi h(x)$ coincides with Clarke's subdifferential of h at x .

Proof. Let $f(x) = \|x\|_1$. Clearly, $f(x)$ is Lipschitz continuous with Lipschitz constant one since

$$\left| \|x\|_1 - \|y\|_1 \right| \leq \|x - y\|_1. \quad (10.19)$$

Moreover, since g is locally Lipschitz at x , it follows that h is locally Lipschitz at x as well. Similarly, given that f is locally integrable, we have that h is locally integrable as well. Also, by definition of $\partial_\Phi(\cdot)$, we have that

$$\begin{aligned} \partial_\Phi h(x) &:= \bigcup_{\varphi} \left(\text{Lim sup}_{m \rightarrow \infty} \{ \nabla(g(x_m) + \lambda \|x_m\|_1^{\varepsilon_m}), x_m \rightarrow x \} \right) \\ &:= \bigcup_{\varphi} \left(\text{Lim sup}_{m \rightarrow \infty} \{ \nabla g(x_m), x_m \rightarrow x \} + \lambda \text{Lim sup}_{m \rightarrow +\infty} \{ \nabla(\|x_m\|_1^{\varepsilon_m}), x_m \rightarrow x \} \right) \\ &:= \nabla g(x) + \lambda \bigcup_{\varphi} \partial_\varphi f(x) \\ &:= \bigcup_{\varphi} \partial_\varphi h(x). \end{aligned}$$

By theorem 10.1.4, since h is locally integrable and locally Lipschitz at x , we have that the last equation is the set of Clarke's subgradients of h at x . \square

Observe that the dissimilarity measure $T(Dx, y)$ as a function of x is bounded—here, D is an $m \times n$ matrix and $y \in \mathbb{R}^m$ is a given vector. Therefore, it is locally integrable over any compact set $K \subset \mathbb{R}^n$. Also, it is differentiable and its gradient is bounded¹, which implies that it is Lipschitz continuous. Given this, Theorems 10.1.3 and 10.1.5 guarantee that the sequence of minimizers x_ε^* of (10.5) converges to a minimizer x^* of (10.1) when ε tends to zero. In other words, $x_\varepsilon^* \rightarrow x^*$ as $\varepsilon \rightarrow 0$.

In the sequel we will use the following family of smoothing Gaussian functions:

$$\hat{\phi}_\varepsilon(x) = \frac{1}{\varepsilon^n} \hat{\phi}\left(\frac{x}{\varepsilon}\right), \quad (10.20)$$

where

$$\hat{\phi}(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{\|x\|_2^2}{2}}. \quad (10.21)$$

It is well known that $\hat{\phi}_\varepsilon$ is a density function, so its integral over \mathbb{R}^n is equal to one, it is smooth, and $\hat{\phi}_\varepsilon$ goes to zero when $\|x\|_2 \rightarrow \infty$. However, this sequence is not a proper

¹Indeed, it is not hard to prove that $|\nabla T(Dx, y)| \leq \frac{\|D^T\|_2}{\sqrt{2\|y\|_2^2 + C - \|y\|_2^2}}$.

family of mollifiers as each element $\hat{\phi}_\varepsilon$ does not have compact support. Nevertheless, it can be proved that, given a function f , the family of smooth functions

$$\hat{f}_\varepsilon(x) = (f * \hat{\phi}_\varepsilon)(x) = \frac{1}{\varepsilon^n} \int_{\mathbb{R}^n} f(x-z) \hat{\phi}\left(\frac{z}{\varepsilon}\right) dz \quad (10.22)$$

epi-converges to f when $\varepsilon \rightarrow 0$.

Theorem 10.1.6. *The family of functions defined as $\hat{f}_\varepsilon(x) := (f * \hat{\phi}_\varepsilon)(x)$ epi-converges to f as $\varepsilon \rightarrow 0$.*

Proof. Let $\{x_m\}$ be a sequence in \mathbb{R}^n such that $x_m \rightarrow x$. Also, let $\{\varepsilon_m\}$ be a sequence of real numbers such that $\varepsilon_m \rightarrow 0$ as $m \rightarrow \infty$. Observe that the sequence of functions $\hat{f}_{\varepsilon_m}(x) = (f * \hat{\phi}_{\varepsilon_m})(x)$ converges to f as $m \rightarrow \infty$. Then, we have that

$$\lim_{m \rightarrow \infty} \hat{f}_{\varepsilon_m}(x_m) = f(x). \quad (10.23)$$

On the other hand, by Jensen's inequality, $f_\varepsilon(x) \geq f(x)$ holds for all $x \in \mathbb{R}^n$ and all $\varepsilon \in \mathbb{R}_+$. Therefore,

$$\hat{f}_{\varepsilon_m}(x_m) \geq f(x_m). \quad (10.24)$$

By taking \liminf of both sides of this inequality over all sequences $\{x_m\}$ such that $x_m \rightarrow x$ we obtain that

$$\liminf_{m \rightarrow \infty} \hat{f}_{\varepsilon_m}(x_m) \geq f(x). \quad (10.25)$$

By definition 10.1.2, we conclude that \hat{f}_ε epi-converges to f as $\varepsilon \rightarrow 0$. \square

10.2 SSIM-based Optimization with Sparsity

In this section, we propose a simple gradient-based method for solving problem (10.5). To develop such a method, we first need to define the following non-linear functional:

$$f_\varepsilon(x) := T(Dx, y) + \lambda \|x\|_1^\varepsilon. \quad (10.26)$$

Its gradient is given by

$$\nabla f_\varepsilon(x) = \frac{2\text{SSIM}(Dx, y) D^T Dx - 2D^T y}{\|Dx\|_2^2 + \|y\|_2^2 + C} + \lambda \int_{\mathbb{R}^n} \|z\|_1 \nabla \hat{\phi}_\varepsilon\left(\frac{x-z}{\varepsilon}\right) dz, \quad (10.27)$$

where $\nabla \hat{\phi}_\varepsilon(x)$ is equal to

$$\nabla \hat{\phi}_\varepsilon \left(\frac{x}{\varepsilon} \right) = \frac{-x}{\sqrt{(2\pi)^n \varepsilon^{n+2}}} e^{-\frac{\|x\|_2^2}{2\varepsilon^2}}. \quad (10.28)$$

By using (10.27), and defining $\mathbf{1} = [1, \dots, 1]^T \in \mathbb{R}^m$, we are now in a position to introduce the following gradient-descent algorithm for solving (10.5):

Algorithm IV: Gradient descent for unconstrained SSIM-L¹ optimization via mollifiers

initialize Choose $x = x_0, \gamma$;
data preprocessing $\bar{y} = \frac{1}{n} \mathbf{1}^T y, y = y - \bar{y} \mathbf{1}$;
repeat
 $x = x - \gamma \nabla f_\varepsilon(x)$;
until stopping criterion is met (e.g., $\|x^{(new)} - x^{(old)}\|_\infty < \delta$);
return $x, y = y + \bar{y} \mathbf{1}$.

The MATLAB code for this algorithm is presented in Appendix C, Section C.4. Also, notice that this algorithm will return a zero-mean optimal solution x^* of problem (10.5). As usual, if it is known that the means of y and Dx^* are equal, then, the non-zero mean optimal solution x^* can be recovered by means of the following equation:

$$x^* = x^* + \bar{y} (D^T D)^{-1} D^T \mathbf{1}, \quad (10.29)$$

provided that the inverse of $D^T D$ exists.

10.3 Experiments

In these experiments we solve the approximate sparse reconstruction problem (10.5) with the proposed gradient-descent algorithm. Its performance is measured by comparing its recovered solutions with the solutions obtained by the Algorithm III for solving (10.1) (see Section 9.2.2) and the solutions of problem (10.2). In all computations a set of Discrete Cosine Transform (DCT) coefficients is to be recovered; therefore, problem (10.2) was solved by means of the soft-thresholding (ST) operator [4, 113].

In all the experiments images were divided into non-overlapping 8×8 pixel blocks. As expected, the means of each block are subtracted prior to processing, which are added after the non-overlapping blocks have been processed. This is also done when problem (10.2) is solved at each pixel block for the sake of a fair comparison between the different methods.

It is worthwhile to mention that for computing the integral of the gradient of (10.26) we performed a Monte Carlo integration (see Eq. (10.27)). This can be done by noticing the fact that calculating

$$\int_{\mathbb{R}^n} \|z\|_1 \nabla \hat{\phi}_\varepsilon \left(\frac{x-z}{\varepsilon} \right) dz \quad (10.30)$$

is equivalent to computing the expected value $\mathbb{E}(\|z\|_1(x-z))$, where z follows a Gaussian distribution of variance ε^2 and mean equal to x .

In the left plot in Figure 10.1 is shown an example of the optimal DCT coefficients that are obtained by the different methods that are being compared. Plots in red and green correspond to the solutions obtained by Algorithm III and ST respectively. The blue plot is the optimal solution that was obtained by the proposed method when $\varepsilon = 0.001$. True sparsity in the solution is not achieved since this occurs in the limit when $\varepsilon \rightarrow 0$; nevertheless, it can be seen that the proposed method gives a good non-sparse approximation of the solution of the non-smooth problem (10.1). This in fact can be useful for providing a good initial guess of a thresholding method that solves (10.1) [114]. In the plot on the right it can be seen how a sequence of optimal solutions of (10.5) gets closer to a solution of (10.1) as ε tends to zero. In this case, the plot in magenta corresponds to the set of optimal DCT coefficients that is obtained by solving problem (10.1).

As for visual results, these are shown in Figure 10.2. In the example shown, a sub-image of the test image *Lena* was employed. In the bottom row the original sub-image and its recovered counterparts can be observed. Regularization was carried out at all non-overlapping pixel blocks in such a way that the number of non-zero DCT coefficients obtained by the Algorithm III and the ST operator is always 19. As for the regularization of the proposed algorithm, the values of the regularization parameter that were used were the same that were employed for the algorithm that solves the non-smooth problem (10.1). This was done in this way since in the limit $\varepsilon \rightarrow 0$, both problems (10.1) and (10.5) are equivalent. In other words, the strength of the regularization tends to be the same for these two methods. Along with the images of the bottom row, the SSIM maps that depict the similarity between the original sub-image and its reconstructions are shown in the top row. The higher the brightness of these maps at a given location, the higher the SSIM at that particular point [122]. As mentioned in Section 9.3, performance of the ST approach and Algorithm III is very similar, however, the average $T(Ax, y)$ of the non-mollified SSIM-based optimization problem (0.9156) is slightly higher than the average $T(Ax, y)$ of the L^2

counterpart (0.9117). As for the proposed approach, the recovered image is visually more appealing than the other two methods, and as expected, the average $T(Ax, y)$ is the highest of the three approaches that are being compared (0.9629). This should not be surprising since several recovered DCT coefficients are not set to zero by the proposed algorithm, which is not always the case for the other two methods.




Figure 10.1: The plot on the left shows an example of the different solutions that were obtained by the three methods that were compared. The plot on the right shows a visual example of how a sequence of minimizers x_ϵ^* of the mollified SSIM-based optimization problem (10.5) converges to a minimizer x^* of the non-smooth problem (10.1).




Figure 10.2: Visual results for a sub-image from the test image *Lena*. In all cases, regularization is carried out to induce the same degree of sparsity for all methods at each non-overlapping pixel block. In the bottom row the original image and its reconstructions are shown. The corresponding SSIM maps can be seen in the top row.

Chapter 11

Conclusions and Future Work

The concept of a function whose range is infinite dimensional has been an object of study in the mathematical community since the third decade of the last century [11]. Nowadays, there is a vast literature on the extension of classical results for real-valued functions to the case of functions that assume values in a Banach space, which are also known as *Banach-valued functions* [40, 97, 23, 86]. As was mentioned before, several fields have benefited from this approach thanks to the works of Bochner [11], Peetre [97], Diestel [40] and many others [126, 23, 85, 69]. Surprisingly, the image processing community has barely considered these contributions.

Given the latter observation, we consider that the main contribution of the first part of this thesis is the connection that has been established between the mathematical theory of Banach-valued functions and image processing, which shows that is possible to carry out image processing tasks without relying on the classical finite dimensional vector-valued approach. For this reason, in Chapter 3, we decided to give a brief review of part of the theory of Banach-valued functions in the context of FVMs, and use these mathematical tools to develop some imaging methods that could be applied in an imaging application in which high dimensionality is always present: hyperspectral images.

Although the experimental results of Chapter 5 show that the simple FVM-based methods that were developed are outperformed by some classical approaches, it should be pointed out that these approaches rely on two important image processing concepts, namely, *sparsity* [4] and *total variation* [15], which we did not explore in the mathematical foundation that we presented. This may imply that if state-of-the-art algorithms are to be developed using the FVM approach, proper mathematical definitions of the notions of sparseness and bounded variation for FVMs should be formulated. In fact, in [86, 100], for-

mulations of total variation for Banach-valued functions are provided, however, to the best of our knowledge, these contributions have not yet been employed in imaging applications. As expected, these venues of research are considered as future work.

Along the lines of further future research, note that in this thesis we have employed just part of the available mathematical theory of Banach-valued functions, therefore, an obvious continuation of this work is to keep investigating in which ways these current mathematical contributions can be applied to image processing.

Also, it is worthwhile to mention that we considered only FVMs whose range is a space of either real or complex-valued functions, nevertheless, it is interesting to explore the notion of vector-FVMs, that is, FVMs whose range is a vector-valued Banach space. In this case, $u : X \subseteq A \rightarrow \mathcal{F}^n(Y)$:

$$u(x) = [u_1(x), \dots, u_n(x)]^T. \quad (11.1)$$

Observe that each u_i , $1 \leq i \leq n$, is a FVM of the form $u_i : X \subseteq A \rightarrow \mathcal{F}(Y)$. The importance of this representation is that it allows to parallelize image processing tasks. For instance, if each u_i is a HS image, a set of n HS datasets can be, say, denoised at the same time with a denoising algorithm that is based on this representation.

Furthermore, the FVM approach seems to provide a natural way of defining the SSIM for Banach-valued functions, and consequently, for vector-valued functions as well. Such a definition is still missing, however, we can conjecture that for, say, FVMs $u : X \subseteq A \rightarrow \mathcal{F}(Y)$ and $v : X \subseteq A \rightarrow \mathcal{F}(Y)$ such that $\int_X u \, dx = \int_X v \, dx = 0$, the SSIM is simply given by

$$\text{SSIM}(u, v) = \frac{\|u - v\|_2^2}{\|u\|_2^2 + \|v\|_2^2}, \quad (11.2)$$

where

$$\|u\|_2^2 = \int_X \|u(x)\|_{\mathcal{F}(Y)}^2 dx. \quad (11.3)$$

Note, however, that we still have to investigate how to define properly the statistics of FVMs to provide a formal definition of the SSIM. Also, observe that the generalization of the SSIM for FVMs provides a natural link between the first part of the thesis and SSIM-based optimization, which suggests that both FVMs and SSIM may be combined to develop imaging tools that take advantage of the benefits of these two approaches.

Regarding the second part of the thesis, a general framework for SSIM-based optimization was proposed along with several algorithms with which a variety of SSIM-based imaging tasks can be carried out. The main contribution of this approach is that it provides the means to include the SSIM as a fidelity term in a wide range of optimization

problems. Moreover, problems in which the L^1 -norm is used either as a constraint or as a regularizing term were addressed, something that, as far as we are concerned, had not been done before.

As it was seen, we employed a simplified version of the SSIM which is an example of the square of a normalized metric. This one is mathematically more tractable than its original counterpart and has nice properties such as quasiconvexity and quasiconcavity. Even though this simplified SSIM does not take into account the luminance component of the original SSIM, it was shown that in many cases is possible to recover the luminance information, as long as the given observation and the solution that is to be found are related by a linear degradation model.

On the other hand, the experimental results show that the SSIM-based methods presented in the second part of the thesis indeed perform moderately better than the classical L^2 approaches with respect to the SSIM. Visually, this translates into an enhanced contrast in some of the visual features of the recovered images. The reason is that the SSIM-based dissimilarity measure $T(x, y)$ takes into account the contrast component of the SSIM, therefore, minimization of $T(x, y)$ leads naturally to an increased contrast in the reconstructions. Although both SSIM and L^2 approaches tend to have a similar performance, this is not the case in all circumstances. For instance, in the second plot from left to right in Figure 9.1, these results suggest that the SSIM-based method outperforms significantly the L^2 approach when sparsity is increased and images are not so regular.

As for future research, given the quasiconvex properties of the SSIM as the square of a normalized metric, it would be interesting to employ different methods from the field of quasiconvex programming [49] for optimizing the SSIM. Also, observe that we have used the SSIM-based dissimilarity measure as a fidelity term, however, such a measure can be employed as a constraint whereas the fidelity term may be a convex functional. For instance, consider the following optimization problem:

$$\begin{aligned} \min_x \quad & \|x\|_1 \\ \text{subject to} \quad & T(Ax, y) \leq \lambda. \end{aligned} \tag{11.4}$$

This is a convex problem subject to a quasiconvex constraint. This approach is not new, similar problems have been considered also in [59]. In general, we consider it is worthwhile to explore the latter and other types of approaches, which may lead to algorithms more efficient than the ones we have proposed in this thesis.

Last but not least, we hope that the contributions of this thesis benefit the image processing community. We consider that these methodologies offer interesting venues of research that are worthy of further investigation.

APPENDICES

Appendix A

Proofs of Chapter 3

Theorem A.0.1. *Let $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ and let x_0 be a point in either X or its boundary $\partial(X)$. Then $\lim_{x \rightarrow x_0} u(x) = f$ if and only if for all $\varepsilon > 0$ there exists a $\delta > 0$ such that for $x \in X$ that satisfies $0 < \|x - x_0\|_X < \delta$, we have that $\|u(x) - f\|_{\mathcal{F}(Y)} < \varepsilon$.*

Proof. Let us suppose first that $\lim_{x \rightarrow x_0} u(x) = f$. Let $\varepsilon > 0$ and consider $N_f = B_\varepsilon(f)$, where

$$B_\varepsilon(f) = \{g \in \mathcal{F}(Y) : \|g - f\|_{\mathcal{F}(Y)} < \varepsilon\}.$$

By definition 3.1.2, f is eventually in $B_\varepsilon(f)$ when $x \rightarrow x_0$ in the $\|\cdot\|_X$ sense, which implies that there exists an open neighbourhood N_{x_0} of x_0 such that $u(x) \in B_\varepsilon(f)$ if $x \in (X \cap N_{x_0})$ and $x \neq x_0$. Since N_{x_0} is open and $x_0 \in N_{x_0}$, there exists $\delta > 0$ such that $B_\delta(x_0) \subset N_{x_0}$, where

$$B_\delta(x_0) = \{x \in A : \|x - x_0\|_X < \delta\}.$$

Thus, $0 < \|x - x_0\|_X < \delta$ and $x \in N_{x_0}$ imply $x \in B_\delta(x_0) \subset N_{x_0}$. Therefore, $u(x) \in B_\varepsilon(f)$, which means that $\|u(x) - f\|_{\mathcal{F}(Y)} < \varepsilon$.

Now we prove the converse. Suppose that for all $\varepsilon > 0$ there exists $\delta > 0$ such that $0 < \|x - x_0\|_X < \delta$ and $x \in X$ implies $\|u(x) - f\|_{\mathcal{F}(Y)} < \varepsilon$. Let N_f be an open neighbourhood of f . Then, there exists $\varepsilon > 0$ such that $B_\varepsilon(f) \subset N_f$. Let $N_{x_0} = B_\delta(x_0)$ be an open neighbourhood of x_0 . According to the hypothesis, $x \in (X \cap N_{x_0})$ and $x \neq x_0$ implies $\|u(x) - f\|_{\mathcal{F}(Y)} < \varepsilon$, therefore, $u(x) \in B_\varepsilon(f) \subset N_f$. \square

Theorem A.0.2. *Let $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$, $v : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$, $f, g \in \mathcal{F}(Y)$, $\alpha \in \mathbb{R}$ and $x_0 \in X$ a point in either X or its boundary $\partial(X)$. Then, the following affirmations hold:*

1. If $\lim_{x \rightarrow x_0} u(x) = f$, then $\lim_{x \rightarrow x_0} \alpha u(x) = \alpha f$, where $\alpha u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ is defined as $x \mapsto \alpha u(x)$.
2. If $\lim_{x \rightarrow x_0} u(x) = f$ and $\lim_{x \rightarrow x_0} v(x) = g$, then $\lim_{x \rightarrow x_0} (u + v)(x) = f + g$, where $(u + v) : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ is defined as $x \mapsto u(x) + v(x)$.

Proof. Regarding affirmation one, let us consider first the case when $\alpha = 0$. If this is so, then one simply has that $\lim_{x \rightarrow x_0} 0 = 0$, which is valid for any $\delta > 0$ such that $0 < \|x - x_0\|_X < \delta$. If $\alpha \neq 0$, since the limit of u is given, for a given $\hat{\varepsilon} = \frac{\varepsilon}{|\alpha|}$ there exists δ such that $0 < \|x - x_0\|_X < \delta$ implies

$$\|u(x) - f\|_{\mathcal{F}(Y)} < \frac{\varepsilon}{|\alpha|} \quad (\text{A.1})$$

$$|\alpha| \|u(x) - f\|_{\mathcal{F}(Y)} < \varepsilon \quad (\text{A.2})$$

$$\|\alpha u(x) - \alpha f\|_{\mathcal{F}(Y)} < \varepsilon \quad (\text{A.3})$$

This proves that $\lim_{x \rightarrow x_0} \alpha u(x) = \alpha f$.

As for the second affirmation, we have that

$$\|u(x) + v(x) - f - g\|_{\mathcal{F}(Y)} \leq \|u(x) - f\|_{\mathcal{F}(Y)} + \|v(x) - g\|_{\mathcal{F}(Y)}. \quad (\text{A.4})$$

Let $\varepsilon > 0$ be any given number. Then, there exists δ_1 such that $0 < \|x - x_0\|_X < \delta_1$ implies $\|u(x) - f\|_{\mathcal{F}(Y)} < \varepsilon/2$. Similarly, there exists δ_2 such that $0 < \|x - x_0\|_X < \delta_2$ implies $\|v(x) - g\|_{\mathcal{F}(Y)} < \varepsilon/2$. Let $\delta = \min\{\delta_1, \delta_2\}$. Then, $0 < \|x - x_0\|_X < \delta$ implies $\|u(x) + v(x) - f - g\|_{\mathcal{F}(Y)} < \varepsilon$, which proves that $\lim_{x \rightarrow x_0} (u + v)(x) = f + g$. \square

Theorem A.0.3. A FVM $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ is continuous at $x_0 \in X$ if and only if for all $\varepsilon > 0$ there exists $\delta > 0$ such that for $x \in X$ that satisfies $0 < \|x - x_0\|_X < \delta$, we have that $\|u(x) - u(x_0)\|_{\mathcal{F}(Y)} < \varepsilon$.

Proof. Suppose first that u is continuous at x_0 . Since $\lim_{x \rightarrow x_0} u(x) = u(x_0)$, by theorem 3.1.1 we have that

$$\forall \varepsilon > 0, \exists \delta > 0 : x \in X \wedge 0 < \|x - x_0\|_X \implies \|u(x) - u(x_0)\|_{\mathcal{F}(Y)} < \varepsilon. \quad (\text{A.5})$$

To prove the converse assume that (A.5) holds. Then, by theorem 3.1.1 we get that $\lim_{x \rightarrow x_0} u(x) = u(x_0)$ and the result follows. \square

Theorem A.0.4. Let $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ and $v : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ be differentiable, and suppose that $(uv)(x) \in \mathcal{F}(Y)$ for all $x \in X$. Also, let $f \in \mathcal{F}(Y)$, $\alpha, \beta \in \mathbb{R}$ and $x \in X$. Then, the following assertions are true:

1. **Derivative of a Constant:** If $u(x) = f$, then $Du(x) = 0$ for all $x \in X$.
2. **Sum Rule:** If $w(x) = \alpha u(x) + \beta v(x)$, then $Dw(x) = \alpha Du(x) + \beta Dv(x)$.
3. **Product Rule:** If $w(x) = (uv)(x)$, then $Dw(x) = Du(x)v(x) + u(x)Dv(x)$.

Proof. The first two assertions follow easily from definition 3.1.5. Regarding the third assertion, if g is a differentiable FVM, then, $g(x+h) = g(x) + Dg(x)h + o(h)$. Therefore,

$$\begin{aligned}
w(x+h) &= u(x+h)v(x+h) \\
&= (u(x) + Du(x)h + o(h))(v(x) + Dv(x)h + o(h)) \\
&= u(x)v(x) + v(x)Du(x)h + u(x)Dv(x)h + (Du(x)h)(Dv(x)h) \\
&\quad + o(h)(u(x) + v(x) + Du(x)h + Dv(x)h + o(h))
\end{aligned}$$

By using this fact, and employing definition 3.1.5, it is easily concluded that

$$Dw(x) = Du(x)v(x) + u(x)Dv(x). \quad (\text{A.6})$$

□

Theorem A.0.5. Let $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ and $v : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ be Gâteaux differentiable, and suppose that $(uv)(x) \in \mathcal{F}(Y)$ for all $x \in X$. Also, let $f \in \mathcal{F}(Y)$, $\alpha, \beta \in \mathbb{R}$ and $x \in X$. Then, the following assertions are true:

1. **Derivative of a Constant:** Let $u(x) = f$, where, $f \in \mathcal{F}(Y)$. Then, $Du(x; h) = 0$.
2. **Sum Rule:** If $w(x) = (u+v)(x)$, then $D(u+v)(x; h) = Du(x; h) + Dv(x; h)$.
3. **Product Rule:** If $w(x) = (uv)(x)$, then $D(uv)(x; h) = Du(x; h)v(x) + u(x)Dv(x; h)$.

Proof. Assertions one and two follow easily from the definition of directional derivative. As for the third assertion, we have that if a FVM g is Gâteaux differentiable, then $g(x+\varepsilon h) = g(x) + \varepsilon Dg(x; h) + o(\varepsilon)$. Then,

$$\begin{aligned}
w(x+\varepsilon h) &= u(x+\varepsilon h)v(x+\varepsilon h) \\
&= (u(x) + \varepsilon Du(x; h) + o(\varepsilon))(v(x) + \varepsilon Dv(x; h) + o(\varepsilon)) \\
&= u(x)v(x) + \varepsilon Du(x; h)v(x) + \varepsilon u(x)Dv(x; h) + \varepsilon^2 Du(x; h)Dv(x; h) \\
&\quad + o(\varepsilon)(u(x) + v(x) + \varepsilon Du(x; h) + \varepsilon Dv(x; h) + o(\varepsilon))
\end{aligned}$$

By employing the latter equation, and using the definition of directional derivative, it follows that

$$Dw(x; h) = Du(x; h)v(x) + u(x)Dv(x; h). \quad (\text{A.7})$$

□

Theorem A.0.6. *Let $u : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ be an integrable FVM. Also, let $D, E \in \Sigma$ and $D \cap E = \emptyset$. Then, we have that*

$$\int_{D \cup E} u \, dx = \int_D u \, dx + \int_E u \, dx. \quad (\text{A.8})$$

Proof. Since u is integrable, we have that

$$\int_{D \cup E} u \, dx = \lim_{n \rightarrow \infty} \int_{D \cup E} \varphi_n \, dx \quad (\text{A.9})$$

$$= \lim_{n \rightarrow \infty} \int \varphi_n \chi_{D \cup E} \, dx. \quad (\text{A.10})$$

Thanks to the properties of the indicator function, the latter equation is equivalent to

$$\int_{D \cup E} u \, dx = \lim_{n \rightarrow \infty} \int \varphi_n (\chi_D + \chi_E) \, dx \quad (\text{A.11})$$

$$= \lim_{n \rightarrow \infty} \int \varphi_n \chi_D \, dx + \lim_{n \rightarrow \infty} \int \varphi_n \chi_E \, dx \quad (\text{A.12})$$

$$= \lim_{n \rightarrow \infty} \int_D \varphi_n \, dx + \lim_{n \rightarrow \infty} \int_E \varphi_n \, dx \quad (\text{A.13})$$

$$= \int_D u \, dx + \int_E u \, dx. \quad (\text{A.14})$$

This completes the proof. □

Theorem A.0.7. *Let $u, v \in \mathcal{M}_{\mathcal{F}}$ such that $u, v \geq 0$ for all $x \in X$ and all $y \in Y$. Then the following affirmations hold:*

1. *If $\mu(Z) = 0$, then $\int_Z u \, dx = 0$.*
2. *If $u = v$ almost everywhere on X , then $\int_X u \, dx = \int_X v \, dx$.*

Proof. Regarding the first affirmation, we have that

$$\int_Z u \, dx \leq \int_Z \sup_{x \in X} \{\|u(x)\|_{\mathcal{F}(Y)}\} \, dx \quad (\text{A.15})$$

$$= \sup_{x \in X} \{\|u(x)\|_{\mathcal{F}(Y)}\} \mu(Z). \quad (\text{A.16})$$

Given that $\mu(Z) = 0$, we have that $\int_Z u \, dx = 0$.

As for the second affirmation, let $Z \subset X$ be the set on which $u \neq v$. Then, we have that

$$\int_X u \, dx = \int_{(X \setminus Z) \cup Z} u \, dx \quad (\text{A.17})$$

$$= \int_{X \setminus Z} u \, dx + \int_Z u \, dx \quad (\text{A.18})$$

$$= \int_{X \setminus Z} v \, dx \quad (\text{A.19})$$

$$= \int_{X \setminus Z} v \, dx + \int_Z v \, dx \quad (\text{A.20})$$

$$= \int_X v \, dx. \quad (\text{A.21})$$

Thus, $\int_X u \, dx = \int_X v \, dx$. □

Theorem A.0.8. *Let $u : X \subset \mathbb{R}^n \rightarrow L^1(Y)$ be an integrable FVM. Then, we have that*

$$\int_Y \left(\int_X u \, dx \right) dy = \int_X \left(\int_Y u \, dy \right) dx. \quad (\text{A.22})$$

Proof. First, let us define the following operator $T : L^1(Y) \rightarrow \mathbb{R}$:

$$Tf := \int_Y f \, dy. \quad (\text{A.23})$$

Clearly, \mathbb{R} is a Banach space under the absolute value as a norm. Therefore, one has that

$$|Tf| = \left| \int_Y f \, dy \right| \quad (\text{A.24})$$

$$\leq \int_Y |f| \, dy \quad (\text{A.25})$$

$$= \|f\|_1. \quad (\text{A.26})$$

Thus, $|T| \leq 1$; that is, T is a bounded operator.

Now, since u is integrable, we have that $\int_X u \, dx$ is measurable and is in $L^1(Y)$. Moreover, by theorem 3.1.12, we have that Tu is integrable and that the following equality holds:

$$T \left(\int_X u \, dx \right) = \int_X Tu \, dx. \quad (\text{A.27})$$

Therefore,

$$\int_Y \left(\int_X u \, dx \right) dy = \int_X \left(\int_Y u \, dy \right) dx. \quad (\text{A.28})$$

□

Theorem A.0.9. *Let (X, Σ, μ) and (Y, T, ν) be finite measure spaces, where $X = [a, b] \in \mathbb{R}$. Also, let $u : X \subset \mathbb{R} \rightarrow \mathcal{F}(Y)$ be a continuous FVM. Then, the FVM U defined as*

$$U(x) = \int_a^x u(z) dz \quad (\text{A.29})$$

is differentiable on (a, b) and $U'(x) = u(x)$ for ν -almost every $y \in Y$.

Proof. For some $h > 0$ we have that

$$\frac{\|U(x+h) - U(x) - DU(x)h\|_{\mathcal{F}(Y)}}{|h|} = \frac{\|\int_x^{x+h} u(z) dz - DU(x)h\|_{\mathcal{F}(Y)}}{|h|} \quad (\text{A.30})$$

$$= \frac{\|\int_x^{x+h} (u(z) - DU(x)) dz\|_{\mathcal{F}(Y)}}{|h|}. \quad (\text{A.31})$$

For Bochner integrals the inequality $\|\int_X u \, dx\| \leq \int_X \|u\| \, dx$ holds for any norm $\|\cdot\|$; therefore,

$$\frac{\|U(x+h) - U(x) - DU(x)h\|_{\mathcal{F}(Y)}}{|h|} \leq \frac{1}{|h|} \int_x^{x+h} \|u(z) - DU(x)\|_{\mathcal{F}(Y)} dz. \quad (\text{A.32})$$

Let $v : X \subset \mathbb{R} \rightarrow \mathbb{R}$ be defined as $v(z) = \|u(z) - DU(x)\|_{\mathcal{F}(Y)}$. We claim that $v(z)$ is a continuous function. This can be seen by noticing the fact that for any given $\varepsilon > 0$ and $z, s \in X$ there exists $\delta > 0$ that satisfies $0 < |z - s| < \delta$ such that $|v(z) - v(s)| < \varepsilon$. This is so since $|v(z) - v(s)|$ can be bounded by $\|u(z) - u(s)\|_{\mathcal{F}(Y)} < \varepsilon$. By choosing $\varepsilon = \delta$ such that the latter assertion holds, it is concluded that $v(z)$ is continuous. Therefore, by the mean value theorem for integrals we have that

$$\frac{1}{|h|} \int_x^{x+h} \|u(z) - DU(x)\|_{\mathcal{F}(Y)} dz = \|u(t) - DU(x)\|_{\mathcal{F}(Y)} \quad (\text{A.33})$$

for some $t \in [x, x + h]$. By taking the limit of h as it tends to zero at both sides of the latter equation, and given that u is continuous on $[a, b]$, we obtain that

$$\lim_{h \rightarrow 0} \frac{1}{|h|} \int_x^{x+h} \|u(z) - DU(x)\|_{\mathcal{F}(Y)} dz = \|u(x) - DU(x)\|_{\mathcal{F}(Y)}. \quad (\text{A.34})$$

The previous expression can be equal to zero if and only if $DU(x)(y) = u(x)(y)$ for ν -almost every $y \in Y$, thus, since x was arbitrary, U is differentiable on (a, b) and its derivative at any $x \in (a, b)$ is $DU(x) = u(x)$. \square

Corollary A.0.9.1. *If $u : X \subset \mathbb{R} \rightarrow \mathcal{F}(Y)$ is a continuous FVM, then*

$$\lim_{h \rightarrow 0} \frac{1}{h} \int_x^{x+h} u(z) dz = u(x) \quad (\text{A.35})$$

for ν -almost all $y \in Y$ and all $x \in (a, b) \subset X$.

Proof. First, notice that for some $h > 0$ Eq. (A.30) can be rewritten as

$$\left\| \frac{1}{h} \int_x^{x+h} u(z) dz - DU(x) \right\|_{\mathcal{F}(Y)}. \quad (\text{A.36})$$

Since u is continuous, the function $U(x) = \int_a^x u(z) dz$ is differentiable and its derivative at $x \in X$ is given by $DU(x) = u(x)$. Thus, we have that for every $x \in (a, b)$

$$\lim_{h \rightarrow 0} \left\| \frac{1}{h} \int_x^{x+h} u(z) dz - u(x) \right\|_{\mathcal{F}(Y)} = \left\| \lim_{h \rightarrow 0} \frac{1}{h} \int_x^{x+h} u(z) dz - u(x) \right\|_{\mathcal{F}(Y)} = 0. \quad (\text{A.37})$$

This completes the proof. \square

Theorem A.0.10. *The functionals $\|u\|_p : L^p(X; \mathcal{F}(Y)) \rightarrow \mathbb{R}$ and $\|u\|_\infty : L^\infty(X; \mathcal{F}(Y)) \rightarrow \mathbb{R}$ are norms.*

Proof. For $1 \leq p \leq \infty$, it is easy to see that $\|u\|_p = 0$ if and only if $u = 0$. Also, absolute homogeneity is clearly fulfilled; that is, $\|\alpha u\|_p = |\alpha| \|u\|_p$ for all $\alpha \in \mathbb{R}$. Verifying that the triangle inequality also holds for $\|\cdot\|_p$ requires some work. Let $1 \leq p < \infty$ and $u, v \in L^p(X; \mathcal{F}(Y))$. Then, we have that

$$\|u + v\|_p = \left(\int_X \|u(x) + v(x)\|_{\mathcal{F}(Y)}^p dx \right)^{\frac{1}{p}}. \quad (\text{A.38})$$

By the property of subadditivity of $\|\cdot\|_{\mathcal{F}(Y)}$ we obtain that

$$\|u + v\|_p \leq \left(\int_X (\|u(x)\|_{\mathcal{F}(Y)} + \|v(x)\|_{\mathcal{F}(Y)})^p dx \right)^{\frac{1}{p}}. \quad (\text{A.39})$$

Finally, thanks to the Minkowski inequality, it follows that

$$\|u + v\|_p \leq \left(\int_X \|u(x)\|_{\mathcal{F}(Y)}^p dx \right)^{\frac{1}{p}} + \left(\int_X \|v(x)\|_{\mathcal{F}(Y)}^p dx \right)^{\frac{1}{p}} \quad (\text{A.40})$$

$$= \|u\|_p + \|v\|_p. \quad (\text{A.41})$$

When $p = \infty$, subadditivity of $\|\cdot\|_{\infty}$ follows thanks to the subadditivity of both $\|\cdot\|_{\mathcal{F}(Y)}$ and the essential supremum:

$$\|u + v\|_{\infty} = \operatorname{ess\,sup}_{x \in X} \{\|u(x) + v(x)\|_{\mathcal{F}(Y)}\} \quad (\text{A.42})$$

$$\leq \operatorname{ess\,sup}_{x \in X} \{\|u(x)\|_{\mathcal{F}(Y)} + \|v(x)\|_{\mathcal{F}(Y)}\} \quad (\text{A.43})$$

$$\leq \operatorname{ess\,sup}_{x \in X} \{\|u(x)\|_{\mathcal{F}(Y)}\} + \operatorname{ess\,sup}_{x \in X} \{\|v(x)\|_{\mathcal{F}(Y)}\} \quad (\text{A.44})$$

$$= \|u\|_{\infty} + \|v\|_{\infty}. \quad (\text{A.45})$$

This completes the proof. \square

Theorem A.0.11. *Let (X, Σ, μ) and (Y, T, ν) be finite measure spaces and let $1 \leq p \leq \infty$. If $\mathcal{F}(Y)$ is a Banach space, then $L^p(X; \mathcal{F}(Y))$ is also a Banach space.*

Proof. Proving completeness of $L^p(X; \mathcal{F}(Y))$ is analogous to the proof of completeness of L^p spaces. Let $1 \leq p < \infty$. Also, let $\sum_k u_k$ be a series such that every $u_k \in L^p(X; \mathcal{F}(Y))$ for all $k \in \mathbb{N}$ and

$$\sum_{k=1}^{\infty} \|u_k\|_p < \infty. \quad (\text{A.46})$$

We must show that such a series converges to an element of $L^p(X; \mathcal{F}(Y))$ with respect to $\|\cdot\|_p$. Let v_n and w_n be defined as

$$v_n(x) := \sum_{k=1}^n u_k(x) \quad \text{and} \quad w_n(x) := \sum_{k=1}^n \|u_k(x)\|_{\mathcal{F}(Y)} \chi_Y(y). \quad (\text{A.47})$$

Observe that

$$\lim_{n \rightarrow \infty} \|w_n\|_p \leq \sum_{k=1}^{\infty} \nu(Y)^{\frac{1}{p}} \|u_k\|_p < \infty, \quad (\text{A.48})$$

thus, $w = \lim_{n \rightarrow \infty} w_n$ is in $L^p(X; \mathcal{F}(Y))$ and the function $\|w(x)\|_{\mathcal{F}(Y)}$ is integrable. Also, for all $n \in \mathbb{N}$ and μ -almost every $x \in X$,

$$\|v_n(x)\|_{\mathcal{F}(Y)} \leq \|w(x)\|_{\mathcal{F}(Y)}, \quad (\text{A.49})$$

which implies that v is in $L^p(X; \mathcal{F}(Y))$ as well. On the other hand, since $\mathcal{F}(Y)$ is complete, we have that for μ -almost all $x \in X$

$$\lim_{n \rightarrow \infty} \|v(x) - v_n(x)\|_{\mathcal{F}(Y)}^p = 0. \quad (\text{A.50})$$

By the dominated convergence theorem we conclude that

$$\lim_{n \rightarrow \infty} \int_X v_n \, dx = \int_X v \, dx \quad (\text{A.51})$$

with respect to the L^p norm; that is,

$$\lim_{n \rightarrow \infty} \int_X \|v(x) - v_n(x)\|_{\mathcal{F}(Y)}^p \, dx = 0. \quad (\text{A.52})$$

As for the case in which $p = \infty$, let $\{u_k\}$ be a Cauchy sequence in $L^\infty(X; \mathcal{F}(Y))$. Then, there exists $N(\varepsilon) \in \mathbb{N}$ such that for all $m, n > N$ one has that

$$\|u_m - u_n\|_p < \varepsilon. \quad (\text{A.53})$$

Then, there exists a set Z of measure zero such that $\|u_m(x) - u_n(x)\|_{\mathcal{F}(Y)} < \delta$ for all $x \in X \setminus Z$ and all $n, m > N(\delta)$. Moreover, there exists a constant K such that $\|u_n(x)\|_{\mathcal{F}(Y)} < K$ for all $x \in X \setminus Z$. Z is independent of n, m and K , therefore, the sequence $\{u_n\}$ converges uniformly to an element of $L^\infty(X; \mathcal{F}(Y))$. This completes the proof. \square

Theorem A.0.12. *If $\mathcal{F}(Y)$ is a Hilbert space with scalar product $\langle \cdot, \cdot \rangle_{\mathcal{F}(Y)}$, then $L^2(X; \mathcal{F}(Y))$ is also a Hilbert space with scalar product defined as*

$$\langle u, v \rangle = \int_X \langle u(x), v(x) \rangle_{\mathcal{F}(Y)} \, dx, \quad (\text{A.54})$$

for all $u, v \in L^2(X; \mathcal{F}(Y))$.

Proof. It can be seen that $\langle \cdot, \cdot \rangle$ fulfills the properties of an inner product, namely, conjugate simetry, linearity and positive definiteness. Also, $\langle u, u \rangle = \|u\|_2^2$; that is, the L^2 norm is induced by the scalar product defined in Eq. (3.53). Completeness of $L^2(X; \mathcal{F}(Y))$ follows from the completeness of the L^p spaces. \square

Theorem A.0.13. [126] *The space $C(X; \mathcal{F}(Y))$ is dense in $L^p(X; \mathcal{F}(Y))$. Even more, the embedding $C(X; \mathcal{F}(Y)) \subseteq L^p(X; \mathcal{F}(Y))$ is continuous.*

Proof. X -step mappings can be arbitrarily closely approximated by continuous FVMs. Since the set of X -step mappings is dense in $L^p(X; \mathcal{F}(Y))$, we have that $C(X; \mathcal{F}(Y))$ is also dense in $L^p(X; \mathcal{F}(Y))$.

As for the embedding, we have the following inequality:

$$\|u\|_p \leq \mu(X)^{\frac{1}{p}} \sup_{x \in X} \{\|u(x)\|_{\mathcal{F}(Y)}\}, \quad (\text{A.55})$$

which implies the continuity of the embedding $C(X; \mathcal{F}(Y)) \subseteq L^p(X; \mathcal{F}(Y))$. \square

Theorem A.0.14. [126] *Let (X, Σ, ν) and (Y, T, ν) be finite measure spaces. If the embedding $\mathcal{F}(Y) \subseteq \mathcal{G}(Y)$ is continuous and $1 \leq q \leq p \leq \infty$, then the embedding $L^p(X; \mathcal{F}(Y)) \subseteq L^q(X; \mathcal{G}(Y))$, is also continuous.*

Proof. Since the embedding $\mathcal{F}(Y) \subseteq \mathcal{G}(Y)$ is continuous, we have that $\|u(x)\|_{\mathcal{G}(Y)} \leq c\|u(x)\|_{\mathcal{F}(Y)}$, where C is a finite positive constant that depends on $\nu(Y) < \infty$. Thus,

$$\left(\int_X \|u(x)\|_{\mathcal{G}(Y)}^q dx \right)^{\frac{1}{q}} \leq C^{\frac{1}{q}} \left(\int_X \|u(x)\|_{\mathcal{F}(Y)}^q dx \right)^{\frac{1}{q}}. \quad (\text{A.56})$$

By using the classical Hölder inequality we obtain that

$$\left(\int_X \|u(x)\|_{\mathcal{G}(Y)}^q dx \right)^{\frac{1}{q}} \leq C^{\frac{1}{q}} \mu(X)^{\frac{1}{r}} \left(\int_X \|u(x)\|_{\mathcal{F}(Y)}^p dx \right)^{\frac{1}{p}}, \quad (\text{A.57})$$

where $\frac{1}{r} + \frac{1}{s} = 1$ and $q \leq p = qs$. This completes the proof. \square

Theorem A.0.15 (Hölder Inequality). [126] *Let $u \in L^p(X; \mathcal{F}(Y))$ and $v \in L^q(X; \mathcal{F}(Y)^*)$, where $\mathcal{F}(Y)$ is the dual space of $\mathcal{F}(Y)$. Also, let $1 < p < \infty$ and $\frac{1}{p} + \frac{1}{q} = 1$. Then, the following Hölder inequality holds for all the elements of both L^p and L^q :*

$$\int_X |\langle u(x), v(x) \rangle_{\mathcal{F}(Y)}| dx \leq \left(\int_X \|u(x)\|_{\mathcal{F}(Y)}^p dx \right)^{\frac{1}{p}} \left(\int_X \|v(x)\|_{\mathcal{F}(Y)^*}^q dx \right)^{\frac{1}{q}}. \quad (\text{A.58})$$

In particular, all the integrals of this inequality exist.

Proof. We have that u and v are measurable on X , then, there exist sequences of X -simple mappings $\{u_n\} : X \rightarrow \mathcal{F}(Y)$ and $\{v_n\} : X \rightarrow \mathcal{F}(Y)^*$ such that $u_n \rightarrow u$ and $v_n \rightarrow v$ as $n \rightarrow \infty$ for μ -almost all $x \in X$. This implies that $\langle u_n(x), v_n(x) \rangle_{\mathcal{F}(Y)} \rightarrow \langle u(x), v(x) \rangle_{\mathcal{F}(Y)}$ as $n \rightarrow \infty$ for μ -almost all $x \in X$. Therefore, the real-valued function $\langle u(x), v(x) \rangle_{\mathcal{F}(Y)} : X \rightarrow \mathbb{R}$ is measurable.

On the other hand, we have that $|\langle u(x), v(x) \rangle_{\mathcal{F}(Y)}| \leq \|u(x)\|_{\mathcal{F}(Y)} \|v(x)\|_{\mathcal{F}(Y)^*}$. Thus,

$$\int_X |\langle u(x), v(x) \rangle_{\mathcal{F}(Y)}| dx \leq \int_X \|u(x)\|_{\mathcal{F}(Y)} \|v(x)\|_{\mathcal{F}(Y)^*} dx. \quad (\text{A.59})$$

The assertion follows by applying the classical Hölder inequality to the above inequality. \square

Theorem A.0.16. *Let (X, Σ, μ) and (Y, T, ν) be finite measure spaces. Also, let $u : X \subseteq \mathbb{R}^n \rightarrow \mathcal{F}(Y)$, $\frac{\partial u}{\partial x_i} : X \subset \mathbb{R}^n \rightarrow \mathcal{G}(Y)$, and assume that the function $\Phi(x) := f(x, u(x), \nabla_x u(x))$ is integrable over X . In addition, suppose that the Fréchet derivatives of $f : X \times \mathcal{F}(Y) \times \mathcal{G}^n(Y) \rightarrow \mathbb{R}$ with respect to all of its arguments are continuous. Define the functional $I(u) : Z(\mathcal{F}(Y), \mathcal{G}(Y)) \rightarrow \mathbb{R}$ as follows:*

$$I(u) := \int_X f(x, u, \nabla_x u) dx, \quad (\text{A.60})$$

where $Z(\mathcal{F}(Y), \mathcal{G}(Y))$ is a Banach space of FVMs that depends on the function spaces $\mathcal{F}(Y)$ and $\mathcal{G}(Y)$. If $u_0 : X \subset \mathbb{R}^n \rightarrow \mathcal{F}(Y)$ is a stationary point of $I(u)$, u_0 is the solution of the equation

$$\frac{\partial f}{\partial u}(u_0) - \nabla \cdot \frac{\partial f}{\partial \nabla_x u}(\nabla_x u_0) = 0. \quad (\text{A.61})$$

where $\frac{\partial f}{\partial u} \in \mathcal{F}(Y)^*$ and $\frac{\partial f}{\partial \nabla_x u} \in \mathcal{G}^n(Y)^*$ are the Fréchet derivatives of f with respect to u and $\nabla_x u$ respectively, $\nabla \cdot$ is the classical divergence operator, and $\mathcal{F}(Y)^*$ and $\mathcal{G}(Y)^*$ are the dual spaces of $\mathcal{F}(Y)$ and $\mathcal{G}(Y)$ respectively.

Proof. Let $\varphi \in C_0^\infty(X)$. As usual, since u_0 is a stationary point, we have that the Gâteaux derivative of $I(u)$ at u_0 in the direction of $\varphi \chi_Y$ is zero, which we denote as $DI(u_0; \varphi \chi_Y)$:

$$DI(u_0; \varphi \chi_Y) = \lim_{\varepsilon \downarrow 0} \frac{1}{\varepsilon} \int_X f(x, u_0 + \varepsilon \varphi \chi_Y, \nabla_x u_0 + \varepsilon \nabla_x(\varphi \chi_Y)) - f(x, u_0, \nabla_x u_0) dx. \quad (\text{A.62})$$

By hypothesis, since the Fréchet derivatives of the integrand are continuous, we have that the integrand is continuous as well; therefore, we can exchange the limit and the integral:

$$DI(u_0; \varphi \chi_Y) = \int_X \lim_{\varepsilon \downarrow 0} \frac{f(x, u_0 + \varepsilon \varphi \chi_Y, \nabla_x u_0 + \varepsilon \nabla_x(\varphi \chi_Y)) - f(x, u_0, \nabla_x u_0)}{\varepsilon} dx. \quad (\text{A.63})$$

Given that the integrand of the previous equation is bounded, we can apply the Lebesgue dominated convergence theorem to obtain

$$DI(u_0; \varphi\chi_Y) = \int_X Df((x, u_0, \nabla_x u_0); (0, \varphi\chi_Y, \nabla_x(\varphi\chi_Y))) dx. \quad (\text{A.64})$$

We have, then, that the integrand of the latter equation is the directional derivative of f at the point $(x, u_0, \nabla_x u_0)$ in the direction of $(0, \varphi\chi_Y, \nabla_x(\varphi\chi_Y))$. Such a derivative is given by

$$Df((x, u_0, \nabla_x u_0); (0, \varphi\chi_Y, \nabla_x(\varphi\chi_Y))) = (\varphi\chi_Y) \frac{\partial f}{\partial u}(u_0) + \nabla_x(\varphi\chi_Y) \cdot \frac{\partial f}{\partial \nabla_x u}(\nabla_x u_0). \quad (\text{A.65})$$

Notice that in this case the partial derivatives $\frac{\partial f}{\partial u}$ and $\frac{\partial f}{\partial \nabla_x u}$ are Fréchet derivatives. Since u and $\nabla_x u$ take values on $\mathcal{F}(Y)$ and $\mathcal{G}^n(Y)$ respectively, we have that $\frac{\partial f}{\partial u} \in \mathcal{F}(Y)^*$ and $\frac{\partial f}{\partial \nabla_x u} \in \mathcal{G}^n(Y)^*$. Substituting Eq. (A.65) into Eq. (A.64) we obtain that

$$DI(u_0; \varphi\chi_Y) = \int_X \left((\varphi\chi_Y) \frac{\partial f}{\partial u}(u_0) + \nabla_x(\varphi\chi_Y) \cdot \frac{\partial f}{\partial \nabla_x u}(\nabla_x u_0) \right) dx = 0. \quad (\text{A.66})$$

Observe that χ_Y does not depend on x , thus the latter equation is equivalent to

$$DI(u_0; \varphi\chi_Y) = \int_X \left(\varphi \frac{\partial f}{\partial u}(u_0) + \nabla \varphi \cdot \frac{\partial f}{\partial \nabla_x u}(\nabla_x u_0) \right) dx = 0. \quad (\text{A.67})$$

By using the product rule of the divergence operator we obtain that

$$DI(u_0; \varphi\chi_Y) = \int_X \varphi \left(\frac{\partial f}{\partial u}(u_0) + \nabla \cdot \left(\varphi \frac{\partial f}{\partial \nabla_x u}(\nabla_x u_0) \right) - \varphi \nabla \cdot \frac{\partial f}{\partial \nabla_x u}(\nabla_x u_0) \right) dx = 0. \quad (\text{A.68})$$

Given that $\varphi(\partial(X)) = 0$, and using the divergence theorem, we have that

$$DI(u_0; \varphi\chi_Y) = \int_X \varphi \left(\frac{\partial f}{\partial u}(u_0) - \nabla \cdot \frac{\partial f}{\partial \nabla_x u}(\nabla_x u_0) \right) dx = 0. \quad (\text{A.69})$$

Finally, by the fundamental lemma of calculus of variations, we obtain that

$$\frac{\partial f}{\partial u}(u_0) - \nabla \cdot \frac{\partial f}{\partial \nabla_x u}(\nabla_x u_0) = 0 \quad (\text{A.70})$$

for μ -almost all $x \in X$. □

Appendix B

Proofs of Chapter 8

Theorem B.0.17. *Let $f : X \subset \mathbb{R}^n \rightarrow \mathbb{R}$ be defined as in Eq. (9.30). Then, its Jacobian is Lipschitz continuous on any open convex set $\Omega \subset X$; that is, there exists a constant $L > 0$ such that for any $x, z \in \Omega$,*

$$\|J_f(x) - J_f(z)\|_F \leq L\|x - z\|_2. \quad (\text{B.1})$$

Here, $\|\cdot\|_F$ denotes the Frobenius norm and

$$L = C_1\|D^T D\|_F + \lambda C_2\|A^T A\|_F, \quad C_1, C_2 > 0. \quad (\text{B.2})$$

Proof. Without loss of generality, and for the sake of simplicity, let the stability constant C of the SSIM be zero. Also, let y be a non-zero vector in \mathbb{R}^m . Let us define

$$s(x) := \frac{2x^T y}{\|Dx\|_2^2 + \|y\|_2^2}, \quad (\text{B.3})$$

and

$$r(x) := \|Dx\|_2^2 + \|y\|_2^2. \quad (\text{B.4})$$

Therefore, we have that $\|J_f(x) - J_f(z)\|_F$ is bounded by

$$\begin{aligned} \|J_f(x) - J_f(z)\|_F &\leq \|D^T D\|_F \|x \nabla s(x)^T - z \nabla s(z)^T\|_F + \\ &\quad \lambda \|A^T A\|_F \|x \nabla r(x)^T - z \nabla r(z)^T\|_F + |s(x) - s(z)| \|D^T D\|_F + \\ &\quad \lambda |r(x) - r(z)| \|A^T A\|_F, \end{aligned} \quad (\text{B.5})$$

To show that J_f is Lipschitz continuous on Ω , we have to show that the terms that multiply the Frobenius norms of both $D^T D$ and $A^T A$ are Lipschitz continuous on Ω as well. Let us begin with the term $|r(x) - r(z)|$. This is equal to

$$|r(x) - r(z)| = \left| \|Dx\|_2^2 - \|Dz\|_2^2 \right| \quad (\text{B.6})$$

$$= (\|Dx\|_2 + \|Dz\|_2) \left| \|Dx\|_2 - \|Dz\|_2 \right|, \quad (\text{B.7})$$

which is bounded above by

$$|r(x) - r(z)| \leq (\|Dx\|_2 + \|Dz\|_2) \|Dx - Dz\|_2 \quad (\text{B.8})$$

$$\leq \|D\|_2^2 (\|x\|_2 + \|z\|_2) \|x - z\|_2 \quad (\text{B.9})$$

$$\leq K_1 \|D\|_2^2 \|x - z\|_2, \quad (\text{B.10})$$

where $K_1 = \max_{x,z \in \Omega} \{\|x\|_2, \|z\|_2\}$.

As for $|s(x) - s(z)|$, by using the mean-value theorem for functions of several variables we have that

$$|s(x) - s(z)| \leq \|\nabla s(\alpha x + (1 - \alpha)z)\|_2 \|x - z\|_2 \quad (\text{B.11})$$

for some $\alpha \in [0, 1]$ and all $x, z \in \Omega$. In fact, it can be shown that for any vector $w \in \mathbb{R}^n$, the norm of the gradient of s is bounded by

$$\|\nabla s(w)\| \leq (\sqrt{2} + 1) \frac{\|D\|_2}{\|y\|_2}. \quad (\text{B.12})$$

Let $K_2 = (\sqrt{2} + 1) \frac{\|D\|_2}{\|y\|_2}$. Thus, $|s(x) - s(z)| \leq K_2 \|x - z\|_2$.

Regarding the term $\|x \nabla s(x)^T - z \nabla s(z)^T\|_F$, we have that the ij -th each entry of the $n \times n$ matrix $x \nabla s(x)^T - z \nabla s(z)^T$ is given by

$$\nabla_j s(x) x_i - \nabla_j s(z) z_i, \quad (\text{B.13})$$

where $\nabla_j s(\cdot)$ is the j -th component of the gradient of $s(\cdot)$. By employing the mean value theorem for functions of one variable we obtain that

$$|\nabla_j s(x) x_i - \nabla_j s(z) z_i| = \left| \frac{\partial}{\partial x_i} (\nabla_j s(x(v))) \right| |x_i - z_i|, \quad (\text{B.14})$$

for some $v \in \mathbb{R}$. Here, $x(v) = [x_1, \dots, x_{i-1}, v, \dots, x_n]$. The partial derivative of the previous equation is bounded, which can be proved using the classical triangle inequality and differential calculus. Given this, we have that

$$\left| \frac{\partial}{\partial x_i} (\nabla_j s(x))(v) \right| \leq (\sqrt{2} + 3) \frac{\|D_i^T\|_2 \|D_j^T\|_2}{\|y\|_2^2} + (2\sqrt{3} + 2) \frac{\|D_j^T\|_2}{\|y\|_2^3} \quad (\text{B.15})$$

$$= K_{ij}, \quad (\text{B.16})$$

where D_k^T is the k -th row of the the transpose of the matrix D . Therefore,

$$|\nabla_j s(x)x_i - \nabla_j s(z)z_i| \leq K_{ij}|x_i - z_i|. \quad (\text{B.17})$$

Using this result, we can conclude that

$$\|x\nabla s(x)^T - z\nabla s(z)^T\|_F \leq K_3\|x - z\|_2, \quad (\text{B.18})$$

where K_3 is equal to

$$K_3 = n \max_{1 \leq i, j \leq n} K_{ij}; \quad (\text{B.19})$$

that is, K_3 is equal to the largest K_{ij} times n . In a similar manner, it can be shown that

$$\|x\nabla r(x)^T - z\nabla r(z)^T\|_F \leq K_4\|x - z\|_2, \quad (\text{B.20})$$

where K_4 is given by

$$K_4 = \max_{1 \leq i, j \leq n} \{2nK_1\|D_j^T\|_2(\|D_i^T\|_2 + \|D\|_2)\}. \quad (\text{B.21})$$

Finally, we obtain that

$$\|J_f(x) - J_f(z)\|_F \leq [(K_2 + K_3)\|D^T D\|_F + \lambda(K_1 + K_4)\|A^T A\|_F]\|x - z\|_2, \quad (\text{B.22})$$

which completes the proof. \square

Theorem B.0.18. *Let $g_{\pm}(x_i) := a_i(x_i) \mp \tau_i(x_i)$, $y \in \mathbb{R}^m$ be a given observation, $x_0 \in \mathbb{R}^n$ an initial condition, and D an $m \times n$ matrix whose columns have L^2 norm equal to one. Then, $g_{\pm}(x_i)$ has a fixed point provided that the regularization parameter λ fulfills the following condition:*

$$\lambda < \min_{x \in \Omega} \left\{ \frac{(SSIM(Dx, y))^2 - \delta\|y\|_2}{2\|D\|_2\|x\|_2} \right\}, \quad (\text{B.23})$$

where δ is a small positive constant, and $\Omega \subset \mathbb{R}^n$ is the set of vectors that lie on the path followed by an algorithm that solves problem (9.39) starting at x_0 .

Proof. The proof amounts to show that g_{\pm} is a contraction provided that the regularization parameter λ is less than a certain value. Consider the following expression:

$$|g_{\pm}(x_i) - g_{\pm}(z_i)| = |g'_{\pm}(t)||x_i - z_i| \quad (\text{B.24})$$

for some $t \in \mathbb{R}$. Thus, we need to see under what conditions the absolute value of $g'_\pm(t)$ is less than one for all $t \in \mathbb{R}$. Let us compute the derivative of g_\pm with respect to x_i :

$$g'_\pm(x_i) = \frac{-D_i^T y (\nabla_i s(x))(x_i) \mp \lambda(2s(x)D_i^T Dx - D_i^T y)}{(s(x))^2}, \quad (\text{B.25})$$

where $s(x) = \text{SSIM}(Dx, y)$ and $\nabla_i s(x)$ is the i -th component of the gradient of $s(x)$. Thus, since $\|D_i^T\|_2 = 1$ for all $1 \leq i \leq n$, we have that the absolute value of $g'_\pm(x_i)$ is bounded above by

$$|g'_\pm(x_i)| = \frac{|-(\nabla_i s(x))(x_i) \pm \lambda(\|y\|_2 + 2\lambda\|D\|_2\|x\|_2)|}{(s(x))^2}, \quad (\text{B.26})$$

Observe that if x_i were zero, we would have that $a_i(0) \in \tau_i(0)[-1, 1]$, which is equivalent to $(\nabla_i s(x))(0) \in \lambda[-1, 1]$. Therefore, if g_\pm is a contraction, we must have that $(\nabla_i s(x))(x_i)$ is approximately equal to λ , if $x_i > 0$, or $-\lambda$, if $x_i < 0$. That is, $(\nabla_i s(x))(x_i)$ is either equal to $\lambda + \delta_1$ or $-\lambda - \delta_2$, where both δ_1 and δ_2 are small positive constants. Let $\delta = \max\{|\delta_1|, |\delta_2|\}$. Thus,

$$|g'_\pm(x_i)| = \frac{\delta\|y\|_2 + 2\lambda\|D\|_2\|x\|_2}{(s(x))^2}, \quad (\text{B.27})$$

Therefore, g_\pm is a contraction if the regularizing parameter λ satisfies the following inequality:

$$\lambda < \frac{(s(x))^2 - \delta\|y\|_2}{2\|D\|_2\|x\|_2}. \quad (\text{B.28})$$

Let Ω be the set of vectors that lie on the path in \mathbb{R}^n that an algorithm may follow to solve problem 9.39 starting at $x_0 \in \mathbb{R}^n$. Then, g_\pm is contractive if the following inequality holds:

$$\lambda < \min_{x \in \Omega} \left\{ \frac{(s(x))^2 - \delta\|y\|_2}{2\|D\|_2\|x\|_2} \right\}. \quad (\text{B.29})$$

□

Appendix C

MATLAB Codes

C.1 Algorithm I

```
function [x,y] = bisection_ssim_lp(A,D,y,K)

% Bisection method for constrained SSIM-based optimization. In this case,
% the following problem is solved:
%
%               min T(Ax,y)
%               subject to \|Dx\|_p<=K.
% Constraints can be changed or added for carrying out the desired
% SSIM-based imaging task. This mehtod employs CVX, which is a
% Matlab-based modeling system for convex optimization. CVX can be
% downloaded from http://cvxr.com/cvx/download/.

n = size(A,2);
x = zeros(n,1);
level_min = 0;
level_max = 2;
epsilon = 1e-5;
maxiter = 50;
i = 0;
p = 1;

average_y = mean(y);
```

```

y = y - average_y;

while ((level_max-level_min) > epsilon) || (cvx_optval ~= 0)
    alpha = (level_min + level_max)/2;
    cvx_begin
        variable x(n);
        minimize 0
        subject to
            (1-alpha)*sum(square(A*x-y)) - ...
            2*alpha*((A*x)'*y) <= 0;
            norm(D*x,p) <= K;
    cvx_end
    if cvx_optval == 0
        level_max = alpha;
    elseif (cvx_optval == inf)&&(alpha == 1)
        break;
    elseif cvx_optval == inf
        level_min = alpha;
    end
    i = i + 1;
    if (i >= maxiter)&&(cvx_optval ~= inf)
        break;
    end
end

y = y + average_y;

end

```


C.2 Algorithm II

```
function [x,y] = ssim_l2(D,A,y,lambda)

% Generalized Newton's Method for unconstrained SSIM-based optimization
% with Tikhonov regularization. This algorithm solves the following
% unconstrained SSIM-base optimization problem:
%           min {T(Dx,y) + \lambda\|Ax\|_2^2}

average_y = mean(y);
y = y - average_y;
Dt = D';
At = A';
x = (Dt*D + lambda*At*A)^(-1)*Dt*y;
C = .03;
epsilon0 = .0000001;
maxiter = 50;
i = 1;
cc0 = x;
cc0(1) = inf;

while (norm(cc0-x,inf))/(norm(x,inf))>epsilon0
    cc0 = x;
    x = x - jacob(x,y,D,Dt,A,At,lambda,C)^(-1)*fx(x,y,D,Dt,A,At,lambda,C);
    if i==maxiter
        break;
    end
    i = i + 1;
end

y = y + average_y;

end

function z = ssim(x,y,epsilon)
    z = (2*x'*y + epsilon)/(norm(x)^2 + norm(y)^2 + epsilon);
end
```

```
function z = rx(x,y,epsilon)
    z = norm(x)^2 + norm(y)^2 + epsilon;
end
```

```
function z = jacob(x,y,D,Dt,A,At,lambda,epsilon)
    DtD = Dt*D;
    AtA = At*A;
    I = eye(size(DtD,1));
    Dx = D*x;
    ss = ssim(Dx,y,epsilon);
    r = rx(Dx,y,epsilon);
    z = DtD*((2/r)*x*(Dt*y - DtD*x*ss)' + ss*I) + ...
        lambda*AtA*(x*(2*DtD*x)' + r*I);
end
```

```
function z = fx(x,y,D,Dt,A,At,lambda,epsilon)
    Dx = D*x;
    z = (ssim(Dx,y,epsilon)*(Dt*D) + lambda*rx(Dx,y,epsilon)*(At*A))*x - ...
        Dt*y;
end
```

C.3 Algorithm III

```
function [x,y] = ssim_l1(D,y,lambda)

% Coordinate Descent algorithm for unconstrained SSIM-based optimization
% with L1 norm regularization. This algorithm solves the following
% unconstrained SSIM-base optimization problem:
%           min {T(Dx,y) + \lambda\|x\|_1}

average_y = mean(y);
y = y - average_y;
Dt = D';
x = (Dt*D)^(-1)*Dt*y;
C = .03;
epsilon0 = 10^(-5);
epsilon1 = 10^(-5);
maxiter = 50;
i = 1;
cc0 = x;
cc0(1) = inf;

while (norm(cc0-x,inf))/(norm(x,inf))>epsilon0
    cc0 = x;
    for j = 1:size(x,1)
        cp = x;
        cp(j) = 0;
        if abs(a(cp,y,cp,D,Dt,j,C)) <= tau(cp,y,D,Dt,j,lambda,C)
            x(j) = 0;
        elseif a(cp,y,cp,D,Dt,j,C) > tau(cp,y,D,Dt,j,lambda,C)
            c0 = inf;
            while abs(x(j)-c0)>epsilon1
                c0 = x(j);
                x(j) = a(x,y,cp,D,Dt,j,C) - tau(x,y,D,Dt,j,lambda,C);
            end
        elseif a(cp,y,cp,D,Dt,j,C) < -tau(cp,y,D,Dt,j,lambda,C)
            c0 = inf;
            while abs(x(j)-c0)>epsilon1
```

```

                c0 = x(j);
                x(j) = a(x,y,cp,D,Dt,j,C) + tau(x,y,D,Dt,j,lambda,C);
            end
        end
    end
    if i==maxiter
        break;
    end
    i = i + 1;
end

y = y + average_y;

end

function z = ssim(x,y,epsilon)
    z = (2*x'*y + epsilon)/(norm(x)^2 + norm(y)^2 + epsilon);
end

function z = a(x,y,xk,A,At,k,epsilon)
    z = (At(k,:)*y)/(ssim(A*x,y,epsilon)*norm(At(k,:))^2) - At(k,:)*A*xk;
end

function z = tau(x,y,A,At,k,lamb,epsilon)
    z = (lamb/2)*((norm(A*x)^2 + norm(y)^2 + epsilon)/...
        (ssim(A*x,y,epsilon)*norm(At(k,:))^2));
end

```

C.4 Algorithm IV

```
function [x,y] = ssim_l1_mollifiers(D,y,lambda)

% Gradient descent for unconstrained SSIM-L1 optimization via mollifiers.
% This algorithm solves the following SSIM-based optimization problem using
% mollifiers:
%           min {T(Dx,y) + \lambda\|x\|_1^\epsilon}

average_y = mean(y);
y = y - average_y;
dim = size(D,2);
epsilon = .001;
gamma = 10^(-3);
n = 10^6;
C = .03;
Dt = D';
DtD = Dt*D;
x = (DtD)^(-1)*Dt*y;
criterion = 10^(-5);
maxiter = 50;
i = 1;
x_0 = x;
x_0(1) = inf;

while (norm(x_0-x,inf))/(norm(x,inf))>criterion
    x_0 = x;
    x = x - gamma*moll_subgrad(D,Dt,DtD,C,dim,epsilon,lambda,n,x,y);
    if i==maxiter
        break;
    end
    i = i + 1;
end

y = y + average_y;

end
```

```

function grad = moll_subgrad(D,Dt,DtD,C,dim,epsilon,lambda,n,x,y)
    grad = grad_ssim(D,Dt,DtD,C,x,y) + ...
        lambda*conv_grad_molli_1(dim,epsilon,n,x);
end

function [ssim_grad] = grad_ssim(D,Dt,DtD,C,x,y)
    Dx = D*x;
    ssim_grad = 2*(ssim(Dx,y,C)*DtD*x - Dt*y)/(norm(Dx)^2 + norm(y)^2 + C);
end

function s = ssim(x,y,C)
    s = (2*x'*y + C)/(norm(x)^2 + norm(y)^2 + C);
end

function [conv] = conv_grad_molli_1(dim,epsilon,n,x)

% Approximates the mollified subgradient of the L1 norm at the point x via
% Monte Carlo integration

[z,x] = gauss_rand_points(dim,epsilon,n,x);

conv = (epsilon^(-2))*mean(func(dim,x,z),2);

end

function y = func(dim,x,z)

% "program" the function that is to be integrated here

y = repmat(sum(abs(z),1),dim,1).*(z-x);

end

```

References

- [1] G. Albuquerque, M. Eisemann, and M.A. Magnor. Perception-based visual quality measures. In *IEEE VAST*, pages 13–20. IEEE, 2011.
- [2] S.K. Alexander, E.R. Vrscay, and S. Tsurumi. A simple, general model for the affine self-similarity of images. In *Image Analysis and Recognition*, pages 192–203. Springer, 2008. 54, 59
- [3] C. D. Aliprantis and K. C. Border. *Infinite Dimensional Analysis: a Hitchhiker’s Guide*. Springer Verlag, 2006. 2, 28, 30, 31, 43
- [4] F. Bach, R. Jenatton, J. Mairal, and G. Obozinski. *Convex Optimization with Sparsity-Inducing Norms*. Optimization for Machine Learning, MIT Press, 2011. 1, 70, 72, 83, 99, 114, 117, 121, 129, 133
- [5] D. Bárcenas. The Radon-Nikodým theorem for reflexive Banach spaces. *Divulgaciones Matemáticas*, 11(1):55–59, 2003. 38
- [6] M.F. Barnsley. *Fractals everywhere*. Academic press, 2014. 11, 58
- [7] M.F. Barnsley and L.P. Hurd. *Fractal image compression*. AK Peters, Ltd., 1993. 11, 55, 59
- [8] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences archive*, 2(1):183–202, January 2009. 1, 70, 72, 83
- [9] G. Beer, R.T. Rockafellar, and R. J-B. Wets. A characterization of epi-convergence in terms of convergence of level sets. *Proceedings of the American Mathematical Society*, 116(3):753–761, 1992. 124

- [10] J.M. Bioucas-Dias and J.M.P. Nascimento. Hyperspectral subspace identification. *Geoscience and Remote Sensing, IEEE Transactions on*, 46(8):2435–2445, Aug 2008. 71
- [11] S. Bochner. Integration von funktionen, deren werte die elemente eines vektorraumes sind. *Fundamenta Mathematicae*, 1(20):262–176, 1933. 28, 133
- [12] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3(1):1–122, January 2011. 1, 71, 72, 75
- [13] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. 1, 85, 86, 92, 93, 95, 96, 97
- [14] A. Braides. *Γ -convergence for Beginners*, volume 22. Oxford University Press, 2002. 124
- [15] X. Bresson and T.F. Chan. Fast dual minimization of the vectorial total variation norm and applications to color image processing. *Inverse Problems and Imaging*, 2(4):455–484, 2008. 1, 71, 73, 133
- [16] H. Brézis. *Functional analysis, Sobolev spaces and partial differential equations*. Universitext. Springer, New York, Heidelberg, London, 2011. 123
- [17] D. Brunet. *A Study of the Structural Similarity Image Quality Measure with Applications to Image Processing*. PhD thesis, University of Waterloo, 2012. 88, 91
- [18] D. Brunet, E.R. Vrscay, and Z. Wang. Structural similarity-based approximation of signals and images using orthogonal bases. In *ICIAR (1)*, volume 6111 of *Lecture Notes in Computer Science*, pages 11–22. Springer, 2010. 2, 84, 107, 111, 118
- [19] D. Brunet, E.R. Vrscay, and Z. Wang. Structural similarity-based affine approximation and self-similarity of images revisited. In *Image Analysis and Recognition*, pages 264–275. Springer, 2011. 54
- [20] D. Brunet, E.R. Vrscay, and Z. Wang. On the mathematical properties of the structural similarity index. *Proc. IEEE Trans. Image Processing*, 21(4):1488–1499, 2012. 85, 91
- [21] A. Buades, B. Coll, and J-M. Morel. A non-local algorithm for image denoising. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 2, pages 60–65. IEEE, 2005. 2

- [22] A. Buades, B. Coll, and J-M. Morel. A review of image denoising algorithms, with a new one. *Multiscale Modeling & Simulation*, 4(2):490–530, 2005. 54
- [23] H. Cartan and H.P. Cartan. *Differential calculus*, volume 1. Hermann, 1971. 2, 26, 27, 133
- [24] A. Chambolle. An algorithm for total variation minimization and applications. *Journal of Mathematical Imaging and Vision*, 20(1-2):89–97, 2004. 1, 9, 70, 72, 73, 101
- [25] T.F. Chan and J Shen. *Image Processing and Analysis*. Society for Industrial and Applied Mathematics, Philadelphia, 2005. 6, 7, 8, 9
- [26] C. Chang. *Hyperspectral Imaging: Techniques for Spectral Detection and Classification*. Kluwer Academic/Plenum Publishers, 2003. 11, 12
- [27] C-I. Chang. *Hyperspectral data exploitation: theory and applications*. John Wiley & Sons, 2007. 55
- [28] Y. Chang, L. Yan, H. Fang, and H. Liu. Simultaneous destriping and denoising for remote sensing images with unidirectional total variation and sparse representation. *Geoscience and Remote Sensing Letters, IEEE*, 11(6):1051–1055, 2014. 70, 72
- [29] S.S. Channappayya, A.C. Bovik, C. Caramanis, and R.W. Jr. Heath. Design of linear equalizers optimized for the structural similarity index. *IEEE Transactions on Image Processing*, 17(6):857–872, 2008. 2, 85
- [30] C. Chapman. *Real Mathematical Analysis*. Springer, New York, 2002. 20, 24, 34
- [31] G. Chen and S-E. Qian. Denoising of hyperspectral imagery using principal component analysis and wavelet shrinkage. *Geoscience and Remote Sensing, IEEE Transactions on*, 49(3):973–980, March 2011. 70
- [32] P.L. Combettes and J-C. Pesquet. Image restoration subject to a total variation constraint. *IEEE Transactions on Image Processing*, 13(9):1213–1222, 2004. 99
- [33] G.P. Crespi, D. La Torre, and M. Rocca. Mollified derivatives and second-order optimality conditions. *Journal of Nonlinear and Convex Analysis*, 4:437–454, 2003. 125
- [34] G.P. Crespi, D. La Torre, and M. Rocca. Second-order mollified derivatives and optimization. *Rendiconti del Circolo Matematico di Palermo*, 52:251–262, 2003. 125

- [35] L. Şendur and I.W. Selesnick. Bivariate shrinkage functions for wavelet-based denoising exploiting interscale dependency. *IEEE Transactions on Signal Processing*, 50(11), November 2002.
- [36] D. Culler, D. Estrin, and M. Srivastava. Guest editors' introduction: Overview of sensor networks. *Computer*, 37(8):41–49, Aug 2004. 17
- [37] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3-D transform-domain collaborative filtering. *Image Processing, IEEE Transactions on*, 16(8):2080–2095, 2007. 54
- [38] Grupo de Inteligencia Computacional de la Universidad del País Vasco. http://www.ehu.es/ccwintco/index.php?title=Hyperspectral_Remote_Sensing_Scenes. 52, 78
- [39] A. de Korvin and C.E. Roberts. Interchange of vector valued integrals when the measures are bochner or pettis indefinite integrals. *Bulletin of the Australian Mathematical Society*, 20(02):199–206, 1979. 2, 48
- [40] J. Diestel and J.J. Uhl. *Vector Measures*. American Mathematical Society, Providence, Rhode Island, 1977. 2, 30, 35, 37, 38, 39, 133
- [41] D.L. Donoho and I.M. Johnstone. Ideal spatial adaptation by wavelet shrinkage. *Biometrika*, 81(3):425–455, August 1994. 115
- [42] D.L. Donoho and I.M. Johnstone. Adapting to unknown smoothness via wavelet shrinkage. *Journal of the American Statistical Association*, 90(432):1220–1224, December 1995.
- [43] D.L. Donoho, I.M. Johnstone, G. Kerkyacharian, and D. Picard. Wavelet shrinkage: Asymptopia? *Journal of the Royal Statistical Society. Series B (Methodological)*, 57(2):pp. 301–369, 1995.
- [44] N. Dunford, J. T Schwartz, W.G. Bade, and R.G. Bartle. *Linear operators*. Wiley-interscience, New York, 1971. 32
- [45] M. Ebrahimi and E.R. Vrscay. Solving the inverse problem of image zooming using “self-examples”. In *Image analysis and Recognition*, pages 117–130. Springer, 2007. 54
- [46] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *The Annals of Statistics*, 32(2):407–451, 2004. 99

- [47] M. Elad and M. Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image Processing*, 15(12):3736–3745, 2006. 84
- [48] M. Elad and D. Datsenko. Example-based regularization deployed to super-resolution reconstruction of a single image. *The Computer Journal*, 52(1):15–30, 2009. 54
- [49] R. Enkhbat. *Quasiconvex Programming and its Applications*. Lambert, 2009. 135
- [50] Y.M. Ermoliev, V.I. Norkin, and R.J.B. Wets. The minimization of semicontinuous functions: mollifier subgradients. *Journal on Control and Optimization*, 33:149–167, 1995. 3, 121, 122, 123, 124, 125, 126
- [51] C. Etemoğlu and V. Cuperman. Structured vector quantization using linear transforms. *Signal Processing, IEEE Transactions on*, 51(6):1625–1631, 2003. 54
- [52] J. Fadili and G. Peyré. Total variation projection with first order schemes. *IEEE Transactions on Image Processing*, 20(3):657–669, 2011. 99
- [53] G.B. Folland. *A Course in Abstract Harmonic Analysis*. CRC Press, 1995. 44
- [54] D.H. Foster. Hyperspectral images of natural scenes 2004, “scene 2”. <http://personalpages.manchester.ac.uk/staff/david.foster/default.htm>. 67
- [55] W.T. Freeman, T.R. Jones, and E.C. Pasztor. Example-based super-resolution. *Computer Graphics and Applications, IEEE*, 22(2):56–65, 2002. 54
- [56] M. Fuhry and L. Reichel. A new Tikhonov regularization method. *Numerical Algorithms*, 59(3):433–445, 2012. 98
- [57] P. Getreuer. Total variation deconvolution using split Bregman. *Image Processing On Line*, 10, 2012. 100
- [58] M. Ghazel, G.H. Freeman, and E.R. Vrscay. Fractal image denoising. *Image Processing, IEEE Transactions on*, 12(12):1560–1578, 2003. 55
- [59] I. Ginchev. Vector optimization problems with quasiconvex constraints. *Journal of Global Optimization*, 44(1):111–130, 2009. 135
- [60] B. Goldlücke, E. Strekalovskiy, and D. Cremers. The natural vectorial total variation which arises from geometric measure theory. *SIAM Journal on Imaging Sciences*, 5:537–563, 2012. 1, 71, 73, 100, 101

- [61] T. Goldstein and S. Osher. The split Bregman method for l^1 -regularized problems. *SIAM J. Img. Sci.*, 2(2), April 2009. 72
- [62] R.C. Gonzalez and R.E. Woods. *Digital image processing*. Prentice-Hall, 2. ed. edition, 2002. 9
- [63] H.F. Grahn and P. Geladi. *Techniques and Applications of Hyperspectral Image Analysis*. John Wiley & Sons, 2007. 12
- [64] R. Haberman. Elementary applied partial differential equations: with Fourier series and boundary value problems. *The American Mathematical Monthly*, 1985. 51
- [65] P. Hagmann, L. Jonasson, P. Maeder, J. P. Thiran, V. J. Wedeen, and R. Meuli. Understanding diffusion MR imaging techniques: From scalar diffusion-weighted imaging to diffusion tensor imaging and beyond. *RadioGraphics*, 26:205–223, 2006. 15
- [66] M.E. Hochstenbach and L. Reichel. An iterative method for Tikhonov regularization with a general linear regularization operator. *J. Integral Equations Applications*, 22(3):465–482, 09 2010. 98
- [67] A.E. Jacquin. Image coding based on a fractal theory of iterated contractive image transformations. *Image Processing, IEEE Transactions on*, 1(1):18–30, 1992. 59
- [68] T.E. Johansen-Berg, H. Behrens. *Diffusion MRI: From quantitative measurement to in-vivo neuroanatomy*. Elsevier, 2009. 14, 15
- [69] V. Keyantuo and C. Lizama. Periodic solutions of second order differential equations in Banach spaces. *Mathematische Zeitschrift*, 253(3):489–514, 2006. 51, 133
- [70] K. Khoshelham and D. Altundag. Wavelet denoising of terrestrial laser scanner data for the characterization of rock surface roughness. In *Proceedings of the Joint International Conference on Theory, Data Handling and Modelling in GeoSpatial Information Science*, volume 38, May 2010.
- [71] M. Kline. *Mathematical Thought from Ancient to Modern Times: Vol. II*. Oxford University Press, New York, 1972. 20
- [72] H. Kunze, D. La Torre, F. Mendivil, and E.R. Vrscay. *Fractal-based methods in analysis*. Springer Science & Business Media, 2011. 2, 59

- [73] D. La Torre and M. Rocca. Remarks on second order generalized derivatives for differentiable functions with Lipschitzian Jacobian. *Applied Mathematics E-Notes [electronic only]*, 3:130–137, 2003. 125
- [74] D. La Torre and E.R. Vrscay. Fourier transforms of measure-valued images, self-similarity and the inverse problem. *Signal Processing*, 101:11–18, 2014.
- [75] D. La Torre, E.R. Vrscay, M. Ebrahimi, and M.F. Barnsley. Measure-valued images, associated fractal transforms, and the affine self-similarity of images. *SIAM Journal on Imaging Sciences*, 2(2):470–507, 2009. 2
- [76] D.A. Landgrebe. *Signal Theory and Methods in Multispectral Remote Sensing*. John Wiley & Sons, 2003. 12
- [77] Jens Lindström. On the origin and early history of functional analysis. Technical report, Technical Report UUDM Project Report 2008: 1, Uppsala Universitet, 2008.
- [78] N. Lu. *Fractal imaging*. Morgan Kaufmann Publishers Inc., 1997. 54, 55, 57, 58, 59, 64, 66, 67
- [79] S. Mallat. *A Wavelet Tour of Signal Processing*. San Diego : Academic Press, 1999. 9, 41, 42, 44
- [80] J.E. Marsden and A. Tromba. *Vector Calculus*. Freeman and Company, New York, 2003. 20, 21, 22
- [81] J. Martín-Herrero. Anisotropic diffusion in the hypercube. *Geoscience and Remote Sensing, IEEE Transactions on*, 45(5):1386–1398, 2007. 1, 10, 70
- [82] G.S. Mayer and E.R. Vrscay. Iterated Fourier transform systems: a method for frequency extrapolation. In *Image Analysis and Recognition*, pages 728–739. Springer, 2007. 55
- [83] R. Mendez-Rial, M. Calvino-Cancela, and J. Martín-Herrero. Anisotropic inpainting of the hypercube. *Geoscience and Remote Sensing Letters, IEEE*, 9(2):214–218, 2012. 10
- [84] O. Michailovich, D. La Torre, and E.R. Vrscay. Function-valued mappings, total variation and compressed sensing for diffusion MRI. In *ICIAR (2)*, volume 7325 of *Lecture Notes in Computer Science*, pages 286–295. Springer, 2012. 2, 10, 55, 56

- [85] M. Milman. Complex interpolation and geometry of Banach spaces. *Annali di Matematica Pura ed Applicata*, 136(1):317–328, 1984. 2, 42, 133
- [86] M. Miranda. Functions of bounded variation on “good” metric spaces. *Journal de mathématiques pures et appliquées*, 82(8):975–1004, 2003. 133
- [87] S. Morigi, L. Reichel, and F. Sgallari. Orthogonal projection regularization operators. *Numerical Algorithms*, 44(2):99–114, 2007. 98
- [88] G. Motta, F. Rizzo, and J.A. Storer. *Hyperspectral data compression*. Springer Science & Business Media, 2006. 55
- [89] D. Mumford and B. Gidas. Stochastic models for generic images. *Quart. Appl. Math.*, 59:85–111, 1999. 8
- [90] V. Namias. The fractional order Fourier transform and its application to quantum mechanics. *IMA Journal of Applied Mathematics*, 25(3):241–265, 1980. 42
- [91] Jet Propulsion Laboratory NASA. Aviris hyperspectral image, “yellowstone calibrated scene 0”. <http://compression.jpl.nasa.gov/hyperspectral/>. 62
- [92] P. Olver and A. Tannenbaum. *Image Analysis: Low and High Level Vision*. Springer-Verlag, 2003. 7
- [93] J.M. Ortega. The Newton-Kantorovich theorem. *The American Mathematical Monthly*, 2(75):658–660, 1968. 112
- [94] S. Osher and J.A. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics*, 79(1):12 – 49, 1988. 9
- [95] H. Othman and S-E. Qian. Noise reduction of hyperspectral imagery using hybrid spatial-spectral derivative-domain wavelet shrinkage. *IEEE Transactions on Geoscience and Remote Sensing*, 44(2), February 2006. 70, 71
- [96] R. Paget and D. Longstaff. Extracting the cliques from a neighbourhood system. *Vision, Image and Signal Processing, IEEE Proceedings*, 144(3):168–170, Jun 1997. 7
- [97] J. Peetre. Sur la transformation de Fourier des fonctions à valeurs vectorielles. *Rendicotti del Seminario Matematico della Università di Padova*, 42:15–26, 1969. 2, 42, 133

- [98] B. Rasti, J.R. Sveinsson, M.O. Ulfarsson, and J.A. Benediktsson. Hyperspectral image denoising using 3-D wavelets. In *Geoscience and Remote Sensing Symposium (IGARSS), 2012 IEEE International*, pages 1349–1352, July 2012. 70
- [99] Smith R.B. Introduction to hyperspectral imaging. Technical report, Microimages, July 2006. 11, 12
- [100] V. Recupero. Bv solutions of rate independent variational inequalities. *Annali della Scuola Normale Superiore di Pisa-Classe di Scienze-Serie V*, 10(2):269, 2011. 133
- [101] A. Rehman, Y. Gao, J. Wang, and Z. Wang. Image classification based on complex wavelet structural similarity. *Sig. Proc.: Image Comm.*, 28(8):984–992, 2013. 86
- [102] A. Rehman, M. Rostami, Z. Wang, D. Brunet, and E.R. Vrscay. SSIM-inspired image restoration using sparse representation. *EURASIP J. Adv. Sig. Proc.*, 2012:16, 2012. 84, 85, 101
- [103] N. Renard, S. Bourennane, and J. Blanc-Talon. Denoising and dimensionality reduction using multilinear tools for hyperspectral images. *Geoscience and Remote Sensing Letters, IEEE*, 5(2):138–142, April 2008. 70
- [104] R.T. Rockafellar and R. J-B. Wets. *Variational Analysis*. Springer-Verlag, Berlin, 1998. 124
- [105] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1-4):259–268, 1992. 9, 70, 72, 73
- [106] Y. Shao, F. Sun, H. Li, and Y. Liu. Structural similarity-optimal total variation algorithm for image denoising. In *Foundations and Practical Applications of Cognitive Systems and Information Processing*, pages 833–843. Springer, 2014. 100
- [107] P. Shippert. Why use hyperspectral imagery? *Photogrammetric engineering and remote sensing*, 70(4):377–380, 2004. xi, 11, 12, 13, 70
- [108] J. Stewart. *Calculus: Early Transcendentals*. Thomson Brooks/Cole, 2007. 20, 21, 22
- [109] H. Thompson. *The Bochner Integral and an Application to Singular Integrals*. PhD thesis, Dalhousie University, 2014. 31, 43
- [110] G. Travaglini. *Number Theory, Fourier Analysis and Geometry Discrepancy*. Cambridge University Press, 2014. 41

- [111] P. Tseng and S. Yun. A coordinate gradient descent method for non-smooth separable minimization. *Journal of Mathematical Programming*, 117(1):387–423, 2009. 114
- [112] M. Tubaishat and S.K. Madria. Sensor networks: an overview. *Potentials, IEEE*, 22(2):20–23, April 2003. 17
- [113] B.A. Turlach. On algorithms for solving least squares problems under an l^1 penalty or an l^1 constraint. In American Statistical Association, editor, *Proceedings of the American Statistical Association, Statistical Computing Section*, pages 2572–2577, 2005. 70, 72, 83, 115, 117, 129
- [114] S. Voronin and D. Yoshida. Gradient based methods for non-smooth regularization via convolution smoothing. *arXiv preprint arXiv:1408.6795*, 2014. 121, 130
- [115] E. R. Vrscay, D. Otero, and D. La Torre. Hyperspectral images as function-valued mappings, their self-similarity and a class of fractal transforms. In *Image Analysis and Recognition*, pages 225–234. Springer, 2013. 10, 11, 54, 55, 60, 62, 67
- [116] E.R. Vrscay, D. Otero, and D. La Torre. A simple class of fractal transforms for hyperspectral images. *Applied Mathematics and Computation*, 231:435–444, 2014. 10, 11
- [117] D.F. Walnut. *An Introduction to Wavelet Analysis*. Applied and Numerical Harmonic Analysis. Birkhäuser Boston, 2002. 9
- [118] S. Wang, A. Rehman, Z. Wang, S. Ma, and W. Gao. SSIM-motivated rate-distortion optimization for video coding. *IEEE Trans. Circuits Syst. Video Techn.*, 22(4):516–529, 2012. 2, 10, 86
- [119] S. Wang, A. Rehman, Z. Wang, S. Ma, and W. Gao. Perceptual video coding based on SSIM-inspired divisive normalization. *IEEE Transactions on Image Processing*, 22(4):1418–1429, 2013.
- [120] Y. Wang, N. Ruiqing, and Y. Xin. Anisotropic diffusion for hyperspectral imagery enhancement. *Sensors Journal, IEEE*, 10(3):469–477, 2010. 10
- [121] Z. Wang and A.C. Bovik. A universal image quality index. *Signal Processing Letters, IEEE*, 9(3):81–84, 2002. 83
- [122] Z. Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *Image Processing, IEEE Transactions on*, 13(4):600–612, 2004. 1, 78, 79, 83, 84, 87, 88, 90, 102, 103, 118, 130

- [123] AVIRIS website. <http://aviris.jpl.nasa.gov/index.html>. 69
- [124] Q. Yuan, L. Zhang, and H. Shen. Hyperspectral image denoising employing a spectral–spatial adaptive total variation model. *Geoscience and Remote Sensing, IEEE Transactions on*, 50(10):3660–3677, 2012. 70, 71, 72
- [125] F. Yuval. Fractal image compression-theory and application, 1994. 54, 55, 57, 58, 59, 64, 66, 67
- [126] E. Zeidler. *Nonlinear Functional Analysis and its Applications*. Springer-Verlag, New York, 1990. 2, 24, 34, 35, 36, 37, 38, 39, 40, 133, 146