Title
Variational Methods in Signal Decomposition and Image Processing

Permalink
https://escholarship.org/uc/item/4zx6s9zs

Author
Dragomiretskiy, Konstantin

Publication Date
2015

Peer reviewed|Thesis/dissertation

University of California
Los Angeles

Variational Methods in Signal Decomposition and Image Processing

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Mathematics

by

Konstantin Vitalyevich Dragomiretskiy

2015

ABSTRACT OF THE DISSERTATION

Variational Methods in Signal Decomposition and Image Processing

by

Konstantin Vitalyevich Dragomiretskiy

Doctor of Philosophy in Mathematics

University of California, Los Angeles, 2015

Professor Andrea Bertozzi, Chair

The work presented in this dissertation is motivated by classical problems in signal and image processing from the perspective of variational and PDE-based methods. Analytically encoding qualitative features of signals into variational energies in conjunction with modern methods in sparse optimization allows for well-founded and robust models, the optimization of which yields meaningful and cohesive signal decomposition.

Part I of this dissertation is based on joint work Variational Mode Decomposition [DZ14] with Dominique Zosso, in which the goal is to recursively decompose a signal into different modes of separate spectral bands, which are unknown beforehand. In the late nineties, Huang [HSL98] introduced the Hilbert-Huang transform, also known as Empirical Mode Decomposition, in order to decompose a signal into so-called intrinsic mode functions (IMF) along with a trend, and obtain instantaneous frequency data. The HHT/EMD algorithm is widely used today, although there is no exact mathematical model corresponding to this algorithm, and, consequently, the exact properties and limits are widely unknown. We propose an entirely non-recursive variational mode decomposition model, where the modes are extracted concurrently. The model looks for a number of modes and their respective center frequencies, such that the modes reproduce the input signal, while being smooth after demodulation into baseband. In Fourier domain, this corresponds to a narrow-band prior. Our model provides a solution to the decomposition problem that is theoretically well-founded, tractable, and motivated. The variational model is efficiently optimized using an alternating direction method of multipliers approach. Preliminary results show excellent performance with respect to existing mode decomposition models.

Part II of this dissertation is the n-dimensional extension of the Variational Mode Decomposition. Decomposing multidimensional signals, such as images, into spatially compact, potentially overlapping modes of essentially wavelike nature makes these components accessible for further downstream analysis such as space-frequency analysis, demodulation, estimation of local orientation, edge and corner detection, texture analysis, denoising, inpainting, and curvature estimation. The model decomposes the input signal into modes with narrow Fourier bandwidth; to cope with sharp region boundaries, incompatible with narrow bandwidth, we introduce binary support functions that act as masks on the narrow-band mode for image re-composition. L^1 and TV-terms promote sparsity and spatial compactness. Constraining the support functions to partitions of the signal domain, we effectively get an image segmentation model based on spectral homogeneity. By coupling several submodes together with a single support function we are able to decompose an image into several crystal grains. Our efficient algorithm is based on variable splitting and alternate direction optimization; we employ Merriman-Bence-Osher-like [MBO92] threshold dynamics to handle efficiently the motion by mean curvature of the support function boundaries under the sparsity promoting terms. The versatility and effectiveness of our proposed model is demonstrated on a broad variety of example images from different modalities. These demonstrations include the decomposition of images into overlapping modes with smooth or sharp boundaries, segmentation of images of crystal grains, and inpainting of damaged image regions through artifact detection.

Part III of this dissertation is based on joint work with Igor Yanovsky of NASA Jet Propulsion Laboratory. We introduce a variational method for destriping data acquired by pushbroom-type satellite imaging systems. The model leverages sparsity in signals and is based on current research in sparse optimization and compressed sensing. It is based on the basic principles of regularization and data fidelity with certain constraints using modern methods in variational optimization - namely total variation (TV), both L^1 and L^2 fidelity, and the alternate direction method of multipliers. The main algorithm in Part III uses sparsity promoting energy functionals to achieve two important imaging effects. The TV term maintains boundary sharpness of content in the underlying clean image, while the L^1 fidelity allows for the equitable removal of stripes without over- or under-penalization, providing a more accurate model of presumably independent sensors with unspecified and unrestricted bias distribution. A comparison is made between the TV-L^1 and TV-L^2 models to exemplify the qualitative efficacy of an L^1 striping penalty. The model makes use of novel minimization splittings and proximal mapping operators, successfully yielding more realistic destriped images in very few iterations.

The dissertation of Konstantin Vitalyevich Dragomiretskiy is approved.

Mark S. Cohen
Luminita A. Vese
Stanley J. Osher
Andrea Bertozzi, Committee Chair

University of California, Los Angeles
2015

To my mother and father for supporting me and providing me with opportunity. And to my younger self, for maintaining discipline and perseverance through stormy waters.

TABLE OF CONTENTS

II Two-Dimensional Compact Variational Mode Decomposition: Spatially compact and spectrally sparse image decomposition and segmentation 42

List of Figures

ACKNOWLEDGMENTS

I would like to thank my advisor Andrea L. Bertozzi for her guidance and support, which made all of the research presented in this dissertation possible. Andrea, thank you for your encouragement in my transition into applied mathematics and flexibility to pursue topics of my interest.

I would like to thank my principal collaborator, co-author and mentor, Dominique Zosso. Ever since our inaugural dialogue at Stan's level set seminar, our collaboration has been both inspiring and rewarding. You have helped motivate and inspire the largest piece of this dissertation, and have exemplified what it means to be a great researcher. I appreciate our earnest conversations and your unbiased advice regarding my path in applied mathematics and future career.

I would like to thank the rest of my graduate committee, Stanley J. Osher, Luminita A. Vese, and Mark S. Cohen for their time and valuable input. They have been a constant inspiration through their prolific work, educational seminars and insightful comments.

Other UCLA mathematicians have made these last five years an educational pleasure. Chris Anderson, Prashant Athavale, Marek Biskup, Christoph Brune, Yves van Gennip, Jerome Gilles, Inwon Kim, Thomas Liggett, Benny Sudokov, Lieven Vandenberghe and Igor Yanovsky have all supported me in their own ways, for which I am very grateful.

A special thanks goes to Xander, Ben, and Nick - some of my best friends who helped keep my spirits up, and made my time in graduate school enjoyable and memorable.

The research presented here was made possible in part by the National Science Foundation through grants DMS-1118971 and DMS-1217239, by the Office of Naval Research through grants N000141210838 and N000141210040, UC Lab Fees Research grant 12-LR-236660, and the W. M. Keck Foundation.

VITA

1988	Born, Odessa, Ukrainian SSR, Soviet Union.
2005	Research Intern at VLSI Group, Sun Microsystems Research Labs.
2006	Colorado School of Mines Award Medal of Achievement in Mathematics.
2010	B.S. Mathematics and B.A. Economics, UCSD, *Magna cum laude*.
2011–2015	Teaching Fellow/Assistant, Department of Mathematics, UCLA.
2012	Academic Mentor, California Research Training Program REU in Applied Mathematics, UCLA.
2012–2015	Graduate Researcher, Department of Mathematics, UCLA.
2014–2015	Graduate Research Assistant, NASA Jet Propulsion Laboratory.
2014	Academic Mentor, Institute for Pure and Applied Mathematics.

PUBLICATIONS

Konstantin Dragomiretskiy and Dominique Zosso. Variational Mode Decomposition. *IEEE Transactions on Signal Processing*, 62(3):531544, 2014.

Konstantin Dragomiretskiy and Dominique Zosso. Two-dimensional Variational Mode Decomposition. *Energy Minimization Methods in Computer Vision and Pattern Recognition*, 197–208, 2015.

Konstantin Dragomiretskiy and Igor Yanovsky. Variational Destriping in Remote Sensing Imagery: Total Variation with L1 Fidelity. *IEEE Transactions on Geoscience and Remote Sensing*, (Submitted 2015).

Dominique Zosso, Konstantin Dragomiretskiy, Andrea L. Bertozzi, Paul S. Weiss. Two-dimensional Compact Variational Mode Decomposition: Spatially compact and spectrally sparse image decomposition and segmentation. *Journal of Mathematical Imaging and Vision*, (Submitted 2015).

Part I

Variational Mode Decomposition

CHAPTER 1

Introduction

Empirical Mode Decompositon (EMD) proposed by Huang et al. [HSL98] is an algorithmic method to detect and decompose a signal into principal "modes" - a signal with mostly compact supported Fourier spectrum. This algorithm recursively detects local minima/maxima in a signal, estimates lower/upper envelopes by spline-interpolation of these extrema, removes the average of the envelopes as "low-pass" centerline, thus isolating the high-frequency oscillations as "mode" of a signal, and continues recursively on the extracted "low-pass" centerline. In some cases, this sifting algorithm does indeed decompose a signal into principal modes, however the resulting decomposition is highly dependent on methods of extremal point finding, interpolation of extremal points into carrier envelopes, and the stopping criteria imposed. The lack of mathematical theory and the aformentioned degrees of freedom reducing the algorithm's robustness all leave room for theoretical development and improvement on the robustness of the decomposition [RFG03, RF08]. In some experiments it has been shown that EMD shares important similarities with wavelets and (adaptive) filter banks [FGR05].

Despite the limited mathematical understanding and some obvious shortcomings, the EMD method, also known as the Hilbert-Huang transform (HHT), has had significant impact and is widely used in a broad variety of time-frequency analysis applications. Applications involve signal decomposition in audio engineering [Klu12], climate analysis [BE11], and various flux, respiratory, and neuromuscular signals found in medicine and biology [AHL05, ANK06, LHG08, MBL09], to name just a few examples.

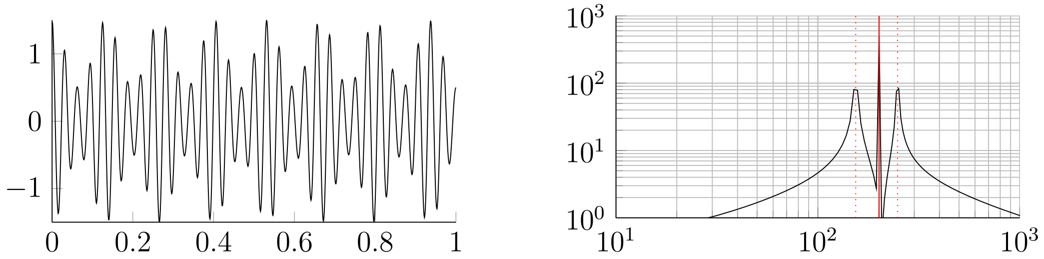

a) AM signal and spectrum. ($\Delta f = 0$)

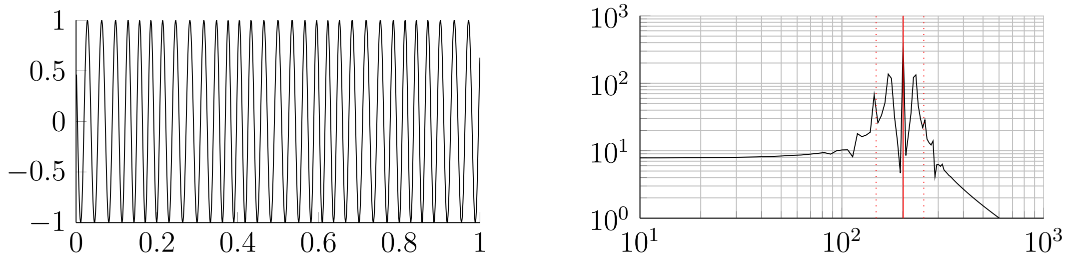

b) FM signal and spectrum. ($f_{\mathrm{FM}} \gg \Delta f$)

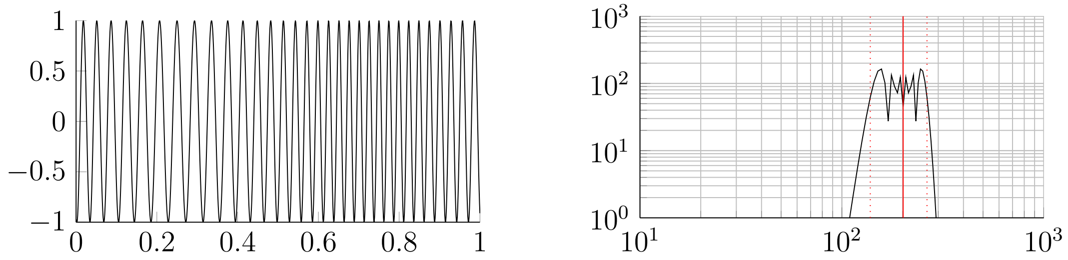

c) FM signal and spectrum. ($f_{\mathrm{FM}} \ll \Delta f$)

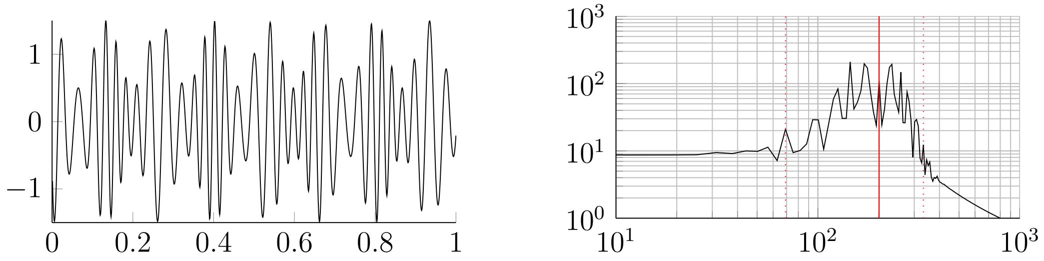

d) AM-FM signal and spectrum. ($f_{\mathrm{AM}} \sim f_{\mathrm{FM}} \sim \Delta f$)

Figure 1.1: AM-FM signals with limited bandwidth. Here, we use a signal $f(t) = (1 + 0.5\cos(2\pi f_{\mathrm{AM}} t)) \cdot \cos(2\pi f_c t + \Delta f / f_{\mathrm{FM}} \cos(2\pi f_{\mathrm{FM}} t))$. a) Pure AM signal. b) Pure FM signal with little but rapid frequency deviations. c) Pure FM signal with slow but important frequency oscillations. d) Combined AM-FM signal. The solid vertical line in the spectrum shows the carrier frequency f_c, the dotted lines correspond to the estimated band limits at $f_c \pm BW/2$, based on (1.2).

1.1 Intrinsic Mode Functions

With EMD, and in all of the previous signals, the core assumption on the individual modes is that they have compact Fourier support. In the original description, in such a mode the number of local extrema and zero-crossings differ at most by one [HSL98]. In most related works, the definition is slightly changed into so-called *Intrinsic Mode Functions* (IMF).

Definition Intrinsic Mode Functions are amplitude-modulated-frequency-modulated (AM-FM) signals, written as:

$$u_k(t) = A_k(t)\cos(\phi_k(t)), \tag{1.1}$$

where the phase $\phi_k(t)$ is a non-decreasing function, $\phi'_k(t) \geq 0$, the envelope is non-negative $A_k(t) \geq 0$, and, very importantly, both the envelope $A_k(t)$ and the instantaneous frequency $\omega_k(t) := \phi'_k(t)$ vary much slower than the phase $\phi_k(t)$ [DLW11, Gil13].

In other words, on a sufficiently long interval $[t-\delta, t+\delta]$, $\delta \approx 2\pi/\phi'_k(t)$, the mode $u_k(t)$ can be considered to be a pure harmonic signal with amplitude $A_k(t)$ and instantaneous frequency $\phi'_k(t)$ [DLW11]. Note that the newer definition of signal components is slightly more restrictive than the original one. The immediate consequence of the IMF assumption is limited bandwidth.

Indeed, if ω_k is the mean frequency of a mode, then its practical bandwidth increases both, with the maximum deviation of the instantaneous frequency, $\Delta f \sim \max(|\omega_k(t) - \omega_k|)$, and with the rate of change of the instantaneous frequency, $f_{\text{FM}} \sim \omega'(t)$, according to Carson's rule: $BW = 2(\Delta f + f_{\text{FM}})$ [Car22]. In addition to this comes the bandwidth of the envelope $A_k(t)$ modulating the amplitude of the FM signal, given by its highest frequency f_{AM}. Hence we estimate the total bandwidth of an IMF as

$$BW = 2(\Delta f + f_{\text{FM}} + f_{\text{AM}}). \tag{1.2}$$

Depending on the actual IMF, either of these terms may be dominant. An illustration of four typical cases is provided in figure 1.1, where the last example is rather extreme in terms of required bandwidth (for illustrational purposes).

Some recent works create a partially variational approach to EMD where the signal is explicitly modeled as an IMF [HS11]. This method still relies on interpolation, selection of a Fourier low-pass filter, and sifting of high-frequency components. Here, the candidate modes are extracted variationally. The signal is recursively decomposed into an IMF with TV3-smooth envelope, and a TV3-smooth residual. The resulting algorithm is very similar to EMD in structure, but somewhat more robust to noise.

A slightly more variational, but still recursive decomposition scheme has been proposed in [Fel06], for the analysis of time-varying vibration. Here, the dominant vibration is extracted by estimating its instantaneous frequency as average frequency after the Hilbert transform. Again, this process is repeated recursively on the residual signal.

An approach based on selecting appropriate wavelet scales, dubbed *synchrosqueezing*, was proposed by Daubechies et al. [DLW11, WFD11]. They remove unimportant wavelet coefficients (both in time and scale) by thresholding of the respective signal energy in that portion. Conversely, locally relevant wavelets are selected as local maxima of the continuous wavelet transform, that are shown to be tuned with the local signals, and from which the current instantaneous frequency of each mode can be recovered.

Other recent work pursuing the same goal is the Empirical Wavelet Transform (EWT) to explicitly build an adaptive wavelet basis to decompose a given signal into adaptive subbands [Gil13]. This model relies on robust preprocessing for peak detection, then performs spectrum segmentation based on detected maxima, and constructs a corresponding wavelet filter bank. The filter bank includes flexibility for some mollification (spectral overlap), but explicit construction of frequency bands still appears slightly strict.

In Part I, we propose a new, fully intrinsic and adaptive, variational method, the minimization of which leads to a decomposition of a signal into its principal modes. Indeed, the current decomposition models are mostly limited by 1) their algorithmic ad-hoc nature lacking mathematical theory (EMD), 2) the recursive sifting in most methods, which does not allow for backward error correction, 3) the inability to properly cope with noise, 4) the hard band-limits of wavelet approaches, and 5) the requirement of predefining filter bank boundaries in EWT. In contrast, we propose a variational model that determines the relevant bands adaptively, and estimates the corresponding modes concurrently, thus properly balancing errors between them. Motivated by the narrow-band properties corresponding to the current common IMF definition, we look for an ensemble of modes that reconstruct the given input signal optimally (either exactly, or in a least-squares sense), while each being band-limited about a center frequency estimated on-line. Here, our variational model specifically can address the presence of noise in the input signal. Indeed, the tight relations to the Wiener filter actually suggest that our approach has some optimality in dealing with noise. The variational model assesses the bandwidth of the modes as H1-norm, after shifting the Hilbert-complemented, analytic signal down into baseband by complex harmonic mixing. The resulting optimization scheme is very simple and fast: each mode is iteratively updated directly in Fourier domain, as the narrow-band Wiener filter corresponding to the current estimate of the mode's center-frequency being applied to the signal estimation residual of all other modes; then the center frequency is re-estimated as the center-of-gravity of the mode's power spectrum. Our quantitative results on tone detection and separation show excellent performance irrespective of harmonic frequencies, in particular when compared to the apparent limits of EMD in this respect. Further, qualitative results on synthetic and real test signals are convincing, also regarding robustness to signal noise.

CHAPTER 2

Tools from Signal Processing

2.1 Gaussian regularizer and Wiener filtering

Let us start with a simple denoising problem. Consider the observed signal $f_0(t)$ to be a copy of the original signal $f(t)$ to be recovered, affected by additive zero-mean Gaussian noise:

$$f_0 = f + \eta \tag{2.1}$$

Recovering the unknown signal f is a typical ill-posed inverse problem [BPT88]. If the original signal is known to vary smoothly, one would typically write the following Tikhonov regularized minimization problem in order to estimate the noise-free signal [Tic63, Mor75]:

$$\min_f \left\{ \|f - f_0\|_2^2 + \alpha \|\partial_t f\|_2^2 \right\} \tag{2.2}$$

This is a standard, Gaussian regularized minimum mean squares, i.e. "L2-H1" problem, of which the Euler-Lagrange equations are easily obtained as

$$f - f_0 = \alpha \partial_t^2 f. \tag{2.3}$$

These EL equations are typically solved in Fourier domain:

$$\hat{f}(\omega) = \frac{\hat{f}_0}{1 + \alpha\omega^2}, \tag{2.4}$$

where $\hat{f}(\omega) := \mathcal{F}\{f(\cdot)\}(\omega) := 1/\sqrt{2\pi} \int_{\mathbb{R}} f(t) e^{-j\omega t} dt$, with $j^2 = -1$, is the Fourier transform of the signal $f(t)$. Clearly, the recovered signal f is a low-pass narrow-band selection of the input signal f_0 around $\omega = 0$. Indeed, the solution corresponds to convolution with a Wiener filter, where α represents the variance of the white noise, and the signal has a lowpass $1/\omega^2$ power spectrum prior [Wie49, GW92].

2.2 Hilbert transform

Here, we cite the definition of the Hilbert transform given in [USV09]:

Definition The 1-D Hilbert transform is the linear, shift-invariant operator $\mathcal{H}$ that maps all 1-D cosine functions into their corresponding sine functions. It is an all-pass filter that is characterized by the transfer function $\hat{h}(\omega) = -j\,\mathrm{sgn}(\omega) = -j\omega/|\omega|$.

Thus, the Hilbert transform is a multiplier operator in the spectral domain. The corresponding impulse response is $h(t) = 1/(\pi t)$. Because $h(t)$ is not integrable the integrals defining the convolution do not converge. Instead, the Hilbert transform $\mathcal{H}f(t)$ of a signal $f(t)$ is therefore obtained as the Cauchy principal value (denoted p.v.) of :

$$\mathcal{H}f(t) = \frac{1}{\pi}\text{p.v.}\int_{\mathbb{R}} \frac{f(v)}{t-v}dv. \tag{2.5}$$

Finally, the inverse Hilbert transform is given by its negative, $\mathcal{H}^{-1} = -\mathcal{H}$, thus:

$$\mathcal{H}^2 f(t) = -f(t). \tag{2.6}$$

For further properties and analysis of the Hilbert transform, we refer e.g. to [Hah96]. The most prominent use of the Hilbert transform is in the construction of an analytic signal from a purely real signal, as proposed by Gabor [Gab46].

2.3 Analytic signal

Definition Let $f(t)$ be a purely real signal. The complex analytic signal is now defined as:

$$f_A(t) = f(t) + j\mathcal{H}f(t) = A(t)e^{j\phi(t)}. \tag{2.7}$$

This analytic signal has the following important properties. The complex exponential term $e^{j\phi(t)}$ is a phasor describing the rotation of the complex signal in time, $\phi(t)$ being the phase, while the amplitude is governed by the real envelope $A(t)$. This representation is particularly useful in the analysis of time-varying amplitude and instantaneous frequency, defined as $\omega(t) = d\phi(t)/dt$. The second property is the unilateral spectrum of the analytic signal, consisting only of non-negative frequencies, hence its use in single-sideband modulation. Finally, we note that from such an analytical signal, the original real signal is easily retrieved as the real part:

$$f(t) = \Re\{f_A(t)\}. \tag{2.8}$$

It is worthwhile highlighting the simple relations between the Fourier spectra of the real signal and its analytic counterpart, as defined by (2.7). First, we recall that the (Fourier) spectrum of a real signal is a Hermitian function:

$$\hat{f}(-\omega) = \overline{\hat{f}(\omega)}. \tag{2.9}$$

In contrast, the spectrum of the analytic signal has only non-negative frequencies. In particular:

$$\hat{f}_A(\omega) = \begin{cases} 0 & \omega < 0 \\ \hat{f}(0) & \omega = 0 \\ 2\hat{f}(\omega) & \omega > 0. \end{cases} \tag{2.10}$$

2.4 Frequency mixing and heterodyne demodulation

The last concept that we wish to recall before introducing the proposed variational mode decomposition, is the principle of frequency mixing. Mixing is the process of combining two signals non-linearly, thus introducing cross-frequency terms in the output. The simplest mixer is multiplication. Multiplying two real signals with frequencies f_1 and f_2, respectively, creates mixed frequencies in the output at $f_1 - f_2$ and $f_1 + f_2$, which is easily illustrated by the following trigonometric identity:

$$2\cos(2\pi f_1 t)\cos(2\pi f_2 t) = \cos(2\pi(f_1 + f_2)t) + \cos(2\pi(f_1 - f_2)t). \tag{2.11}$$

Typical applications are the heterodyne downmixing of the modulated high-frequency carrier signal with a local (heterodyne) oscillator in a radio receiver. In such devices, the selection of either of the two output terms is achieved by filtering. Here, instead of filtering the output, we mix the two respective analytic signals:

$$e^{j2\pi f_1 t} e^{j2\pi f_2 t} = e^{j2\pi(f_1+f_2)t}, \tag{2.12}$$

i.e., the mixed signal is automatically "mono-tone" (constituted of a single frequency only). In Fourier terms, this is well known as the following transform pair:

$$f_A(t)e^{-j\omega_0 t} \overset{\mathcal{F}}{\longleftrightarrow} \hat{f}_A(\omega) * \delta(\omega + \omega_0) = \hat{f}_A(\omega + \omega_0), \tag{2.13}$$

where δ is the Dirac distribution and $*$ denotes convolution. Thus, multiplying an analytic signal with a pure exponential results in simple frequency shifting.

CHAPTER 3

Variational Mode Decomposition

3.1 Constrained model

In this section we introduce our proposed model for variational mode decomposition, essentially based on the three concepts outlined in the previous section.

The goal of VMD is to decompose an input signal into a discrete number of sub-signals (modes), that have specific sparsity properties while reproducing the input. Here, the sparsity prior of each mode is chosen to be its bandwidth in spectral domain. In other words, we require each mode k to be mostly compact around a center pulsation ω_k, which is to be determined along with the decomposition.

In order to assess the bandwidth of a mode, we propose the following scheme: 1) for each mode u_k, compute the associated analytic signal by means of the Hilbert transform in order to obtain a unilateral frequency spectrum. 2) for each mode, shift the mode's frequency spectrum to "baseband", by mixing with an exponential tuned to the respective estimated center frequency. 3) The bandwidth is now estimated through the H^1 Gaussian smoothness of the demodulated signal, i.e. the squared L^2-norm of the gradient. The resulting constrained variational problem is the following:

$$\min_{u_k,\omega_k} \left\{ \sum_k \left\| \partial_t \left[\left(\delta(t) + \frac{j}{\pi t} \right) * u_k(t) \right] e^{-j\omega_k t} \right\|_2^2 \right\} \quad \text{s.t.} \quad \sum_k u_k = f \tag{3.1}$$

3.2 Unconstrained model

The reconstruction constraint can be addressed in different ways. Here, we suggest making use of both a quadratic penalty term and Lagrangian multipliers in order to render the problem unconstrained. Therefore, we introduce the augmented Lagrangian $\mathcal{L}$ as follows [Ber76, NW06]:

$$\mathcal{L}(u_k, \omega_k, \lambda) = \\ \alpha \sum_k \left\| \partial_t \left[\left(\delta(t) + \frac{j}{\pi t} \right) * u_k(t) \right] e^{-j\omega_k t} \right\|_2^2 + \left\| f - \sum u_k \right\|_2^2 + \left\langle \lambda, f - \sum u_k \right\rangle. \quad (3.2)$$

The solution to the original minimization problem (1.1) is now found as the saddle point of the augmented Lagrangian $\mathcal{L}$ in a sequence of iterative sub-optimizations called alternate direction method of multipliers (ADMM). In the next paragraphs, we detail how the respective sub-problems can be solved.

3.3 Minimization with respect to the modes u_k

To update the modes u_k, we first rewrite the subproblem as the following equivalent minimization problem:

$$u_k^{n+1} = \operatorname*{arg\,min}_{u_k \in \mathbb{R}} \left\{ \alpha \left\| \partial_t \left[\left(\delta(t) + \frac{j}{\pi t} \right) * u_k(t) \right] e^{-j\omega_k t} \right\|_2^2 + \left\| f - \sum u_i + \frac{\lambda}{2} \right\|_2^2 \right\}. \quad (3.3)$$

Making use of the Parseval/Plancherel Fourier isometry under the L^2 norm, this problem can be solved in spectral domain:

$$\hat{u}_k^{n+1} = \operatorname*{arg\,min}_{\hat{u}_k, \hat{u}_k = \hat{u}_k^*} \left\{ \alpha \left\| j\omega \left[(1 + \operatorname{sgn}(\omega + \omega_k)) \hat{u}_k(\omega + \omega_k) \right] \right\|_2^2 + \left\| \hat{f} - \sum \hat{u}_i + \frac{\hat{\lambda}}{2} \right\|_2^2 \right\}. \quad (3.4)$$

We now perform a change of variables $\omega \to \omega + \omega_k$ in the first term:

$$\hat{u}_k^{n+1} = \operatorname*{arg\,min}_{\hat{u}_k, \hat{u}_k = \hat{u}_k^*} \left\{ \alpha \left\| j(\omega - \omega_k) \left[(1 + \operatorname{sgn}(\omega)) \hat{u}_k(\omega) \right] \right\|_2^2 + \left\| \hat{f} - \sum \hat{u}_i + \frac{\hat{\lambda}}{2} \right\|_2^2 \right\}. \quad (3.5)$$

Exploiting the Hermitian symmetry of the real signals in the reconstruction fidelity term, we can write both terms as half-space integrals over the non-negative frequencies:

$$\hat{u}_k^{n+1} = \underset{\hat{u}_k, \hat{u}_k = \hat{u}_k^*}{\arg\min} \left\{ \int_0^\infty 4\alpha(\omega - \omega_k)^2 |\hat{u}_k(\omega)|^2 + 2 \left(\hat{f} - \sum \hat{u}_i + \frac{\hat{\lambda}}{2} \right)^2 d\omega \right\}. \tag{3.6}$$

The solution of this quadratic optimization problem is readily found by letting the first variation vanish for the positive frequencies:

$$\hat{u}_k^{n+1} = \left(\hat{f} - \sum_{i \neq k} \hat{u}_i + \frac{\hat{\lambda}}{2} \right) \frac{1}{1 + 2\alpha(\omega - \omega_k)^2} \tag{3.7}$$

which is clearly identified as a Wiener filtering of the current residual, with signal prior $1/(\omega - \omega_k)^2$. The full spectrum of the real mode is then simply obtained by Hermitian symmetric completion. Conversely, the mode in time domain is obtained as the real part of the inverse Fourier transform of this filtered analytic signal.

3.4 Minimization with respect to the center frequencies ω_k

The center frequencies ω_k do not appear in the reconstruction fidelity term, but only in the bandwidth prior. The relevant problem thus writes:

$$\omega_k^{n+1} = \underset{\omega_k}{\arg\min} \left\{ \left\| \partial_t \left[\left(\delta(t) + \frac{j}{\pi t} \right) * u_k(t) \right] e^{-j\omega_k t} \right\|_2^2 \right\}. \tag{3.8}$$

As before, the optimization can take place in Fourier domain, and we end up optimizing:

$$\omega_k^{n+1} = \underset{\omega_k}{\arg\min} \left\{ \int_0^\infty (\omega - \omega_k)^2 |\hat{u}_k(\omega)|^2 d\omega \right\}, \tag{3.9}$$

This quadratic problem is easily solved as:

$$\omega_k^{n+1} = \frac{\int_0^\infty \omega |\hat{u}_k(\omega)|^2 d\omega}{\int_0^\infty |\hat{u}_k(\omega)|^2 d\omega}, \tag{3.10}$$

which puts the new ω_k at the center of gravity of the corresponding mode's power spectrum. This mean carrier frequency is the frequency of a least squares linear regression to the instantaneous phase observed in the mode.

3.5 Complete VMD algorithm

Plugging the solutions of the sub-optimizations into the ADMM algorithm and directly optimizing in Fourier domain where appropriate, we get the complete algorithm for variational mode decomposition, summarized in algorithm 1.

Algorithm 1 Complete optimization of VMD

```
Initialize $\hat{u}_k^1$, $\omega_k^1$, $\hat{\lambda}^1$, $n \leftarrow 0$
repeat
    $n \leftarrow n + 1$
    for $k = 1 : K$ do
        Update $\hat{u}_k$ for all $\omega \geq 0$:
            $$\hat{u}_k^{n+1} \leftarrow \frac{\hat{f} - \sum_{i<k} \hat{u}_i^{n+1} - \sum_{i>k} \hat{u}_i^n + \frac{\hat{\lambda}^n}{2}}{1 + 2\alpha(\omega - \omega_k^n)^2} \tag{3.11}$$
        Update $\omega_k$:
            $$\omega_k^{n+1} \leftarrow \frac{\int_0^\infty \omega |\hat{u}_k^{n+1}(\omega)|^2 d\omega}{\int_0^\infty |\hat{u}_k^{n+1}(\omega)|^2 d\omega} \tag{3.12}$$
    end for
    Dual ascent for all $\omega \geq 0$:
        $$\hat{\lambda}^{n+1} \leftarrow \hat{\lambda}^n + \tau \left( \hat{f} - \sum_k \hat{u}_k^{n+1} \right) \tag{3.13}$$
until convergence: $\sum_k \|\hat{u}_k^{n+1} - \hat{u}_k^n\|_2^2 / \|\hat{u}_k^n\|_2^2 < \epsilon$.
```

3.6 Inexact reconstruction and denoising

Here, the role of the Lagrangian multiplier is to enforce the constraint, while the quadratic penalty improves convergence. If exact reconstruction is not required, but some slack is to be allowed, using the quadratic penalty only while dropping the Lagrangian multiplier would be the appropriate choice. Indeed, the quadratic penalty on its own represents the least-squares fidelity prior associated with additive Gaussian noise.

3.7 On boundaries, periodicity, and windowing

Up until now, the signals f and the modes u_k have been considered continuous over the whole axis $t \in \mathbb{R}$. However, in signal processing we are much more likely to be working with signals that are both finite in time and resolution. Let us say we restrict the time window to $t \in [0, 1]$. Luckily the results presented so far equally hold for discrete, finite time signals, where simply the continuous Fourier transform is replaced by its discrete counterpart. The only problems arise at the boundaries of the signal.

Indeed, when considering short-time signals, the implicit assumption here is that the signal considered is just a one-period extract of an infinitely long, periodic signal. Consequently, the spectrum of a seemingly simple "general trend"-function on a short interval, say $f : [0, 1] \mapsto \mathbb{R} : f(t) = t$, contains an important amount of high-frequency harmonics, since we are effectively looking at the spectrum of the periodic sawtooth function. Conversely in time domain, we realize that at the endpoints of the domain, the periodized function is discontinuous, thus severely affecting the H^1 smoothing term.

There are two remedies to this. Ideally, one should exclude the boundaries of the domain in the evaluation of the smoothness, i.e. restrict its evaluation to the open interval $(0, 1)$. However, this clearly breaks the Parseval/Plancherel Fourier isometry and the whole beauty of the spectral solution is lost. Therefore, we suggest a less far-reaching remedy, that is classically used in short-term Fourier analysis: smooth windowing. This approach is particularly useful in cases, where the variational mode decomposition is anyway performed on short chunks of a much longer time series signal.

For simplicity, in the following examples, we will use a Gaussian window. This window is simply applied to the input signal f prior to performing the VMD algorithm. After decomposition, the individual modes can be "unwindowed" by simple division. This, however, will largely affect reconstruction fidelity close to the window borders. This is particularly apparent in the single frame decomposition. In a sliding window short-time analysis of a larger time series signal, however, instead of window division, the modes can be stitched together by simple addition without error amplification.

CHAPTER 4

Experiments and Results

In this chapter, we apply the proposed VMD algorithm to a series of test signals in order to assess the validity of our approach. First, we focus on a few problems that have been successfully employed for highlighting the strengths and shortcomings of the EMD / Hilbert-Huang-Transform, namely tones versus sampling, and tones separation [RFG03]. Then we briefly investigate noise robustness of VMD. Finally, we shift our attention to more complex signals, which have already been used in [HS11] and [Gil13].

4.1 Tones and sampling

When the input signal $f = f_\nu(t) = \cos(2\pi\nu t)$ is composed of a pure harmonic, then the mode decomposition is expected to output exactly this harmonic. As reported in [RFG03], this does not happen to be the case with EMD, since the local extrema can suffer from important jittering with increasing frequency. In [RFG03], the relative error

$$e(\nu) = \|f_\nu(t) - u_1(t)\|_2 / \|f_\nu(t)\|_2 \tag{4.1}$$

was introduced, and a quadratic increase with frequency of an upper bound to this relative error was reported for EMD. Further, EMD has pronounced spikes of near-perfect reconstruction when the sampling frequency is an even multiple of the tone's frequency.

Here, we perform this analysis for the proposed VMD model. We refrain from windowing and consider exactly the same signals as for the EMD analysis. The results for different convergence tolerance levels ϵ are shown in figure 4.1. It can be clearly seen that the relative reconstruction error is largely independent of the harmonic's frequency. Moreover, the relative error is nicely controlled by the tolerance level ϵ.

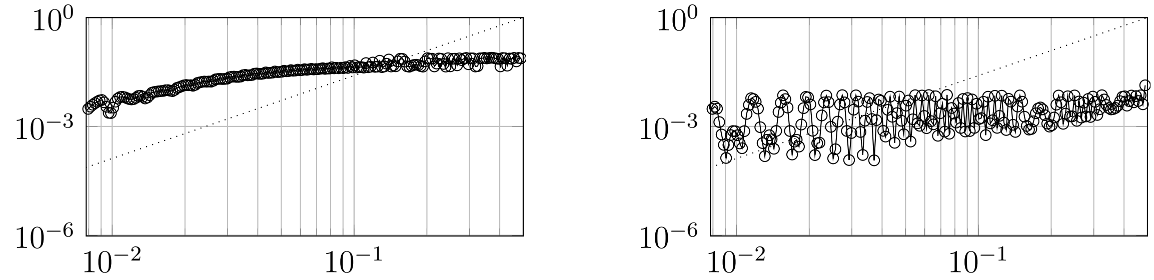

a) $\epsilon = 1e - 3$ b) $\epsilon = 1e - 5$

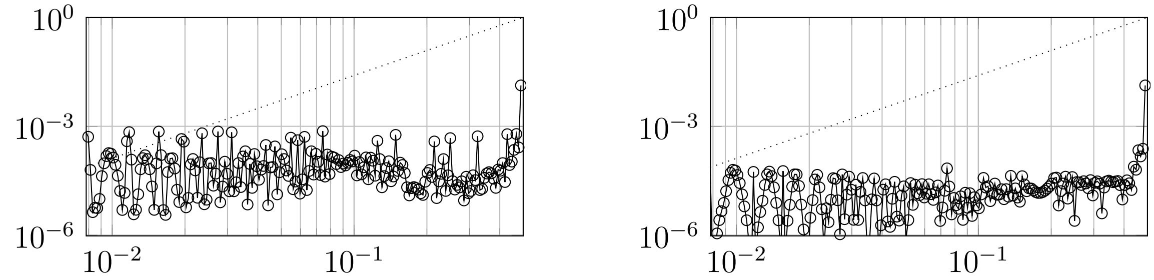

c) $\epsilon = 1e - 7$ d) $\epsilon = 1e - 9$

Figure 4.1: Mode decomposition of a pure harmonic: Relative error for a range of 257 frequencies, for different convergence tolerance levels ϵ. The relative error does not correlate with the tone frequency. Further, reconstruction error can be controlled by decreasing the stopping criterion's convergence tolerance, except for frequencies very close to the Nyquist frequency. In contrast, EMD's relative tone reconstruction error is bounded by a quadratic increase with frequency (dotted line) [RFG03].

4.2 Tones separation

The next slightly more complicated challenge is the separation of two different superimposed tones [RFG03]. Here, the input signal is composed of two different, pure harmonics:

$$f_{\nu_1,\nu_2}(t) = a_1 \cos(2\pi\nu_1 t) + a_2 \cos(2\pi\nu_2 t), \tag{4.2}$$

with $\nu_2 < \nu_1 < f_s/2$, and $a_{1,2}$ two possibly different amplitudes. As a function of the amplitude ratio $\rho = a_1/a_2$, EMD exhibits different, important regions of confusion, where the two signals are too close in frequency to be separated correctly, as reported in [RFG03] and illustrated in figure 4.2.

Again, we apply the same analysis to the proposed VMD model, and again we do not employ any windowing. The results for varying amplitude ratios $\rho \in \{1/4, 1, 4\}$ are shown in figure 4.2 along with the corresponding EMD results. As can be clearly seen, the proposed VMD achieves good tones separation over the whole domain except at the Nyquist frequency. In particular, the decomposition quality is not significantly worse for close harmonics.

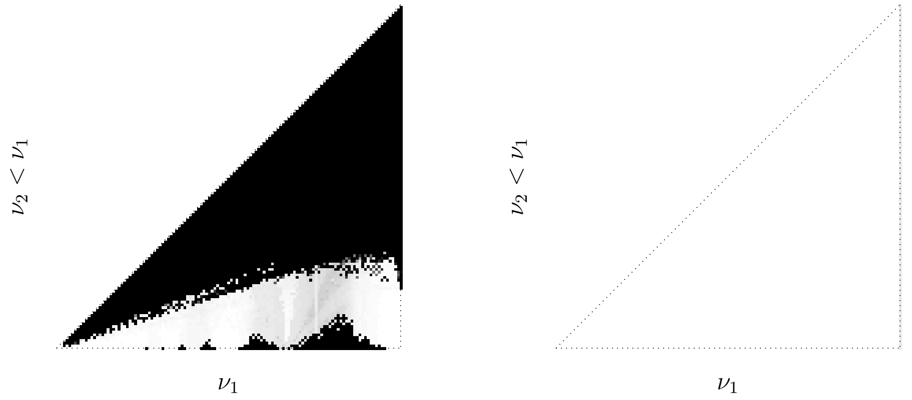

a) EMD, $\rho = 1/4$ b) VMD, $\rho = 1/4$

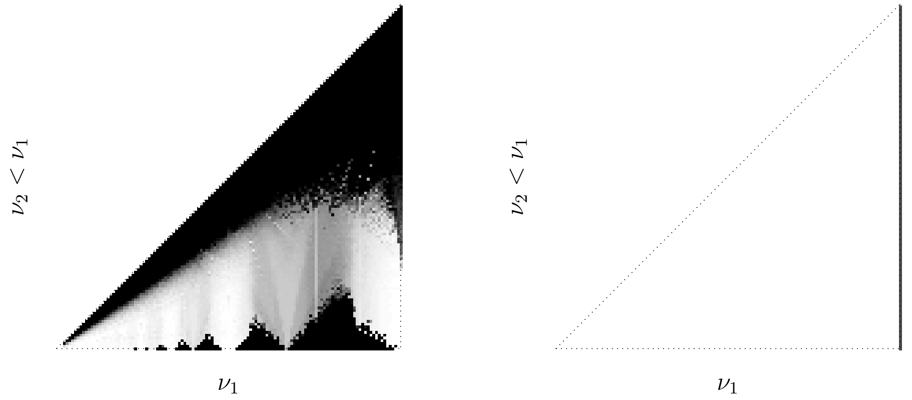

c) EMD, $\rho = 1$ d) VMD, $\rho = 1$

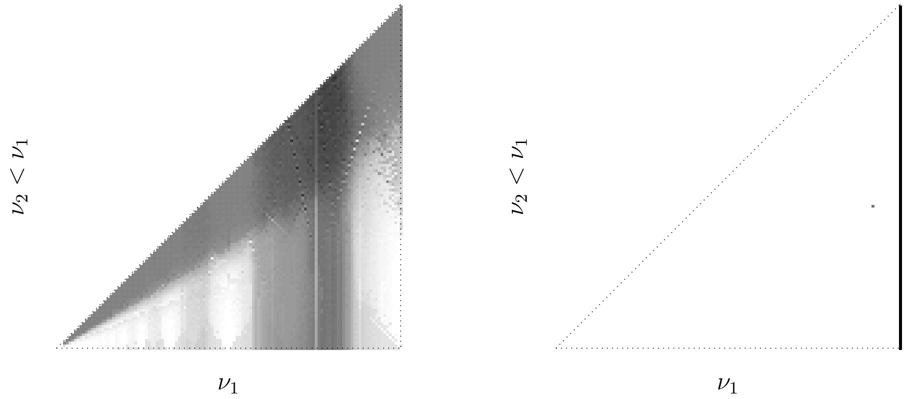

e) EMD, $\rho = 4$ f) VMD, $\rho = 4$

Figure 4.2: Tones separation. In a superposition of two tones of frequencies $\nu_2 < \nu_1 < f_s/2$ and equal amplitudes, the mode decompositions between EMD and VMD vary significantly. The plot indicates relative error, with values between 0 (white) and 0.5 (black). a,c,e) EMD has important areas of confusion (dark), where the tones cannot be separated correctly [RFG03]. b,d,f) In contrast, VMD achieves good tones separation almost everywhere but for ν_1 too close to the Nyquist frequency.

4.3 Noise robustness

To illustrate the VMD robustness with respect to noise in the input signal, we test using the following tri-harmonic signal, affected by noise:

$$f_n(t) = \cos(4\pi t) + \frac{1}{4}\cos(48\pi t) + \frac{1}{16}\cos(576\pi t) + 0.1\eta, \tag{4.3}$$

where $\eta \sim \mathcal{N}(0,1)$ represents the Gaussian additive noise. The noise level is quite important with respect to the amplitude of the highest harmonic. We perform variational modes decomposition into three modes, without Lagrangian multipliers in order to remove the noise. The signal, and the three components estimated using VMD are shown in Fig. 4.3. The strong, lowest frequency signal is recovered almost flawlessly. The medium-strength medium-frequency signal is still detected at acceptable quality. The weak, high-frequency signal, however, is difficult. The VMD algorithm correctly tunes the third center-frequency on this harmonic, but the recovered mode is highly affected by the noise. Here, decreasing the bandwidth by increasing α comes at the risk of not properly capturing the correct center frequency, while too low an α includes more noise in the estimated mode. The mode could, however, be cleaned further in post-processing. For reference, we note that the estimated VMD center frequencies are off by 0.27%, 1.11% and 0.18% only.

We provide a comparison with EMD[1] based on exactly the same signal in Fig. 4.4. The EMD produces 7 estimated modes. The first two modes contain the highest-frequency harmonic, and important amounts of noise. The forth mode comes closest to the middle harmonic, however important features have been attributed to the third and fifth mode. The sixth mode picks up most of the low frequency harmonic, but is severely distorted.

---

[1] Implementation by Gabriel Rilling, available at `http://perso.ens-lyon.fr/patrick.flandrin/emd.html`

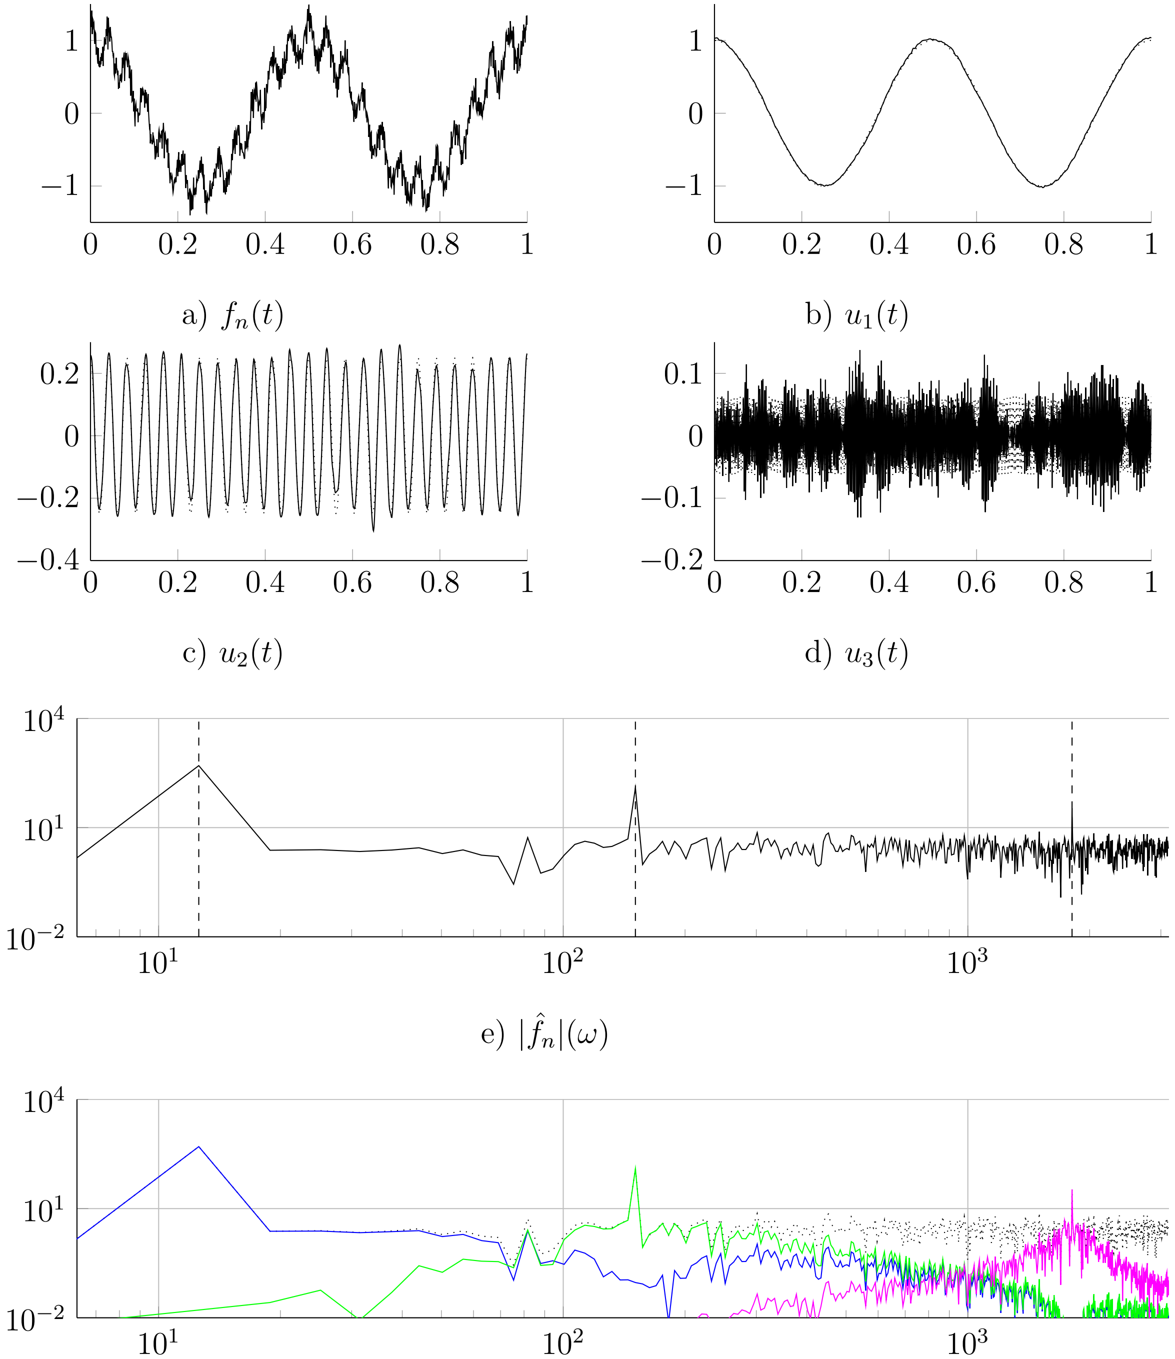

a) $f_n(t)$

b) $u_1(t)$

c) $u_2(t)$

d) $u_3(t)$

e) $|\hat{f}_n|(\omega)$

e) $|\hat{u}_k|(\omega)$

Figure 4.3: VMD decomposition of noisy tri-harmonic. (a) The noisy input signal. (b)-(d) The three modes extracted by denoising VMD, and the theoretical mode (dotted). (e) The spectrum of the input signal, and (f) its distribution over the three modes.

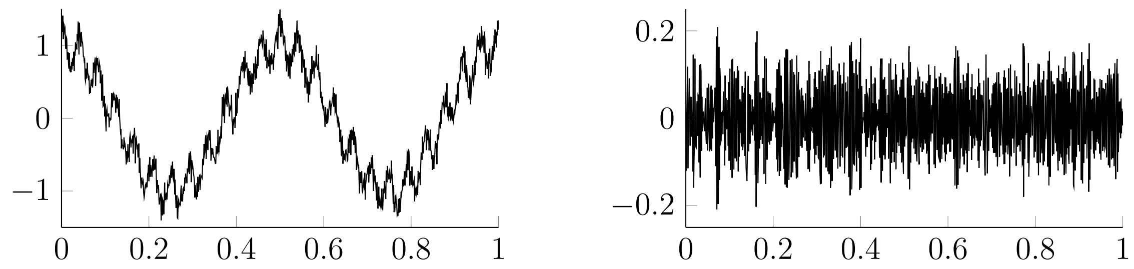

a) $f_n(t)$

b) $u_1(t)$

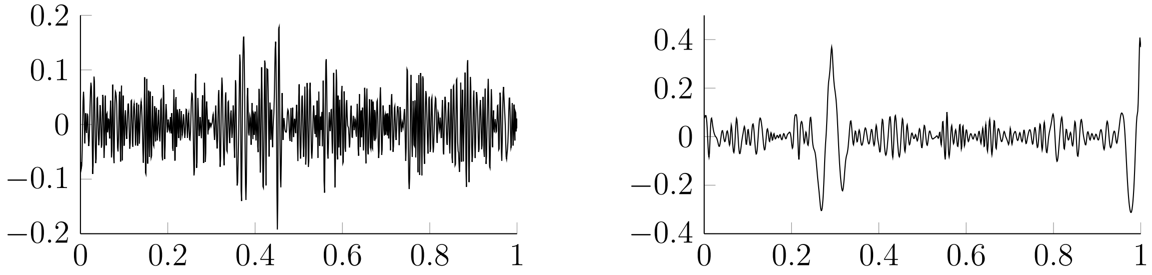

c) $u_2(t)$

d) $u_3(t)$

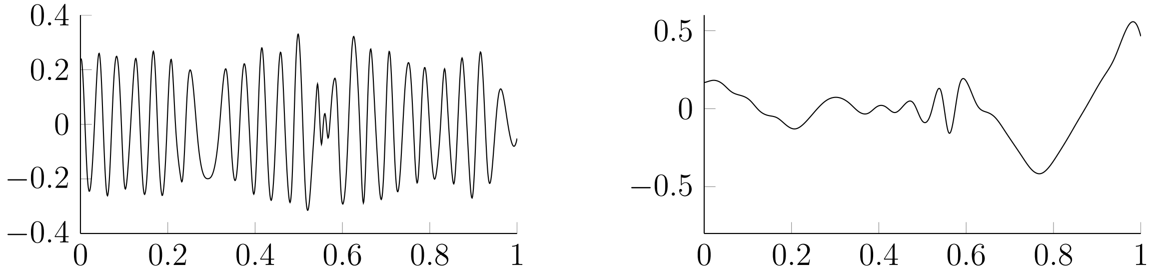

e) $u_4(t)$

f) $u_5(t)$

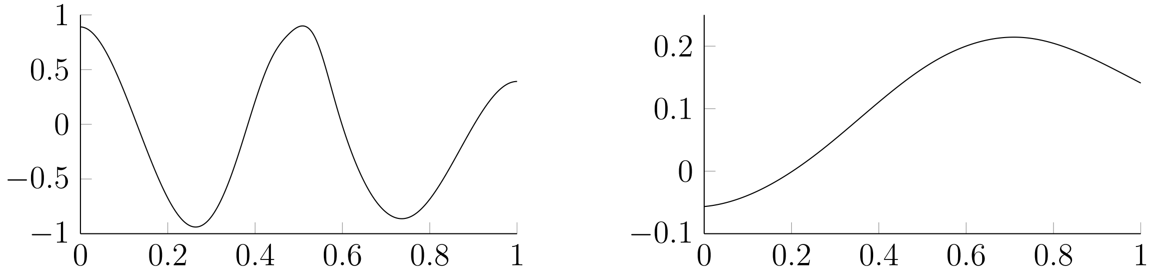

g) $u_6(t)$

h) $u_7(t)$

Figure 4.4: EMD decomposition of noisy tri-harmonic. (a) noisy input signal. (b)-(h) The seven modes extracted by EMD. None of the modes corresponds to a pure harmonic.

4.4 Complex multimode signals

Now we look at slightly more complex signals to be decomposed. In particular, we consider the same test signals that were previously suggested in [HS11] and also used in [Gil13], with the purpose of increased comparability.

4.4.1 General Linear Growth with Two Harmonics

The first signal is a composition of three simple components, namely a general linear trend and two different harmonics:

$$f_{Sig1}(t) = 6t + \cos(8\pi t) + 0.5\cos(40\pi t). \tag{4.4}$$

The signal, its three constituent modes, and the composite Fourier spectrum are shown in figure 4.5. The main challenge of this signal is the linear growth term. Without windowing, the higher order harmonics of the periodized sawtooth signal spread over the whole spectrum.

In order to reduce the effects of periodization, we apply Gaussian windowing. The corresponding windowed signal, and the respective VMD results are illustrated in detail in figure 4.6. In particular, we show how the two non-zero center frequencies ω_2 and ω_3 quickly converge towards the exact harmonics. The corresponding modes constitute a nice partition of the input spectrum, with each mode being clearly dominant around its respective center frequency. The three modes in time domain show nice separation into three distinct signals of characteristic oscillations. After “unwindowing” by pointwise division of the estimated modes by the Gaussian window, we recover good estimates of the true underlying modes (dotted lines), valid on the central 60% of the signal.

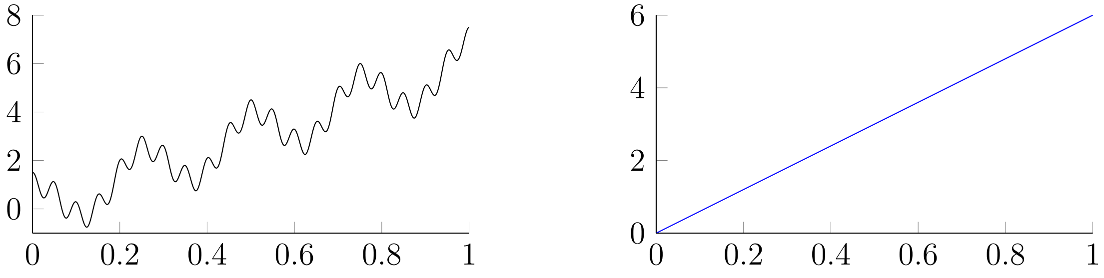

a) $f_{Sig1}(t)$ b) $6t$

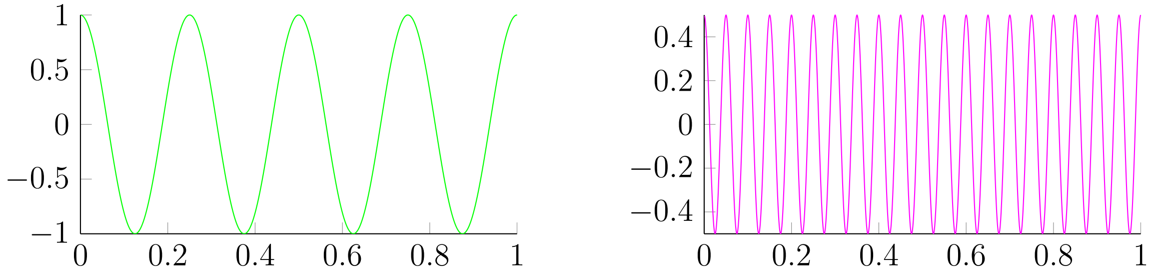

c) $\cos(8\pi t)$ d) $0.5\cos(40\pi t)$

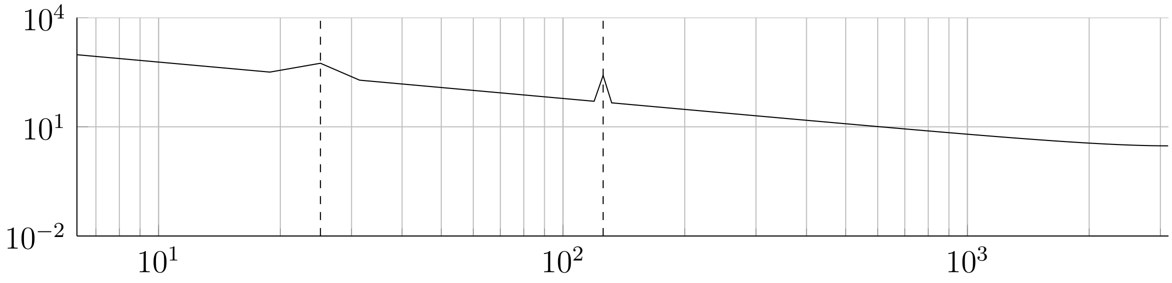

e) $|\hat{f}_{Sig1}|(\omega)$

Figure 4.5: a) $f_{Sig1}(t)$, b–d) its constituent modes. e) The signal's spectrum.

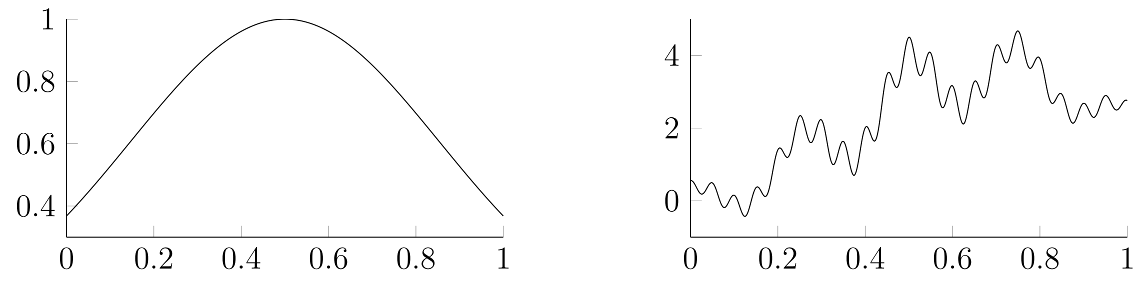

a) $w(t)$

b) $f^{w}_{Sig1}(t)$

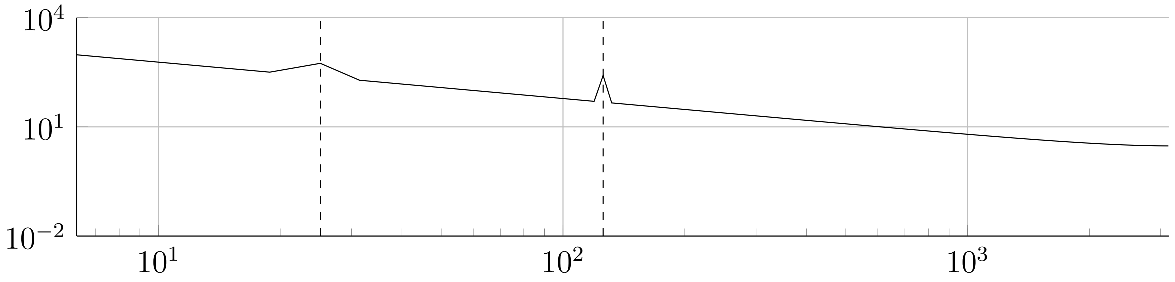

c) $|\hat{f}^{w}_{Sig1}|(\omega)$

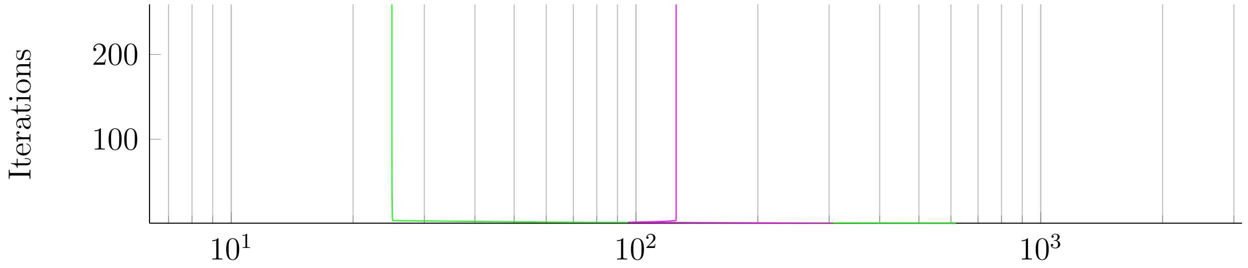

d) ω_k

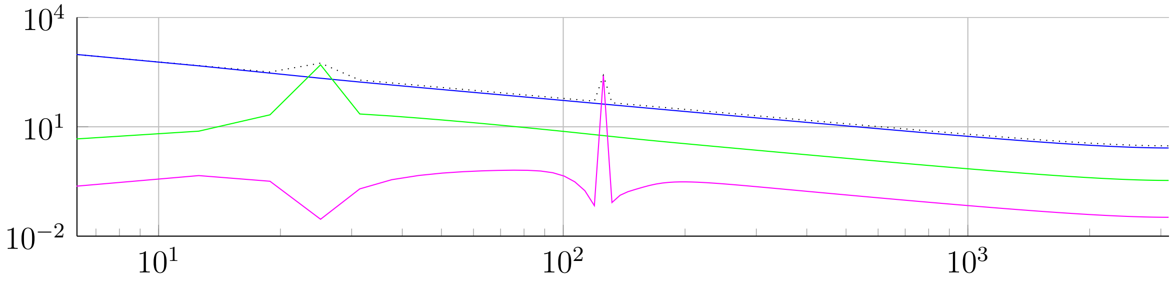

e) $|\hat{u}^{w}_{k}|(\omega)$

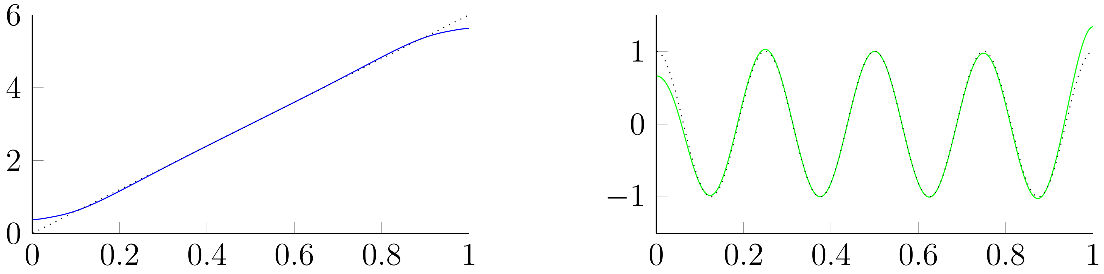

f) $u_1^w(t)$ g) $u_2^w(t)$

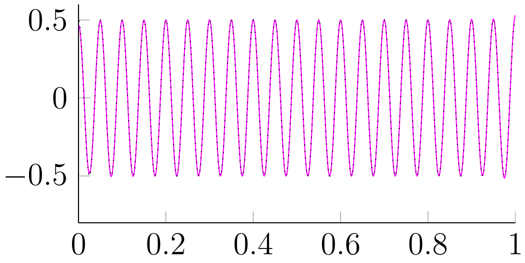

h) $u_3^w(t)$

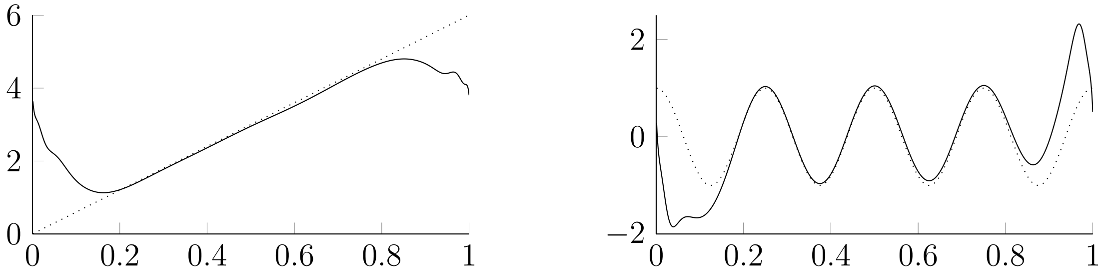

i) $u_1(t)$ j) $u_2(t)$

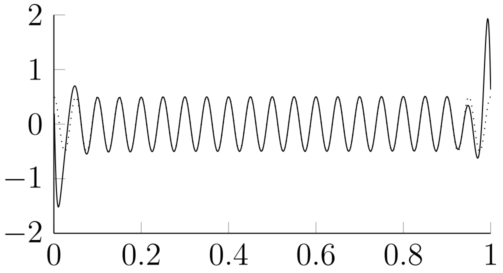

k) $u_3(t)$

Figure 4.6: VMD decomposition of f_{Sig1}. a) The applied window, b) the windowed signal, and c) its spectrum. d) Evolution of the detected center frequencies, and e) the corresponding spectrum decomposition. f–h) the reconstructed modes prior to, and i–k) after Gaussian window removal.

4.4.2 General Quadratic Growth with Chirp Signal and Piecewise Harmonics

The second example uses a quadratic trend, a chirp signal, and a third mode with sharp transition between two constant frequencies[2]:

$$f_{Sig2}(t) = 6t^2 + \cos(10\pi t + 10\pi t^2) + \begin{cases} \cos(60\pi t) & t \leq 0.5 \\ \cos(80\pi t - 10\pi) & t > 0.5 \end{cases} \tag{4.5}$$

The signal, its three constituent modes, and the composite Fourier spectrum are shown in figure 4.7. The instantaneous frequency of the chirp is given by the time derivative of its phase:

$$\omega(t) := \partial_t \phi(t) = 10\pi + 20\pi t. \tag{4.6}$$

Thus, for $t \in [0, 1]$ the instantaneous frequency varies linearly between 10π an 30π. Consequently, the theoretical center frequency of the mode is located at 20π. The piecewise-constant bi-harmonic has spectral peaks expected at 60π and 80π.

Here, too, we employ Gaussian windowing to alleviate periodization artifacts. Indeed, the windowed signal has a much cleaner spectrum, and the expected peaks of the signal's components become more prominent, as illustrated in figure 4.8. Again, the estimated center frequencies ω_k converge to the expected frequencies precisely. Here, we chose to decompose into four modes, thus assigning each half of the piecewise-constant frequency signal to a separate mode. The spectral partitioning can be nicely appreciated in the spectral plot of the different modes. The unwindowed mode estimates fit well the theoretical signals, except again for boundary issues.

---

[2] Here, we changed the phase shift in the third component, with piecewise-constant frequency, from 15π to 10π, in order to have a continuous signal.

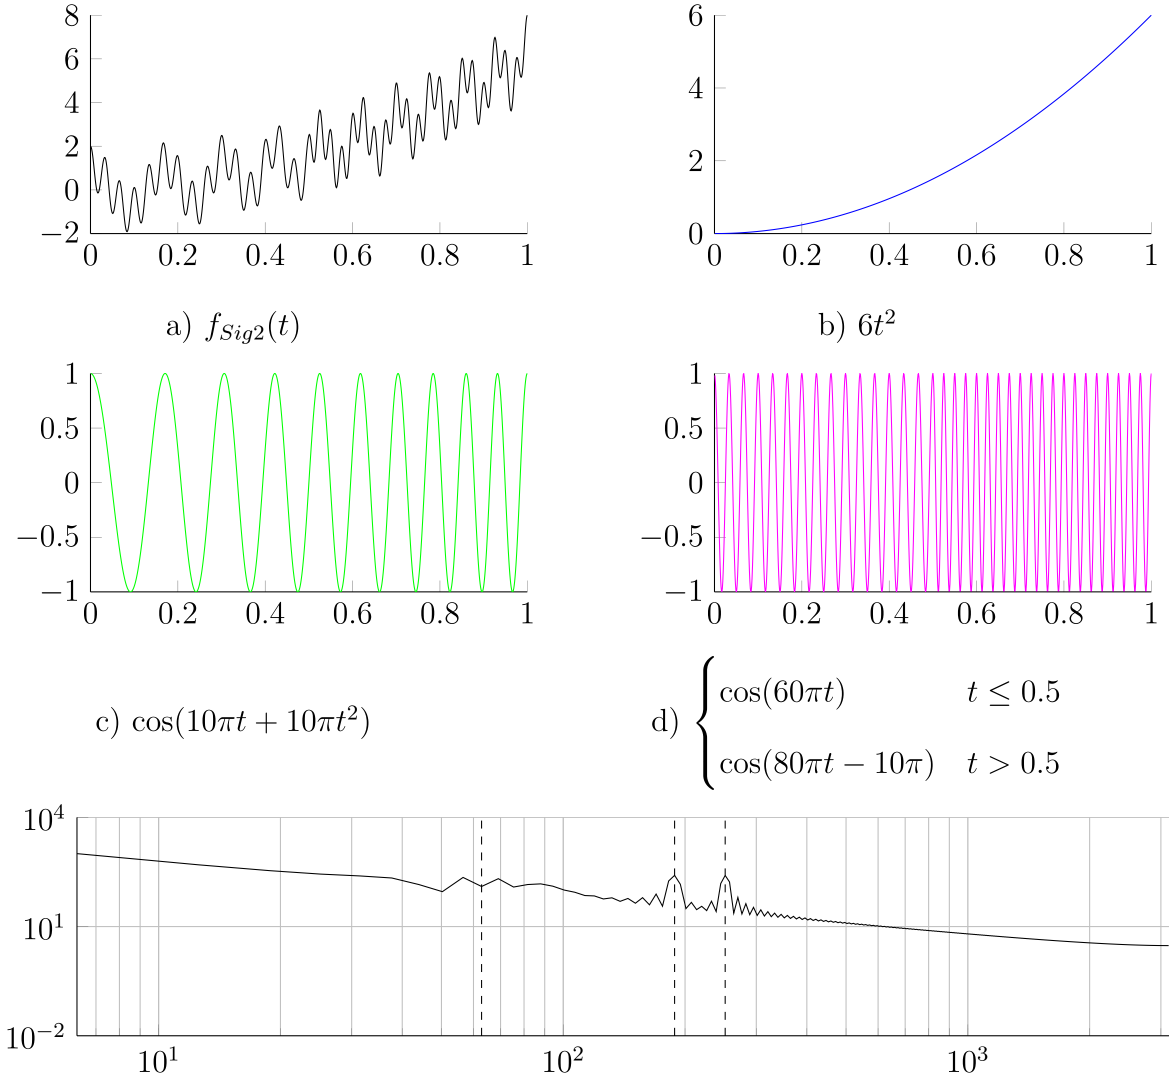

a) $f_{Sig2}(t)$

b) $6t^2$

c) $\cos(10\pi t + 10\pi t^2)$

d) $\begin{cases} \cos(60\pi t) & t \leq 0.5 \\ \cos(80\pi t - 10\pi) & t > 0.5 \end{cases}$

e) $|\hat{f}_{Sig2}|(\omega)$

Figure 4.7: a) f_{Sig2}, b–d) its constituent modes. e) The signal's spectrum.

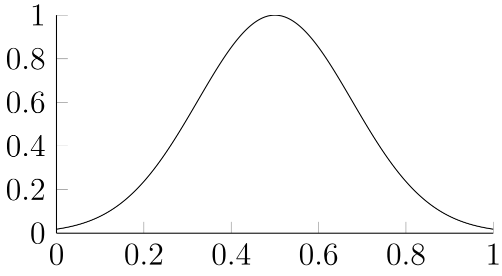

a) $w(t)$

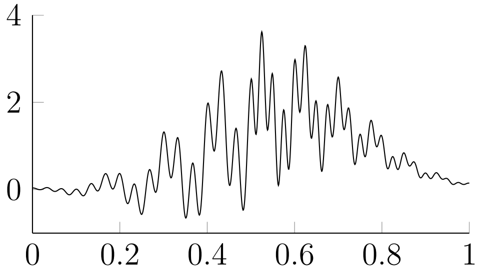

b) $f^{w}_{Sig2}(t)$

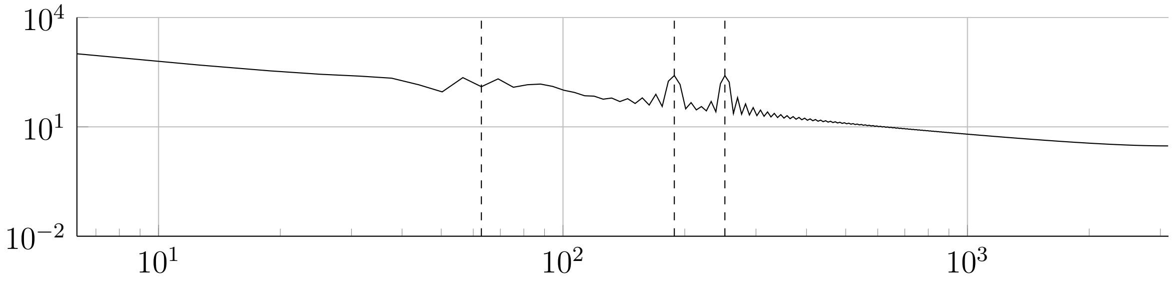

c) $|\hat{f}^{w}_{Sig2}|(\omega)$

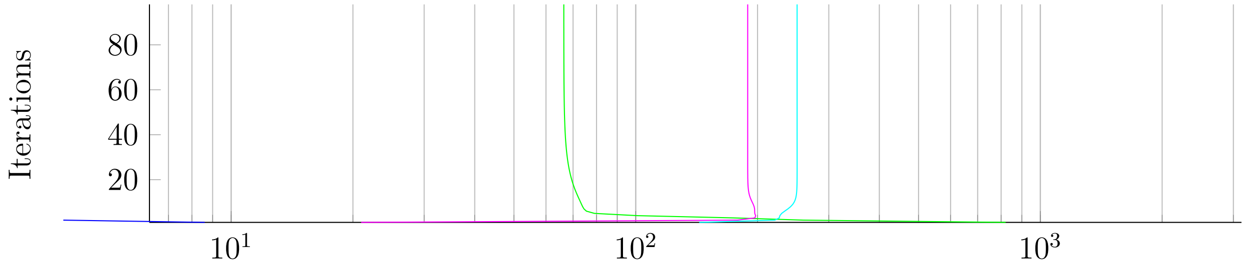

d) ω_k

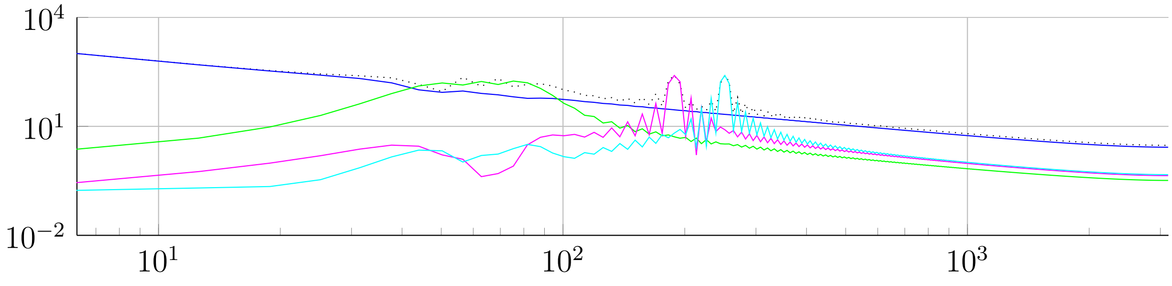

e) $|\hat{u}^{w}_{k}|(\omega)$

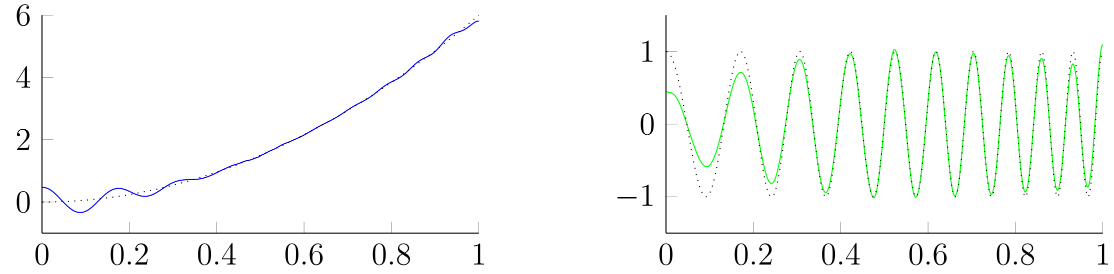

f) $u_1^w(t)$ g) $u_2^w(t)$

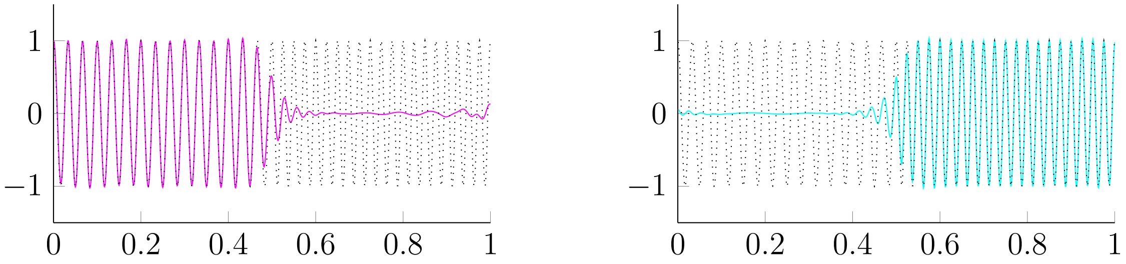

h) $u_3^w(t)$ i) $u_4^w(t)$

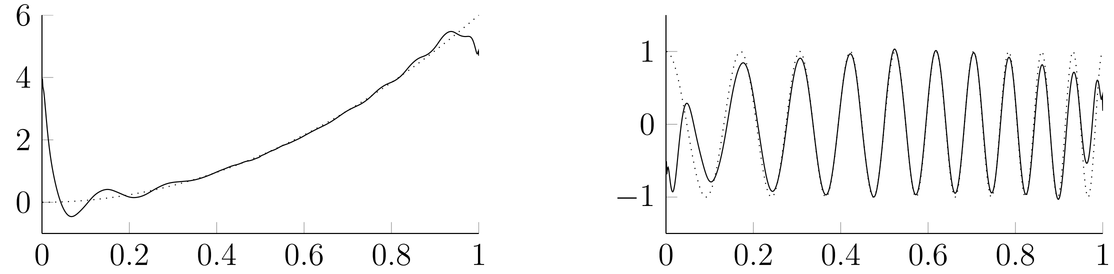

j) $u_1(t)$ k) $u_2(t)$

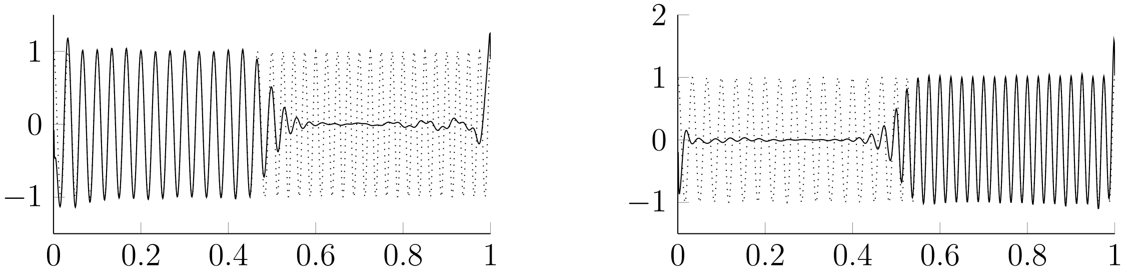

l) $u_3(t)$ m) $u_4(t)$

Figure 4.8: Results of VMD on f_{Sig2}. a) The applied window, b) the windowed signal, and c) its spectrum. d) Evolution of the detected center frequencies, and e) the corresponding spectrum decomposition. f–i) the reconstructed modes prior to, and j–m) after Gaussian window removal.

4.4.3 Intrawave Frequency Modulation

The third synthetic signal has intrawave frequency modulation:

$$f_{Sig3}(t) = \frac{1}{1.2 + \cos(2\pi t)} + \frac{\cos(32\pi t + 0.2\cos(64\pi t))}{1.5 + \sin(2\pi t)}. \tag{4.7}$$

The signal, its three constituent modes, and the composite Fourier spectrum are shown in figure 4.9. While the first, bell-shaped component has mostly low-pass content, the second mode's main peak is clearly identified at 32π. However, due to the non-linear intrawave frequency modulation, an important amount of higher-order harmonics are also observed at $32\pi + 64\pi = 96\pi$, $32\pi + 2 \cdot 64\pi = 160\pi$ and $32\pi + 3 \cdot 64\pi = 224\pi$, respectively. This second component obviously violates the narrowband assumption, and one would naturally expect some difficulties recovering this mode using VMD. Indeed, by Carson's rule, the mode's bandwidth here is dominantly controlled by the relatively high frequency of the modulating term $\cos(64\pi t)$, essentially spreading the mode over the whole practical spectrum.

The slightly windowed signal and the corresponding VMD results are illustrated in figure 4.10. The non-zero ω_2 quickly converges to the correct main frequency 32π. The higher order harmonics are not uniquely attributed to the second mode, but shared between both modes. Consequently, the intrawave frequency modulation is shared by both modes, creating some ripples in the otherwise low-frequency mode. Nonetheless, the reconstructed estimated modes fit well the constituent signals (dotted lines). Most of the error occurs at the boundaries, and at the very center of the signal, where the low-frequency mode has a sharp peak, involving some higher frequency features wrongly attributed to the higher-frequency mode.

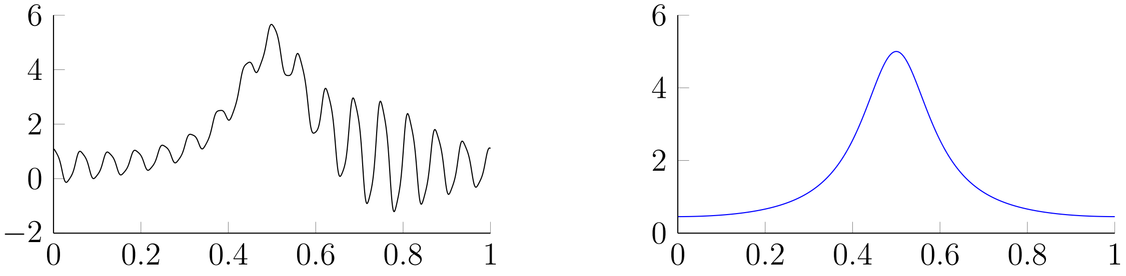

a) $f_{Sig3}(t)$ b) $(1.2+\cos(2\pi t))^{-1}$

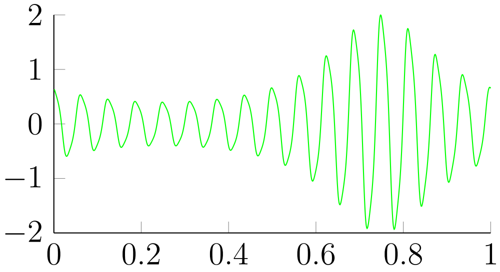

c) $(1.5+\sin(2\pi t))^{-1}\cos(32\pi t+0.2\cos(64\pi t))$

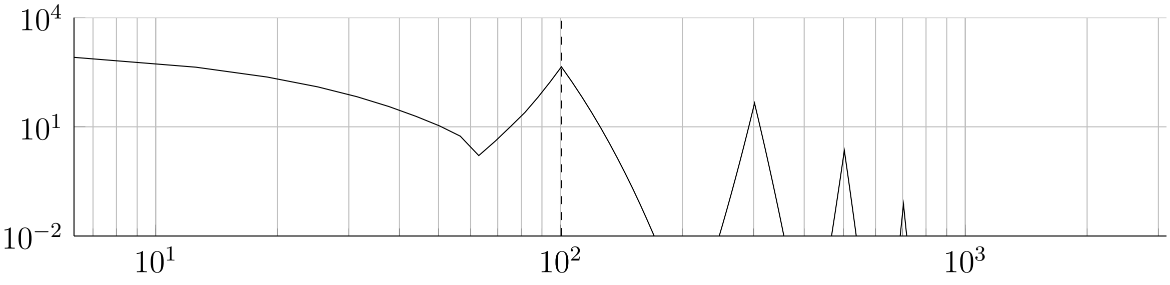

d) $|\hat{f}_{Sig2}|(\omega)$

Figure 4.9: a) f_{Sig3}, b–c) its constituent modes. d) The signal's spectrum.

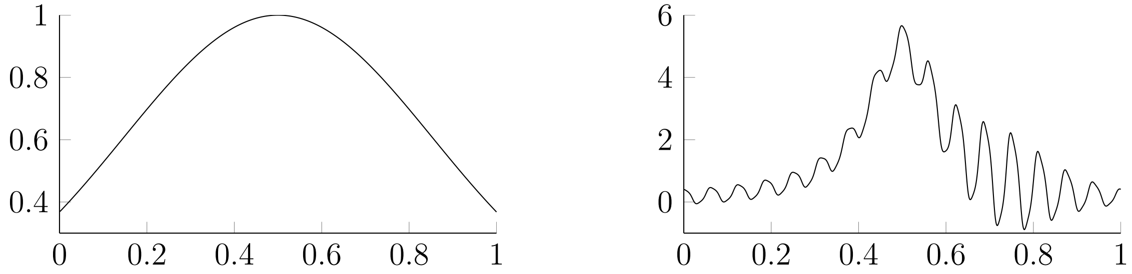

a) $w(t)$ b) $f^{w}_{Sig3}(t)$

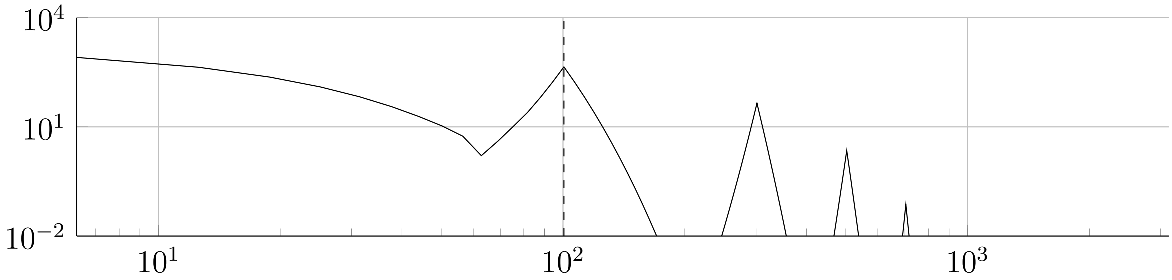

c) $|\hat{f}^{w}_{Sig3}|(\omega)$

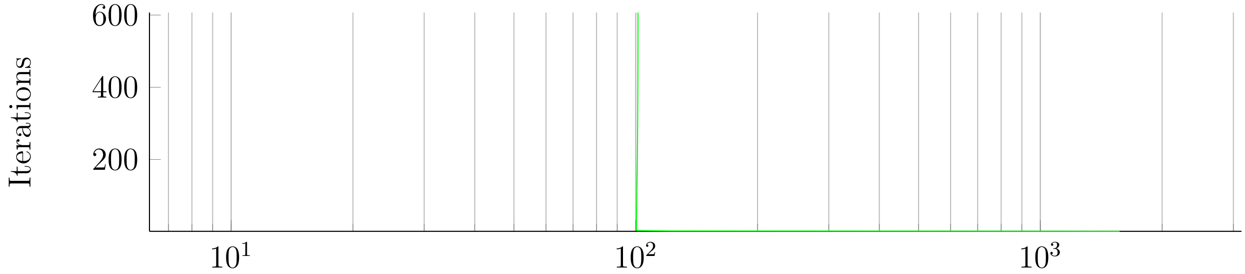

d) ω_k

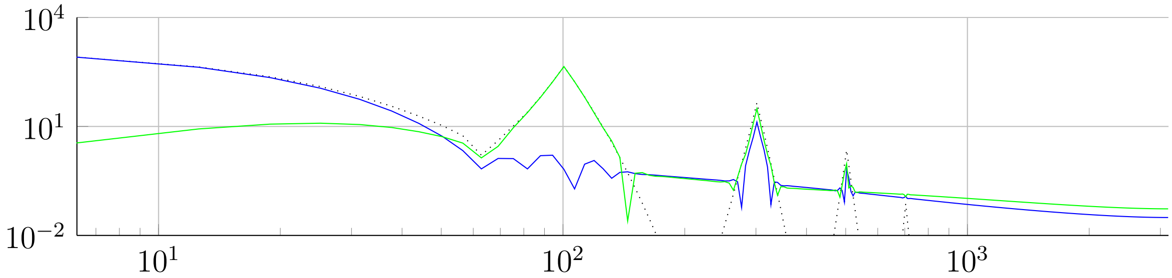

e) $|\hat{u}^{w}_{k}|(\omega)$

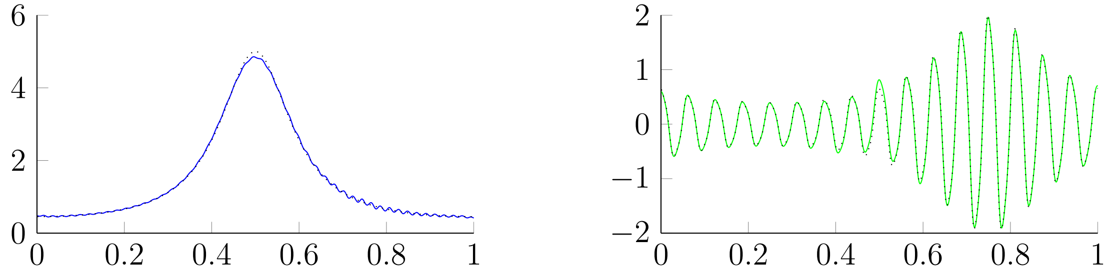

f) $u_1^w(t)$

g) $u_2^w(t)$

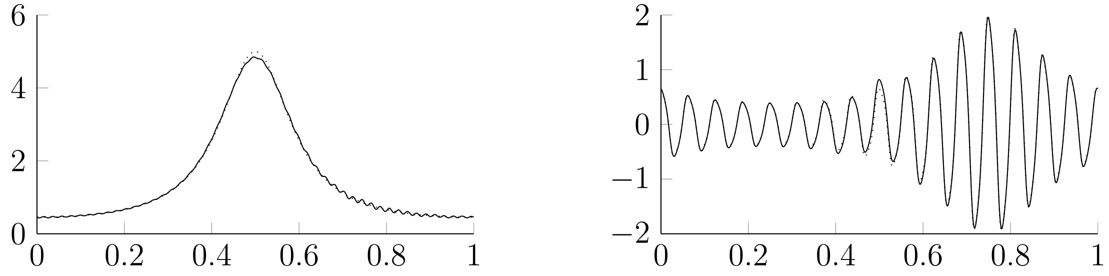

h) $u_1(t)$

i) $u_2(t)$

Figure 4.10: Results of VMD on f_{Sig3}. a) The applied window, b) the windowed signal, and c) its spectrum. d) Evolution of the detected center frequencies, and e) the corresponding spectrum decomposition. f–g) the reconstructed modes prior to, and h–i) after Gaussian window removal.

4.4.4 Electrocardiogram Signal

The fourth example is a real signal from an electrocardiogram (ECG), data shared by [Gil13]. These data present numerous components, as seen in figure 4.11. Beyond the expected spikes-train driven by the rhythm of the heartbeat, one can clearly see an oscillating low-frequency pattern. At the other end of the spectrum, there is distinct high-frequency noise at a single high-pitch harmonic, most likely the electric power-line frequency. The distinct spikes of the ECG signal create important higher-order harmonics.

The spectrum after slight Gaussian windowing, and the results of VMD are depicted in figure 4.12. We chose a high-number of 10 modes to be detected, to accommodate the numerous higher-order harmonics of the spikes. The respective center frequencies nicely converge to these spectral peaks. The first, low-frequency mode captures the low-frequency oscillation of the baseline. The highest frequency mode contains the most noise. The first actual ECG specific mode oscillates precisely at the frequency of the heartbeat. The higher ECG modes then contain the higher-order wave-packages around the highly non-sinusoidal spikes. A "clean" ECG signal can be reconstructed by summing all but the first and last VMD modes, thus discarding the low-frequency baseline oscillation and most of the high-frequency noise.

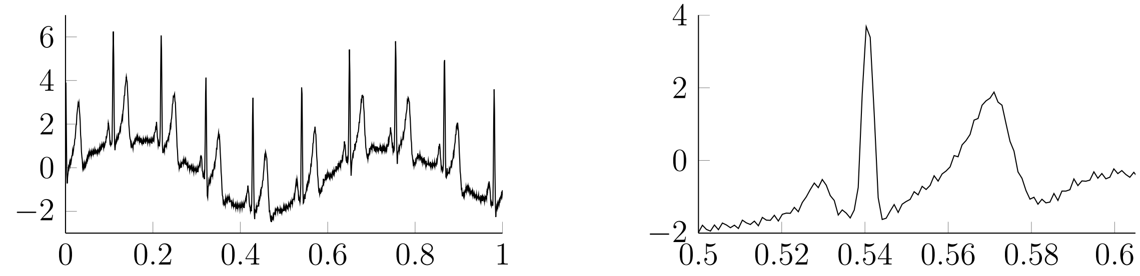

a) $f_{ECG}(t)$ b) $f_{ECG}(t)$ (Detail)

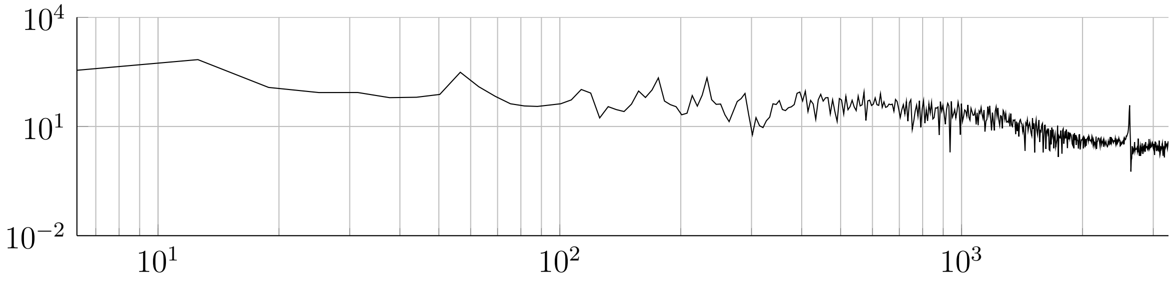

c) $|\hat{f}_{ECG}|(\omega)$

Figure 4.11: a) ECG signal 7. b) Detail. c) The signal's spectrum.

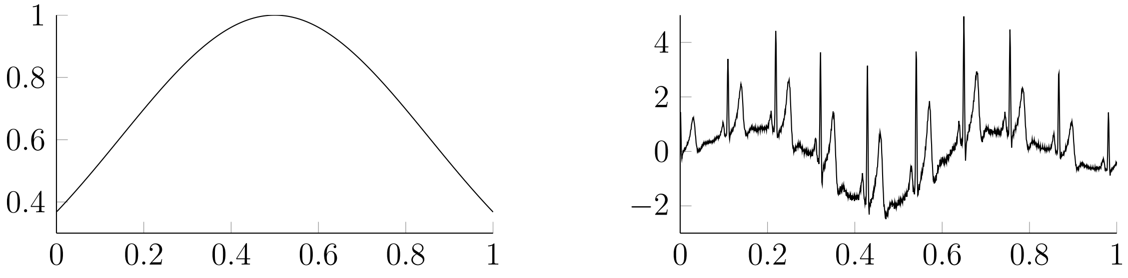

a) $w(t)$ b) $f^w_{ECG}(t)$

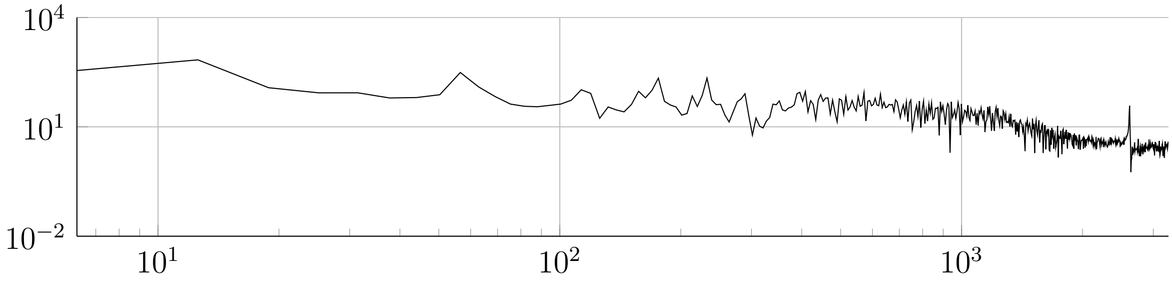

c) $|\hat{f}^w_{ECG}|(\omega)$

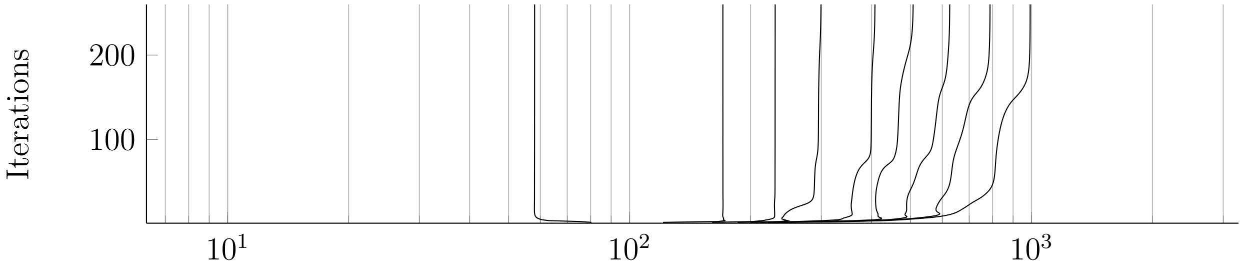

d) ω_k

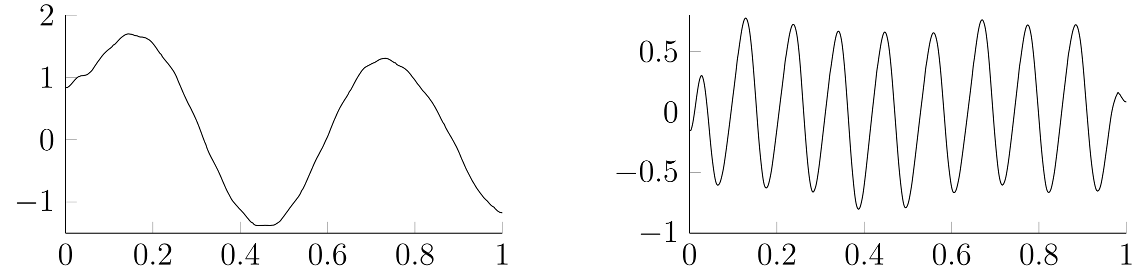

e) $u_1^w(t)$ f) $u_2^w(t)$

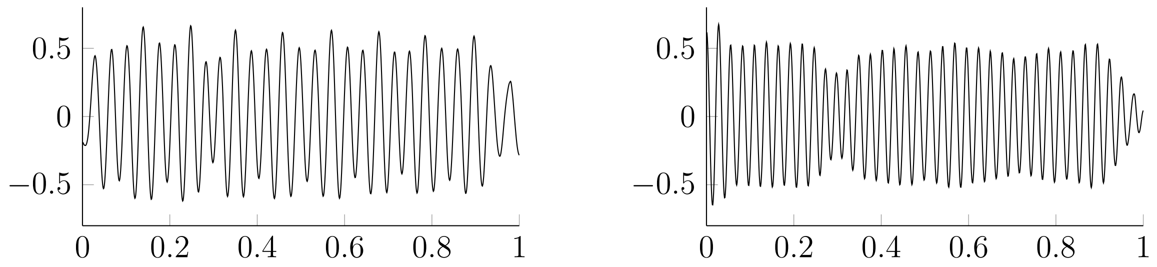

g) $u_3^w(t)$ h) $u_4^w(t)$

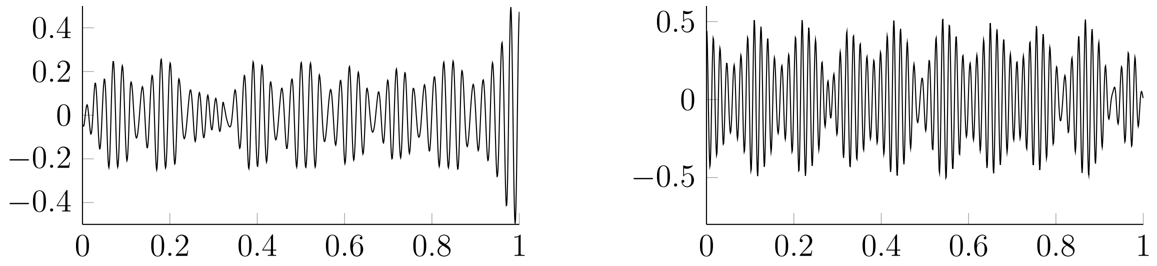

i) $u_5^w(t)$ j) $u_6^w(t)$

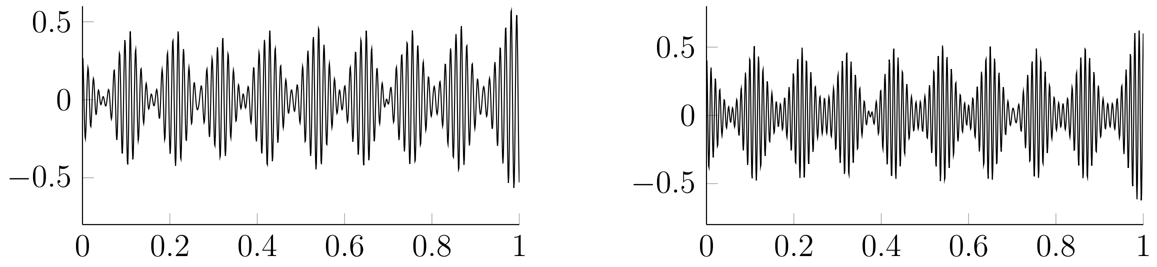

k) $u_7^w(t)$ l) $u_8^w(t)$

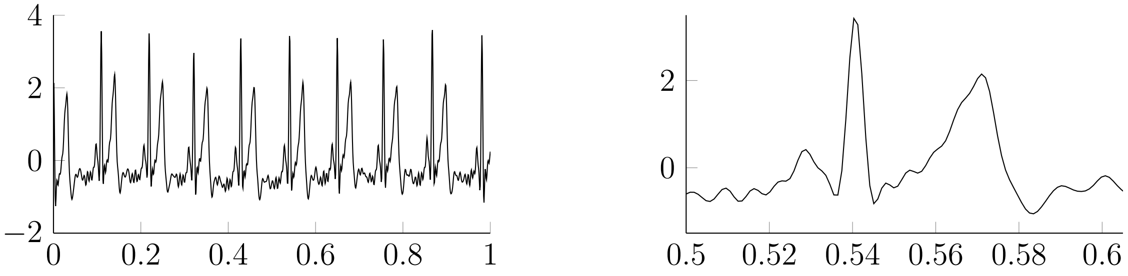

m) $\sum_{k=2}^{8} u_k(t)$

n) $\sum_{k=2}^{8} u_k(t)$ (Detail)

Figure 4.12: Results of ECG signal 7. a) The applied window, b) the windowed signal, and c) its spectrum. d) Evolution of the detected center frequencies. e–l) The reconstructed modes prior to Gaussian window removal. m) Cleaned ECG, and n) detail.

CHAPTER 5

Conclusions and Outlook

In Part I, we have presented a novel variational method for decomposing a signal into an ensemble of band-limited intrinsic mode functions, that we call Variational Mode Decomposition, (VMD). In contrast to existing decomposition models, like the empirical mode decomposition (EMD), we refrain from modeling the individual modes as signals with explicit IMFs. Instead, we replace the most recent definition of IMFs, namely their characteristic description as AM-FM signals, by the corresponding narrow-band property. Indeed, we provide a formula that relates the parameters of the explicit AM-FM descriptors to the estimated signal bandwidth.

Our decomposition model solves the inverse problem as follows: decompose a signal into a given number of modes, either exactly or in a least squares sense, such that each individual mode has limited bandwidth. We assess the mode's bandwidth as the squared H^1 norm of its Hilbert complemented analytic signal with only positive frequencies, shifted to baseband by mixing with a complex exponential of the current center frequency estimate. The variational problem is solved very efficiently in a classical ADMM approach: The modes are updated by simple Wiener filtering, directly in Fourier domain with a filter tuned to the current center frequency, then the center frequencies are updated as the center of gravity of the mode's power spectrum, and finally the Lagrangian multiplier enforcing exact signal reconstruction is updated as dual ascent.

In our experiments, we show that the proposed VMD scheme clearly outperforms EMD with regards to tone detection, tone separation, and noise robustness. Further, we apply our model to more complicated signals for comparison with other state-of-the-art methods, and can show successful decomposition.

The most important limitation of the proposed VMD is with boundary effects, and sudden signal onset in general. This is strongly related to the use of an L^2-based smoothness term, that overly penalizes jumps at the domain borders and within; conversely, this is also reflected by implicit periodicity assumptions when optimizing in Fourier domain, and by the narrow-band violation caused by discontinuous envelopes in such AM-FM signals. Another point that critics might highlight, is the required explicit (manual) selection of the number of active modes in the decomposition, like in EWT but as opposed to EMD. The work in Part II addresses these shortcomings and extends to signals on domains of dimension greater than one.

Part II

Two-Dimensional Compact Variational Mode Decomposition: Spatially compact and spectrally sparse image decomposition and segmentation

CHAPTER 1

Introduction

In Part II, we are interested in decomposing images $f\colon \mathbb{R}^n \to \mathbb{R}$ into ensembles of constituent modes (components) that have specific directional and oscillatory characteristics. Simply put, the goal is to retrieve a small number K of modes $u_k\colon \mathbb{R}^n \to \mathbb{R}$, that each have a very limited bandwidth around their characteristic center frequency ω_k. These modes are called intrinsic mode functions (IMF) and can be seen as amplitude- and frequency-modulated (AM-FM) n-D signals ("plane"-waves). Such a mode can have limited spatial support, its local (instantaneous) frequency and amplitude vary smoothly, several modes can overlap in space, and together the ensemble of modes should reconstruct the given input image up to noise and singular features.

Many fields use signal decomposition as a fundamental tool for quantitative and technical analysis. In remote sensing, decomposing images based on frequency content and signal priors, such as housing lattices and terrain structures, is useful for segmentation, identification, and classification [DLL14]. In oceanography, a combination of baroclinic modes helps model density profiles of seasonal cycles, and other geophysical phenomena such as thermal or solar variation [FPS08, SBC12]. Similarly, in seismology, modes with differing frequency components help highlight different geological and stratigraphic information [HB13]. In holography, mode decomposition allows reducing speckle [LPD14]. In the fields of energy and power engineering, mode decompositions are used for vibration analysis and fault detection, e.g., [GLS13, TZY11]. Multivariate mode decomposition and mode entropy analysis are useful tools in neural data analysis [HL11]. In crystallography, because the crystal lattice exhibits multiple spatial periodicities, interpretable as a superposition of multiple different cosine-waves, we wish to couple several “sub-modes” into a single phase. This coupled-mode decomposition enables robust estimates of mesoscopic properties such as crystal defects, rotations, and grain boundaries. Recent work in crystal orientation detection includes variational methods based on tensor maps in conjunction with a regularization scheme [EW14] and 2D synchrosqueezed transforms [YLY14]. In nanoscale imaging, segmentation enables analyses and comparisons of surface regions of different structures as well as directed measurements of function, spectra, and dynamics [SMH05, TSH15]. Ultimately, efficient segmentation will enable directed data acquisition and parsing acquisition time between different modalities to assemble and to converge complementary structural, functional, and other information.

Independent of the scientific discipline, sparse signal decomposition provides expansive utility and a more advanced podium from which to elucidate greater understanding.

1.1 Recent and related work

The problem is inspired by the one-dimensional empirical mode decomposition (EMD) algorithm [HSL98] and its more recent derivates, such as [FGR05, HS11, HS13, HS15, MP07, RF08, RF09, RFG03, SV05, THS14, WH09]. We are interested in the two-dimensional (2D) analogs and extensions of such decomposition problems. The 2D extension of EMD [NBD03] similarly uses recursive sifting of 2D spatial signals by means of interpolating upper and lower envelopes, median envelopes, and thus extracting image components in different “frequency” bands. This 2D-EMD, however, suffers from the same drawbacks in robustness as the original EMD in extremal point finding, interpolation of envelopes, and stopping criteria imposed. More recent work, such as the Prony-Huang Transform [SPB14], has only partially improved on some of these drawbacks using modern variational and transform methods.

Classical decomposition methods include the discrete Fourier transform (DFT) and the continuous wavelet transform (CWT), where a fixed basis can be used to find a sparse representation. Using more general bases or frames, extended methods such as matching pursuit decomposition (MP), method of frames, best orthogonal basis (BOB), and basis pursuit (BP) are more robust and, in principle, decompose a signal into an "optimal" superposition of dictionary elements. Though these methods have had success with simple signals, they are still not fully robust to non-stationary waves and require a large, redundant dictionary of elements, which are not reflective of the specifics of the given signal.

More specific methods for directional image decomposition work by mostly rigid frames, decomposing the Fourier spectrum into fixed, mostly or strictly disjoint, (quasi-)orthogonal basis elements. Examples include Gabor filters [Tai96], wavelets [Dau88, DV01, Mal89], curvelets [CD99], or shearlets [GL07, LLK05]. These methods are not adaptive relative to the signal, and can attribute principle components of the image to different bands, as well as contain several different image components in the same band. Adaptivity and tuned sparsity concerns have been addressed through synchrosqueezed wavelet transforms [COP12, DLW11, WFD11, YY13], where unimportant wavelet coefficients are removed by thresholding based on energy content. In pursuit of the same goal, the 2D empirical wavelet transform (EWT) [Gil13, GTO14] decomposes an image by creating a more adaptive wavelet basis.

In previous work [DZ14], Dragomiretskiy and Zosso defined a fully variational model for mode decomposition of 1D signals. The so-called *variational mode decomposition* (VMD) in 1D is essentially based on well-established concepts such as Wiener filtering, the 1D Hilbert transform and the analytic signal, and heterodyne demodulation. The goal of 1D-VMD is to decompose an input signal into a discrete number of sub-signals (modes), where each mode has limited bandwidth in the spectral domain. In other words, one requires each mode $u_k\colon \mathbb{R} \to \mathbb{R}$ to be mostly compact around a center pulsation ω_k, which is to be determined along with the decomposition. In order to assess the bandwidth of a mode, the following scheme was proposed: 1) for each mode u_k, compute the associated analytic signal by means of the Hilbert transform in order to obtain a unilateral frequency spectrum. 2) For each mode, shift the mode's frequency spectrum to "baseband", by mixing with an exponential tuned to the respective estimated center frequency. 3) The bandwidth is now estimated through the H^1 smoothness (Dirichlet energy) of the demodulated signal. The resulting constrained variational problem is the following:

$$\min_{u_k\colon \mathbb{R}\to\mathbb{R},\, \omega_k} \left\{ \sum_k \left\| \partial_t \left[\left\{ \left(\delta(\cdot) + \frac{j}{\pi \cdot} \right) * u_k(\cdot) \right\} (t) e^{-j\omega_k t} \right] \right\|_2^2 \right\} \quad \text{s.t.} \quad \forall t \in \mathbb{R}\colon \sum_k u_k(t) = f(t). \tag{1.1}$$

In [DZ14], it was shown that this variational model can be minimized efficiently and it outperforms empirical mode decomposition algorithms in various respects, most notably regarding noise robustness and mode cleanliness.

1.2 Proposed method

Here we propose a natural two-dimensional extension of the (1D) variational mode decomposition algorithm [DZ14] in the context of image segmentation and directional decomposition. The 2D-VMD algorithm is a non-recursive, fully adaptive, variational method that sparsely decomposes images in a mathematically well-founded manner.

Here, we are interested in making the advantages of the variational model accessible for the 2D case (and higher dimensions equally so). The first order of business is thus to generalize the 1D-VMD model to the multidimensional case, as sketched in [DZ15]. Second, we want to address an intrinsic conflict of the VMD model, namely the inverse relation between spatial and frequency support: in 1D-VMD it was noted that the algorithm had difficulties whenever signals exhibited sudden onset and amplitude changes, since these effectively represent a violation of the assumptions of Bedrosian's theorem, a key element of the VMD model. In this work, we address this issue by further introducing a separate amplitude function that masks the underlying mode spatially, which allows decoupling spatial from spectral support. In 2D, this approach allows extraction of modes with sharp boundaries. We then introduce various priors on the shape of the amplitude function. Requiring the amplitude function to be binary and penalizing its total variation regularizes the mode boundaries. Restricting the ensemble of amplitude functions associated with the various modes to the probability simplex at each pixel leads to non-overlapping modes effectively segmenting the image. Coupling several modes to share a single support function further allows extraction of multi-wave textures, such as hexagonal lattice patterns.

CHAPTER 2

Two-dimensional Variational Mode Decomposition

We design the 2D model analogously to its 1D predecessor, minimizing the constituent sub-signals bandwidth while maintaining data fidelity. While derivatives in higher dimensions are simply generalized by gradients, and modulation is also straightforward, the generalization of the analytic signal is less obvious. To complete the analogy, we must first define the appropriate "analytic signal"-equivalent in the n-D context.

2.1 n-D Hilbert transform / Analytic signal

In the 1D time domain, the analytic signal is achieved by adding the Hilbert transformed copy of the original signal $f\colon \mathbb{R} \to \mathbb{R}$ as imaginary part [Gab46]:

$$\begin{aligned} f_{AS}\colon \mathbb{R} &\to \mathbb{C} \\ f_{AS}(t) &\mapsto f(t) + j\mathcal{H}\{f\}(t), \end{aligned} \tag{2.1}$$

where $j^2 = -1$, and the 1D Hilbert transform is defined as:

$$\mathcal{H}\{f\}(t) := \left\{ \frac{1}{\pi s} * f(s) \right\}(t) = \frac{1}{\pi}\text{p.v.} \int_{\mathbb{R}} \frac{f(s)}{t-s} ds, \tag{2.2}$$

where $*$ denotes convolution. We note that the real signal is recovered simply by taking the real component of the analytic signal.

In the spectral domain, this definition of analytic signal corresponds to suppressing the negative frequencies, thus giving it a unilateral spectrum:

$$\hat{f}_{AS}(\omega) = \begin{cases} 2\hat{f}(\omega), & \text{if} \quad \omega > 0, \\ \hat{f}(\omega), & \text{if} \quad \omega = 0, \\ 0, & \text{if} \quad \omega < 0, \end{cases} \tag{2.3}$$

where

$$\hat{f}(\omega) := \mathcal{F}\{f(\cdot)\}(\omega) = 1/\sqrt{2\pi} \int_{\mathbb{R}} f(t) e^{-j\omega t} dt$$

is the unitary Fourier transform in 1D.

Single-sidedness of the analytic signal spectrum was the key property motivating its use in the 1D case, since this property allowed for easy frequency shifting to base-band by complex exponential mixing. Therefore, to mimic this spectral property in 2D, one half-plane of the frequency domain must effectively be set to zero;[1] this half-plane is chosen relative to a vector, in our case to $\vec{\omega}_k$. Thus the 2D analytic signal of interest can first be defined in the frequency domain by generalizing the concept of half-space spectrum suppression:

$$\begin{aligned} \hat{f}_{AS}(\vec{\omega}) &= \begin{cases} 2\hat{f}(\omega), & \text{if} \quad \langle \vec{\omega}, \vec{\omega}_k \rangle > 0, \\ \hat{f}(\omega), & \text{if} \quad \langle \vec{\omega}, \vec{\omega}_k \rangle = 0, \\ 0, & \text{if} \quad \langle \vec{\omega}, \vec{\omega}_k \rangle < 0, \end{cases} \\ &= (1 + \operatorname{sgn}(\langle \vec{\omega}, \vec{\omega}_k \rangle)) \hat{f}(\vec{\omega}) \end{aligned} \tag{2.4}$$

where the n-D Fourier transform is defined as

$$\hat{f}(\vec{\omega}) := \mathcal{F}\{f(\cdot)\}(\vec{\omega}) = (2\pi)^{-n/2} \int_{\mathbb{R}^n} f(\vec{x}) e^{-j\langle \vec{\omega}, \vec{x} \rangle} d\vec{x}.$$

The 2D analytic signal in the time domain with the aforementioned Fourier property is given in [BS99]. It is easy to see how the generalized analytic signal reduces to the classical definition in 1D.

---

[1] Similarly, in higher dimensions, a half-space of the frequency domain needs to be suppressed.

2.2 n-D VMD functional

We are now able to put all the generalized VMD ingredients together to define the two-dimensional extension of variational mode decomposition. The functional to be minimized, stemming from this definition of n-D analytic signal, is:

$$\min_{u_k\colon \mathbb{R}^n\to\mathbb{R},\, \vec{\omega}_k\in\mathbb{R}^n} \left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j\langle \vec{\omega}_k, \vec{x}\rangle} \right] \right\|_2^2 \right\} \quad \text{s.t.} \quad \forall \vec{x}\in\mathbb{R}^n\colon\ \sum_k u_k(\vec{x}) = f(\vec{x}), \tag{2.5}$$

where $u_{AS,k}$ denotes the generalized analytic signal obtained from the mode u_k according to (2.4) using its associated center frequency ω_k. We thus minimize the Dirichlet energy of the modes after half-space spectrum suppression ($u_k \to u_{AS,k}$) and demodulation to baseband ($e^{-j\langle \vec{\omega}_k, \vec{x}\rangle}$), subject to collective signal fidelity. This model specifically includes the desired two-dimensional case $n = 2$, and reduces to the earlier 1D-VMD for $n = 1$.

Entirely analogous to the 1D-VMD model, the reconstruction constraint is addressed through the introduction of a quadratic penalty and Lagrangian multiplier (the *augmented Lagrangian*, AL, method), and we proceed by alternate direction minimization (ADMM) for optimization [Ber76, DZ14, NW06].

2.3 Augmented Lagrangian and ADMM Optimization

To render the constrained minimization problem (2.5) unconstrained, we include both a quadratic penalty and a Lagrangian multiplier to enforce the fidelity constraint. We thus define the augmented Lagrangian:

$$\begin{aligned} &\mathcal{L}(\{u_k\}, \{\omega_k\}, \lambda) := \\ &\quad \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j\langle \vec{\omega}_k, \vec{x}\rangle} \right] \right\|_2^2 + \left\| f(\vec{x}) - \sum u_k(\vec{x}) \right\|_2^2 + \left\langle \lambda(\vec{x}), f(\vec{x}) - \sum u_k(\vec{x}) \right\rangle . \end{aligned} \tag{2.6}$$

where $\lambda\colon \mathbb{R}^n \to \mathbb{R}$ is the Lagrangian multiplier. We can now solve the unconstrained saddle point problem instead of (2.5):

$$\min_{u_k\colon \mathbb{R}^n\to\mathbb{R},\, \vec{\omega}_k\in\mathbb{R}^n} \; \max_{\lambda\colon \mathbb{R}^n\to\mathbb{R}} \mathcal{L}(\{u_k\}, \{\omega_k\}, \lambda) \tag{2.7}$$

The solution to the original constrained minimization problem (2.5) is now found as the saddle point of the augmented Lagrangian $\mathcal{L}$ in a sequence of iterative sub-optimizations called alternate direction method of multipliers (ADMM) [Ber76, Hes69, Roc73a]. The idea is to iterate the following sequence of variable updates:

$$u_k^{t+1} \leftarrow \operatorname*{arg\,min}_{u_k\colon \mathbb{R}^n\to\mathbb{R}} \mathcal{L}\left(\left\{u_{i<k}^{t+1}\right\}, u_k, \left\{u_{i>k}^{t}\right\}, \left\{\omega_i^t\right\}, \lambda^t\right) \tag{2.8a}$$

$$\vec{\omega}_k^{t+1} \leftarrow \operatorname*{arg\,min}_{\vec{\omega}_k\in\mathbb{R}^n} \mathcal{L}\left(\left\{u_i^{t+1}\right\}, \left\{\vec{\omega}_{i<k}^{t+1}\right\}, \vec{\omega}_k, \left\{\vec{\omega}_{i>k}^{t}\right\}, \lambda^t\right) \tag{2.8b}$$

$$\lambda^{t+1} \leftarrow \lambda^t + \tau\left(f - \sum u_k^{t+1}\right) \tag{2.8c}$$

for $1 > \tau \geq 0$. For simplified notation while considering the subminimization problems (2.8a) and (2.8b) in the following paragraphs, we incorporate the Lagrangian multiplier term λ into the quadratic penalty term, and rewrite the objective expression slightly different:

$$\mathcal{L}(\{u_k\}, \{\omega_k\}, \lambda) = \sum_k \alpha_k \left\|\nabla\left[u_{AS,k}(\vec{x})e^{-j\langle\vec{\omega}_k,\vec{x}\rangle}\right]\right\|_2^2 + \left\|f(\vec{x}) - \sum u_k(\vec{x}) + \frac{\lambda(\vec{x})}{2}\right\|_2^2 - \left\|\frac{\lambda(\vec{x})^2}{4}\right\|_2^2 \tag{2.9}$$

2.4 Minimization w.r.t. the modes u_k

The relevant update problem derived from (2.9) is

$$u_k^{n+1} = \operatorname*{arg\,min}_{u_k\colon \mathbb{R}^n\to\mathbb{R}} \left\{\alpha_k \left\|\nabla\left[u_{AS,k}(\vec{x})e^{-j\langle\vec{\omega}_k,\vec{x}\rangle}\right]\right\|_2^2 + \left\|f(\vec{x}) - \sum_i u_i(\vec{x}) + \frac{\lambda(\vec{x})}{2}\right\|_2^2\right\} \tag{2.10}$$

Since we are dealing with L^2-norms, we can make use of the L^2 Fourier isometry and rewrite the subminimization problem in spectral domain (thus implicitly assuming periodic boundary conditions):

$$\hat{u}_k^{n+1} = \operatorname*{arg\,min}_{\hat{u}_k | u_k\colon \mathbb{R}^n\to\mathbb{R}} \left\{\alpha_k \left\|j\vec{\omega}\left[\hat{u}_{AS,k}(\vec{\omega}+\vec{\omega}_k)\right]\right\|_2^2 + \left\|\hat{f}(\vec{\omega}) - \sum_i \hat{u}_i(\vec{\omega}) + \frac{\hat{\lambda}(\vec{\omega})}{2}\right\|_2^2\right\}. \tag{2.11}$$

The ω_k-term in the spectrum of the analytic signal is due to the modulation with the complex exponential, and justified by the well-known transform pair:

$$f(\vec{x})e^{-j\langle\vec{\omega}_0,\vec{x}\rangle} \overset{\mathcal{F}}{\longleftrightarrow} \hat{f}(\vec{\omega}) * \delta(\vec{\omega}+\vec{\omega}_0) = \hat{f}(\vec{\omega}+\vec{\omega}_0), \tag{2.12}$$

where δ is the Dirac distribution and $*$ denotes convolution. Thus, multiplying an analytic signal with a pure exponential results in simple frequency shifting. Further, we can push the frequency shift out of the analytic signal spectrum through a change of variables, to obtain:

$$\hat{u}_k^{n+1} = \underset{\hat{u}_k | u_k\colon \mathbb{R}^n \to \mathbb{R}}{\arg\min} \left\{ \alpha_k \left\| j(\vec{\omega}-\vec{\omega}_k)\left[\hat{u}_{AS,k}(\vec{\omega})\right]\right\|_2^2 + \left\| \hat{f}(\vec{\omega}) - \sum_i \hat{u}_i(\vec{\omega}) + \frac{\hat{\lambda}(\vec{\omega})}{2} \right\|_2^2 \right\}. \tag{2.13}$$

We now plug in the spectral definition of the n-D analytic signal (2.4),

$$\hat{u}_{AS,k}(\vec{\omega}) = (1+\operatorname{sgn}(\langle\vec{\omega},\vec{\omega}_k\rangle))\hat{u}_k(\vec{\omega}).$$

Also, the spectra in the second term have Hermitian symmetry, since they correspond to real signals. Let

$$\Omega_k \subset \mathbb{R}^n\colon \Omega_k := \{\vec{\omega} \mid \langle\vec{\omega},\vec{\omega}_k\rangle \geq 0\}$$

denote the frequency domain half-space to which the n-D analytic signal is restricted. We rewrite both terms as integrals over these frequency domain half-spaces:

$$\hat{u}_k^{n+1} = \underset{\hat{u}_k | u_k\colon \mathbb{R}^n \to \mathbb{R}}{\arg\min} \left\{ 2\alpha_k \int_{\Omega_k} |\vec{\omega}-\vec{\omega}_k|^2 |\hat{u}_k(\vec{\omega})|^2 d\vec{\omega} + \int_{\Omega_k} \left| \hat{f}(\vec{\omega}) - \sum_i \hat{u}_i(\vec{\omega}) + \frac{\hat{\lambda}(\vec{\omega})}{2} \right|^2 d\vec{\omega} \right\}. \tag{2.14}$$

This subminimization problem is now solved by letting the first variation w.r.t. $\hat{u}_k$ vanish[2]. The optimal mode spectrum thus satisfies:

$$0 = 2\alpha_k|\vec{\omega}-\vec{\omega}_k|^2\hat{u}_k - \Big(\hat{f}(\vec{\omega}) - \sum_i \hat{u}_i(\vec{\omega}) + \frac{\hat{\lambda}(\vec{\omega})}{2}\Big), \quad \forall\vec{\omega}\in\Omega_k. \tag{2.15}$$

With this optimality condition, solving for $\hat{u}_k$ yields the following Wiener-filter update:

$$\hat{u}_k^{n+1}(\vec{\omega}) = \Big(\hat{f}(\vec{\omega}) - \sum_{i\neq k} \hat{u}_i(\vec{\omega}) + \frac{\hat{\lambda}(\vec{\omega})}{2}\Big)\frac{1}{1+2\alpha_k|\vec{\omega}-\vec{\omega}_k|^2}, \quad \forall\vec{\omega}\in\Omega_k. \tag{2.16}$$

---

[2] Note that the spectrum of u_k is complex valued so the process of "taking the first variation" is not self-evident. However, the functional is analytic in $\hat{u}_k$ and complex-valued equivalents to the standard derivatives do indeed apply.

The full spectrum $\hat{u}_k^{n+1}$ can then be obtained by symmetric (Hermitian) completion. Equivalently, we can decide to update the half-space analytic signal of the mode, $\hat{u}_{AS,k}^{n+1}$, on the entire frequency domain, instead:

$$\hat{u}_{AS,k}^{n+1}(\vec{\omega}) = \left(\hat{f}(\vec{\omega}) - \sum_{i \neq k} \hat{u}_i(\vec{\omega}) + \frac{\hat{\lambda}(\vec{\omega})}{2}\right) \frac{1 + \operatorname{sgn}(\langle \vec{\omega}, \vec{\omega}_k \rangle)}{1 + 2\alpha_k |\vec{\omega} - \vec{\omega}_k|^2}, \quad \forall \vec{\omega} \in \mathbb{R}^n, \tag{2.17}$$

from which the actual mode estimate is recovered as the real part after inverse Fourier transform. The term in parentheses is the signal's k-th residual, where $\hat{f}(\vec{\omega}) - \sum_{i \neq k} \hat{u}_i(\vec{\omega})$ is the explicit current residual, and $\hat{\lambda}$ accumulates the reconstruction error over iterations (see below). The second term is identified as a frequency filter tuned to the current estimate of the mode's center pulsation, $\vec{\omega}_k$, and whose bandwidth is controlled by the parameter α_k.

2.5 Minimization w.r.t. the center frequencies $\vec{\omega}_k$

Optimizing for $\vec{\omega}_k$ is even simpler. Indeed, the respective update goal derived from (2.9) is

$$\vec{\omega}_k^{n+1} = \operatorname*{arg\,min}_{\vec{\omega}_k \in \mathbb{R}^n} \left\{ \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j\langle \vec{\omega}_k, \vec{x} \rangle} \right] \right\|_2^2 \right\}. \tag{2.18}$$

Or, again we may consider the equivalent problem in the Fourier domain:

$$\begin{aligned} \vec{\omega}_k^{n+1} &= \operatorname*{arg\,min}_{\vec{\omega}_k \in \mathbb{R}^n} \left\{ \alpha_k \left\| j(\vec{\omega} - \vec{\omega}_k)(1 + \operatorname{sgn}(\langle \vec{\omega}_k, \vec{\omega} \rangle)) \hat{u}_k(\vec{\omega}) \right\|_2^2 \right\} && (2.19) \\ &= \operatorname*{arg\,min}_{\vec{\omega}_k \in \mathbb{R}^n} \left\{ 4\alpha_k \int_{\Omega_k} |\vec{\omega} - \vec{\omega}_k|^2 \, |\hat{u}_k(\vec{\omega})|^2 \, d\vec{\omega} \right\}. && (2.20) \end{aligned}$$

The minimization is solved by letting the first variation w.r.t. $\vec{\omega}_k$ vanish, leading to:

$$\int_{\Omega_k} (\vec{\omega} - \vec{\omega}_k^{n+1}) \, |\hat{u}_k(\vec{\omega})|^2 \, d\vec{\omega} = 0. \tag{2.21}$$

The resulting solutions are the centers of gravity of the modes' power spectra, $|\hat{u}_k(\vec{\omega})|^2$, restricted to the half-space Ω_k:

$$\vec{\omega}_k^{n+1} = \frac{\int_{\Omega_k} \vec{\omega} |\hat{u}_k(\vec{\omega})|^2 d\vec{\omega}}{\int_{\Omega_k} |\hat{u}_k(\vec{\omega})|^2 d\vec{\omega}} = \frac{\int_{\mathbb{R}^n} \vec{\omega} |\hat{u}_{AS,k}(\vec{\omega})|^2 d\vec{\omega}}{\int_{\mathbb{R}^n} |\hat{u}_{AS,k}(\vec{\omega})|^2 d\vec{\omega}}, \tag{2.22}$$

where the second form is given for implementation purposes, based on the analytic signal spectrum and involving the entire frequency domain.

2.6 Maximization w.r.t. the Lagrangian multiplier λ

Maximizing the λ is the simplest step in the algorithm. The first variation for λ is just the data reconstruction error, $f(\vec{\omega}) - \sum_k u_k^{n+1}(\vec{\omega})$. We use a standard gradient ascent with fixed time step $1 > \tau \geq 0$ to achieve this maximization:

$$\lambda^{n+1}(\vec{x}) = \lambda^n(\vec{x}) + \tau \left(f(\vec{x}) - \sum_k u_k^{n+1}(\vec{x}) \right). \tag{2.23}$$

It is important to note that choosing $\tau = 0$ effectively eliminates the Lagrangian update and thus reduces the algorithm to the penalty method for data fidelity purposes. Doing so is useful when exact data fidelity is not appropriate, such as in (high) noise scenarios, we reconstruction error actually allows capturing noise separately.

Note also that the linearity of the Euler-Lagrange equation allows an impartial choice in which space to update the Lagrangian multiplier, either in the time domain or in the frequency domain. In our implementation, we perform our dual ascent update in the frequency domain, since the other appearance of the Lagrangian multiplier in (2.17) is in spectral terms, as well. Thus:

$$\hat{\lambda}^{n+1}(\vec{\omega}) = \hat{\lambda}^n(\vec{\omega}) + \tau \left(\hat{f}(\vec{\omega}) - \sum_k \hat{u}_k^{n+1}(\vec{\omega}) \right). \tag{2.24}$$

2.7 Complete 2D VMD algorithm

The entire proposed algorithm for the 2D-VMD functional optimization problem (2.5) is summarized in algorithm 2. Variables are trivially initialized at 0, except for the center frequencies, $\vec{\omega}_k$, for which smart initialization is of higher importance; initial $\vec{\omega}_k^0$ can, e.g., be spread randomly, radially uniform, or initialized by user input. Further, we choose to assess convergence in terms of the normalized rate of change of the modes. Typical thresholds $\epsilon > 0$ range in orders of magnitude from 10^{-4} (fast) down to 10^{-7} (very accurate). An example of image decomposition achieved with 2D VMD according to algorithm 2 is shown in figure 7.1.

Algorithm 2 2D-VMD

```
Input: signal $f(\vec{x})$, number of modes $K$, parameters $\alpha_k$, $\tau$, $\epsilon$.
Output: modes $u_k(\vec{x})$, center frequencies $\vec{\omega}_k$.
Initialize $\{\omega_k^0\}$, $\{\hat{u}_k^0\} \leftarrow 0$, $\hat{\lambda}^0 \leftarrow 0$, $n \leftarrow 0$
repeat
    $n \leftarrow n+1$
    for $k = 1 : K$ do
        Create 2D mask for analytic signal Fourier multiplier:
            $\mathcal{H}_k^{n+1}(\vec{\omega}) \leftarrow 1 + \operatorname{sgn}(\langle \vec{\omega}_k^n, \vec{\omega} \rangle)$
        Update $\hat{u}_{AS,k}$:
            $\hat{u}_{AS,k}^{n+1}(\vec{\omega}) \leftarrow \mathcal{H}_k^{n+1}(\vec{\omega}) \left[ \frac{\hat{f}(\vec{\omega}) - \sum_{i<k} \hat{u}_i^{n+1}(\vec{\omega}) - \sum_{i>k} \hat{u}_i^n(\vec{\omega}) + \frac{\hat{\lambda}^n(\vec{\omega})}{2}}{1 + 2\alpha_k |\vec{\omega} - \vec{\omega}_k^n|^2} \right]$
        Update $\vec{\omega}_k$:
            $\vec{\omega}_k^{n+1} \leftarrow \frac{\int_{\mathbb{R}^2} \vec{\omega} |\hat{u}_{AS,k}^{n+1}(\vec{\omega})|^2 d\vec{\omega}}{\int_{\mathbb{R}^2} |\hat{u}_{AS,k}^{n+1}(\vec{\omega})|^2 d\vec{\omega}}$
        Retrieve $u_k$:
            $u_k^{n+1}(\vec{x}) \leftarrow \Re\left(\mathcal{F}^{-1}\left\{\hat{u}_{AS,k}^{n+1}(\vec{\omega})\right\}\right)$
    end for
    Dual ascent (optional):
        $\hat{\lambda}^{n+1}(\vec{\omega}) \leftarrow \hat{\lambda}^n(\vec{\omega}) + \tau \left( \hat{f}(\vec{\omega}) - \sum_k \hat{u}_k^{n+1}(\vec{\omega}) \right)$
until convergence: $\sum_k \|\hat{u}_k^{n+1} - \hat{u}_k^n\|_2^2 / \|\hat{u}_k^n\|_2^2 < \epsilon$.
```

CHAPTER 3

VMD with Compact Spatial Support

A main assumption regarding the intrinsic mode functions considered so far is that their amplitude (spatially) varies much more slowly than the wavelength of the carrier. Indeed, IMFs can be defined as signals (in time or space) that are both amplitude and frequency modulated [DLW11]. In [DZ14], we have defined the *total practical IMF bandwidth* of such an AM-FM signal, as an extension to Carson's rule for FM-signal bandwidth [Car22]:

$$BW_{\text{AM-FM}} := 2(\Delta f + f_{\text{FM}} + f_{\text{AM}}), \tag{3.1}$$

where Δf and f_{FM} represent the frequency swing and modulation bandwidth, respectively, of the FM part, while f_{AM} denotes the bandwidth of the amplitude modulation. The last, AM bandwidth, conflicts with signals composed of modes having sudden signal onset, in particular those with compact spatial support. Indeed, this inverse relation between spatial and spectral compactness is well known and stated by the Heisenberg uncertainty principle.

3.1 Introducing binary support functions A_k

To make our "modes have limited bandwidth"-prior compatible with signals of limited spatial support, it is thus necessary to deal with the spatial and spectral compactness of the modes, separately. To this end, we introduce a binary support function for each mode, in order to capture the signal onset and offset disconnected from the smooth AM-FM modulations.

We consider signals and modes $f, u_k \colon \mathbb{R}^n \to \mathbb{R}$ (thus including both the 1D-VMD and higher dimensional signals such as 2D-VMD stated above). Let

$$A_k \colon \mathbb{R}^n \to \{0, 1\}$$

denote the binary support functions for each mode u_k. The mode decomposition problem can then formally be stated as

$$\text{find } u_k, A_k \quad \text{s.t.} \quad f = \sum_k A_k \cdot u_k,$$

i.e., we want the modes u_k, now masked by their binary support function A_k, to reproduce collectively the given input signal. Note that the modes u_k can extend arbitrarily into their inactive regions where $A_k = 0$; in particular, they can decay smoothly or oscillate *ad infinitum*, thus keeping small spectral bandwidth.

3.2 Sparsity promoting VMD functional

It is important to introduce sparsity promoting regularity constraints on the support function to achieve reasonable compact local support. Here, we consider both total variation (TV) and L^1 penalties on A_k, thus effectively penalizing support area and boundary length (through the co-area formula).

We incorporate the binary support functions A_k and their regularizers in the n-D VMD functional as follows:

$$\min_{u_k \colon \mathbb{R}^n \to \mathbb{R},\, A_k \colon \mathbb{R}^n \to \{0,1\},\, \vec{\omega}_k \in \mathbb{R}^n} \left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j \langle \vec{\omega}_k, \vec{x} \rangle} \right] \right\|_2^2 + \beta_k \|A_k\|_1 + \gamma_k \|\nabla A_k\|_1 \right\}$$

$$\text{s.t.} \quad \forall \vec{x} \in \mathbb{R}^n \colon \sum_k A_k(\vec{x}) u_k(\vec{x}) = f(\vec{x}). \quad (3.2)$$

The L^1 penalties on A_k and ∇A_k ensure that an individual mode is only active in places where it is "sufficiently justified" (i.e., the increased data fidelity outweighs the incurred friction cost), and represent the prior on modes to have limited spatial support and regular outlines.

3.3 Model relaxation

Due to the introduction of the binary support functions A_k in the fidelity constraint, and the L^1-based prior terms, the functional is no longer directly translatable to the spectral domain. Moreover, the L^1-terms do not lend themselves to standard calculus of variations methods, directly. Instead, we propose an ensemble of splitting techniques [Coh96, CP11, GL89] that have been applied to L^1-based and related optimization problems with great success, such as [EZB14, GO09, ZBT14].

First, we would like to restore spectral solvability of the modes u_k. Currently, the masks A_k prevent this, since in the quadratic penalty addressing the reconstruction constraint, the spatial multiplication translates to spectral convolution. Spectral solvability for u_k is restored by introducing a splitting of the modes $u_k = v_k$, and applying spectral bandwidth penalty and reconstruction over the separate copies:

$$\min_{u_k\colon \mathbb{R}^n\to\mathbb{R},\, A_k\colon \mathbb{R}^n\to\{0,1\},\, \vec{\omega}_k\in\mathbb{R}^n} \left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j\langle \vec{\omega}_k, \vec{x}\rangle} \right] \right\|_2^2 + \beta_k \|A_k\|_1 + \gamma_k \|\nabla A_k\|_1 \right\}$$

$$\text{s.t.} \quad \forall \vec{x} \in \mathbb{R}^n\colon \begin{cases} u_k(\vec{x}) = v_k(\vec{x}), \\ \sum_k A_k(\vec{x}) v_k(\vec{x}) = f(\vec{x}). \end{cases} \tag{3.3}$$

The splitting constraint can be addressed with a quadratic penalty (proximal splitting, [CP11]), or using an augmented Lagrangian [GL89]. As an intermediate illustration, and since the latter includes the former, we give the full saddle-point functional (augmented Lagrangian) incorporating both equality constraints through quadratic penalty and Lagrangian multipliers, in analogy to (2.6):

$$\begin{aligned}&\mathcal{L}(\{u_k\},\{v_k\},\{A_k\},\{\omega_k\},\lambda,\{\lambda_k\}) := \\ &\Bigg\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j\langle\vec{\omega}_k,\vec{x}\rangle}\right] \right\|_2^2 + \beta_k \|A_k\|_1 + \gamma_k \|\nabla A_k\|_1 + \rho \left\| f(\vec{x}) - \sum A_k(\vec{x}) v_k(\vec{x}) \right\|_2^2 \\ &\quad + \left\langle \lambda(\vec{x}), f(\vec{x}) - \sum A_k(\vec{x}) v_k(\vec{x}) \right\rangle + \sum_k \rho_k \left\| u_k(\vec{x}) - v_k(\vec{x}) \right\|_2^2 + \langle \lambda_k(\vec{x}), u_k(\vec{x}) - v_k(\vec{x})\rangle \Bigg\},\end{aligned} \tag{3.4}$$

where λ_k are the Lagrangian multipliers associated with the K equality constraints $u_k = v_k$, and ρ, ρ_k are parameters weighting the different quadratic penalties. All terms involving u_k translate nicely into the spectral domain, while all terms in v_k lend themselves to efficient point-wise optimization in time domain. Before actually looking at the specific subminimization problems, we want to study the L^1-terms further by recognizing them as essentially balloon and motion-by-mean-curvature forces acting on the binary support functions A_k.

3.4 Excursion on MBO

The first variation associated with the TV-term is proportional to $\operatorname{div}(\nabla A_k/|\nabla A_k|)$. One can expect difficulties with this term, for example in flat regions where $|\nabla A_k| \to 0$. Moreover, if the gradient descent PDE is integrated explicitly, then the time step is also heavily limited by the Courant-Friedrich-Lewy (CFL) condition [CFL28].

An important contribution stems from the diffusion-threshold scheme for approximating motion by mean curvature proposed by Merriman, Bence, and Osher (MBO) [MBO94]. The fundamental idea is to reproduce the motion by mean curvature due to the boundary-length term $TV(A_k)$ by more efficient means than direct gradient descent.

Since A_k is binary we opt for alternative schemes other than split-Bregman/shrinkage or dual minimization [Cha04, GO09, ZWC10]. As a preliminary, motivational step, let us replace the total variation of the support function A_k, by the real Ginzburg-Landau (GL, also known as Allen-Cahn) functional [Mod87]:

$$E^{\epsilon}_{\mathrm{GL}}(A_k) := \epsilon \int_{\Omega} |\nabla A_k(\vec{x})|^2 d\vec{x} + \frac{1}{\epsilon} \int_{\Omega} W(A_k(\vec{x}))d\vec{x}, \quad \epsilon > 0, \tag{3.5}$$

where $W(s)$ is a double-well potential with two equal minima at $s = 0$ and $s = 1$, for example $W(s) := s^2(1-s)^2$. Minimizing this functional yields a phase field that is smooth and tends to be binary. In particular, it has been shown [Mod87] that the GL-functional Γ-converges to the total variation functional of binary phase-fields $A_k \in \{0, 1\}$ as $\epsilon \to 0$:

$$E^{0}_{\mathrm{GL}}(A_k) = \sigma(W) \int_{\Omega} |\nabla A_k|, \tag{3.6}$$

where $\sigma(W)$ is a surface tension term depending on the double well potential. The minimizing flow of this functional for $\epsilon \to 0^+$ produces motion by mean curvature of the interface, which is exactly what one needs in the spatially sparse VMD model minimization. However, now, the PDE associated with the GL-functional minimization is

$$\frac{\partial A_k}{\partial t} = 2\epsilon \nabla^2 A_k - \frac{1}{\epsilon} W'(A_k), \tag{3.7}$$

and this PDE is conveniently solved in a discrete-time two step time-splitting approach:

1. Propagate A_k according to the heat equation,

   $$\frac{\partial A_k}{\partial t} = 2\epsilon \nabla^2 A_k$$

2. Propagate A_k according to the double well potential gradient descent,

   $$\frac{\partial A_k}{\partial t} = -\frac{1}{\epsilon} W'(A_k).$$

The heat equation is efficiently solved, e.g., based on convolution or spectral transforms [Ruu98].

Now, the MBO-scheme [MBO94] improves on this time-split GL-optimization in that the ODE is recognized as essentially performing thresholding. While the first step is reduced to propagation according to the standard heat equation, the second step in MBO is actual thresholding (projection onto the binary set $\{0,1\}$):

1. Propagate A_k according to the heat equation,

$$\frac{\partial A_k}{\partial t} = \nabla^2 A_k$$

2. Rectify A_k by thresholding:

$$A_k(\vec{x}) = \begin{cases} 0 & \text{if } A_k(\vec{x}) \leq \frac{1}{2} \\ 1 & \text{if } A_k(\vec{x}) > \frac{1}{2} \end{cases} \qquad \forall \vec{x} \in \mathbb{R}^n$$

These MBO threshold dynamics have already been successfully integrated with imaging data terms, such as [ET06, ZAS15], where in addition to the heat diffusion and thresholding steps, a data-driven gradient descent step is included in the iterations. We propose a similar structure here, to account for the balloon force and reconstruction fidelity term contributions to the A_k minimization.

3.5 n-D-TV-VMD Minimization

Based on the preparatory steps of the preceding sections, we now propose to solve the constraint, sparsity promoting n-D VMD functional (3.2) through its augmented Lagrangian (3.4). Consider the following saddle point problem:

$$\min_{u_k, v_k : \mathbb{R}^n \to \mathbb{R},\, A_k : \mathbb{R}^n \to \{0,1\},\, \vec{\omega}_k \in \mathbb{R}^n} \quad \max_{\lambda, \lambda_k : \mathbb{R}^n \to \mathbb{R}} \left\{ \mathcal{L}(\{u_k\}, \{v_k\}, \{A_k\}, \{\omega_k\}, \lambda, \{\lambda_k\}) \right\}. \tag{3.8}$$

This saddle point problem is an extended version of the 2D VMD saddle point problem (2.7) (without spatial sparsity promoting terms), and is again efficiently solved through alternate direction minimization and dual ascent (ADMM):

$$u_k^{t+1} \leftarrow \underset{u_k\,:\,\mathbb{R}^n\to\mathbb{R}}{\arg\min} \mathcal{L}\left(\left\{u_{i<k}^{t+1}\right\}, u_k, \left\{u_{i>k}^t\right\}, \left\{v_i^t\right\}, \left\{A_i^t\right\}, \left\{\omega_i^t\right\}, \lambda^t, \left\{\lambda_i^t\right\}\right) \tag{3.9a}$$

$$v_k^{t+1} \leftarrow \underset{v_k\,:\,\mathbb{R}^n\to\mathbb{R}}{\arg\min} \mathcal{L}\left(\left\{u_i^{t+1}\right\}, \left\{v_{i<k}^{t+1}\right\}, v_k, \left\{v_{i>k}^t\right\}, \left\{A_i^t\right\}, \left\{\omega_i^t\right\}, \lambda^t, \left\{\lambda_i^t\right\}\right) \tag{3.9b}$$

$$A_k^{t+1} \leftarrow \underset{A_k\,:\,\mathbb{R}^n\to\{0,1\}}{\arg\min} \mathcal{L}\left(\left\{u_i^{t+1}\right\}, \left\{v_i^{t+1}\right\}, \left\{A_{i<k}^{t+1}\right\}, A_k, \left\{A_{i>k}^t\right\}, \left\{\omega_i^t\right\}, \lambda^t, \left\{\lambda_i^t\right\}\right) \tag{3.9c}$$

$$\vec{\omega}_k^{t+1} \leftarrow \underset{\vec{\omega}_k\in\mathbb{R}^n}{\arg\min} \mathcal{L}\left(\left\{u_i^{t+1}\right\}, \left\{v_i^{t+1}\right\}, \left\{A_i^{t+1}\right\}, \left\{\vec{\omega}_{i<k}^{t+1}\right\}, \vec{\omega}_k, \left\{\vec{\omega}_{i>k}^t\right\}, \lambda^t, \left\{\lambda_i^t\right\}\right) \tag{3.9d}$$

$$\lambda^{t+1} \leftarrow \lambda^t + \tau\left(f - \sum A_k^{t+1} v_k^{t+1}\right) \tag{3.9e}$$

$$\lambda_k^{t+1} \leftarrow \lambda_k^t + \tau_k\left(u_k^{t+1} - v_k^{t+1}\right) \tag{3.9f}$$

We provide details on the individual sub-minimization problems in the following paragraphs. The complete algorithm for n-D-TV-VMD functional (with spatial sparsity promoting terms) is then easily derived in analogy to algorithm 2.

3.5.1 Subminimization w.r.t. u_k

The relevant minimization problem (3.9a) with respect to the modes u_k reads

$$u_k^{t+1} = \underset{u_k\,:\,\mathbb{R}^n\to\mathbb{R}}{\arg\min}\left\{\alpha_k\left\|\nabla\left[u_{AS,k}(\vec{x})e^{-j\langle\vec{\omega}_k,\vec{x}\rangle}\right]\right\|_2^2 + \rho_k\left\|u_k(\vec{x}) - v_k(\vec{x}) + \frac{\lambda_k(\vec{x})}{\rho_k}\right\|_2^2\right\}. \tag{3.10}$$

In full analogy to the problem without spatial sparsity terms, (2.10), the update is most easily computed in spectral domain, like (2.16). Unsurprisingly, the update rule on the frequency halfspace $\Omega_k = \{\omega \mid \langle\omega,\omega_k\rangle \geq 0\}$ is found to be:

$$\hat{u}_k^{t+1}(\vec{\omega}) = (\rho_k\hat{v}_k - \hat{\lambda}_k)\frac{1}{\rho_k + 2\alpha_k|\omega-\omega_k|^2}, \quad \forall\omega\in\Omega_k. \tag{3.11}$$

From this half-space update, the full spectrum can again be obtained by Hermitian completion; Or by updating the mode's half-space analytic signal instead:

$$\hat{u}_{AS,k}^{t+1}(\vec{\omega}) = (\rho_k \hat{v}_k - \hat{\lambda}_k)\frac{1 + \operatorname{sgn}(\langle \omega, \omega_k \rangle)}{\rho_k + 2\alpha_k |\omega - \omega_k|^2}. \tag{3.12}$$

3.5.2 Subminimization w.r.t. v_k

The update (3.9b) of v_k reduces to the following minimization problem:

$$v_k^{t+1} = \operatorname*{arg\,min}_{v_k \colon \mathbb{R}^n \to \mathbb{R}} \left\{ \rho \left\| f(\vec{x}) - \sum A_i(\vec{x}) v_i(\vec{x}) + \frac{\lambda(\vec{x})}{\rho} \right\|_2^2 + \rho_k \left\| u_k(\vec{x}) - v_k(\vec{x}) + \frac{\lambda_k(\vec{x})}{\rho_k} \right\|_2^2 \right\} \tag{3.13}$$

This problem admits the following pointwise Euler-Lagrange equations:

$$-\rho A_k(\vec{x}) \left(f(\vec{x}) - \sum A_i(\vec{x}) v_i(\vec{x}) + \frac{\lambda(\vec{x})}{\rho} \right) - \rho_k \left(u_k(\vec{x}) - v_k(\vec{x}) + \frac{\lambda_k(\vec{x})}{\rho_k} \right) = 0, \quad \forall \vec{x} \in \mathbb{R}^n \tag{3.14}$$

yielding the simple update rule

$$v_k^{t+1}(\vec{x}) = \frac{\rho A_k(\vec{x}) \left(f(\vec{x}) - \sum_{i \neq k} A_i(\vec{x}) v_i(\vec{x}) + \frac{\lambda(\vec{x})}{\rho} \right) + \rho_k u_k(\vec{x}) + \lambda_k(\vec{x})}{\rho A_k(\vec{x})^2 + \rho_k}, \quad \forall \vec{x} \in \mathbb{R}^n. \tag{3.15}$$

This update is interpreted as a balance between fidelity to the split mode u_k (enforced through Lagrangian multiplier λ_k), and the reconstruction fidelity constraint where A_k is active (enforced through λ).

3.5.3 Subminimization w.r.t. A_k

As outlined above, the minimization problem with respect to the binary support functions A_k involves the L^1-based priors:

$$A_k^{t+1} = \operatorname*{arg\,min}_{A_k \colon \mathbb{R}^n \to \{0,1\}} \left\{ \beta_k \|A_k\|_1 + \gamma_k \|\nabla A_k\|_1 + \rho \left\| f(\vec{x}) - \sum A_i(\vec{x}) v_i(\vec{x}) + \frac{\lambda(\vec{x})}{\rho} \right\|_2^2 \right\}. \tag{3.16}$$

Motivated by successful implementation for image segmentation problems, for example, we want to employ diffusion and threshold dynamics for the efficient solution of this problem. In analogy to the image segmentation scheme, we devise a three-fold time-split gradient descent iteration: The first step is gradient descent based on the support area and reconstruction-fidelity penalty. The second step is diffusion by the heat equation, followed by thresholding, to deal with the boundary length term and the projection on the admissible set $\{0,1\}$.

Since A_k is non-negative, it is safe to drop the absolute value and relax the L^1-area term to $\beta_k \int_{\mathbb{R}^n} A_k$. This makes the functional smoothly differentiable in the area and reconstruction term.

We thus propose to update the binary support functions A_k^{t+1} in MBO-like fashion [ET06, MBO94, ZAS15] by iterating over the following three evolution equations:

1. Area penalty and reconstruction fidelity ODE:

$$\frac{\partial A_k(\vec{x})}{\partial t} = -\beta_k + 2\rho v_k(\vec{x}) \left(f(\vec{x}) - \sum A_i(\vec{x}) v_i(\vec{x}) + \frac{\lambda(\vec{x})}{\rho} \right), \tag{3.17}$$

2. Heat equation PDE for diffusion:

$$\frac{\partial A_k(\vec{x})}{\partial t} = \gamma_k \nabla^2 A_k(\vec{x}), \tag{3.18}$$

3. Rectification by thresholding:

$$A_k(\vec{x}) = \begin{cases} 0 & \text{if } A_k(\vec{x}) \leq \frac{1}{2} \\ 1 & \text{if } A_k(\vec{x}) > \frac{1}{2} \end{cases} \qquad \forall \vec{x} \in \mathbb{R}^n. \tag{3.19}$$

Note that the ODE problem can be addressed through an implicit (backward) Euler scheme, and the heat equation PDE is efficiently solved spectrally.

3.5.4 Subminimization w.r.t. ω_k

The last, remaining sub-problem of the saddle-point problem (3.8) is the update of the mode's central frequency, ω_k. The relevant portion of the functional (3.4) is identical to the non-sparse 2D-VMD model (2.6). Therefore the corresponding subminimization problem here is identical to (2.18), and thus the update is equally given by (2.22).

The complete algorithm for the ADMM optimization of the 2D-TV-VMD model is shown in algorithm 3, and illustrative examples of its use are given in figures 7.1 and 7.2.

Algorithm 3 2D-TV-VMD (sparsity promoting)

```
**Input:** signal $f(\vec{x})$, number of modes $K$, parameters $\alpha_k$, $\beta_k$, $\gamma_k$, $\rho$, $\rho_k$, $t$, $\tau$, $\tau_k$, $\epsilon$.
**Output:** modes $u_k(\vec{x})$, support functions $A_k(\vec{x})$, center frequencies $\vec{\omega}_k$.
Initialize $\{\omega_k^0\}$, $\{u_k^0\} \leftarrow 0$, $\{v_k^0\} \leftarrow 0$, $\{A_k^0\} \leftarrow 1$, $\{\lambda_k\}^0 \leftarrow 0$, $\lambda^0 \leftarrow 0$, $n \leftarrow 0$
**repeat**
    $n \leftarrow n+1$
    **for** $k = 1 : K$ **do**
        Create 2D mask for analytic signal Fourier multiplier:
            $$\mathcal{H}_k^{n+1}(\vec{\omega}) \leftarrow 1 + \operatorname{sgn}(\langle \vec{\omega}_k^n, \vec{\omega} \rangle)$$
        Update $\hat{u}_{AS,k}$:
            $$\hat{u}_{AS,k}^{n+1}(\vec{\omega}) \leftarrow \mathcal{H}_k^{n+1}(\vec{\omega}) \left[ \frac{\rho_k \hat{v}_k^n(\vec{\omega}) - \hat{\lambda}_k^n(\vec{\omega})}{\rho_k + 2\alpha_k |\vec{\omega} - \vec{\omega}_k^n|^2} \right]$$
        Retrieve $u_k$:
            $$u_k^{n+1}(\vec{x}) \leftarrow \Re\left(\mathcal{F}^{-1}\left\{\hat{u}_{AS,k}^{n+1}(\vec{\omega})\right\}\right)$$
        Update $v_k$:
            $$v_k^{n+1}(\vec{x}) \leftarrow \frac{\rho A_k^n(\vec{x}) \left( f(\vec{x}) - \sum_{i<k} A_i^n(\vec{x}) v_i^{n+1}(\vec{x}) - \sum_{i>k} A_i^n(\vec{x}) v_i^n(\vec{x}) + \frac{\lambda^n(\vec{x})}{\rho} \right) + \rho_k u_k^{n+1}(\vec{x}) + \lambda_k^n(\vec{x})}{\rho A_k^n(\vec{x})^2 + \rho_k}$$
```

Algorithm 3 2D-TV-VMD (sparsity promoting) (continued)

```
        Update $A_k$ through modified MBO:
            $A_k^{n+1/3}(\vec{x}) \leftarrow \frac{1}{1+2t\rho(v_k^{n+1}(\vec{x}))^2}\Big(A_k^n(\vec{x}) + t\Big(-\beta_k + 2\rho v_k^{n+1}(\vec{x})\Big(f(\vec{x}) - \sum_{i<k} A_i^{n+1}(\vec{x})v_i^{n+1}(\vec{x}) - \sum_{i>k} A_i^n(\vec{x})v_i^n(\vec{x}) + \frac{\lambda^n(\vec{x})}{\rho}\Big)\Big)\Big)$
            $\hat{A}_k^{n+2/3}(\vec{\omega}) \leftarrow \frac{\hat{A}_k^{n+1/3}(\vec{\omega})}{1+t\gamma_k|\vec{\omega}|^2}$
            $A_k^{n+1}(\vec{x}) \leftarrow \begin{cases} 0 & \text{if} \quad A_k^{n+2/3}(\vec{x}) \le \frac{1}{2} \\ 1 & \text{if} \quad A_k^{n+2/3}(\vec{x}) > \frac{1}{2} \end{cases}$
        Update $\vec{\omega}_k$:
            $\vec{\omega}_k^{n+1} \leftarrow \frac{\int_{\mathbb{R}^2} \vec{\omega}|\hat{u}_{AS,k}^{n+1}(\vec{\omega})|^2 d\vec{\omega}}{\int_{\mathbb{R}^2} |\hat{u}_{AS,k}^{n+1}(\vec{\omega})|^2 d\vec{\omega}}$
        Dual ascent $u$-$v$ coupling:
            $\lambda_k^{n+1}(\vec{x}) \leftarrow \lambda_k^n(\vec{x}) + \tau_k\left(u_k^{n+1}(\vec{x}) - v_k^{n+1}(\vec{x})\right)$
    end for
    Dual ascent data fidelity:
        $\lambda^{n+1}(\vec{x}) \leftarrow \lambda^n(\vec{x}) + \tau\left(f(\vec{x}) - \sum_k A_k^{n+1}(\vec{x})v_k^{n+1}(\vec{x})\right)$
until convergence
```

CHAPTER 4

Spectral Image Segmentation

Up to now we have considered modes whose spatial support was mutually independent. In particular, this means that VMD and TV-VMD modes can be spatially overlapping, and conversely, that not all parts of a signal are covered by an active mode. Here, we want to consider the case where modes are restricted to be non-overlapping while covering the entire signal domain. In other words, the modes' support functions A_k form a partition of the signal domain. For example, such a model includes the image segmentation problem.

In terms of the binary support functions, $A_k\colon \mathbb{R}^n \to \{0,1\}$, this means imposing the following constraint:

$$\sum_k A_k(\vec{x}) = 1, \qquad \forall \vec{x} \in \mathbb{R}^n. \tag{4.1}$$

In return, the area penalty $\beta_k \|A_k\|_1$ is obsolete, of course, unless not all modes incur the same area penalty due to different size priors, corresponding to $\beta_i \neq \beta_j$ for at least some $(i,j) \in \{1,\dots,K\}^2$.

We propose the following spatially disjoint n-D-TV-VMD model, as a modification of (3.2):

$$\min_{u_k\colon \mathbb{R}^n \to \mathbb{R},\, A_k\colon \mathbb{R}^n \to \{0,1\},\, \vec{\omega}_k \in \mathbb{R}^n} \left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j\langle \vec{\omega}_k, \vec{x}\rangle} \right] \right\|_2^2 + \beta_k \|A_k\|_1 + \gamma_k \|\nabla A_k\|_1 \right\}$$

$$\text{s.t.} \quad \forall \vec{x} \in \mathbb{R}^n\colon \begin{cases} \sum_k A_k(\vec{x}) u_k(\vec{x}) = f(\vec{x}), \\ \sum_k A_k(\vec{x}) = 1. \end{cases} \tag{4.2}$$

Next we outline two different strategies to accommodate this extra constraint on the support functions in the minimization scheme. The first strategy incorporates the partitioning constraint through another augmented Lagrangian to be included in the saddle point problem. The second model deals with the restricted solution space through projection, more precisely by modifying the current rectification step included in the MBO-like diffusion and threshold-dynamics.

4.1 Augmented Lagrangian method

In the first approach, we incorporate the segmentation constraint as a third augmented Lagrangian term. Based on the AL (3.4) of the spatially overlapping compact VMD functional (3.2), we write:

$$\begin{aligned}
&\mathcal{L}(\{u_k\},\{v_k\},\{A_k\},\{\omega_k\},\lambda,\{\lambda_k\}) := \\
&\Bigg\{\sum_k \alpha_k \left\|\nabla\left[u_{AS,k}(\vec{x})e^{-j\langle\vec{\omega}_k,\vec{x}\rangle}\right]\right\|_2^2 + \beta_k\|A_k\|_1 + \gamma_k\|\nabla A_k\|_1 + \rho\left\|f(\vec{x}) - \sum A_k(\vec{x})v_k(\vec{x})\right\|_2^2 \\
&+ \left\langle \lambda(\vec{x}), f(\vec{x}) - \sum A_k(\vec{x})v_k(\vec{x})\right\rangle + \sum \rho_k\|u_k(\vec{x}) - v_k(\vec{x})\|_2^2 + \sum\langle\lambda_k(\vec{x}), u_k(\vec{x}) - v_k(\vec{x})\rangle \\
&\qquad + \rho'\left\|\sum A_k(\vec{x}) - 1\right\|_2^2 + \left\langle \lambda'(\vec{x}), \sum A_k(\vec{x}) - 1\right\rangle\Bigg\}, \qquad (4.3)
\end{aligned}$$

where $\lambda'\colon \mathbb{R}^n \to \mathbb{R}$ is the newly introduced Lagrangian multiplier, and ρ' the weight of the corresponding quadratic penalty term. Sticking to the alternate direction gradient descent and dual ascent scheme (3.8) for optimization, we realize that all sub-optimization problems remain unchanged, except for the A_k update and an additional dual ascent step.

The heat diffusion and thresholding steps are not affected by the extra terms in the functional. Instead, the corresponding first variation is incorporated in the first, ODE step (3.17):

$$\frac{\partial A_k(\vec{x})}{\partial t} = -\beta + 2\rho v_k(\vec{x})\left(f(\vec{x}) - \sum A_i(\vec{x})v_i(\vec{x}) + \frac{\lambda(\vec{x})}{\rho}\right) - 2\rho'\left(\sum A_k(\vec{x}) - 1\right) - \lambda'(\vec{x}). \quad (4.4)$$

4.2 Projection: Multiphase MBO and rearrangement

Instead of the additional penalty and Lagrangian multiplier term, the partitioning constraint can be dealt with by the rectification step in the MBO-like part. Indeed, the partitioning problem corresponds to a multiphase interface problem. The fundamental idea is to propagate the data-ODE (3.17) and the heat-diffusion PDE (3.18) on each support function A_k individually, but to replace the individual thresholding step (3.19) by a single, common "winner-takes-it-all" rectification. This idea has been discussed more rigorously in [EO15], and is related to the rearrangement algorithm for the discrete graph partitioning problem [OWO14].

The projection-based partitioning update for A_k becomes:

1. Area penalty and reconstruction fidelity ODE propagation for each mode k, according to (3.17).

2. Heat diffusion PDE for each mode k according to (3.18).

3. "Winner-takes-it-all" rectification; Projection of the intermediate A_k on the feasible set $A_k \in \{0,1\} \cap \sum A_k = 1$:

$$A_k^{t+1} = \begin{cases} 1 & \text{if} \quad k = \arg\max_i A_i, \\ 0 & \text{otherwise.} \end{cases} \tag{4.5}$$

For an application of the same strategy to graph-based image processing, see [GMB14, HSB14]. The modified 2D-TV-VMD algorithm with segmentation constraint is given in algorithm 4, while illustrative examples are shown in figures 7.3 *et seqq.*

Algorithm 4 2D-TV-VMD with segmentation constraint

```
Input: signal $f(\vec{x})$, number of modes $K$, parameters $\alpha_k$, $\beta_k$, $\gamma_k$, $\rho$, $\rho_k$, $t$, $\tau$, $\tau_k$, $\epsilon$.
Output: modes $u_k(\vec{x})$, domain partitioning support functions $A_k(\vec{x})$, center frequencies $\vec{\omega}_k$.
Initialize $\{\omega_k^0\}$, $\{u_k^0\} \leftarrow 0$, $\{v_k^0\} \leftarrow 0$,$\{A_k^0\} \leftarrow 1$, $\{\lambda_k\}^0 \leftarrow 0$, $\lambda^0 \leftarrow 0$, $n \leftarrow 0$
repeat
    $n \leftarrow n + 1$
    for $k = 1 : K$ do
        Create 2D mask for analytic signal Fourier multiplier:
            $\mathcal{H}_k^{n+1}(\vec{\omega}) \leftarrow 1 + \operatorname{sgn}(\langle \vec{\omega}_k^n, \vec{\omega} \rangle)$
        Update $\hat{u}_{AS,k}$:
            $\hat{u}_{AS,k}^{n+1}(\vec{\omega}) \leftarrow \mathcal{H}_k^{n+1}(\vec{\omega}) \left[ \frac{\rho_k \hat{v}_k^n(\vec{\omega}) - \hat{\lambda}_k^n(\vec{\omega})}{\rho_k + 2\alpha_k |\vec{\omega} - \vec{\omega}_k^n|^2} \right]$
        Retrieve $u_k$:
            $u_k^{n+1}(\vec{x}) \leftarrow \Re\left(\mathcal{F}^{-1}\left\{\hat{u}_{AS,k}^{n+1}(\vec{\omega})\right\}\right)$
        Update $v_k$:
            $v_k^{n+1}(\vec{x}) \leftarrow \frac{\rho A_k^n(\vec{x}) \left( f(\vec{x}) - \sum_{i<k} A_i^n(\vec{x}) v_i^{n+1}(\vec{x}) - \sum_{i>k} A_i^n(\vec{x}) v_i^n(\vec{x}) + \frac{\lambda^n(\vec{x})}{\rho} \right) + \rho_k u_k^{n+1}(\vec{x}) + \lambda_k^n(\vec{x})}{\rho A_k^n(\vec{x})^2 + \rho_k}$
        Update $\vec{\omega}_k$:
            $\vec{\omega}_k^{n+1} \leftarrow \frac{\int_{\mathbb{R}^2} \vec{\omega} |\hat{u}_{AS,k}^{n+1}(\vec{\omega})|^2 d\vec{\omega}}{\int_{\mathbb{R}^2} |\hat{u}_{AS,k}^{n+1}(\vec{\omega})|^2 d\vec{\omega}}$
        Dual ascent $u$-$v$ coupling:
            $\lambda_k^{n+1}(\vec{x}) \leftarrow \lambda_k^n(\vec{x}) + \tau_k \left( u_k^{n+1}(\vec{x}) - v_k^{n+1}(\vec{x}) \right)$
    end for
```

Algorithm 4 2D-TV-VMD with segmentation constraint (continued)

```
    for $k = 1 : K$ do
        Update $A_k$ through time split ODE and PDE propagation:
        $A_k^{n+1/3}(\vec{x}) \leftarrow \frac{1}{1+2t\rho(v_k^{n+1}(\vec{x}))^2}\Big(A_k^n(\vec{x}) + t\Big(-\beta_k + 2\rho v_k^{n+1}(\vec{x})\Big((f(\vec{x}) - \sum_{i<k} A_i^{n+2/3}(\vec{x})v_i^{n+1}(\vec{x}) - \sum_{i>k} A_i^n(\vec{x})v_i^{n+1}(\vec{x}) + \frac{\lambda^n}{\rho}\Big)\Big)\Big)$
        $\hat{A}_k^{\,n+2/3}(\vec{\omega}) \leftarrow \frac{\hat{A}_k^{n+1/3}(\vec{\omega})}{1+t\gamma_k|\vec{\omega}|^2}$
    end for
    for $k = 1 : K$ do
        Rectify $A_k$ through winner-takes-it-all:
        $A_k^{n+1}(\vec{x}) = \begin{cases} 1 & \text{if} \quad k = \arg\max_i A_i^{n+2/3}(\vec{x}) \\ 0 & \text{otherwise} \end{cases}$
    end for
    Dual ascent data fidelity:
    $\lambda^{n+1}(\vec{x}) \leftarrow \lambda^n(\vec{x}) + \tau\left(f(\vec{x}) - \sum_k A_k^{n+1}(\vec{x})v_k^{n+1}(\vec{x})\right)$
until convergence
```

CHAPTER 5

Lattice Segmentation

Until now, our decomposition associates one spatial characteristic support function, A_k, with only one intrinsic mode function, u_k. This results in a simple decomposition where each spatial region has exactly one simple oscillation. Let us now consider a case where the image is composed of regions not corresponding to plane waves, but combinations of simple oscillatory patterns, such as a checkerboard or hexagonal pattern. Microscopy of single-molecule layers, colloids, and crystal grains have such patterns. In biochemistry and nanoscience, the decomposition of such microscopy images into regions of homogeneity provides a necessary mechanic for further downstream analyses.

In microscopy, a crystal image contains different mesoscopic grains, where each grain typically can be a homogeneous, lattice region. Each grain has different spatial periodicities, depending on the crystal lattice structure. These structures are modelled by Bravais lattices, which, depending on the 2D crystalline arrangement, come in five forms: oblique, rectangular, centered rectangular, hexagonal, and square. Thus a grain's Fourier spectrum has several distinct peaks, associated with the various cosine waves that constitute the pattern, which share a common spatial support (function). For example, a grain in a homogeneously hexagonal lattice patch would have three coupled peaks in the spectral half-space. Grains differ by orientation, so it is interesting to find the grain supports, their boundaries and defects, and the Fourier peaks associated with each grain. A crystal image composed of such grains can be considered as an assemblage of 2D general intrinsic mode type functions with non-overlapping supports, specified propagating directions and smoothly varying local wave vectors. A recent state-of-the-art method uses 2D synchrosqueezed transforms together with slow-oscillating, global-structure providing functions, known as shape functions, in order to model atomic crystal images [YLY14]. In general, knowing the Bravais lattice structure yields strong priors on the relative positions of the frequency peaks; here, however, we only make use of the known number of peaks, but not their relative positions.

To accommodate such regions, our spectral image segmentation needs to be adapted to allow for multiple single-Fourier-peak modes to be joined together through a single binary support function. Let $\{u_{kj}\}_j$ denote the set of modes associated with the single binary support function A_k. Each of these modes needs to be individually of small bandwidth, but they contribute to the signal reconstruction jointly through their single support function A_k. This simple modification allows us to *segment* signals into meaningful pieces.

To this end, we modify the spatially disjoint n-D-TV-VMD model (4.2) as follows:

$$\min_{u_{ki}\colon \mathbb{R}^n \to \mathbb{R},\, A_k\colon \mathbb{R}^n \to \{0,1\},\, \vec{\omega}_{ki} \in \mathbb{R}^n} \left\{ \sum_{k,i} \alpha_{ki} \left\| \nabla \left[u_{AS,ki}(\vec{x}) e^{-j\langle \vec{\omega}_{ki}, \vec{x} \rangle} \right] \right\|_2^2 + \sum_k \beta_k \|A_k\|_1 + \sum_k \gamma_k \|\nabla A_k\|_1 \right\}$$

$$\text{s.t.} \quad \forall \vec{x} \in \mathbb{R}^n \colon \begin{cases} \sum_k A_k(\vec{x}) \sum_i u_{ki}(\vec{x}) = f(\vec{x}), \\ \sum_k A_k(\vec{x}) = 1. \end{cases} \tag{5.1}$$

We call this the n-D-TV-VMD lattice segmentation model. The model can be optimized in much the same way as the simpler model (4.2). The only significant difference is in the ODE propagation step of the A_k update: Here, all associated modes u_{ki} (resp. their copies v_{ki}) jointly influence the update of the single A_k. Indeed, (3.17) now becomes:

$$\frac{\partial A_k(\vec{x})}{\partial t} = -\beta + 2\rho \left(\sum_i v_{ki}(\vec{x}) \right) \left(f(\vec{x}) - \sum_l A_l(\vec{x}) \sum_j v_{lj}(\vec{x}) + \frac{\lambda(\vec{x})}{\rho} \right). \tag{5.2}$$

Explicitly modifying the previous algorithms to incorporate this submode coupling is fairly straightforward and left as an exercise to the reader. Examples of image decomposition with submode coupling are shown in figures 7.7–7.10.

CHAPTER 6

Outlier Detection: Artifact Detection and Inpainting

As a final complication regarding crystallography images, we now wish to deal with image features that cannot be explained by the VMD model thus far, such as defects and artifacts. While artifacts can be due to acquisition noise or sample impurities (accidental or intended), defects are irregularities in the regular crystal structure, within crystal grains, or more frequently at the grain boundaries. In imaging terms, these are characterized by a stark deviation from the regular spatial pattern modeled by the band-limited modes of the VMD model. In the presence of imaging noise, one naturally relaxes the data-fidelity constraint by just a quadratic penalty, i.e., not making use of a Lagrangian multiplier. Therefore, unless otherwise accounted for, such defects and artifacts appear in the data-fidelity residual, but due to their non-Gaussian nature as strong outliers will also affect and deteriorate the mode decomposition. It is imperative, therefore, to address these features more specifically beyond making Gaussian noise assumptions.

6.1 Artifact indicator function

Recently, a dynamic artifact detection model was introduced in the framework of classical Chan-Vese image segmentation [ZAS15]. There, individual pixels were eliminated from the region-based segmentation terms to prevent skewing and misleading the segmentation. This method is related to similar approaches in occlusion detection in optical flow [ARS11] and salt-and-pepper denoising [Yan13]. Here, the goal is to isolate defects and artifacts from interfering with the regular modes.

We introduce an artifact indicator function,

$$\chi \colon \mathbb{R}^n \to \{0, 1\},$$

where for each pixel a 1 denotes an artifact, and 0 absence thereof. We use this artifact indicator function to limit the data-fidelity constraint to non-artifact regions, only, e.g.,

$$\forall \vec{x} \in \mathbb{R}^n \mid \chi(\vec{x}) = 0 \colon \sum_k A_k(\vec{x}) u_k(\vec{x}) = f(\vec{x}). \tag{6.1}$$

This is equivalent to

$$\forall \vec{x} \in \mathbb{R}^n \colon \sum_k (1 - \chi(\vec{x})) A_k(\vec{x}) u_k(\vec{x}) = (1 - \chi(\vec{x})) f(\vec{x}), \tag{6.2}$$

where $(1 - \chi(\vec{x})) = 1$ in regions not classified as artifacts, which is where data fidelity is to be enforced. A similar modification can be made to all data-fidelity constraints of the previous models.

6.2 Defect and artifact detection and inpainting

We have not described, so far, how the values of the binary defect and artifact indicator function χ are to be determined, in the first place. While there are reasonable grounds to believe that these defect and artifact locations could be heuristically identified from images in preprocessing, we want to integrate this detection process into the very same decomposition model.

At this point, we do not have a concise and simple characterization of the shape and appearance of defects and artifacts, and for the general case we even want to avoid including too many such priors. Instead, we characterize lattice defects and image artifact locations by what they are not; indeed, at these locations the image simply fails to be sufficiently well modeled by the band-limited modes extracted nearby. We thus decide to classify a certain pixel $f(\vec{x})$ as an artifact or defect, $\chi(\vec{x}) = 1$, if the incurred data-fidelity cost would be too large, locally, otherwise. This is most simply achieved by including an L^1-term on χ.

We modify the constrained n-D-TV-VMD cost functional (3.2) to become the n-D-TV-XVMD (with artifact detection) functional as follows:

$$\min_{u_k\colon \mathbb{R}^n\to\mathbb{R},\, A_k,\chi\colon \mathbb{R}^n\to\{0,1\},\, \vec{\omega}_k\in\mathbb{R}^n} \left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j\langle\vec{\omega}_k,\vec{x}\rangle} \right] \right\|_2^2 + \beta_k \|A_k\|_1 + \gamma_k \|\nabla A_k\|_1 + \delta \|\chi\|_1 \right\}$$
$$\text{s.t.} \quad \forall \vec{x}\in\mathbb{R}^n\colon \sum_k (1-\chi(\vec{x})) A_k(\vec{x}) u_k(\vec{x}) = (1-\chi(\vec{x})) f(\vec{x}). \quad (6.3)$$

The corresponding unconstrained saddle point problem (without Lagrange multiplier on the data-fidelity) then becomes:

$$\mathcal{L}(\{u_k\},\{v_k\},\{A_k\},\{\omega_k\},\chi,\{\lambda_k\}) :=$$
$$\left\{ \sum_k \alpha_k \left\| \nabla \left[u_{AS,k}(\vec{x}) e^{-j\langle\vec{\omega}_k,\vec{x}\rangle} \right] \right\|_2^2 + \beta_k \|A_k\|_1 + \gamma_k \|\nabla A_k\|_1 + \delta \|\chi\|_1 \right.$$
$$\left. +\rho \left\| (1-\chi(\vec{x}))(f(\vec{x}) - \sum A_k(\vec{x}) v_k(\vec{x})) \right\|_2^2 + \sum_k \rho_k \|u_k(\vec{x}) - v_k(\vec{x})\|_2^2 + \langle \lambda_k(\vec{x}), u_k(\vec{x}) - v_k(\vec{x}) \rangle \right\}. \quad (6.4)$$

It is important to note that the masking only impacts the data-fidelity evaluation domain, while all other terms are not affected. Indeed, only two sub-minimization steps will be altered by the introduction of the $(1-\chi)$-term:

1. the area penalty and reconstruction fidelity ODE (3.17) will collapse to just $\partial_t A_k(\vec{x}) = -\beta_k$ whenever $\chi(\vec{x}) = 1$ (and remain unchanged, otherwise). In particular, the TV- and L^1-terms on the binary support functions A_k will now exclusively drive the evolution of the latter whenever a location is marked as artifact, since the data-fidelity constraint is the only link between modes and support functions.

2. Similarly, the update (3.15) of v_k collapses to $v_k^{t+1}(\vec{x}) = u_k(\vec{x}) + \lambda_k(\vec{x})/\rho_k$ when $\chi(\vec{x}) = 1$, which effectively unlinks the local mode estimate from the observed data and simply *in-paints the artifact regions* by Fourier interpolation of the modes.

On the other hand, the estimation of the artifact indicator function χ itself also leads to a straightforward optimization step. The binary optimization can be carried out independently for each pixel, and the optimal $\chi^*(\vec{x})$ chooses between paying data-fidelity penalty versus artifact cost δ, as follows:

$$\chi^*(\vec{x}) = \begin{cases} 0 & \text{if} \quad \rho(f(\vec{x}) - \sum A_k(\vec{x}) v_k(\vec{x}))^2 \leq \delta \\ 1 & \text{otherwise} \end{cases} \tag{6.5}$$

This thresholding scheme has an immediate interpretation from a hypothesis-testing perspective. Indeed, if we consider the data-fidelity weight ρ to be the precision of the implicitly assumed Gaussian noise distribution, then the expression $\rho(f(\vec{x}) - \sum A_k(\vec{x}) v_k(\vec{x}))^2$ represents the squared z-score (standard score) of the local image intensity under such a noise distribution. This squared z-score is compared against the threshold δ. The artifact classification is effectively a concealed statistical hypothesis z-test of the pixel intensity with a Gaussian distribution

$$p(f(\vec{x})) = \mathcal{N}(f(\vec{x}) \mid \sum A_k(\vec{x}) v_k(\vec{x}), \rho^{-1})$$

as null-hypothesis H_0, and a pixel is classified as an artifact (H_1) if the z-score of its intensity is more extreme than $\sqrt{\delta}$. The model parameter δ is thus intimately related to the level of statistical significance attached to the artifact classification and its expected false positives rate.

Again, in the interest of conciseness, we leave the modification of the algorithms to include the artifact detection and inpainting terms as an exercise for the reader. An inpainting example is illustrated in figure 7.6.

CHAPTER 7

Experiments and Results

We have implemented the above three algorithms 2–4, including the submode coupling of section §5 and the artifacts detection and inpainting (§6) extensions, in MATLAB®. The algorithms can be implemented in a single code file, because they are mostly generalizations of each other.

In the implementation, we make two deliberate choices that have not been discussed, so far. The first choice is with respect to initialization of the center frequencies, where we include four options:

1. initialization of frequencies uniformly spread on a circle (deterministic),
2. random initialization on the positive half-space,
3. user selection through graphical user interface, and
4. user input as parameters.

Unless otherwise noted, all the examples shown below make use of the deterministic radial frequency initialization scheme.

The second particularity is with respect to model selection 2D-VMD, 2D-TV-VMD, and 2D-SEG-VMD. Indeed it is useful in practice to initialize the TV-VMD model by some iterations of unrestricted 2D-VMD, in order to settle the center frequencies close to the optimal location; and similarly, the segmentation model is best initialized based on the outcome of 2D-TV-VMD optimization. We will thus always start optimizing in 2D-VMD mode, and over the iterations, switch to the two more complicated models at user-defined time-points (which may be set to infinity, thereby producing results of simpler models as final output).

Our implementation is publicly available for download at `http://www.math.ucla.edu/~zosso/code.html`, and on MATLAB Central.

7.1 Synthetic overlapping texture decomposition

The first, synthetic image is a composition of spatially overlapping basic shapes, more precisely six ellipses and a rectangle, with frequency patterns varying in both periodicity and direction, courtesy of J. Gilles [Gil12]. The spectrum is ideal for segmentation due to modes being deliberately both well spectrally isolated and narrow-banded. The resolution of the synthetic image is 256×256.

We feed the synthetic image to our models and show the resulting decompositions for both 2D-VMD and 2D-TV-VMD models in figure 7.1. The parameters are[1]:

K	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	t
5	1000	0.5	500	∞	10	10	2.5	2.5	1.5

In addition, the center frequency of the first mode is held fixed at $\omega_1 = 0$ to account for the DC component of the image. As a result, the first mode contains the solid ellipse and rectangle, while the four remaining decompositions in figure 7.1 show clear separation of the patterned ellipses.

In the simple 2D-VMD model of figure 7.1(e), due to the solid pieces having sharp edges, their

---

[1] Of course, the simpler 2D-VMD model only uses a subset of these parameters, for the support functions are fixed at $A_k = 1$ uniformly.

spectra are not band-limited and only smoothed versions are recovered. This is naturally paired with the two lower frequency modes absorbing residual boundary artifacts of the DC component, and ghost contours appearing in these modes.

The spatially compact 2D-TV-VMD model, figure 7.1(f)–(h), however, can handle sharp boundaries through the support functions A_k, while the modes u_k can smoothly decay. The resulting masked modes, $A_k u_k$, are thus clean and sharp.

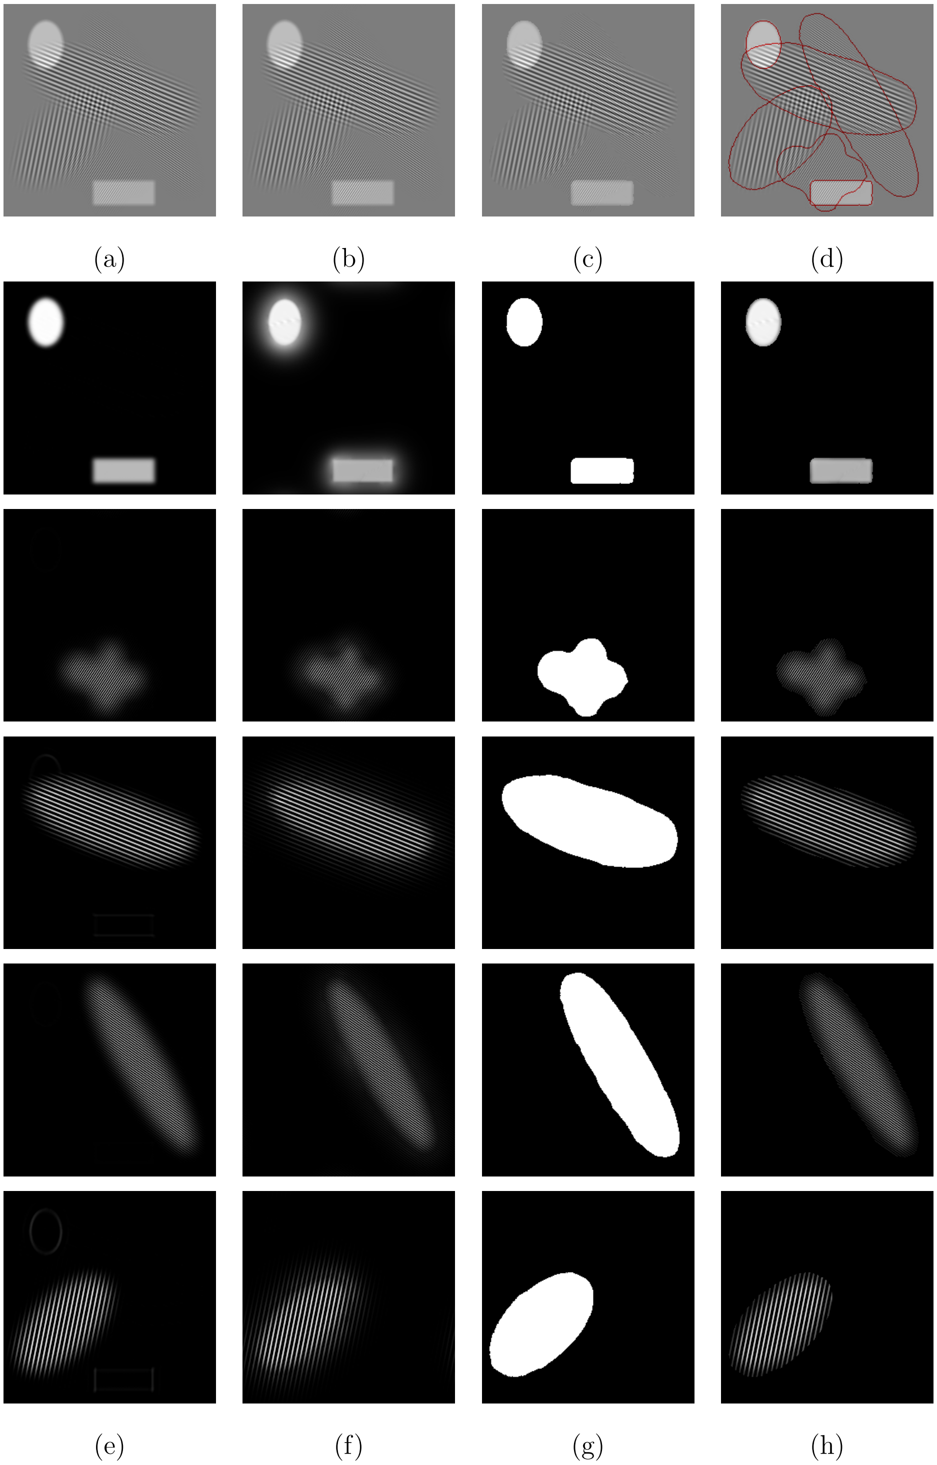

(a) (b) (c) (d)

(e) (f) (g) (h)

Figure 7.1: Synthetic overlapping texture (a) Input image f. (b) 2D-VMD reconstruction $\sum_k u_k$. (c) Compactly supported 2D-TV-VMD reconstruction $\sum_k A_k u_k$. (d) Support boundaries overlaid onto original image. (e) 2D-VMD modes u_k. (f) 2D-TV-VMD modes u_k. (g) Detected supports A_k. (h) Masked modes $A_k u_k$.

7.2 Overlapping chirps

The second example problem is still synthetic, but the modes have non-trivial Fourier support. More precisely, the synthetic image is a superposition of three compactly supported yet spatially overlapping 2D chirps (see figure 7.2). Starting from radial initialization, we let our algorithm determine the correct support and appropriate center frequencies for this problem, based on the following parameters:

K	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	t
3	2000	1	1000	∞	7	10	1	1	1

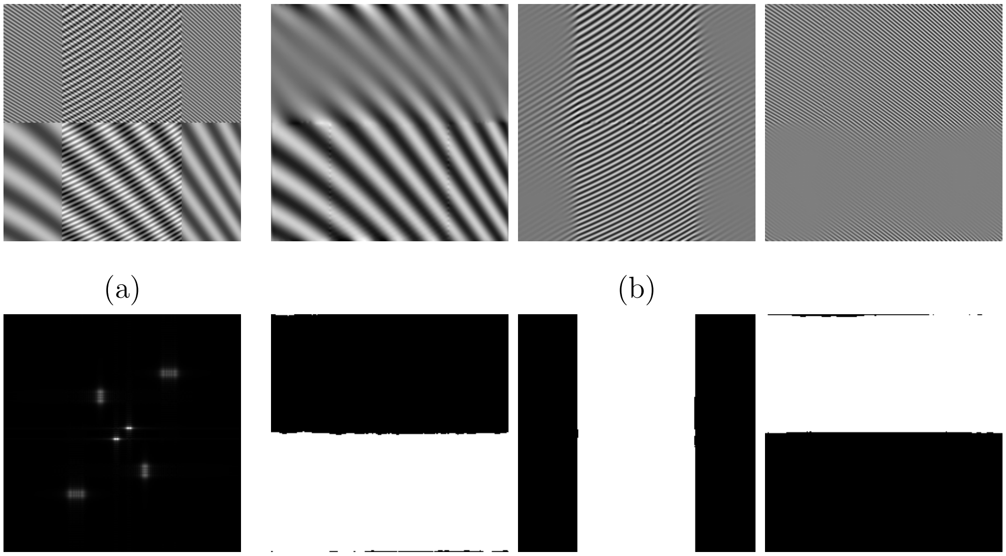
(a) (b)

(c) (d)

Figure 7.2: Chirp decomposition. (a) Input signal f. (b) 2D-TV-VMD modes u_k. (c) Fourier spectrum $\hat{f}$. (d) Determined supports A_k. See §7.2.

The resulting decomposition is accurate with only little error on the true support functions. The modes are spectrally clean. It is interesting to observe how our model extrapolates the modes outside their rectangular domain boundaries. Note that the decay distance correlates with the wave-length of the mode.

7.3 Textural segmentation for denoising

The two examples encountered so far were noise-free and perfect reconstruction was possible through the use of Lagrange multipliers ($\tau, \tau_k > 0$). In the presence of noise, however, enforcing strict data fidelity may be inappropriate, and instead relying on just the quadratic penalty to promote data-fidelity is the proper way to go. This is easily achieved by preventing the Lagrangian multipliers from updating: $\tau, \tau_k = 0$. As a result, the noise can be handled with a residual slack between the splitting variables. In particular, the quadratic penalty term corresponds to a Gaussian noise assumption, where the penalty coefficients ρ, ρ_k relate to the noise precision.

Here, we explore the idea of using the slack in the absence of Lagrangian multipliers for denoising based on spectral sparsity. To this end, we construct a four-quadrant, non-overlapping unit-amplitude cosine-texture image with different levels of noise, shown in figure 7.3. Because the quadrants are non-overlapping, we are interested in the output of the 2D-SEG-VMD model using the following parameters:

K	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	t
4	3500	1.5	750	∞	7	10	0	0	1

Without the Lagrangian multipliers active, it is important to realize that the two copies of the modes, u_k and v_k, may be different; and that u_k is the potentially cleaner copy of the two.

In figure 7.3, we can see that even for important noise levels, the partition is recovered with good precision (red contours). In addition, the recovered composite of the four masked modes is very clean, seemingly irrespective of the degrading noise level.

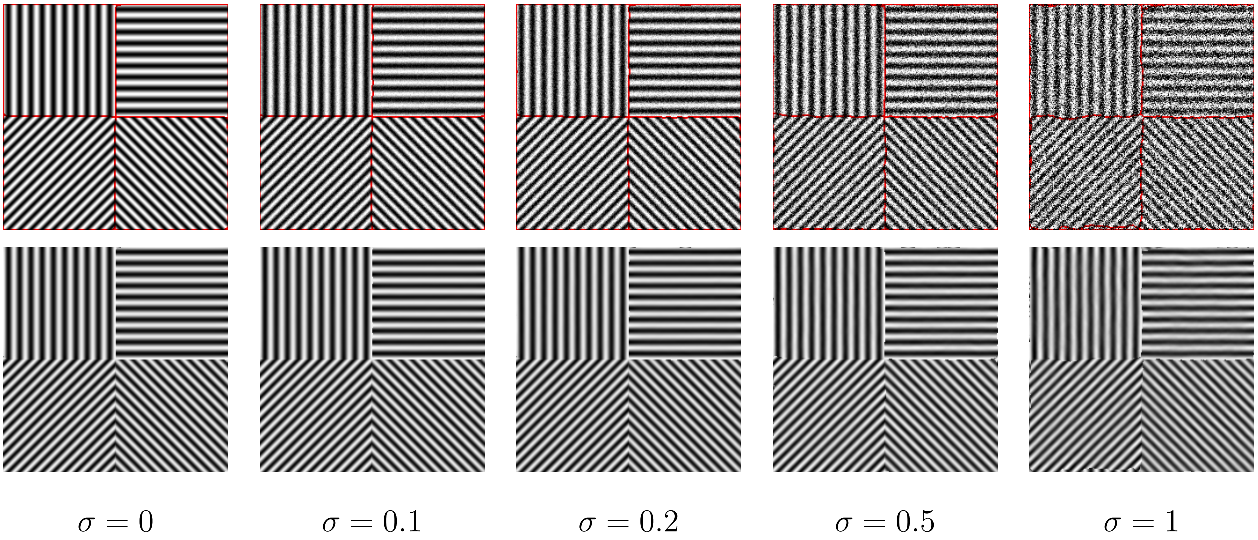
$\sigma = 0$ $\sigma = 0.1$ $\sigma = 0.2$ $\sigma = 0.5$ $\sigma = 1$

Figure 7.3: Denoising. Noise standard deviation σ. Top: noisy f with detected phase borders (red). Bottom: denoised signal $\sum_k A_k u_k$. See §7.3.

7.4 Segmentation of peptide β-sheets

The next test case are two scanning tunneling microscopy (STM) images of peptide β-sheets bonding on a graphite base, courtesy of the Weiss group at the California NanoSystems Institute (CNSI) at UCLA, [CTS13]. The peptide sheets grow in regions of directional homogeneity and form natural spatial boundaries where the regions meet. It is important to scientists to have accurate segmentation for their dual interests in complementary analysis of the homogeneous regions and their boundaries. Identifying regions of homogeneity enables the subsequent study of isolated peptide sheets of one particular bonding class. For these types of scans, manually finding the boundaries is a tedious problem that demands the attention of a skilled scientist on a rote task. In addition to speed and automation, the proposed 2D-VMD is superior in accuracy to manual boundary identification due to regions potentially having very similar patterns, of which the orientation differs by only a few degrees, difficult to discern by eye.

Nanoscale images such as these are a useful testbed since data are often oversampled relative to the smallest observable features, atoms and molecular parts. Also, segmentation in one imaging modality can be used to guide segmentation or data acquisition in a complementary imaging mode [BBB12, CSW11, HKG09, MYY10, TSH15].

The first example, shown in figure 7.4, is a 512 × 512 false-color image, of which we only consider the average intensity across color channels as a proxy, *in lieu* of the actual raw data produced by the microscope. Also, as classical pre-processing step, we apply a *Laplacian of Gaussians* (LoG) band-pass filter to the image in order to remove both some noise and the DC component. Expert inspection suggests that there are six different grain orientations represented in this image. We perform 2D-VMD, 2D-TV-VMD, and 2D-SEG-VMD using these parameters:

K	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	t
6	2000	1	250	∞	7	10	0	0	2.5

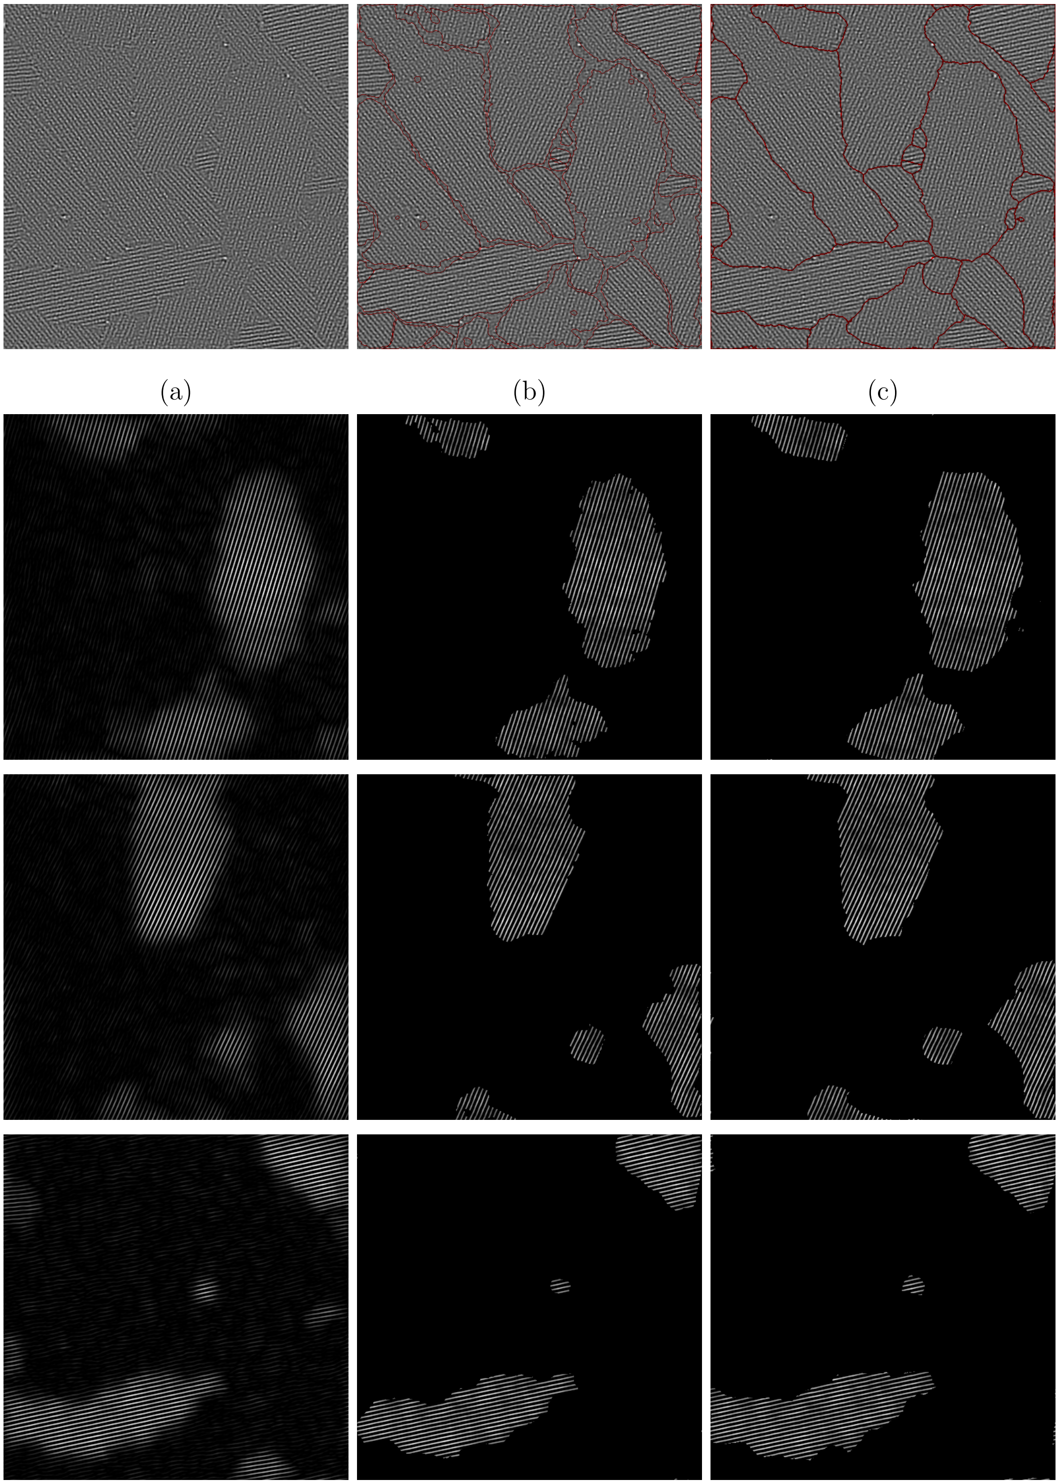

(a) (b) (c)

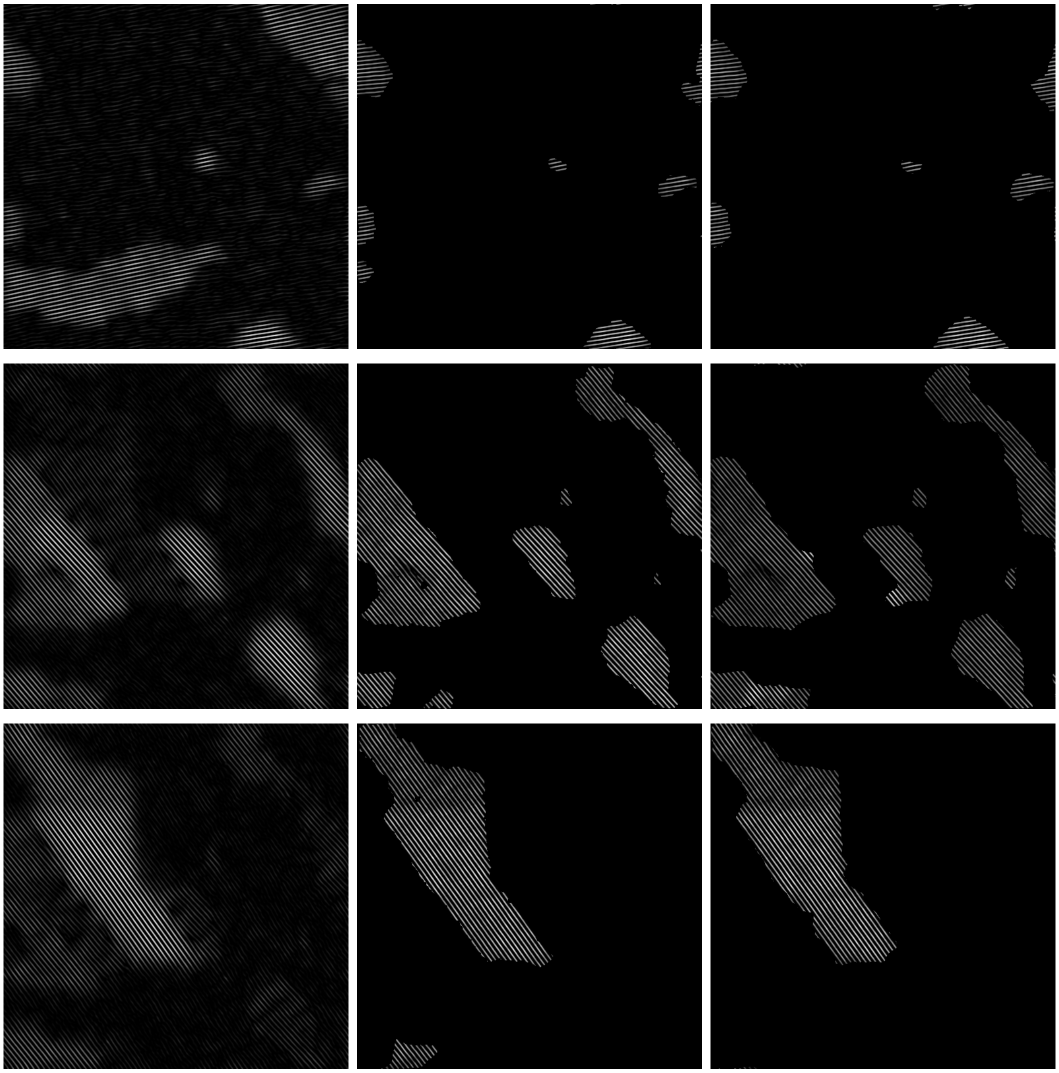

(d) (e) (f)

Figure 7.4: Scanning tunneling microscopy (STM) image of peptide β-sheets, 512×512 (I). (a) Input f. (b) 2D-TV-VMD boundaries (red). (c) 2D-SEG-VMD partition (red). (d) 2D-VMD modes u_k. (e) 2D-TV -VMD modes $A_k u_k$. (f) 2D-SEG-VMD modes $A_k u_k$. See §7.4 in the text for details and discussion.

The recovered modes are shown in figure 7.4(d)–(f). The unconstrained 2D-VMD model produces overly smooth modes without clear boundaries. The compactly supported 2D-TV-VMD model yields modes with sharp delineation. As can be seen from the grain boundaries overlaid to the input image, in figure 7.4(b), the modes are not overlapping, but do not cover the entire image domain, leaving unaccounted space at the grain boundaries. This problem is effectively addressed by the addition of the segmentation constraint, as seen by the boundaries in 7.4(c).

The second example, shown in figure 7.5, is believed to consist of only three main grain orientations. This 512×512 image is of the same type as the previous example and pre-processed in the same way. The image exhibits strong singular spots due to additional material deposition on the sample surface. In order to address these outliers, we make use of the artifact detection and inpainting extension, for δ finite:

K	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	t
3	2000	1	75	3.5	7	10	0	0	2.5

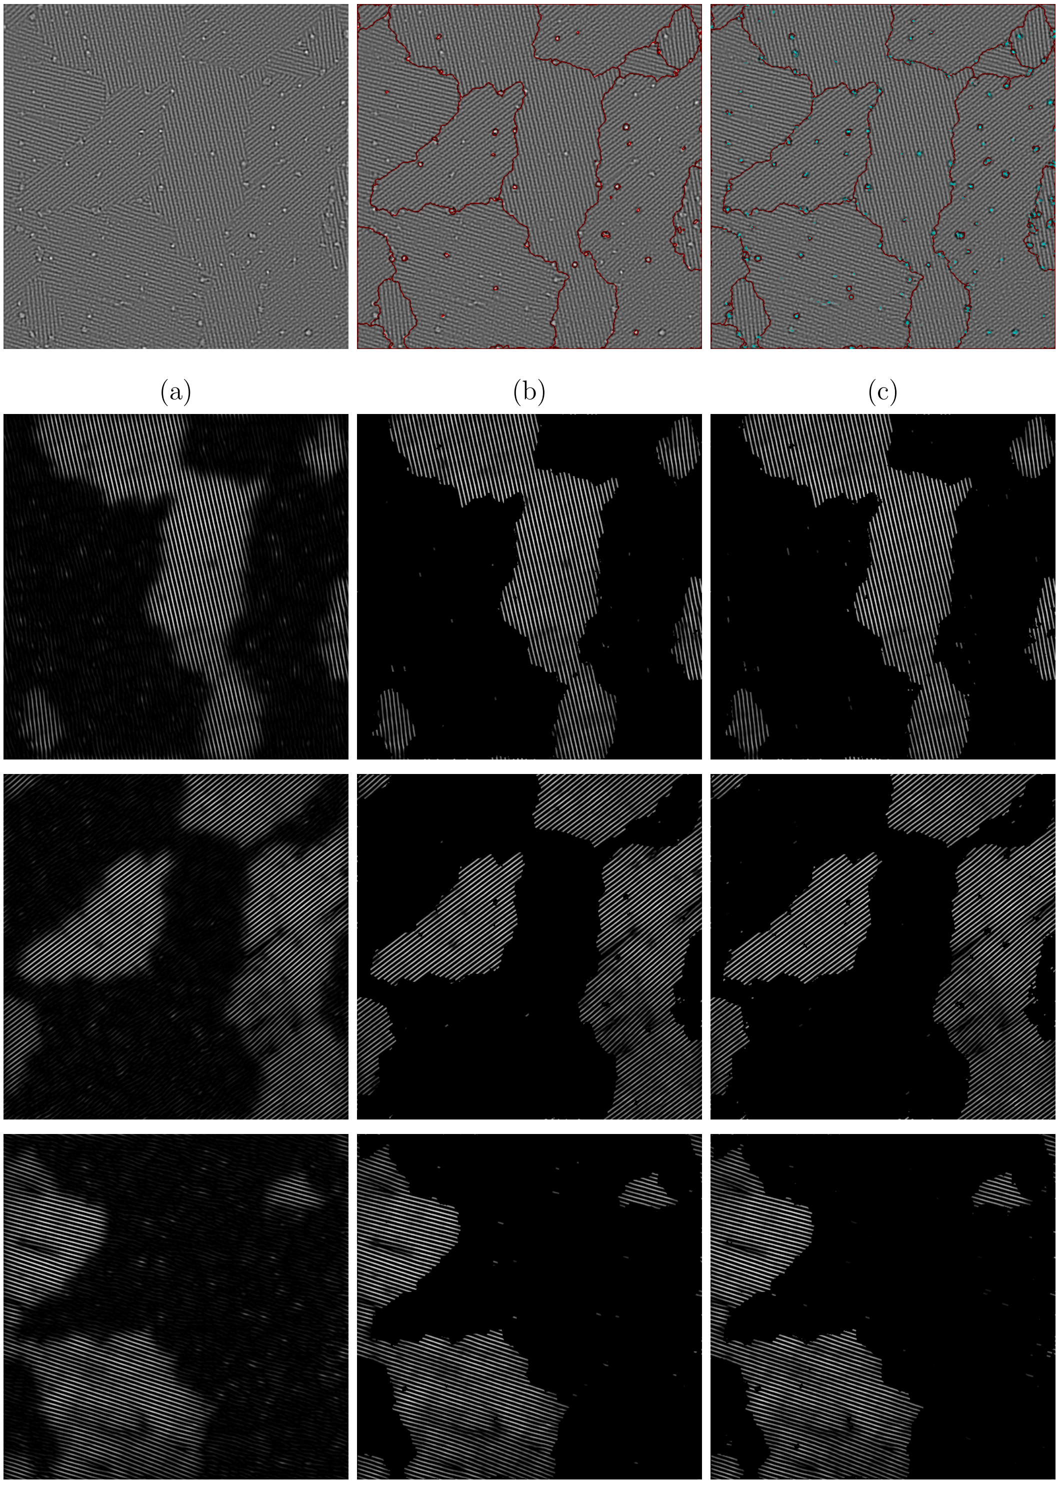
(a) (b) (c)

(d) (e) (f)

Figure 7.5: STM of peptide β-sheets, 512×512 (II). (a) Input f. (b) 2D-SEG-VMD partition (red). (c) Partition (red) with enabled artifact detection (cyan). (d) 2D-VMD modes u_k. (e) 2D-SEG-VMD modes $A_k u_k$. (f) Modes obtained with artifact detection enabled. See §7.4 for details.

While the singular deposits ("artifacts") negatively impact the mode purity for both 2D-VMD and 2D-TV-VMD (figure 7.5(d)–(e)), this effect is partially alleviated by the automatic detection and inpainting capability of the artifacts-extension (figure 7.5(f))[2]. In addition to the outlined grain boundaries (red), the location of the detected artifacts is highlighted in cyan, in figure 7.5(c). Note that the artifact detection also allows spotting at least some of the grain defects, in addition to the deposits.

7.5 Inpainting

Here, we are interested in exploiting the model's capability of intrinsically inpainting the modes (and therefore the input image) in regions that are labeled as artifacts/outliers. To this end, we construct a simple checkerboard image, which essentially corresponds to a superposition of two cosine-waves with full support each. In addition, portions of the image are corrupted by "pencil-scribble", as shown in figure 7.6(a). We set up the model as a two-modes 2D-VMD image decomposition problem, with a finite artifact detection threshold. The data-fidelity Lagrangian is inactive in order to allow some slack (Gaussian noise assumption) and artifact detection, while we maintain an active Lagrangian multiplier on the $u - v$ splitting:

K	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	t
2	1500	n/a	n/a	30	150	20	0	1	n/a

---

[2] Lower artifact threshold δ and higher TV-weight γ_k might increase the mode cleanliness even further.

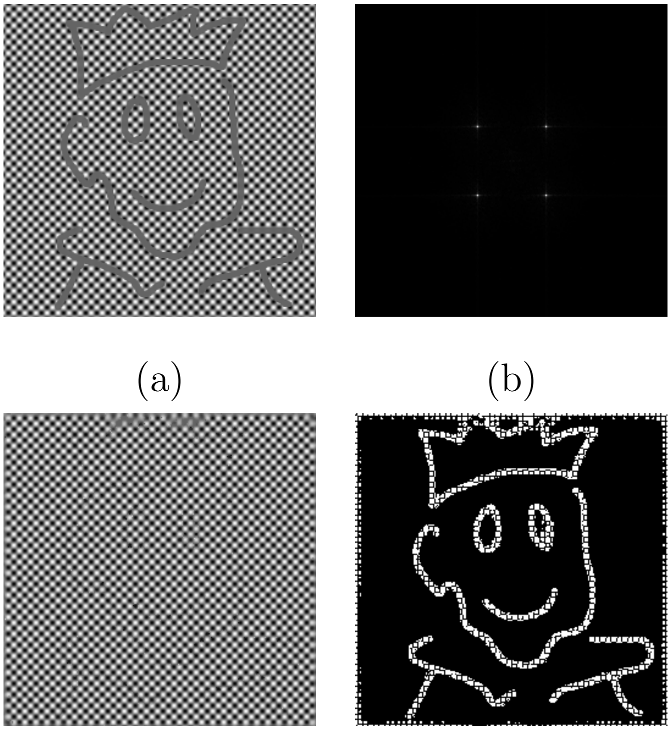
(a) (b) (c) (d)

Figure 7.6: 2D-VMD inpainting. (a) Input image f. (b) Fourier spectrum $\hat{f}$. (c) Recovered modes $\sum_k u_k$. (d) Detected artifacts χ. See §7.5.

As can be seen in figure 7.6(c)–(d), the model succeeds well in detecting the scribble as outliers. In the artifact-labeled image portions, the submodes are inpainted by intrinsic Fourier-interpolation, and as a result, a full checkerboard can be recovered from the decomposition.

7.6 Textural segmentation: Lattices

We finally turn our attention to the segmentation of images with lattice texture, as observed, for example in crystallography and microscopy images of crystalloid samples. The fundamental assumed property of such images is that they consist of K different domains (grains) forming a partition of the image, such that each grain has a distinct lattice texture composed of a superposition of M different essentially wavelike sub-bands. As seen earlier, a checkerboard lattice would consist of a superposition of $M = 2$ orthogonal cosine waves, while a hexagonal lattice consist of $M = 3$ modes differing by 60° rotation. Our model allows for multiple sub-modes u_{ki} to share a common support function A_k, and thus be spatially coupled.

7.6.1 Checkerboard: 2 phases with 2 sub-modes

As a first simple example, we consider the composite of two checkerboard halves, of which one is slightly rotated, as shown in figure 7.7(a). The goal is to find the support of two phases, partitioning the 256×256 image domain, and the respective two sub-modes for each such grain. We run the 2D-SEG-VMD model with the following parameters:

K	M	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	t
2	2	2000	1	250	∞	7	10	0	0	2.5

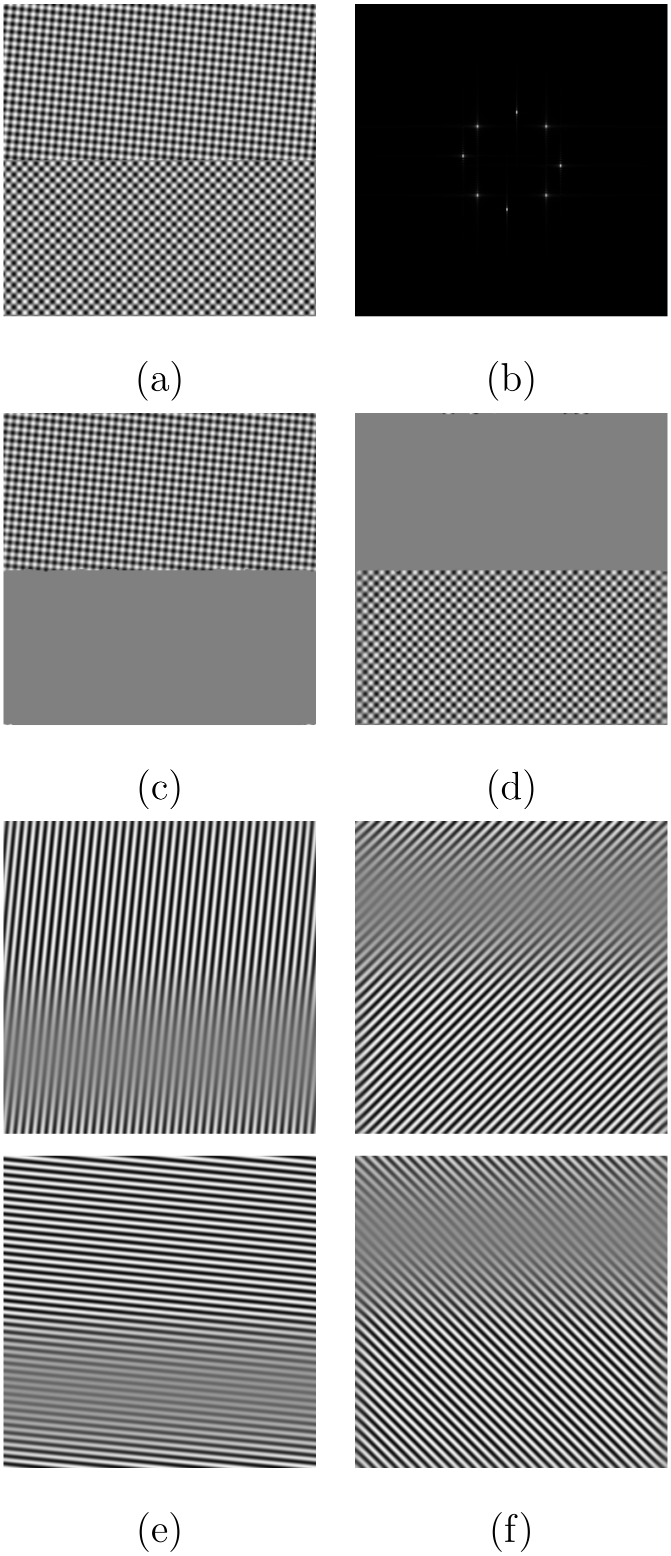

(a) (b)

(c) (d)

(e) (f)

Figure 7.7: Lattice decomposition. (a) Input f. (b) Fourier spectrum $\hat{f}$. (c)–(d) Recovered phases $\sum_i A_k u_{ki}$. (e)–(f) Submodes u_{ki}. See §7.6.1.

The resulting decomposition into the two checkerboard phases, $A_k \sum_i u_{ki}$, is shown in figure 7.7(c)–(d), while the constituting two sub-modes per phase, u_{ki}, are illustrated in figure 7.7(e)–(f).

7.6.2 Hexagonal lattice: 3 phases with 3 sub-modes

A slightly more complicated problem is illustrated in figure 7.8. We start with a tripartite 256×256 image, where each domain consists of an artificial hexagonal lattice pattern, obtained by superposing three cosine waves rotated by $60\,^\circ$ against each other. Each domain has a slightly different lattice orientation $(0\,^\circ, 15\,^\circ, 45\,^\circ)$. Like the previous example, this is a 2D-SEG-VMD problem, this time with three phases and three sub-modes, each. The other parameters remain unchanged:

K	M	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	t
3	3	2000	1	250	∞	7	10	0	0	2.5

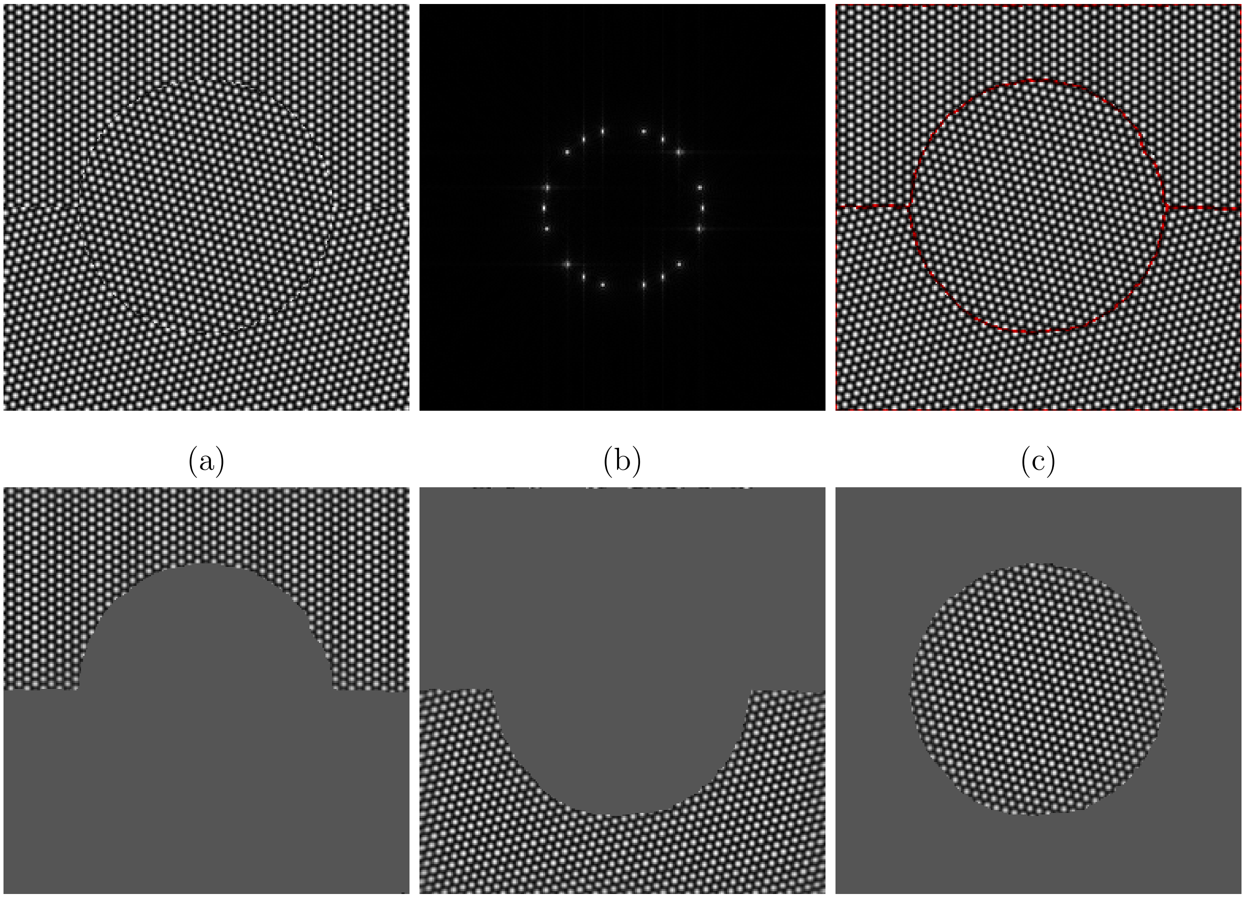

(a) (b) (c)

(d)

Figure 7.8: 3 phase 3 modes. (a) Input f. (b) Fourier spectrum $\hat{f}$. (c) 2D-SEG-VMD partition (red). (d) Phases $A_k \sum_i u_{ki}$. See §7.6.2.

As can be seen in figure 7.8(c)–(d), the recovered phases and their boundaries are very precise. Note that this decomposition involves the identification of nine center frequencies and associated wave functions, and the delineation of three support functions partitioning the image domain.

7.6.3 Simulated hexagonal crystal

The 3-phase-3-waves hexagonal lattice image of the previous sub-section was an idealized synthetic version of what real world acquired images of hexagonally arranged crystal structures might look like. In an attempt to make the problem more realistic, we created a more complicated synthetic lattice image as follows: We predefine a 5-partition of the 256 × 256 image domain. In each domain, individual pixels corresponding to approximate "bubble locations" of the crystal lattice are activated. The exact center position is affected by discretization noise (the pixel locations are obviously limited to the Cartesian grid) as well as additional, controllable jitter. The resulting "nail board" is then convoluted with a circular point spread function designed to mimic the approximate appearance of an individual lattice element, and Gaussian white noise is added. An example is shown in figure 7.9(a). Due to this construction the grain boundaries exhibit very irregular defects. All of these complications make the resulting image much more interesting and challenging to segment.

In a first, simple attempt, we configure the 2D-SEG-VMD algorithm as follows:

K	M	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	t
5	3	2000	1	250	∞	7	10	1	1	2.5

In contrast to the actually noise-free preceding examples, here, we enforce data-fidelity strictly by picking $\tau = \tau_k = 1$, so as to make sure the phases and modes pick-up the relevant center frequencies and do not lazily get stuck in local minima (see a discussion in [DZ14] for the role of the Lagrangian multipliers in low-noise regimes). The model is thus obliged to over-explain all image noise (jitter and Gaussian noise) in terms of mode decomposition. As a result, the obtained partition captures the five phases largely, but suffers from strong noise, as shown in the middle of figure 7.9. Most importantly, though, this procedure found the correct 5 × 3 center frequencies.

These correctly identified center frequencies can now be used as a very strong prior when running the 2D-SEG-VMD model a second time, in a different regime with inactive Lagrangian multipliers to allow noise-slack. To this end, we use the obtained center frequencies as user initialization for a second run, with parameters as follows:

K	M	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	t
5	3	$2e4$	1	500	∞	7	10	0	0	2.5

Now, the increased α_k renders the modes more pure, and also keeps the center frequencies from drifting too much, while the partition regularity is regularized slightly stronger (increased γ_k). The main difference are the inactivated Lagrangian multipliers, relaxing the data-fidelity constraint considerably. The resulting decomposition is shown in figure 7.9. In the correctly initialized denoising regime we obtain a very accurate partition and much cleaner crystal grain estimates.

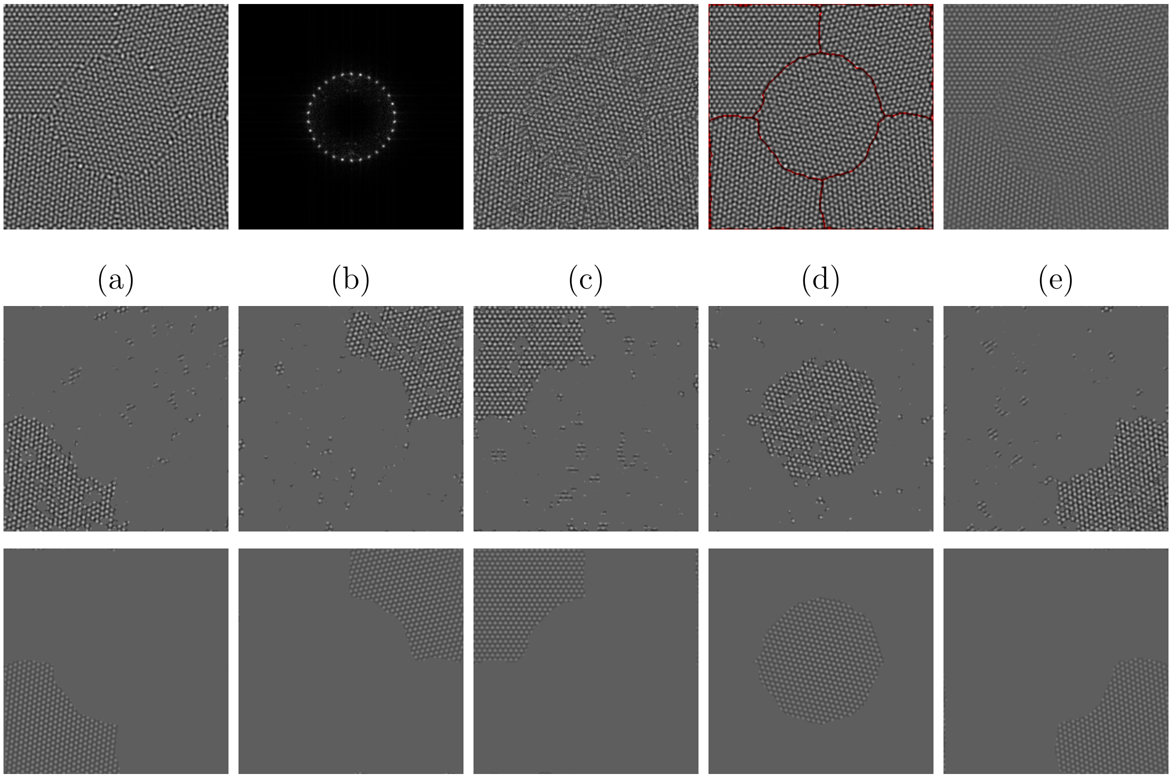

(a) (b) (c) (d) (e)

Figure 7.9: Simulated crystal lattice. 2D-SEG-VMD decomposition in two runs, first with, then without Lagrangian multipliers. See §7.6.3 in text for details and discussion. (a) Input image f. (b) Fourier spectrum $\hat{f}$. (c) First run reconstruction $\sum_{k,i} A_k u_{ki}$. (d) Partition (red) of second run. (e) Reconstruction of second run. Middle row: Phases obtained in first run with with $\tau, \tau_k > 0$ to find correct ω_{ki}. Bottom row: Clean phases $A_k \sum_i u_{ik}$ of second run with $\tau = \tau_k = 0$ and well-initialized ω_{ki}.

7.6.4 Colloidal image

As a last example problem, we consider a bright-field light microscopy image of 10 μm-sized spherical glass particles suspended in water[3]. These glass particles form a collection of small 2D colloidal crystals with grain boundaries between them. These grains have a hexagonal lattice structure similar to the previously considered examples. For our purposes, the original image is cropped, band-pass filtered with a LoG-filter, and downsampled to a final dimension of 256×256. The effective input image is shown in figure 7.10(a).

Visual inspection of the Fourier spectrum suggests that there are probably four different grain orientations to be found in the image (see figure 7.10(b)). We thus configure the 2D-SEG-VMD model with the following parameter choice:

K	M	α_k	β_k	γ_k	δ	ρ	ρ_k	τ	τ_k	t
4	3	2000	1	250	∞	10	50	0.1	0.1	2.5

The resulting grain boundaries shown in figure 7.10(c) should be compared to computationally determined lattice irregularities (grain boundaries, defects) in figure 7.10(d)[4].

---

[3] Image used with permission, courtesy by Richard Wheeler, Sir William Dunn School of Pathology, University of Oxford, UK.

[4] *Ibid.*

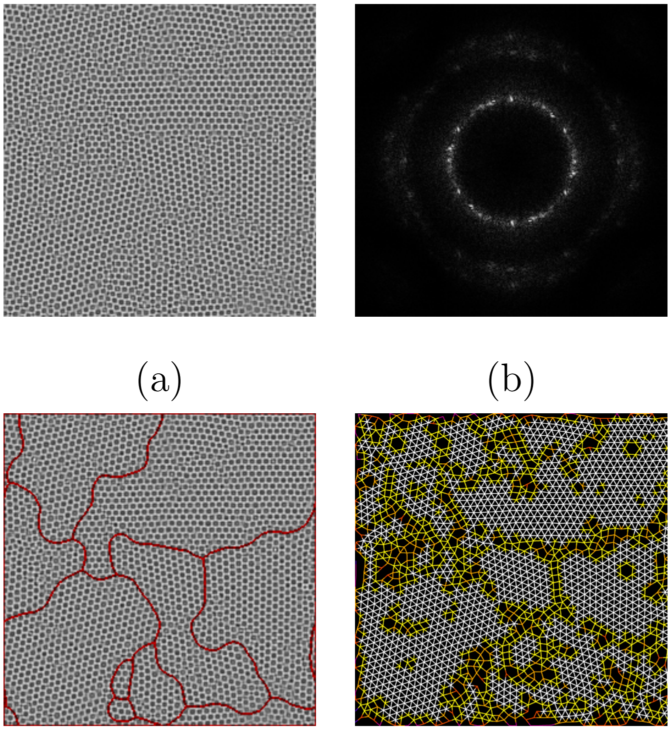

(a) (b)

(c) (d)

Figure 7.10: Bright-field microscopy image of colloidal crystal and its segmentation. Individual beads are $10\,\mu$m in diameter. See §7.6.4. (a) Cropped, LoG-filtered, and downsampled input image f. (b) Fourier spectrum $\hat{f}$. (c) 2D-SEG-VMD 4-partition (red) overlaid on input image. (d) Colloidal connectivity graph for comparison: white edges indicate hexagonal alignment (six equally spaced neighbors) and that a particle is therefore part of a crystalline domain (grain), while colored edges indicate grain boundaries and defects.

CHAPTER 8

Conclusions and Outlook

We have presented a variational method for decomposing a multidimensional signal, $f\colon \mathbb{R}^n \to \mathbb{R}$, (images for $n = 2$) into ensembles of constituent modes, $u_k\colon \mathbb{R}^n \to \mathbb{R}$, intrinsic mode functions which have specific directional and oscillatory characteristics. This multidimensional extension of the variational mode decomposition (VMD) method [DZ14] yields a sparse representation with band-limited modes around a center frequency ω_k, which reconstructs the initial signal, exactly or approximately.

In addition to generalizing the 1D-VMD model to higher dimensions, we introduce a binary support function $A_k\colon \mathbb{R}^n \to \{0, 1\}$ for each mode u_k, such that the signal decomposition obeys $f \approx \sum_k A_k \cdot u_k$. In order to encourage compact spatial support, an L^1 and a TV-penalty term on A_k are introduced. After appropriate variable splitting, we present an ADMM scheme for efficient optimization of this model. This includes MBO-like threshold dynamics to tackle the motion by mean curvature stemming from the support-function regularizing TV-term.

In this general setting, our model allows for spatially compact modes that may be spatially overlapping. By restricting the support functions on the probability simplex, $\sum_k A_k = 1$, the modes have mutually exclusive spatial support and actually form a partition of the signal domain. In this fashion, we obtain an image segmentation model that can be seen as a Chan-Vese-like region-based model, where the homogeneity is assessed through spectral bandwidth. Our variable splitting and the handling of region boundaries through the binary support functions elegantly overcomes the usual tradeoff between spatial and spectral compactness/bandwidth.

In order to deal with images of crystal grains, each region being more complicated than a simple cosine-wave, we introduce the coupling of sub-modes with a single binary support function. This allows the segmentation of crystal grain images, e.g., from microscopy, into respective grains of different lattice orientation. Further, non-Gaussian image noise, outliers, and lattice defects are efficiently addressed by the introduction of an artifact indicator function, $\chi\colon \mathbb{R}^n \to \{0,1\}$.

The work presented in Part II greatly extends the precursor model in [DZ14]. The models and algorithms allow decomposing a signal/image into modes that may:

- have smooth or sharp boundaries (with or without TV/L^1 terms on A_k),
- overlap or form a partition of the domain (image segmentation),
- be essentially wavelike (single mode) or crystalline (coupled sub-modes),
- reconstruct the input image exactly or up to Gaussian noise,
- identify outlier pixels/regions and inpaint them.

Without any doubt, these skills are important to numerous image and signal analysis applications.

Part III

Variational Image Destriping: Stripe Removal via Total Variation with L^1 Fidelity in Remote Sensing Imagery

CHAPTER 1

Introduction

Image striping is a well-known phenomenon that arises in multi-detector imaging systems ranging from pushbroom-type instruments, such as the Airborne Multi-angle Spectro Polarimetric Imager (AirMSPI), to atomic force mircroscopy (AFM). Biases in lateral detection occur due to response variation in spatial detectors, such as in satellite imaging systems, or temporal changes, such as in raster scans. Though these systems are optimally pre-calibrated, post-processing, such as destriping, of data is prerequisite for accurate and valid analyses. Striping removal has been traditionally performed using either statistically based methods, [CI10], [SZ09], or low-pass filtering in the frequency domain [SW88], [Cri89], [HH92], [SL98], [CZ03]. This method, however, does not remove stripes completely and has an effect of blurring the image. More recently, wavelet-based filtering methods have been proposed [TI01], [CY06], [RY07]. However, such methods also blur the images and produce ringing effects in reconstruction.

We follow the pedigree of variational and PDE-based methods applied to images [DB08], [CV01] in order to construct a well-defined, optimizable model yielding fast and quality destriping. During our research, we have come across a similar work achieved by a total variation and framelet regularization model [CL13]. Our model and results were found independently, but share a similar foundation. The focus here is not about creating a sparse wavelet representation of the destriped image, but rather on how to remove the optimal striping mask while preserving high image fidelity. We include detailed derivations and a motivated evolution of the optimization problem with pedagogy in mind so that these novel variational methods can be accessible to all academic disciplines involved with image processing.

Our research is robust to both isotropic and anisotropic versions of total variation, whereas [CL13] argue that the anisotropic case is the only appropriate one. While it is true that the anisotropic case uses a decoupled energy for measure of smoothness and is therefore easier to minimize, isotropic total variation is not selective in which direction smoothness is penalized. Image content smoothness (or lack thereof) is not known *a priori*, thus no preference should be immediately given to certain directions for evaluating smoothness.

Additionally, our research considers both L^2 *and* L^1 penalties for striping size, and compares the two, ultimately favoring the L^1 due to a wider yet tighter distribution of the striping mask. Using the L^1 penalty, and depending on the data, the isotropic total variation, which theoretically uses more local information, allows for a qualitatively better, less invasive and more intelligent destriping.

We construct a variational model that is well-defined, qualitatively motivated, and easily minimized. The constructed energy uses sparsity promoting energy functionals, based on total variation and L^1 energy, to achieve minimally invasive destriping. Both isotropic and anisotropic total variation, along with L^1 and L^2 energies, are considered in our variational model. The alternating direction method of multipliers (ADMM) (split-Bregman) is used in conjunction with non-linear proximal operators to efficiently optimize the energy, yielding quick and quality results.

CHAPTER 2

Variational Formulation

2.1 Striping structure

Let $U(x,y)$ be a stripe-free image of size R by C, and let $G(y)$ be a multiplicative stripe noise of length R. Then the observed image, F can be written as $F(x,y) = G(x,y)U(x,y)$. Taking logarithms yields an additive structure, more suitable for energy minimization methods. The model can now be written as $f(x,y) = g(y) + u(x,y)$ where $f(x,y) = \ln(F(x,y))$, $g(x,y) = \ln(G(x,y))$, and $u(x,y) = \ln(U(x,y))$.

Striping in images can be viewed as a structured noise, of which variations are mainly concentrated along one axis. This can be mathematically encoded as $\|\nabla_x G\| \ll \|\nabla_y G\|$, or with the logarithmic terms, $\|\nabla_x g\| \ll \|\nabla_y g\|$.

2.2 Tikhonov minimization

A classical Tikhonov minimization problem would consist of a smoothness regularizer and a data fidelity term, both easily differentiable, with the striping constraint [TA]:

$$\min_u \left\{ \|\nabla u(x,y)\|_2^2 + \frac{\lambda}{2}\|u(x,y) - f(x,y)\|_2^2 \right\} \quad \text{s.t.} \quad \|\nabla_x g\| \ll \|\nabla_y g\| \tag{2.1}$$

This constraint can be simplified by the approximation that $\nabla_x G(x,y) = 0 \quad \forall (x,y)$, which would make $G(x, y) = G(y)$, and $g(x, y) = g(y)$, functions of only one variable. Using the additive identity between f, g, and u along with the constraint approximation, the new unconstrained minimization problem is:

$$\min_{g} \Big\{ \|\nabla_y (f(x, y) - g(x, y))\|_2^2 + \frac{\lambda}{2} \|g(x, y)\|_2^2 \Big\}$$

By taking the first variation of the energy and setting it to zero, closed form solution to this minimization problem is

$$g(x, y) = (\nabla_y \cdot \nabla_y + \lambda I)^{-1} (\nabla_y \cdot \nabla_y f(x, y)) = (\frac{\partial^2}{\partial y^2} + \lambda I)^{-1} (f_{yy}(x, y)) \tag{2.2}$$

However, this solution would cause g to become bivariate, in contradiction to the constraint. Instead, using the Cartesian regularity of our rectangular domain $\Omega = I_x \times I_y$ and using $g(x, y) = g(y)$, we can come to a solution that is in agreement with the constraint by integrating with respect to x:

$$\int_{\Omega} g(x, y) dx = \int_{I_x} g(y) dx = g(y) \int_{I_x} dx = g(y) \mu(I_x) = \int_{I_x} (\frac{\partial^2}{\partial y^2} + \lambda I)^{-1} (f_{yy}(x, y)) dx \quad \Rightarrow \tag{2.3}$$

$$g(y) = \frac{1}{\mu(I_x)} \int_{I_x} (\frac{\partial^2}{\partial y^2} + \lambda I)^{-1} (f_{yy}(x, y)) dx$$

2.3 Fourier interpretation

Utilizing the Plancherel Fourier isometry, the solution can be interpreted in spectral form as:

$$\hat{g}(\omega_y) = \frac{1}{\mu(I_x)} \Big(\frac{\omega_y^2}{\lambda + \omega_y^2} \Big) \hat{f} = \Big(\frac{\omega_y^2}{\lambda + \omega_y^2} \Big) \bar{\hat{f}}$$

For a specific x, the stripping $g(y)$ is constant of higher frequency, whereas the underlying clean image varies more slowly (has more low-frequency content) and while for each x having a somewhat different frequencies. Therefore, the average frequencies of the clean image are low in magnitude and of lower-frequency, while the average frequencies of the stripping are high in magnitude and of higher-frequency. Therefore, the average frequencies (averaged over ω_x) of the cleaned image are simply the average frequencies of the original image multiplied by a one dimensional Low-Pass filter $\frac{\lambda}{\lambda+z^2}$. Likewise, the striping mask on the spectral side, $\hat{g}$, is obtained analogously with a one dimensional High-Pass filter $\frac{z^2}{\lambda+z^2}$.

Though this minimization problem is readily solvable in closed form and has a motivated physical interpretation, we must abandon the quadratic energy terms so that we may have less penalization for heavier striping and to allow for less smooth solutions. Though the differentiability of terms is nice, enough optimization machinery has been developed that we may tread forward. We now investigate and outline some tools from signal processing in order to refine our model. Stripe and ring artifact removal from this frequency perspective has been accomplished using wavelet and Fourier filtering [MS09].

CHAPTER 3

Tools from Signal Processing

3.1 Total Variation

The idea of using total variation as a regularizer and denoiser that promotes sparsity and piecewise constant smoothness dates back to Rudin, Osher, and Fatemi [RF92, RO94]. We begin with defining the notion of total variation, which will be used as a regularizer in the model.

Definition The *total variation* of a function $f \in L^1(\Omega)$ is

$$V(f,\Omega) := \sup\Big\{ \int_\Omega f(x)\mathrm{div}\phi(x)\,\mathrm{d}x \colon \phi \in C_c^1(\Omega,\mathbb{R}^n), \quad \|\phi\|_{L^\infty(\Omega)} \leq 1 \Big\}$$

For a differentiable function $f \in \Omega$, with $\Omega \subseteq \mathbb{R}^n$, the total variation of f can be written as

$$V(f,\Omega) = \int\limits_\Omega |\nabla f(x)|\,\mathrm{d}x$$

The choice of vectorial norm inside the integral yields two different types of total variation.

3.1.1 Isotropic

Definition *Isotropic* total variation: $|\cdot|$ denotes the l_2-norm, in which case

$$V(f,\Omega) = \int\limits_\Omega (\sum_i^n f_{x_i}^2(x))^{\frac{1}{2}}\mathrm{d}x$$

3.1.2 Anisotropic

Definition *Anisotropic* total variation: $|\cdot|$ denotes the l_1-norm, in which case

$$V(f,\Omega) = \int\limits_{\Omega} \sum_{i}^{n} |f_{x_i}(x)| \mathrm{d}x$$

The isotropic and anisotropic cases differ in terms of the geometries they each preserve. While the decoupled anisotropic total variation preserves piecewise constant orthogonal structures such as rectangular roofs, the coupled isotropic total variation preserves piecewise constant radial structures such as silos. Our model will be robust with respect to either choice of total variation and dual derivations of variable updates will be shown.

3.2 Shrinkage Proximal Operator

We will introduce a splitting variable and quadratic penalty into the model. The solution to the l_1-regularized least squares problem

$$\arg\min_{\vec{x}} \mu\|\vec{x}\|_1 + \frac{1}{2}\|\vec{x}-\vec{y}\|_2^2$$

is given by the soft threshold proximal mapping operator, shrinkage [Don95, WZ07]:

Definition

$$Shrink(\vec{x},\mu) = S_\mu(\vec{x}) = \frac{\vec{x}}{|\vec{x}|}\max\{|\vec{x}|-\mu, 0\}$$

If $\|\vec{x}\|_1 = \|x_1\| + \|x_2\|$, as in the anisotropic case of total variation, the shrinkage is decoupled and done component wise. On the other hand, if $\|\vec{x}\|_1 = \sqrt{\|x_1\|^2 + \|x_2\|^2}$, as in the isotropic case, the terms are coupled and both components are updated simultaneously. Both variants have their merits, while the former is computationally simpler, the latter has the advantage of using more local information and may be more conformant to certain image processing application.

CHAPTER 4

TV-L^1 ADMM Optimization

We will now make some technical modifications to our model, while preserving the qualitative ideas and motivations. The two energy components of the minimization problem are the smoothness regularizer and the data fidelity term. The energy of the data fidelity term, $\frac{\lambda}{2}\|f(x,y) - u(x,y)\|_2^2 = \frac{\lambda}{2}\|g(y)\|_2^2$, can be interpreted as the size of the striping mask. The L_2 fidelity overly penalizes stripes of large magnitude, and likewise under-exaggerates the significance of stripes of small magnitude. In areas of no striping, we intend our (logarithm of the) striping mask to be very close to zero, while in areas of heavy striping, we wish to remove said striping and thus will yield a larger magnitude of our striping mask in that region. Using the L^1 fidelity gives us a smaller striping mask in areas of no striping, leaving enough energy to remove the heavier striping in localized areas of the image. Because there is no prior knowledge of the distribution of the stripes, and qualitatively we may wish to remove deep striping effects while preserving sharp geometry, we believe it is better to update the model with an L_1 striping penalty, $\|g\|_1$.

An L^2 gradient term would cause over-smoothing of the retrieved clean image $u(x,y)$. This could cause a loss in boundary sharpness of elements in the image (e.g. lakes, rooftops, etc.), which seems important in the pursuit and usage of destriped images. Implementing a total variation based regularizer would act similarly to the L^2 gradient but maintains boundary sharpness more natural to the underlying image. Though these terms are not differentiable, impeding a closed form solution, state-of-the-art nonlinear optimization algorithms are available for fast convergence to qualitatively meaningful minimizers. The unconstrained total variation L^1 model (TV-L^1) is:

$$\min_u \left\{ V(u(x,y), I_x \times I_y) + \lambda |u(x,y) - f(x,y)\|_1 \right\}$$

or equivalently, minimizing with respect to the striping mask g:

$$\min_g \left\{ V(f(x,y) - g(y), I_x \times I_y) + \lambda \|g(y)\|_1 \right\}$$

4.1 Discretization

For the purpose of application and computation, we shall now move the problem into a discrete setting. Let $\Omega = \{x_1, ..., x_C\} \times \{y_1, ..., y_R\}$ be an $R \times C$ matrix. First variations will be approximated via forward differences, so that $\frac{\partial f}{\partial y}(x_i, y_j) \approx f(x_i, y_{j+1}) - f(x_i, y_j) := \delta_y f_{i,j}$ for $j = 1, ..., R-1$ and analogously for $\frac{\partial f}{\partial x} \approx f(x_{i+1}, y_j) - f(x_i, y_j) := \delta_x f_{i,j}$ for $i = 1, ..., C-1$. We will take Neumann boundary conditions, so that on the forward boundary ($i = C$ or $j = R$), the derivative is set to zero.

$$D = \begin{bmatrix} -1 & 1 & 0 & \cdots & 0 & 0 \\ 0 & -1 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & -1 & 1 \\ 0 & 0 & 0 & \cdots & 0 & 0 \end{bmatrix} \in M_{R\times R}(\mathbb{R})$$

so that $Du(x_{i_0}, y) = (\delta_y u_{i_0,1}, \cdots, \delta_y u_{i_0,R-1}, 0)^T$

Isotropic total variation:

$$|\nabla^I f(x_i, y_j)| \approx \sqrt{(\delta_x f_{i,j})^2 + (\delta_y f_{i,j})^2}$$

Anisotropic total variation:

$$|\nabla^{AI} f(x_i, y_j)| \approx |\delta_x f_{i,j}| + |\delta_y f_{i,j}|$$

With these discrete operators defined, the discrete unconstrained TV-L^1 minimization problem is:

$$\min_g \Big\{ \sum_{i,j} \|\langle \delta_x f_{i,j}, \delta_y (f_{i,j} - g_j) \rangle\|_1 + \lambda \|g(y)\|_1 \Big\}$$

The two flavors of the minimization problem are:

Anisotropic

$$\min_g \Big\{ \sum_{i,j} |\delta_y (f_{i,j} - g_j)| + \lambda \|g(y)\|_1 \Big\}$$

Isotropic

$$\min_g \Big\{ \sum_{i,j} \sqrt{(\delta_x f_{i,j})^2 + (\delta_y (f_{i,j} - g_j))^2} + \lambda \|g(y)\|_1 \Big\}$$

4.2 Augmented Lagrangian

4.2.1 Anisotropic

With the discrete forward difference approximation matrix defined above, we can rewrite the minimization problem as:

$$\begin{aligned}
\textit{Point Form:}\quad & \min_{\vec{g}}\Big\{\sum_{i,j}|\delta_y(f_{i,j}-g_j)|+\lambda\|g(y)\|_1\Big\}\\
\textit{Vector Form:}\quad & \min_{\vec{g}}\Big\{\sum_{i}\|D(f_i-\vec{g})\|_1+\lambda\|g(y)\|_1\Big\}\\
\textit{Matrix Form:}\quad & \min_{\vec{g}}\Big\{\|D(f-\vec{g}\otimes\vec{1_C})\|_{1,1}+\lambda\|g(y)\|_1\Big\}
\end{aligned}$$

To render the constrained minimization problem unconstrained, we introduce auxiliary variables, Lagrangian multipliers (split Bregman), and quadratic penalty terms, so that the augmented Lagrangian is defined as:

$$\begin{aligned}
&\mathcal{L}_{\alpha,\lambda}(b_i,h,g,q_i,r)= \qquad (4.1)\\
&\sum_i\Big(\|b_i\|_1+\frac{\alpha}{2}\|b_i-D(g-f_i)\|_2^2+\langle q_i,b_i-D(g-f_i)\rangle\Big)+\lambda\Big(\|h\|_1+\frac{\alpha}{2}\|h-g\|_2^2+\langle r,h-g\rangle\Big)=\\
&\sum_i\Big(\|b_i\|_1+\frac{\alpha}{2}\|b_i-D(g-f_i)+\frac{q_i}{\alpha}\|_2^2\Big)+\lambda\Big(\|h\|_1+\frac{\alpha}{2}\|h-g+\frac{r}{\alpha}\|_2^2\Big)+\mathcal{O}(q_i^2,r^2)
\end{aligned}$$

We now solve the unconstrained saddle point problem.

$$\min_{b_i,h,g}\max_{q_i,r}\mathcal{L}_{\alpha,\lambda}(b_i,h,g,q_i,r)$$

The solution to the original constrained minimization problem is now found as the saddle point of the augmented Lagrangian $\mathcal{L}$ in a sequence of iterative sub-optimizations called alternating direction method of multipliers (ADMM) [Roc73b, GM75, GM76, Hes69, WZ08, GO08].

The splitting variables b_i, and h are updated by the proximal mapping operator:

$$\begin{aligned} b_i^{k+1} &= \arg\min_{b_i} \mathcal{L}_{\alpha,\lambda}(b_i, h^k, g^k, q_i^k, r^k) \\ &= \arg\min_{b_i} \left\{ \|b_i\|_1 + \frac{\alpha}{2}\|b_i - D(g - f_i) + \frac{q_i}{\alpha}\|_2^2 \right\} \\ &= S_{\frac{1}{\alpha}}\left(D(g - f_i) - \frac{q_i}{\alpha} \right) \end{aligned} \tag{4.2}$$

$$\begin{aligned} h^{k+1} &= \arg\min_{h} \mathcal{L}_{\alpha,\lambda}(b_i^{k+1}, h, g^k, q_i^k, r^k) \\ &= \arg\min_{h_i} \left\{ \|h\|_1 + \frac{\alpha}{2}\|h - g\|_2^2 \right\} \\ &= S_{\frac{1}{\alpha}}\left(g - \frac{r}{\alpha} \right) \end{aligned} \tag{4.3}$$

Due to the introduction of the splitting variables, g is only contained in quadratic terms, and thus easily solved for:

$$\begin{aligned} g^{k+1} &= \arg\min_{g} \mathcal{L}_{\alpha,\lambda}(b_i^{k+1}, h^{k+1}, g, q_i^k, r^k) \\ &= \arg\min_{g} \left\{ \frac{\alpha}{2}\sum_i \|b_i - D(g - f_i) + \frac{q_i}{\alpha}\|_2^2 + \frac{\lambda\alpha}{2}\|h - g + \frac{r}{\alpha}\|_2^2 \right\} \quad \Rightarrow \end{aligned} \tag{4.4}$$

$$\begin{aligned} \frac{\delta\mathcal{L}}{\delta g} &= \alpha\sum_i -D^T(b_i - D(g - f_i) + \frac{q_i}{\alpha}) - \lambda\alpha(h - g + \frac{r}{\alpha}) \\ &= \alpha(CD^TD + \lambda I)g - \alpha D^T(\sum_i b_i + Df_i + \frac{q_i}{\alpha}) - \lambda\alpha h + \lambda r = 0 \quad \Rightarrow \end{aligned} \tag{4.5}$$

$$g = (CD^TD + \lambda I)^{-1}\left(D^T(\sum_i b_i + Df_i + \frac{q_i}{\alpha}) + \lambda(h - \frac{r}{\alpha}) \right) \tag{4.6}$$

The Langrangian multipliers (split Bregman variables) are updated through gradient ascent.

$$
\begin{aligned}
q_i^{k+1} &= q_i^k + \tau\alpha(b_i^{k+1} - D(g^{k+1} - f_i)) \\
r^{k+1} &= r^k + \tau\lambda\alpha(h^{k+1} - g^{k+1}))
\end{aligned}
$$

4.2.2 Isotropic

Due to the coupling of the terms in this version of the minimization problem, we cannot compactly write the problem with matrices as above; however, the solution is just as readily available. Here the $\diamond$ denotes the Hadamard matrix power operator, which acts pointwise on the matrix.

Point Form:

$$\min_{\vec{g}} \Big\{ \sum_{i,j} \sqrt{(\delta_x f_{i,j})^2 + (\delta_y (f_{i,j} - g_j))^2} + \lambda \|g(y)\|_1 \Big\}$$

Matrix Form:

$$\min_{\vec{g}} \Big\{ \| [(fD^T)^{\diamond 2} + (D(f - \vec{g} \otimes \vec{1_C}))^{\diamond 2}]^{\diamond \frac{1}{2}} \|_{1,1} + \lambda \|g(y)\|_1 \Big\}$$

Just as before, we introduce splitting variables and Lagrangian multipliers to form the augmented Lagrangian:

$$\mathcal{L}_{\alpha,\lambda}(a_{i,j}, b_{i,j}, h, g, p_{i,j}, q_{i,j}, r) = \tag{4.7}$$

$$\sum_{i,j} \sqrt{|a_{i,j}|^2 + |b_{i,j}|^2} + \frac{\alpha}{2}\|a_{i,j} - \delta_x f_{i,j}\|_2^2 + \langle p_{i,j}, a_{i,j} - \delta_x f_{i,j}\rangle + \frac{\alpha}{2}\|b_{i,j} - \delta_y(f_{i,j} - g_j)\|_2^2$$
$$+ \langle q_{i,j}, b_{i,j} - \delta_y(f_{i,j} - g_j)\rangle + \lambda(\|h\|_1 + \frac{\alpha}{2}\|h - g\|_2^2 + \langle r, h - g\rangle) =$$

$$\sum_{i,j} \sqrt{|a_{i,j}|^2 + |b_{i,j}|^2} + \frac{\alpha}{2}\|a_{i,j} - \delta_x f_{i,j} + \frac{p_{i,j}}{\alpha}\|_2^2 + \frac{\alpha}{2}\|b_{i,j} - \delta_y(f_{i,j} - g_j) + \frac{q_{i,j}}{\alpha}\|_2^2$$
$$+ \lambda\big(\|h\|_1 + \frac{\alpha}{2}\|h - g + \frac{r}{\alpha}\|_2^2\big) + \mathcal{O}(q_{i,j}^2, p_{i,j}^2, r^2)$$

The splitting variables $a_{i,j}$, $b_{i,j}$ are updated by the vectorial proximal mapping operator:

$$\begin{aligned}
\langle a_{i,j}, b_{i,j}\rangle &= \arg\min_{\langle a_{i,j}, b_{i,j}\rangle} \mathcal{L}_{\alpha,\lambda}(a_{i,j}, b_{i,j}, h^k, g^k, p_i^k, q^k) \qquad (4.8)\\
&= \arg\min_{\langle a_{i,j}, b_{i,j}\rangle} \left\{ \sqrt{|a_{i,j}|^2 + |b_{i,j}|^2} + \frac{\alpha}{2}\|a_{i,j} - \delta_x f_{i,j} + \frac{p_{i,j}}{\alpha}\|_2^2 + \frac{\alpha}{2}\|b_{i,j} - \delta_y(f_{i,j} - g_j) + \frac{q_{i,j}}{\alpha}\|_2^2 \right\}\\
&= \arg\min_{\langle a_{i,j}, b_{i,j}\rangle} \left\{ \|\langle a_{i,j}, b_{i,j}\rangle\| + \frac{\alpha}{2}\|\langle a_{i,j}, b_{i,j}\rangle - \langle \delta_x f_{i,j} + \frac{p_{i,j}}{\alpha}, \delta_y(f_{i,j} - g_j) + \frac{q_{i,j}}{\alpha}\rangle\|^2 \right\}\\
&= S_{\frac{1}{\alpha}}(\langle \delta_x f_{i,j} + \frac{p_{i,j}}{\alpha}, \delta_y(f_{i,j} - g_j) + \frac{q_{i,j}}{\alpha}\rangle)
\end{aligned}$$

Each component of the vector is updated via shrinkage as follows:

$$a_{i,j} = \frac{\delta_x f_{i,j} + \frac{p_{i,j}}{\alpha}}{s} \cdot \max(s - \frac{1}{\alpha}, 0), \quad b_{i,j} = \frac{\delta_y(f_{i,j} - g_j) + \frac{q_{i,j}}{\alpha}}{s} \cdot \max(s - \frac{1}{\alpha}, 0)$$

$$s = \sqrt{(\delta_x f_{i,j} + \frac{p_{i,j}}{\alpha})^2 + (\delta_y(f_{i,j} - g_j) + \frac{q_{i,j}}{\alpha})^2}$$

The splitting variable h, the striping mask g, and the Lagrangian multipliers are updated as before due to the common structure between the two models.

$$p_{i,j}^{k+1} = p_{i,j}^k + \tau\alpha(\delta_x f_{i,j} + \frac{p_{i,j}}{\alpha}))$$

4.3 TV-L^2 Comparison

We introduce a slight variant to the TV-L^1 model where the norm on the striping mask is replaced by an L_2 norm, this is called the TV-L^2 model.

$$\min_g \Big\{ \sum_{i,j} \|\langle \delta_x f_{i,j}, \delta_y(f_{i,j} - g_j)\rangle\|_1 + \frac{\lambda}{2}\|g(y)\|_2^2 \Big\}$$

In both the isotropic and anisotropic cases, the energy terms with g in the augmented Lagrangians are the same:

$$\sum_{i,j} \frac{\alpha}{2}|b_{i,j} - \delta_y(f_{i,j} - g_j) + \frac{q_{i,j}}{\alpha}|_2^2 + \frac{\lambda}{2}\|g\|_2^2 = \sum_i \frac{\alpha}{2}\|b_i - D(g - f_i) + \frac{q_i}{\alpha}\|_2^2 + \frac{\lambda}{2}\|g\|_2^2 \tag{4.9}$$

The optimization of g is done analogously.

$$\begin{aligned} g^{k+1} &= \arg\min_g \mathcal{L}_{\alpha,\lambda}(b_i^{k+1}, h^{k+1}, g, p_i^k, q_i^k, r^k) \qquad (4.10) \\ &= \arg\min_g \left\{ \frac{\alpha}{2} \sum_i \|b_i - D(g - f_i) + \frac{q_i}{\alpha}\|_2^2 + \frac{\lambda}{2}\|g\|_2^2 \right\} \Rightarrow \end{aligned}$$

$$\begin{aligned} \frac{\delta \mathcal{L}}{\delta g} &= \alpha \sum_i -D^T(b_i - D(g - f_i) + \frac{q_i}{\alpha}) + \lambda g \qquad (4.11) \\ &= (\alpha C D^T D + \frac{\lambda}{\alpha} I)g - \alpha D^T(\sum_i b_i + D f_i + \frac{q_i}{\alpha}) = 0 \Rightarrow \end{aligned}$$

$$g = (C D^T D + \frac{\lambda}{\alpha} I)^{-1} \left(D^T(\sum_i b_i + D f_i + \frac{q_i}{\alpha}) \right) \tag{4.12}$$

The quadratic penalty on the size of the striping mask is included for comparison with the L_1 penalty term. Given the same parameters α, and λ, the TV-L^1 should be able to remove deeper stripes (of higher magnitude) while preserving small fluctuations (of lesser magnitude) by not classifying them as stripes. We compare both the TV-L^1 and TV-L^2 models in our experiments.

The first algorithm (anisotropic) is presented in vector fashion. The second algorithm (isotropic) is presented in matrix fashion.

4.4 Complete TV-L^1 ADMM algorithm

Algorithm 5 ADMM optimization of TV-L^1

```
Initialize: $A^0, B^0, Q^0 \leftarrow 0 \in \mathbb{R}^{R\times C}$ , $\vec{g^0}$, $\vec{h^0}$, $\vec{r^0} \leftarrow 0 \in \mathbb{R}^{R\times 1}$, $n \leftarrow 0 \in \mathbb{R}$
$f = (\vec{f_i}) \leftarrow \ln(F)$, $D \leftarrow 0 \in \mathbb{R}^{R\times R}$, $(D)_{i,i} = -1$, $(D)_{i,i+1} = 1$ for $i = 1, ..., R-1$
repeat
    $n \leftarrow n + 1$
    case Anisotropic:
        for $i = 1 : C$ do
Update splitting variable for smoothness regularizer term via shrinkage:
$$\vec{b_i}^{n+1} \leftarrow S_{\frac{1}{\alpha}}\left(D(\vec{g}^n - \vec{f_i}) - \frac{\vec{q_i}^n}{\alpha}\right)$$
Update Lagrangian multiplier for regularizer term via dual ascent:
$$\vec{q_i}^{n+1} \leftarrow \vec{q_i}^n + \tau\alpha(\vec{b_i}^{n+1} - D(\vec{g}^n - \vec{f_i}))$$
    case Isotropic:
        for $i = 1 : C$, $j = 1 : R$ do
Update splitting variables for smoothness regularizer term via vectorial shrinkage:
$$\langle a_{i,j}^{n+1}, b_{i,j}^{n+1}\rangle \leftarrow S_{\frac{1}{\alpha}}(\langle \delta_x f_{i,j} + \frac{p_{i,j}^n}{\alpha}, \delta_y(f_{i,j} - g_j^n) + \frac{q_{i,j}^n}{\alpha}\rangle)$$
Update Lagrangian multipliers for regularizer term via dual ascent:
$$p_{i,j}^{n+1} \leftarrow p_{i,j}^n + \tau\alpha(\delta_x f_{i,j} + \frac{p_{i,j}}{\alpha})), \quad q_{i,j}^{n+1} \leftarrow q_{i,j}^n + \tau\alpha(b_{i,j}^{n+1} - \delta_y(g_j^n - f_{i,j}))$$
```

Algorithm 5 ADMM optimization of TV-L^1 (continued)

```
    case TV-$L^1$:
Update splitting variable for data fidelity term via shrinkage:
    $\vec{h}^{n+1} \leftarrow S_{\frac{1}{\alpha}}\left(\vec{g}^n - \frac{\vec{r}^n}{\alpha}\right)$
Update striping mask:
    $\vec{g}^{n+1} \leftarrow (CD^TD + \lambda I)^{-1}\left(D^T(\sum_i \vec{b}_i^{n+1} + D\vec{f}_i + \frac{\vec{q}_i^{n+1}}{\alpha}) + \lambda\vec{h}^{n+1} - \frac{\lambda}{\alpha}\vec{r}^n\right)$
Update Lagrangian multiplier for data fidelity term via dual ascent:
    $\vec{r}^{n+1} \leftarrow \vec{r}^n + \tau\lambda\alpha(\vec{h}^{n+1} - \vec{g}^{n+1}))$
    case TV-$L^2$:
        Update striping mask:
            $\vec{g}^{n+1} \leftarrow (CD^TD + 2\frac{\lambda}{\alpha}I)^{-1}\left(D^T(\sum_i \vec{b}_i^{n+1} + D\vec{f}_i + \frac{\vec{q}_i^{n+1}}{\alpha})\right)$
        Update energy terms:
            $E_1^{n+1} \leftarrow \sum_i \|D(\vec{f}_i - \vec{g}^{n+1})\|_1, \quad E_2^{n+1} \leftarrow \lambda\|\vec{g}^{n+1}\|_1$
            $E^{n+1} \leftarrow E_1^{n+1} + E_2^{n+1}$
until convergence:
$\|\vec{g}^{n+1} - \vec{g}^n\|_2^2/\|\vec{g}^n\|_2^2 < \epsilon_g$ and $|E^{n+1} - E^n|^2/|E^n|^2 < \epsilon_E$.
Retrieve clean image:
    $u \leftarrow f - g^{n+1} \otimes [1_1, 1_2, ..., 1_C]$
    $U \leftarrow \exp(u)$
```

CHAPTER 5

Experiments and Results

In our experiments, we used data acquired by the Airborne Multi-angle Spectro Polarimetric Imager (AirMSPI). AirMSPI is an airborne prototype instrument similar to that of the future satellite-borne MSPI instrument for obtaining multi-angle polarization imagery [DM13]. The instrument was built for NASA by the Jet Propulsion Laboratory in Pasadena, California and has been flying aboard the NASA ER-2 high altitude aircraft since October 2010.

AirMSPI is an eight-band (355, 380, 445, 470, 555, 660, 865, 935 nm) pushbroom camera, measuring polarization in the 470, 660, and 865 nm bands, mounted on a gimbal to acquire multiangular observations over a ±67° along-track range. Two principal observing modes are employed: step-and-stare, in which 11 km x 11 km targets are observed at a discrete set of view angles with a spatial resolution of ~ 10 m; and continuous sweep, in which the camera slews back and forth along the flight track between ±67° to acquire wide area coverage (11 km swath at nadir, target length 108 km) with ~ 25 m spatial resolution. Step-and-stare provides more angles, but continuous sweep gives greater coverage. Multiple observing modes can be programmed into the instrument and activated under cockpit control. Multi-angle radiance and polarization imagery from AirMSPI will provide 3-D scene context where clouds and aerosol plumes are present. It will also enable retrieval of aerosol and cloud macrophysical properties (distribution, height), microphysical properties (size distribution, single scattering albedo, shape), and optical depth.

5.1 TV-L^1 vs. TV-L^2

We first compare destriping results generated using TV-L^1 and TV-L^2 models. Fig. 5.1 shows the 355 nm UV channel image with stripes captured by the AirMSPI instrument from Nadir angle at Mojave, California. The image is destriped using TV-L^2 and TV-L^1 models. As we see from the destriped images and corresponding differences between captured images and destriped images, TV-L^2 model does not preserve radiometric intensities in regions where no stripes are present. Fig. 5.2(a,left) shows plots of recovered function g for TV-L^1 and TV-L^2 destriped images. TV-L^1 recovered function g is closer to identity especially at the rows where with no stripes, suggesting TV-L^1 reconstruction is more accurate than TV-L^2 reconstruction. Fig. 5.2(a,right) shows plots of sums over all rows of the original image with stripes (from Fig. 5.1), as well as sums of rows for TV-L^1 and TV-L^2 destriped images. These plots indicate that TV-L^1 model preserves radiometric intensities of the original images better, which TV-L^2 model produces more artificial smoothing throughout the image. Figures 5.2 (b) and (c) show histograms of function g for TV-L^2 and TV-L^1 reconstructions, respectively. TV-L^1 reconstruction is pointier than TV-L^2 reconstruction. It is also centered at 1, as opposed to TV-L^2 reconstruction, which further indicates better accuracy of TV-L^1 model. Note that the actual stripes, at around $g \approx 0.95$ are represented in the histograms by small bumps.

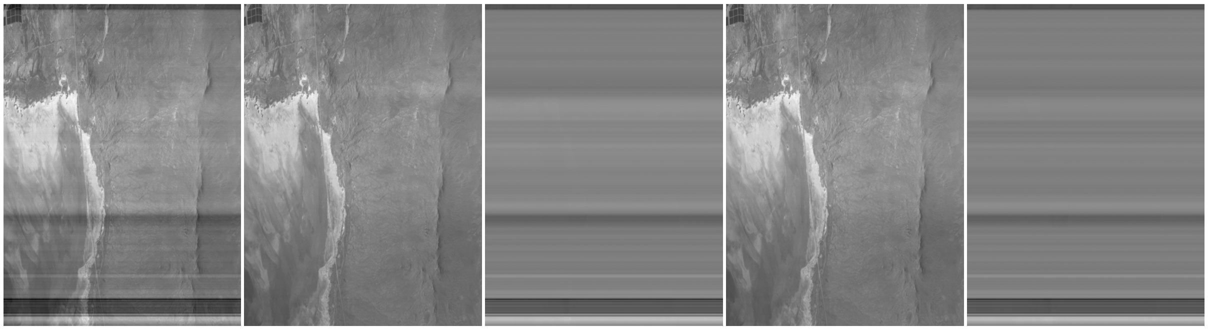

(a) (b) (c) (d) (e)

Figure 5.1: (a) The 355 nm channel image with stripes captured by AirMSPI instrument from Nadir angle at Mojave, California. (b) Destriped image using TV-L^2 model. (c) Difference between captured image from (a) and TV-L^2 destriped image from (b). (d) Destriped image using TV-L^1 model. (e) Difference between captured image from (a) and TV-L^1 destriped image from (d).

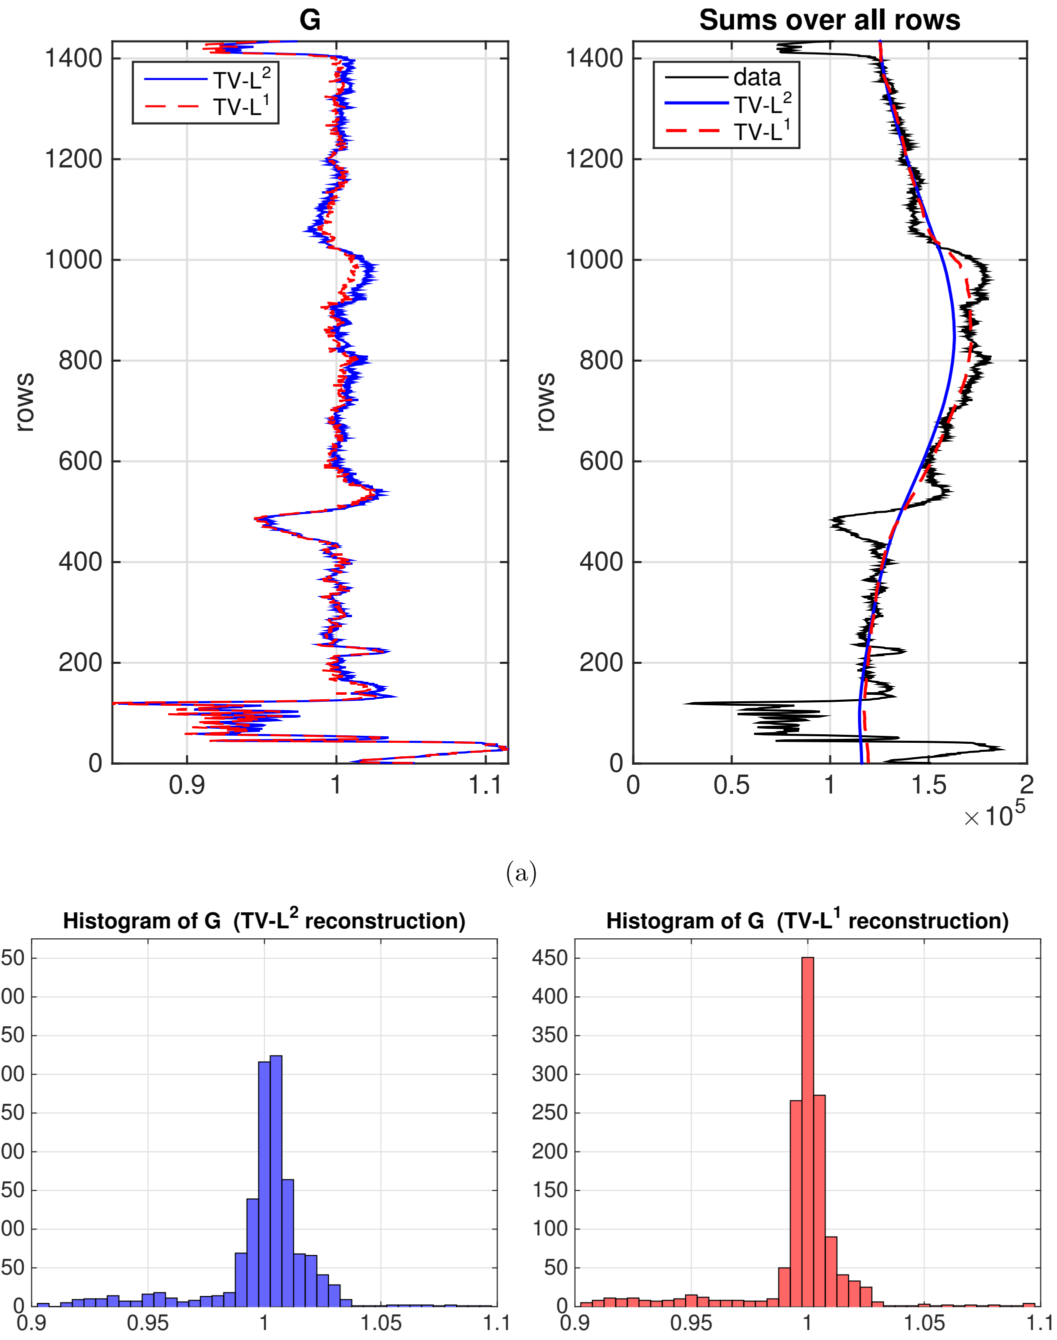

(a)

(b)

(c)

Figure 5.2: Comparisons of TV-L^2 and TV-L^1 destriping for the results shown in Fig. 5.1. (a) Left: plots of recovered function G for TV-L^2 destriped image (blue) and TV-L^1 destriped image (red) are shown. Right: plots of sums over all rows of original image with stripes (black), TV-L^2 destriped image (blue), and TV-L^1 destriped image (red) are shown. (b) Histogram of function G for TV-L^2 reconstruction. (c) Histogram of function G for TV-L^1 reconstruction.

Fig. 5.3 and Fig. 5.4 display similar results as in Fig. 5.1 and 5.2 for the 355 nm channel image with stripes depicting clouds over the Pacific Ocean captured by AirMSPI instrument from 66.0°F angle.

Figures 5.5, 5.6, and 5.7 show more examples of TV-L^1 reconstruction of images captured using continuous sweep observing mode. Figures 5.8 and 5.9 display images captured using the step-and-stare observing mode as well as destriped results using TV-L^1 model.

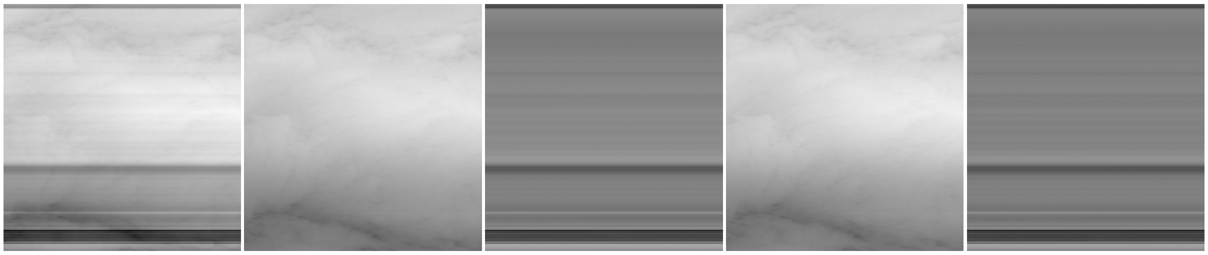
(a) (b) (c) (d) (e)

Figure 5.3: (a) The 355 nm channel image with stripes depicting clouds over the Pacific Ocean captured by AirMSPI instrument from 66.0°F angle. (b) Destriped image using TV-L^2 model. (c) Difference between captured image from (a) and TV-L^2 destriped image from (b). (d) Destriped image using TV-L^1 model. (e) Difference between captured image from (a) and TV-L^1 destriped image from (d).

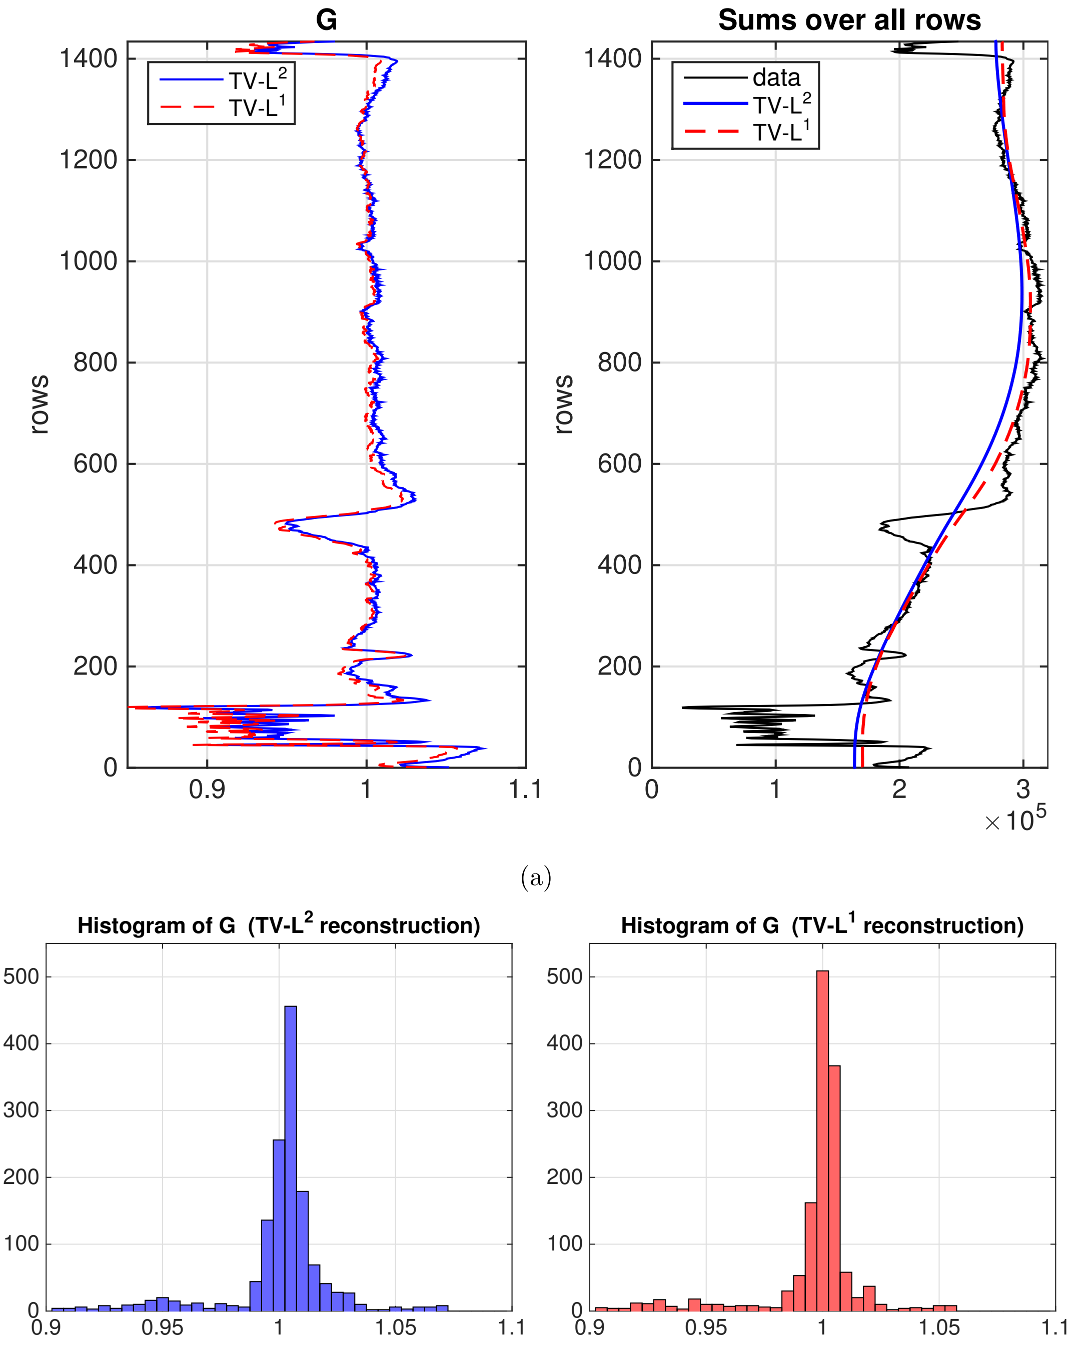

(a)

(b)

(c)

Figure 5.4: Comparisons of TV-L^2 and TV-L^1 destriping for the results shown in Fig. 5.3. (a) Left: plots of recovered function G for TV-L^2 destriped image (blue) and TV-L^1 destriped image (red) are shown. Right: plots of sums over all rows of original image with stripes (black), TV-L^2 destriped image (blue), and TV-L^1 destriped image (red) are shown. (b) Histogram of function G for TV-L^2 reconstruction. (c) Histogram of function G for TV-L^1 reconstruction.

5.2 Mojave

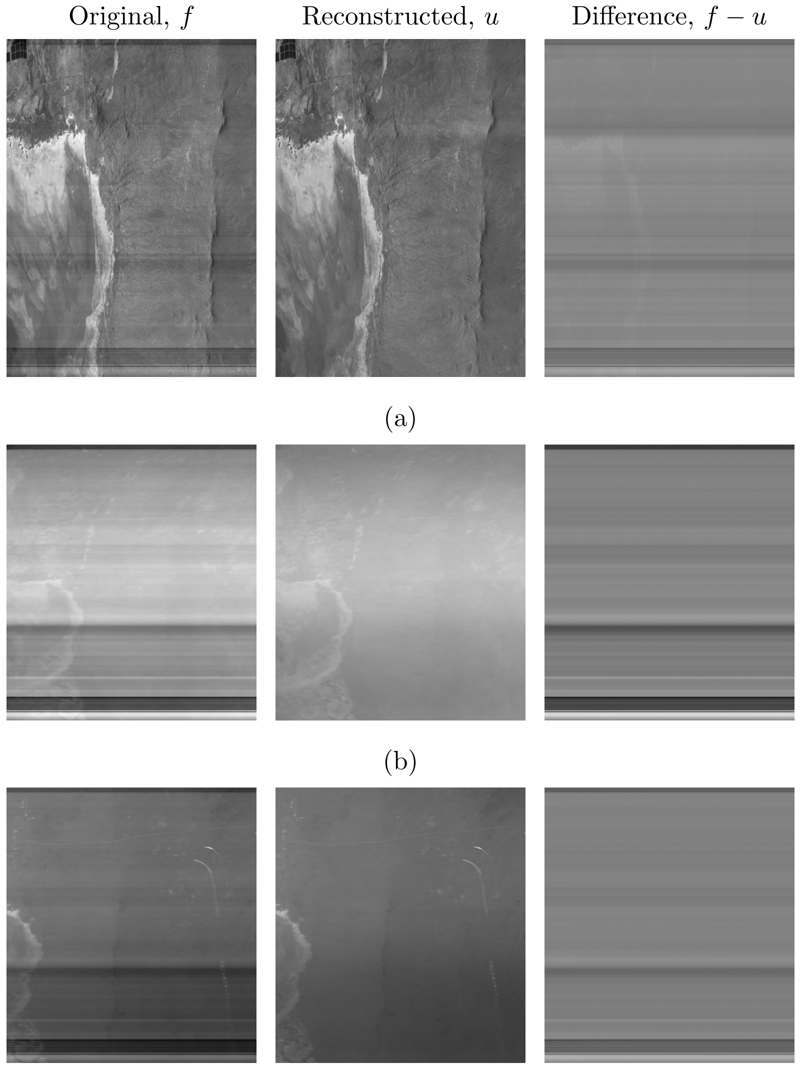

(a)

(b)

(c)

Figure 5.5: Images with stripes captured by AirMSPI instrument at Mojave, California (left), destriped images using TV-L^1 model (center), and differences between captured and destriped images (right) are shown. The bands and viewing angles are: (a) 380 nm band at Nadir angle; (b) 355 nm band at 66.1°F angle; (c) 355 nm band at 66.1°A angle. 355 nm band at Nadir angle is shown in Fig. 5.1.

5.3 Pacific Ocean

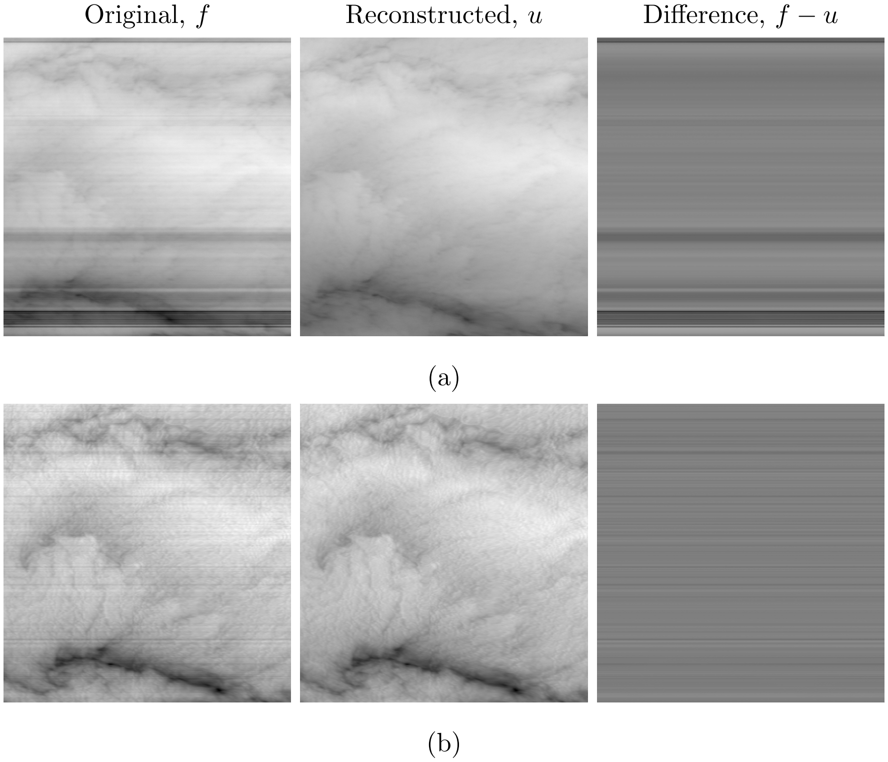

Figure 5.6: Images with stripes depicting clouds over the Pacific Ocean captured by AirMSPI instrument (left), destriped images using TV-L^1 model (center), and differences between captured and destriped images (right) are shown. The bands are: (a) 380 nm, and (b) 660 nm all at 66.0°F. 355 nm band at 66.0°F angle is shown in Fig. 5.3.

5.4 Ivanpah

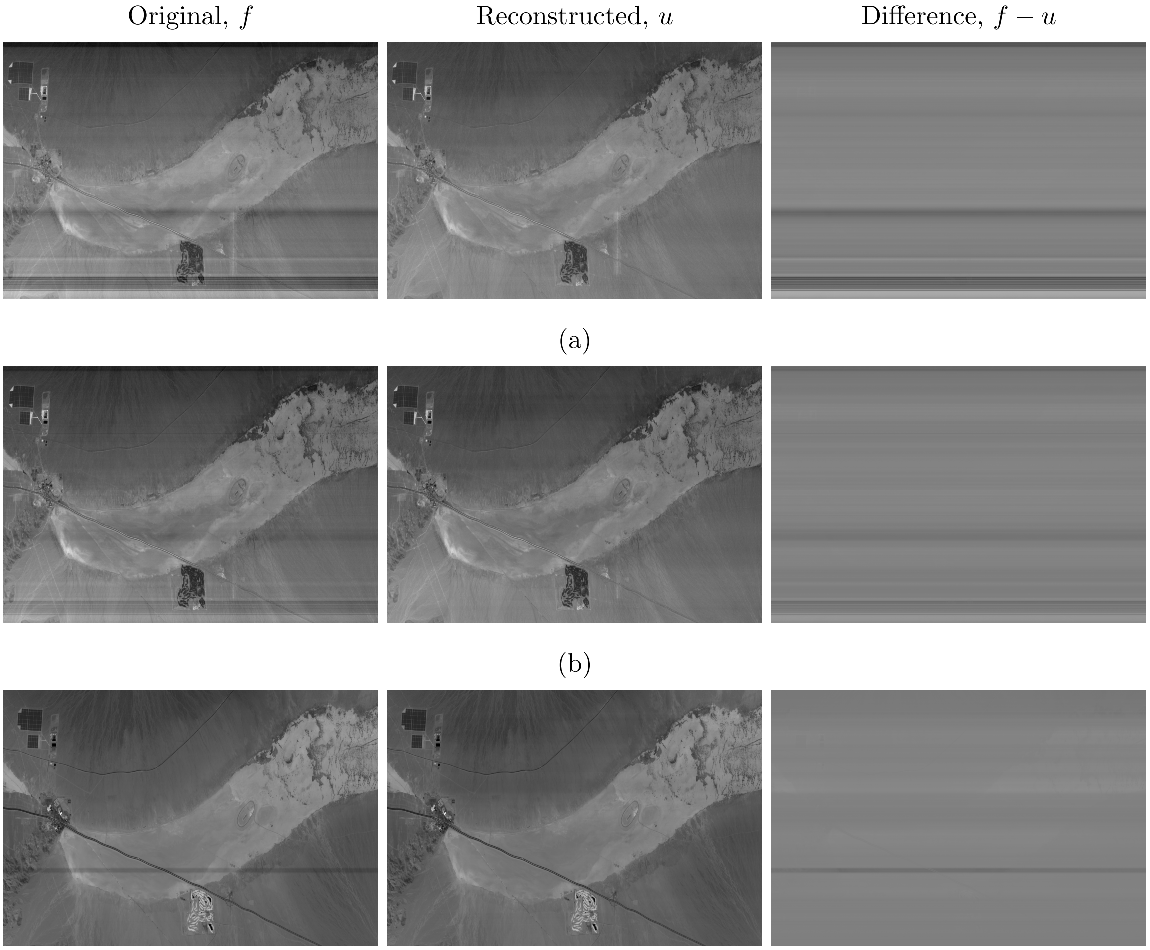

Original, f Reconstructed, u Difference, $f - u$

(a)

(b)

(c)

Figure 5.7: Images with stripes of dry lake Ivanpah, California captured by AirMSPI instrument (left), destriped images using TV-L^1 model (center), and differences between captured and destriped images (right) are shown. The bands are: (a) 355 nm, (b) 380 nm, and (c) 865 nm bands, all at Nadir angle.

5.5 Avalon

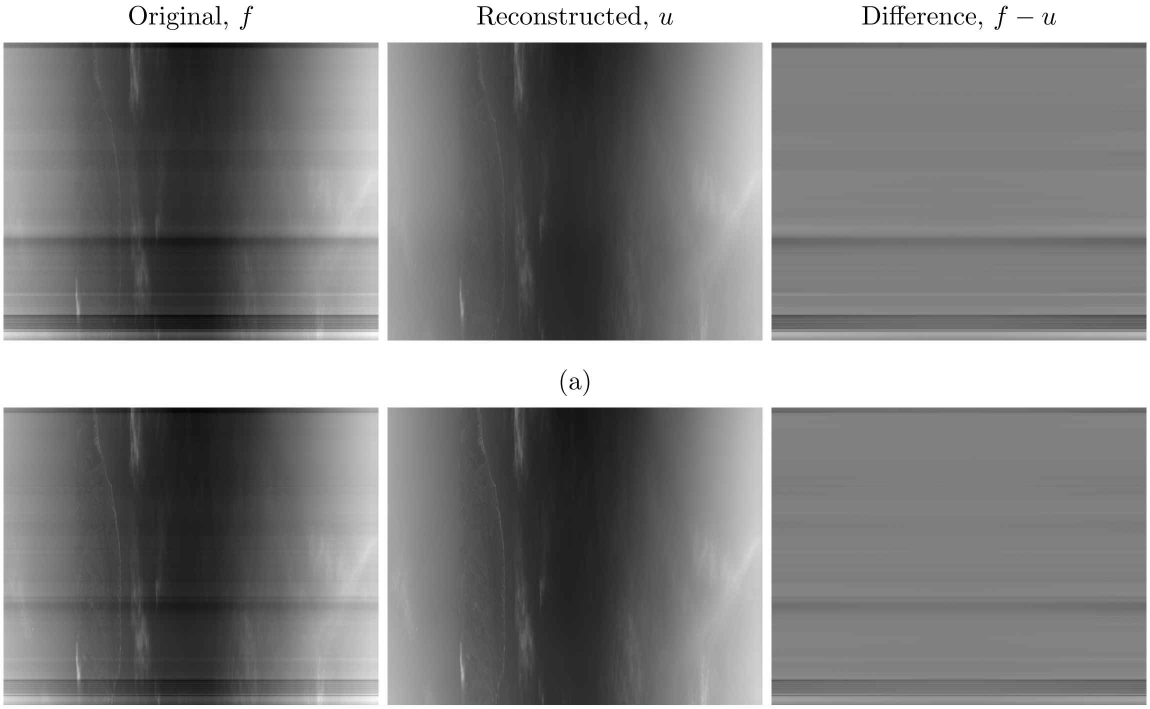

(a)

(b)

Figure 5.8: Images captured by AirMSPI instrument at Avalon, California (left), destriped images using TV-L^1 model (center), and differences between captured and destriped images (right) are shown. The bands are: (a) 355 nm and (b) 380 nm, both captured using the continuous sweep observing mode.

5.6 Fallbrook

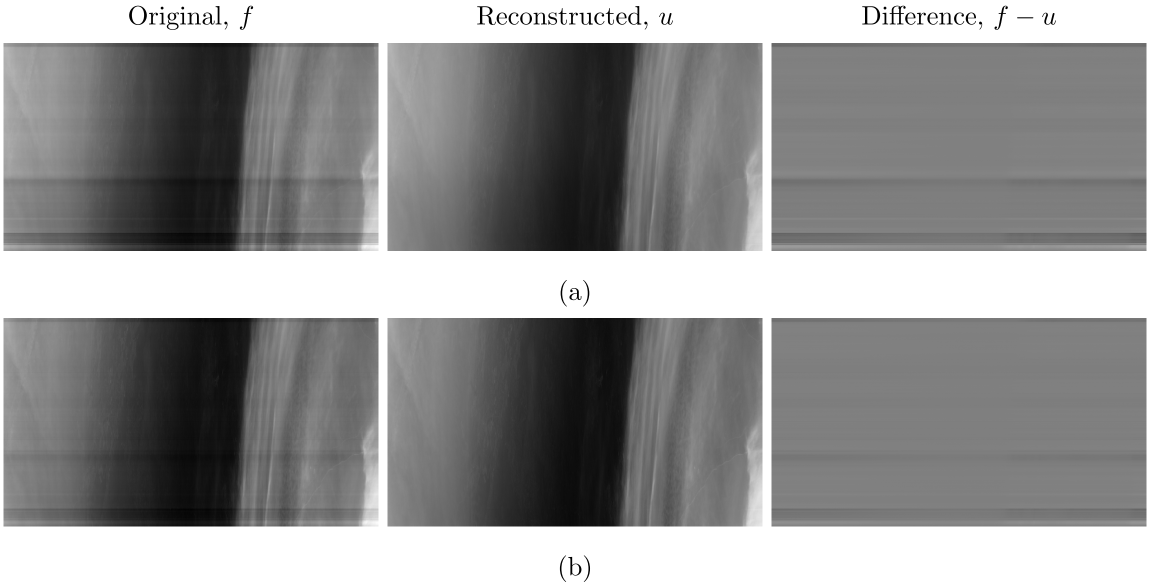

(a)

(b)

Figure 5.9: Images with stripes captured by AirMSPI instrument at Fallbrook, California (left), destriped images using TV-L^1 model (center), and differences between captured and destriped (right) are shown. The bands are: (a) 355 nm and (b) 380 nm, both captured using the continuous sweep observing mode.

CHAPTER 6

Conclusions and Outlook

In Part III, we have presented a novel variational method for image destriping through fast minimization techniques of appropriately modelled energy functionals - namely total variation and L^1 data fidelity term. In contrast to existing destriping models, such as statistical estimation models, we simplify the calculations while achieving excellent qualitative results quickly and with few explicit parameters.

Our destriping model solves the inverse problem as follows: minimally remove a univariate multiplicative striping mask from the data, such that the clean image is somewhat smooth and the removed stripe has low energy. We assess the smoothness of the clean image using total variation, which maintains sharp image features and preserves definition and contrast. We address both isotropic and anisotropic total variation, each having their respective strengths and weaknesses. We use L^1, and for comparison, L^2 energy to measure the removed striping, ensuring minimal data removal and thus a clean image of high fidelity.

The variational problem is solved very efficiently in an ADMM approach: introduce splitting variables and quadratic penalties for deviations from said splitting variables to allow efficient optimization via proximal shrinkage operators, explicit quadratic solutions, and simple gradient ascent for the Lagrangian multipliers. In our experiments, we have shown that the proposed method yields qualitatively good results, removes very minimal masking, and does so quickly in both iterations and time. From the histogram distributions of G, we observe a narrower spread around 1, yet a wider, more equidistributed support, suggesting that most of the time, there is minimal masking removal (multiplier close to 1), yet in areas of heavy striping, the destriping effect is more prevalent and of greater magnitude.

Applications of this algorithm are not limited to satellite imagery, and may be analogized to other fields such as raster scans in microscopy. Any scientific measurements (of images) made mostly along a curve – parameterizable by a single dimension – may be susceptible to such striping biases, and may be a candidate for similar destriping. Future work will expand this model to multi-modal images, color images, and may incorporate other specific priors on the data.

References

[AHL05] Said Assous, Anne Humeau, and J-P L'huillier. "Empirical mode decomposition applied to laser Doppler flowmetry signals : diagnosis approach." *IEEE Engineering in Medicine and Biology Conference (EMBC)*, **2**:1232–5, January 2005.

[ANK06] Adriano O. Andrade, Slawomir Nasuto, Peter Kyberd, Catherine M. Sweeney-Reed, and F.R. Van Kanijn. "EMG signal filtering based on empirical mode decomposition." *Biomedical Signal Processing and Control*, **1**(1):44–55, January 2006.

[ARS11] Alper Ayvaci, Michalis Raptis, and Stefano Soatto. "Sparse occlusion detection with optical flow." *International Journal of Computer Vision*, **97**(3):322–338, October 2011.

[BBB12] Dawn A. Bonnell, D. N. Basov, Matthias Bode, Ulrike Diebold, Sergei V. Kalinin, Vidya Madhavan, Lukas Novotny, Miquel Salmeron, Udo D. Schwarz, and Paul S. Weiss. "Imaging physical phenomena with local probes: from electrons to photons." *Reviews of Modern Physics*, **84**(3):1343–1381, September 2012.

[BE11] B.L. Barnhart and W.E. Eichinger. "Empirical mode decomposition applied to solar irradiance, global temperature, sunspot number, and CO2 concentration data." *Journal of Atmospheric and Solar-Terrestrial Physics*, **73**(13):1771–1779, August 2011.

[Ber76] Dimitri P Bertsekas. "Multiplier methods: A survey." *Automatica*, **12**(2):133–145, 1976.

[BPT88] Mario Bertero, Tomaso A Poggio, and Vincent Torre. "Ill-Posed Problems in Early Vision." In *Proceedings of the IEEE*, volume 76, pp. 869–889, 1988.

[BS99] Thomas Bülow and Gerald Sommer. "A novel approach to the 2D analytic signal." In *Computer Analysis of Images and Patterns*, pp. 25–32, 1999.

[Car22] J.R. Carson. "Notes on the theory of modulation." *Proceedings of the IRE*, **10**(1):57–64, February 1922.

[CD99] Emmanuel J. Candes and David L. Donoho. "Curvelets: A surprisingly effective nonadaptive representation for objects with edges." In *Curve and Surface Fitting*, pp. 105–120, 1999.

[CFL28] R Courant, K Friedrichs, and H Lewy. "Über die partiellen Differenzengleichungen der mathematischen Physik." *Mathematische Annalen*, **100**(1):32–74, 1928.

[Cha04] Antonin Chambolle. “An algorithm for total variation minimization and applications.” *Journal of Mathematical Imaging and Vision*, **20**:89–97, 2004.

[CI10] H. Carfantan and J. Idier. “Statistical Linear Destriping of Satellite-Based Pushbroom-Type Images.” *IEEE Transactions on Geoscience and Remote Sensing*, **48**(4):1860–1871, March 2010.

[CL13] Fang H. Yan L. Chang, Y. and H. Liu. “Robust destriping method for unidirectional total variation and framelet regularization.” *Optics Express*, **21**(20):23307–23323, October 2013.

[Coh96] Laurent D. Cohen. “Auxiliary variables and two-step iterative algorithms in computer vision problems.” *J. Math. Imaging Vis.*, **6**(1):59–83, January 1996.

[COP12] Marianne Clausel, Thomas Oberlin, and Valérie Perrier. “The monogenic synchrosqueezed wavelet transform: A tool for the decomposition/demodulation of AM-FM images.”, November 2012.

[CP11] Patrick L. Combettes and Jean-Christophe Pesquet. “Proximal splitting methods in signal processing.” In Heinz H. Bauschke, Regina S. Burachik, Patrick L. Combettes, Veit Elser, D. Russell Luke, and Henry Wolkowicz, editors, *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, volume 49 of *Springer Optimization and Its Applications*, pp. 185–212. Springer New York, New York, NY, 2011.

[Cri89] R.E. Crippen. “A simple spatial filtering routine for the cosmetic removal of scan-line noise from LANDSAT TM P-tape imagery.” *Photogrammetric Engineering and Remote Sensing*, **55**(3):327–331, March 1989.

[CSW11] Shelley A. Claridge, Jeffrey J. Schwartz, and Paul S. Weiss. “Electrons, photons, and force: quantitative single-molecule measurements from physics to biology.” *ACS Nano*, **5**(2):693–729, February 2011.

[CTS13] Shelley A Claridge, John Christopher Thomas, Miles A Silverman, Jeffrey J Schwartz, Yanlian Yang, Chen Wang, and Paul S Weiss. “Differentiating amino acid residues and side chain orientations in peptides using scanning tunneling microscopy.” *Journal of the American Chemical Society*, **135**(49):18528–18535, November 2013.

[CV01] T F Chan and L A Vese. “Active contours without edges.” *IEEE Transactions on Image Processing*, **10**(2):266–277, 2001.

[CY06] Lin H. Shao Y. Chen, J. and L. Yang. “Oblique striping removal in remote sensing imagery based on wavelet transform.” *International Journal of Remote Sensing*, **27**(8):1717–1723, April 2006.

[CZ03] Shao Y. Guo H. Wang W. Chen, J. and B. Zhu. "Destriping CMODIS data by power filtering." *IEEE Transactions on Geoscience and Remote Sensing*, **41**(9):2119–2124, September 2003.

[Dau88] Ingrid Daubechies. "Orthonormal bases of compactly supported wavelets." *Communications on Pure and Applied Mathematics*, **41**(7):909–996, October 1988.

[DB08] J.A. Dobrosotskaya and A.L. Bertozzi. "A Wavelet-Laplace Variational Technique for Image Deconvolution and Inpainting." *IEEE Transactins on Image Processing*, **17**(5):657–663, April 2008.

[DLL14] Weihua Dong, Xian'en Li, Xiangguo Lin, and Zhilin Li. "A bidimensional empirical mode decomposition method for fusion of multispectral and panchromatic remote sensing images." *Remote Sensing*, **6**(9):8446–8467, September 2014.

[DLW11] Ingrid Daubechies, Jianfeng Lu, and Hau-Tieng Wu. "Synchrosqueezed wavelet transforms: An empirical mode decomposition-like tool." *Applied and Computational Harmonic Analysis*, **30**(2):243–261, March 2011.

[DM13] Xu F. Garay M.J. Martonchik J.V. Rheingans B.E. Geier S. Davis A. Hancock B.R. Jovanovic V.M. Bull M.A. Capraro K. Chipman R.A. Diner, D.J. and S.C. McClain. "The Airborne Multiangle SpectroPolarimetric Imager (AirMSPI): A new tool for aerosol and cloud remote sensing." *Atmospheric Measurement Techniques*, **6**(8):2007–2025, 2013.

[Don95] D.L. Donoho. "De-noising by soft-thresholding." *IEEE Transactions on Information Theory*, **41**(3):613–627, May 1995.

[DV01] M.N. Do and M. Vetterli. "Pyramidal directional filter banks and curvelets." In *Proceedings 2001 International Conference on Image Processing*, volume 2, pp. 158–161. IEEE, 2001.

[DZ14] Konstantin Dragomiretskiy and Dominique Zosso. "Variational mode decomposition." *IEEE Transactions on Signal Processing*, **62**(3):531–544, February 2014.

[DZ15] Konstantin Dragomiretskiy and Dominique Zosso. "Two-dimensional variational mode decomposition." In Xue-Cheng Tai, Egil Bae, Tony F. Chan, and Marius Lysaker, editors, *Energy Minimization Methods in Computer Vision and Pattern Recognition*, volume 8932 of *Lecture Notes in Computer Science*, pp. 197–208. Springer International Publishing, 2015.

[EO15] Selim Esedoglu and Felix Otto. "Threshold dynamics for networks with arbitrary surface tensions." *Communications on Pure and Applied Mathematics*, **68**(5):808–864, 2015.

[ET06] Selim Esedoglu and Yen Hsi Richard Tsai. "Threshold dynamics for the piecewise constant Mumford-Shah functional." *Journal of Computational Physics*, **211**:367–384, 2006.

[EW14] Matt Elsey and Benedikt Wirth. "Fast automated detection of crystal distortion and crystal defects in polycrystal images." *Multiscale Modeling & Simulation*, January 2014.

[EZB14] Virginia Estellers, Dominique Zosso, Xavier Bresson, and Jean-Philippe Thiran. "Harmonic active contours." *IEEE Transactions on Image Processing*, **23**(1):69–82, January 2014.

[Fel06] Michael Feldman. "Time-varying vibration decomposition and analysis based on the Hilbert transform." *Journal of Sound and Vibration*, **295**(3-5):518–530, August 2006.

[FGR05] Patrick Flandrin, Paulo Gonçalvès, and Gabriel Rilling. "EMD equivalent filter banks, from interpretation to applications." In *Hilbert-Huang Transform and Its Applications*, pp. 57–74. 2005.

[FPS08] N. Fauchereau, G. G. S. Pegram, and S. Sinclair. "Empirical mode decomposition on the sphere: application to the spatial scales of surface temperature variations." *Hydrology and Earth System Sciences*, **12**(3):933–941, June 2008.

[Gab46] D. Gabor. "Theory of communication." *Journal of the Institution of Electrical Engineers - Part III: Radio and Communication Engineering*, **93**(26):429–457, 1946.

[Gil12] Jérôme Gilles. "Multiscale texture separation." *Multiscale Modeling & Simulation*, **10**(4):1409–1427, December 2012.

[Gil13] Jérôme Gilles. "Empirical wavelet transform." *IEEE Transactions on Signal Processing*, **61**(16):3999–4010, August 2013.

[GL89] Roland Glowinski and Patrick Le Tallec. *Augmented Lagrangian and operator-splitting methods in nonlinear mechanics.* Society for Industrial and Applied Mathematics (SIAM), Philadelphia, 1989.

[GL07] Kanghui Guo and Demetrio Labate. "Optimally sparse multidimensional representation using shearlets." *SIAM Journal on Mathematical Analysis*, **39**(1):298–318, January 2007.

[GLS13] George Georgoulas, Theodore Loutas, Chrysostomos D. Stylios, and Vassilis Kostopoulos. "Bearing fault detection based on hybrid ensemble detector and empirical mode decomposition." *Mechanical Systems and Signal Processing*, **41**(1-2):510–525, December 2013.

[GM75] R. Glowinski and A. Marrocco. "Sur l'approximation, par elements finis d'ordre un, et la resolution, par, penalisation-dualité, d'une classe de problems de Dirichlet non lineares." *Revue Française d'Automatique, Informatique, et Recherche Opéationelle*, **9**:41–76, January 1975.

[GM76] D. Gabay and B. Mercier. "A dual algorithm for the solution of nonlinear variational problems via finite element approximations." *Computers and Mathematics with Applications*, **2**:17–40, January 1976.

[GMB14] Cristina Garcia-Cardona, Ekaterina Merkurjev, Andrea L Bertozzi, Arjuna Flenner, and Allon G. Percus. "Multiclass data segmentation using diffuse interface methods on graphs." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **36**(8):1600–1613, 2014.

[GO08] T. Goldstein and S. Osher. "The split bregman method for L1 regularized problems." *SIAM Journal of Imaging Science*, **2**(2):323–343, 2008.

[GO09] Tom Goldstein and Stanley Osher. "The split Bregman method for L1-regularized problems." *SIAM Journal on Imaging Sciences*, **2**(2):323–343, January 2009.

[GTO14] Jérôme Gilles, Giang Tran, and Stanley Osher. "2D empirical transforms. Wavelets, ridgelets, and curvelets revisited." *SIAM Journal on Imaging Sciences*, **7**(1):157–186, January 2014.

[GW92] Rafael C. Gonzalez and Richard E. Woods. *Digital Image Processing*. Addison Wesley, 1992.

[Hah96] Stefan L. Hahn. *Hilbert transforms in signal processing*. Artech House, Inc., 1996.

[HB13] Jiajun Han and Mirko van der Baan. "Empirical mode decomposition for seismic time-frequency analysis." *GEOPHYSICS*, **78**(2):O9–O19, February 2013.

[Hes69] Magnus R. Hestenes. "Multiplier and gradient methods." *Journal of Optimization Theory and Applications*, **4**(5):303–320, 1969.

[HH92] Quirk B.K. Helder, D.L. and J.J. Hood. "A technique for the reduction of banding in LANDSAT Thematic Mapper images." *Photogrammetric Engineering and Remote Sensing*, **58**(10):1425–1431, October 1992.

[HKG09] Patrick Han, Adam R. Kurland, Andrea N. Giordano, Sanjini U. Nanayakkara, Meaghan M. Blake, Chris M. Pochas, and Paul S. Weiss. "Heads and tails: simultaneous exposed and buried interface imaging of monolayers." *ACS Nano*, **3**(10):3115–3121, October 2009.

[HL11] Meng Hu and Hualou Liang. "Intrinsic mode entropy based on multivariate empirical mode decomposition and its application to neural data analysis." *Cognitive*

Neurodynamics, **5**(3):277–284, September 2011.

[HS11] Thomas Y. Hou and Zuoqiang Shi. "Adaptive data analysis via sparse time-frequency representation." *Advances in Adaptive Data Analysis*, **03**(1 & 2):1–28, April 2011.

[HS13] Thomas Y. Hou and Zuoqiang Shi. "Data-driven timefrequency analysis." *Applied and Computational Harmonic Analysis*, **35**(2):284–308, September 2013.

[HS15] Thomas Y. Hou and Zuoqiang Shi. "Sparse time-frequency decomposition for multiple signals with same frequencies." July 2015.

[HSB14] Huiyi Hu, Justin Sunu, and Andrea L. Bertozzi. "Multi-class graph Mumford-Shah model for plume detection using the MBO scheme.", 2014.

[HSL98] N. E. Huang, Z. Shen, S. R. Long, M. C. Wu, H. H. Shih, Q. Zheng, N.-C. Yen, C. C. Tung, and H. H. Liu. "The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis." *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, **454**(1971):903–995, March 1998.

[Klu12] Niklas Klügel. "Practical empirical mode decomposition for audio synthesis." In *Int. Conference on Digital Audio Effects (DAFx-12)*, number 2, pp. 15–18, 2012.

[LHG08] Shaopeng Liu, Qingbo He, Robert X Gao, and Patty Freedson. "Empirical mode decomposition applied to tissue artifact removal from respiratory signal." *IEEE Engineering in Medicine and Biology Conference (EMBC)*, pp. 3624–3627, January 2008.

[LLK05] Demetrio Labate, Wang-Q Lim, Gitta Kutyniok, and Guido Weiss. "Sparse multidimensional representation using shearlets." In Manos Papadakis, Andrew F. Laine, and Michael A. Unser, editors, *Optics & Photonics 2005*, pp. 1–9. International Society for Optics and Photonics, August 2005.

[LPD14] Marco Leo, Roberta Piccolo, Cosimo Distante, Pasquale Memmolo, Melania Paturzo, and Pietro Ferraro. "Multilevel bidimensional empirical mode decomposition: a new speckle reduction method in digital holography." *Optical Engineering*, **53**(11):112314, June 2014.

[Mal89] S.G. Mallat. "A theory for multiresolution signal decomposition: the wavelet representation." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **11**(7):674–693, July 1989.

[MBL09] Isar Mostafanezhad, Olga Boric-Lubecke, Victor Lubecke, and Danilo P Mandic. "Application of empirical mode decomposition in removing fidgeting interference in doppler radar life signs monitoring devices." *IEEE Engineering in Medicine and Biology Conference (EMBC)*, pp. 340–343, January 2009.

[MBO92] B Merriman, J K Bence, and S Osher. "Diffusion generated motion by mean curvature." *UCLA CAM Report 92-18*, pp. 1–12, 1992.

[MBO94] Barry Merriman, James K Bence, and Stanley J Osher. "Motion of multiple junctions: a level set approach." *Journal of Computational Physics*, **112**(2):334–363, June 1994.

[Mod87] Luciano Modica. "The gradient theory of phase transitions and the minimal interface criterion." *Archive for Rational Mechanics and Analysis*, **98**:123–142, 1987.

[Mor75] VA A Morozov. "Linear and nonlinear ill-posed problems." *Journal of Mathematical Sciences*, **II**(6):706–736, 1975.

[MP07] S. Meignen and V. Perrier. "A new formulation for empirical mode decomposition based on constrained optimization." *IEEE Signal Processing Letters*, **14**(12):932–935, December 2007.

[MS09] Trtik P. Marone F. Muench, B. and M. Stampanoni. "Stripe and ring artifact removal with combined wavelet–Fourier filtering." *Optics Express*, **17**(10):8567–8591, May 2009.

[MYY10] Amanda M. Moore, Sina Yeganeh, Yuxing Yao, Shelley A. Claridge, James M. Tour, Mark A. Ratner, and Paul S. Weiss. "Polarizabilities of adsorbed and assembled molecules: measuring the conductance through buried contacts." *ACS Nano*, **4**(12):7630–7636, December 2010.

[NBD03] J.C Nunes, Y Bouaoune, E Delechelle, O Niang, and Ph Bunel. "Image analysis by bidimensional empirical mode decomposition." *Image and Vision Computing*, **21**(12):1019–1026, November 2003.

[NW06] J Nocedal and S J Wright. *Numerical optimization.* Springer, Berlin, 2nd edition, 2006.

[OWO14] Braxton Osting, Chris D. White, and Édouard Oudet. "Minimal Dirichlet energy partitions for graphs." *SIAM Journal on Scientific Computing*, **36**(4):A1635–A1651, August 2014.

[RF92] Osher S. Rudin, L.I. and E. Fatemi. "A Nonlinear total variation based noise removal algorithms." *Physica D*, **60**:256–268, January 1992.

[RF08] G. Rilling and P. Flandrin. "One or two frequencies? The empirical mode decomposition answers." *IEEE Transactions on Signal Processing*, **56**(1):85–95, January

2008.

[RF09] Gabriel Rilling and Patrick Flandrin. “Sampling effects on the empirical mode decomposition.” *Advances in Adaptive Data Analysis*, **01**(01):43–59, January 2009.

[RFG03] Gabriel Rilling, Patrick Flandrin, and Paulo Gonçalvès. “On empirical mode decomposition and its algorithms.” In *IEEE-EURASIP workshop on Nonlinear Signal and Image Processing*, 2003.

[RO94] L Rudin and S Osher. “Total Variation based Image Restoration with Free Local Constraints.” *Proc. IEEE Int. Conf. Image Process.*, **1**:31–35, 1994.

[Roc73a] R. Tyrrell Rockafellar. “A dual approach to solving nonlinear programming problems by unconstrained optimization.” *Mathematical Programming*, **5**(1):354–373, December 1973.

[Roc73b] R.T. Rockafeller. “A dual approach to solving nonlinear programming problems by unconstrained optimization.” *Mathematical Programming*, **5**(1):354–373, January 1973.

[Ruu98] SJ Ruuth. “Efficient algorithms for diffusion-generated motion by mean curvature.” *Journal of Computational Physics*, **625**:603–625, 1998.

[RY07] Takeuchi W. Rakwatin, P. and Y. Yasuoka. “Stripe noise reduction in MODIS data by combining histogram matching with facet filter.” *IEEE Transactions on Geoscience and Remote Sensing*, **45**(6):1844–1856, June 2007.

[SBC12] Z. B. Szuts, J. R. Blundell, M. P. Chidichimo, and J. Marotzke. “A vertical-mode decomposition to investigate low-frequency internal motion across the Atlantic at 26° N.” *Ocean Science*, **8**(3):345–367, June 2012.

[SL98] Stitt J.R. Simpson, J.J. and D.M. Leath. “Improved finite impulse response filters for enhanced destriping of geostationary satellite data.” *Optical Engineering*, **66**(3):235–249, December 1998.

[SMH05] E. Charles H. Sykes, Brent A. Mantooth, Patrick Han, Zachary J. Donhauser, and Paul S. Weiss. “Substrate-mediated intermolecular interactions: a quantitative single molecule analysis.” *Journal of the American Chemical Society*, **127**(19):7255–7260, May 2005.

[SPB14] Jérémy Schmitt, Nelly Pustelnik, Pierre Borgnat, Patrick Flandrin, and Laurent Condat. “2-D Prony-Huang transform: A new tool for 2-D spectral analysis.” p. 24, April 2014.

[SV05] Robert C. Sharpley and Vesselin Vatchev. “Analysis of the intrinsic mode functions.” *Constructive Approximation*, **24**(1):17–47, August 2005.

[SW88] R. Srinivasan and M.C.J. White. “Landsat data destriping using power spectral filtering.” *Optical Engineering*, **27**(11):939–943, November 1988.

[SZ09] H. Shen and L. Zhang. “A MAP-based algorithm for destriping and inpainting of remotely sensed images.” *IEEE Transactions on Geoscience and Remote Sensing*, **47**(5):1492–1502, May 2009.

[TA] A.N. Tikhonov and V.Y. Arsenin. Winston and Sons, Washington.

[Tai96] Tai Sing Lee. “Image representation using 2D Gabor wavelets.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **18**(10):959–971, 1996.

[THS14] Peyman Tavallali, Thomas Y. Hou, and Zuoqiang Shi. “Extraction of intrawave signals using the sparse time-frequency representation method.” *Multiscale Modeling & Simulation*, **12**(4):1458–1493, October 2014.

[TI01] J. Torres and S.O. Infante. “Wavelet analysis for the elimination of striping noise in satellite images.” *Optical Engineering*, **40**(7):1309–1314, July 2001.

[Tic63] A N Tichonov. “Solution of incorrectly formulated problems and the regularization method.” *Soviet Mathematics*, **4**:1035–1038, 1963.

[TSH15] John C. Thomas, Jeffrey J. Schwartz, J. Nathan Hohman, Shelley A. Claridge, Harsharn S. Auluck, Andrew C. Serino, Alexander M. Spokoyny, Giang Tran, Kevin F. Kelly, Chad A. Mirkin, Jerome Gilles, Stanley J. Osher, and Paul S. Weiss. “Defect-tolerant aligned dipoles within two-dimensional plastic lattices.” *ACS Nano*, **9**(5):4734–4742, May 2015.

[TZY11] Jian Tang, Lijie Zhao, Heng Yue, Wen Yu, and Tianyou Chai. “Vibration analysis based on empirical mode decomposition and partial least square.” *Procedia Engineering*, **16**:646–652, 2011.

[USV09] Michael Unser, Daniel Sage, and Dimitri Van De Ville. “Multiresolution Monogenic Signal Analysis Using the Riesz Laplace Wavelet Transform.” *IEEE Transactions on Image Processing*, **18**(11):2402–2418, 2009.

[WFD11] Hau-Tieng Wu, Patrick Flandrin, and Ingrid Daubechies. “One or two frequencies? The synchrosqueezing answers.” *Advances in Adaptive Data Analysis*, **03**(01n02):29–39, April 2011.

[WH09] Zhaohua Wu and Norden E. Huang. “Ensemble empirical mode decomposition: A noise-assisted data analysis method.” *Advances in Adaptive Data Analysis*, **01**(01):1–41, January 2009.

[Wie49] Norbert Wiener. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series.* 1949.

[WZ07] Yin W. Wang, Y. and Y. Zhang. "A fast algorithm for image deblurring with total variation regularization." *CAAM Technical Report*, (1):1–19, June 2007.

[WZ08] Yang J. Yin W. Wang, Y. and Y. Zhang. "A new alternating minimization algorithm for total variation image reconstruction." *SIAM Journal of Imaging Science*, **1**(3):248–272, 2008.

[Yan13] Ming Yan. "Restoration of images corrupted by impulse noise and mixed gaussian impulse noise using blind inpainting." *SIAM Journal on Imaging Sciences*, **6**(3):1227–1245, July 2013.

[YLY14] Haizhao Yang, Jianfeng Lu, and Lexing Ying. "Crystal image analysis using 2D synchrosqueezed transforms." p. 27, February 2014.

[YY13] Haizhao Yang and Lexing Ying. "Synchrosqueezed wave packet transform for 2D mode decomposition." *SIAM Journal on Imaging Sciences*, **6**(4):1979–2009, January 2013.

[ZAS15] Dominique Zosso, Jing An, James Stevick, Nicholas Takaki, Morgan Weiss, Liane S. Slaughter, Huan H. Cao, Paul S. Weiss, and Andrea L. Bertozzi. "Image segmentation with dynamic artifacts detection and bias correction." *submitted to: AIMS J. Inverse Problems and Imaging*, (submitted), 2015.

[ZBT14] Dominique Zosso, Xavier Bresson, and Jean-Philippe Thiran. "Fast geodesic active fields for image registration based on splitting and augmented Lagrangian approaches." *IEEE Transactions on Image Processing*, **23**(2):673–683, March 2014.

[From Duplicate 1 (.]

[ZWC10] Mingqiang Zhu, Stephen J. Wright, and Tony F. Chan. "Duality-based algorithms for total-variation-regularized image restoration." *Computational Optimization and Applications*, **47**(3):377–400, December 2010.