UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

HIGH-FIDELITY AND PERFECT RECONSTRUCTION TECHNIQUES FOR SYNTHESIZING MODULATION DOMAIN FILTERED IMAGES

A DISSERTATION

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

DOCTOR OF PHILOSOPHY

By

PATRICK ADRIAN CAMPBELL
Norman, Oklahoma
2016

Downloaded from SHAREOK. Persistent Handle URL: https://hdl.handle.net/11244/47112

HIGH-FIDELITY AND PERFECT RECONSTRUCTION TECHNIQUES
FOR SYNTHESIZING MODULATION DOMAIN FILTERED IMAGES

A DISSERTATION APPROVED FOR THE
SCHOOL OF ELECTRICAL AND COMPUTER ENGINEERING

BY

Joseph P. Havlicek

Ronald D. Barnes

Monte P. Tull

Mohammed Atiquzzaman

Tomasz Przebinda

©Copyright by PATRICK ADRIAN CAMPBELL 2016
All Rights Reserved.

This dissertation is dedicated to my mother,

Sarah Jo Campbell,

for her gift to me of a beautiful life -

a gift requiring so much more than birth,

possible only through unending belief and support.

Acknowledgements

This dissertation would not have been possible without the support of many others. I would like to thank each of the members of my committee for their invaluable contributions to my education. Over the course of my studies, I have learned a great deal from each of them. In particular, I am very grateful to my committee chair, Dr. Joseph "Joebob" Havlicek. He has provided me with extraordinary vision and guidance ever since I first took a class under him, and he has taught me a tremendous amount concerning methods of image processing. I am also grateful to have had the opportunity to both support my education and positively impact the lives of Oklahoma citizens by contributing to the projects at the OU Center for Intelligent Transportation Systems.

I also very much wish to thank my friends and family, who have always been with me as I continued and continued my research. Most especially, my mother has always provided both unconditional support and an anchoring voice of reason. It is entirely thanks to her that I am a person capable of this achievement. With such a foundation of support, I look forward to whatever new challenges may lie in the future.

Table of Contents

Acknowledgements **iv**

List of Tables **vii**

List of Figures **viii**

Abstract **xiii**

Chapter 1. Introduction **1**
1.1 Dissertation Organization 5
1.2 Notation and Nomenclature 6

Chapter 2. Background **8**
2.1 Inspiration for AM-FM Analysis 8
2.1.1 Image Composition and Representation 8
2.1.2 Analysis Informed by Study of Biological Vision 11
2.1.3 Signal Reconstruction 13
2.2 AM-FM Demodulation and Analysis 13
2.2.1 Gabor Filterbank Decomposition 18
2.2.2 Demodulation Via Energy Separation 19
2.2.3 Analytic Image Based Demodulation 22
2.2.4 Component Analysis Paradigms 31
2.2.5 Practical Analysis Applications 34
2.3 Orthogonal Filterbank Decomposition 39
2.3.1 Orthogonal Filterbanks for Reconstruction 40
2.3.2 The Steerable Pyramid Filterbank 41
2.4 Spline Representation and Computation 42
2.4.1 Spline Representation 43
2.4.2 Spline Differentiation and Integration 46
2.5 Phase Unwrapping 49
2.6 Perfect AM-FM Reconstruction 54

2.7 Recent Application of the AM-FM Model 56
2.7.1 Infrared Target Tracking 57
2.7.2 Texture Component Decomposition 58
2.8 Modulation Domain Filtering 58
2.8.1 AM Filters 59
2.8.2 FM Filters 61
2.9 Summary 65

Chapter 3. Phase Unwrapping for Modulation Domain Filtering 67
3.1 Challenges Presented by Modulation Domain Filtering 68
3.2 Frequency Guided Phase Unwrapping 70
3.3 Improved Stability in the AM-FM Transform 75
3.4 Summary 77

Chapter 4. Image Reconstruction From Processed Modulations 78
4.1 Reconstruction Techniques 79
4.1.1 Image-Wide Integration 79
4.1.2 Subwindow Integration from Multiple Points 85
4.1.3 Phase Growing from Multiple Points 88
4.2 Avoidance of Phase Discontinuities During Interpolation 91
4.3 Modulation Domain Scaling, Rotation, and Translation 96
4.4 Filtering via Attenuation of Amplitude Modulation 100
4.5 Rotation and Scaling of Target Image Texture 103
4.6 Cross-Component Amplitude Attenuation 111
4.7 Cross-Component Scaling of Target Image Texture 116
4.8 Summary 119

Chapter 5. Conclusion and Future Work 121
5.1 Dissertation Summary 121
5.2 Contributions to the Field 123
5.3 Future Work 124

Bibliography 127

List of Tables

3.1 Queue-based phase region growing algorithm 72
3.2 Mean squared reconstruction error for *Chirp*, *Lines*, and *Lena* images using different phase unwrapping algorithms for AM-FM reconstruction. 76

List of Figures

2.1 Fourier transform example. (a) Original *Lena* image. (b) Fourier transform of (a). 9

2.2 From [27], the envelope of the elementary signal in both the time and frequency domains. 10

2.3 From [89] (a) Input *Lena* image. (b) A Gabor decomposition channel at 30 degrees. 13

2.4 From [88] (a) Input *Lena* image. (b) Reconstruction from 12626 5x5 Gabor elementary functions. (c) Reconstruction from 4062 GEF's. 14

2.5 From [9] (a) Input image of juxtaposed textures. (b) Needle diagram showing computed emergent image frequencies. (c) Computed amplitude modulation. 17

2.6 From [30], 40 channel Gabor filterbank used for decomposition. 19

2.7 From [66], DESA applied to speech. (a) Signal s(n) from speech vowel /E/. (b) Estimated amplitude envelope using DESA, from (a) filtered by Gabor filter centered around 1580 Hz. (c) Estimated instantaneous frequency, smoothed by an 11-point median filter. Dotted line is Gabor filter center. 23

2.8 From [63], DESA applied to a texture image. (a) Original *wood* texture image. (b) Estimated amplitude envelope using DESA, from (a) filtered by Gabor filter centered around horizontal and vertical frequencies of 35.5 and 14.7 cycles per image. (c) Estimated instantaneous frequency vectors, superimposed upon the bandpass image of (a). 24

2.9 From [39], reconstruction from Tracked Components Analysis. (a) Original *Reptile* image. (b) Reconstruction of (a) from six tracked AM-FM components. (c) Original *Raffia* image. (d) Reconstruction of (c) from nine tracked AM-FM components. . 33

2.10 From [42], reconstruction from Dominant Component Analysis. (a) Original *Tree* image. (b) Reconstruction of (a) from a single dominant AM-FM component. 34

2.11 From [42], reconstruction from Channelized Components Analysis. (a) Original *Mandrill* image. (b) Reconstruction of (a) from a 43 channelized AM-FM components. (c) Original *Peppers* image. (d) Channelized Components Analysis Reconstruction of (c). 35

2.12 From [1], texture inpainting applied to a texture image. (a) Original *Bark* texture image. (b) Occluded image from (a). (c) Image from (b) after texture inpainting by AM-FM reaction-diffusion. 38

2.13 From [97], the discrete wavelet transform filterbank developed by Tay. The gray levels in this representation have been scaled to improve display. 41

2.14 From [73], frequency domain representation of example Steerable Pyramid filterbank. 43

2.15 Example image components obtained via Steerable Pyramid decomposition. (a) Original *Lena* image. (b-m) Components at orientations: 7.5° (column 1), 60° (column 2), 120° (column 3). 44

2.16 Example modulation functions for image components. (a) Original *Lena* image. (b-j) Example components (column 1) and corresponding amplitude modulations (column 2) and frequency modulations (column 3). 55

2.17 From [84], example of modulation domain filtering on amplitude modulation functions. (a) Original *Lena* texture image. (b) Amplitude modulation attenuation within the black rectangle of (a) to remove the bands of Lena's hat. 59

2.18 From [76], example of modulation domain filtering on amplitude modulation functions. (a),(b) Input clock images. (c) AM-FM fusion of (a) and (b). 60

2.19 From [74], modulation domain filtering on frequency modulation functions. (a) Original *Chirp* image. (b) Rotation of frequency vectors in (a) using least squares based reconstruction. (c) Flattening of frequency vectors in (a) using integration based reconstruction. (d) Rotation and inversion of frequency vectors in (a) using integration based reconstruction. 62

2.20 From [74], modulation domain filtering on frequency modulation functions in a natural image. (a) Original *Barbara* image. Processing was performed within the black rectangle. (b) Detail view of original image. (c) Result of FM processing to rotate the stripes on the pants. 63

2.21 From [85], modulation domain scaling and rotation. (a) *Barbara* image zoomed in by 2.0. (d) *Boat* image rotated by 45 degrees. 64

3.1 Phase unwrapping example. (a) Real part of component t_{20} of *Chirp* image. (b) Unwrapped phase from (a) using frequency guided phase unwrapping. (c) FM function $\nabla\varphi_{20}$ obtained by differentiating (b). 74

3.2 Perfect reconstruction examples. (a) Reconstructed *Chirp* image. (b) Reconstructed *Lines* image. (c) Reconstructed *Lena* image. 75

4.1 Perfect reconstruction from unmodified modulation functions using standard image-wide spline integration. 80
4.2 Rotated *Lena* image reconstructed directly from rotated AM and PM functions. 81
4.3 Perfect reconstruction from unmodified modulation functions using spline integration split into four subwindows. 83
4.4 Estimated reconstruction from unmodified modulation functions using spline integration split into four subwindows. Integration is computed from a single initial phase value, and other required initial phase values are estimated. 84
4.5 Estimated reconstruction from rotated modulation functions using spline integration split into four subwindows. (a) Reconstruction of entire rotated image. (b) Reconstruction of single component of image. 85
4.6 Reconstruction using windowed integration from multiple points without phase growing step. 86
4.7 Estimated reconstruction using windowed integration from multiple points with phase growing step. (a) Result of phase growing algorithm run with requeuing. (b) Result of phase growing algorithm run without requeuing. 88
4.8 Estimated reconstruction from unmodified modulation functions using phase growing only algorithm. 90
4.9 Estimated reconstruction from rotated modulation functions using phase growing only algorithm. (a) Reconstruction of entire rotated image. White arrow points to interpolation artifact. (b) Reconstruction of phase from rotated frequency modulation for a single component. The component displays a phase discontinuity that is visible as artifacts in (a). 91
4.10 Comparison of interpolation methods for reconstruction of rotated mid-frequency channel image. (a) Unrotated region from zoomed in section of component image. (b) Reconstruction of rotated region from (a) with interpolation across phase discontinuities. (c) Reconstruction of rotated region from (a) with special interpolation algorithm used near phase discontinuities. (d) Reconstruction of entire channel image with interpolation across phase discontinuities. White arrow points to interpolation artifacts. (e) Reconstruction of entire channel image with special interpolation near phase discontinuities. 93
4.11 Comparison of interpolation methods for reconstruction of rotated low-frequency channel image. (a) Interpolation across phase discontinuities. White arrow points to interpolation artifact. (b) Special interpolation near phase discontinuities. . . . 94
4.12 Comparison of interpolation methods for reconstruction of rotated natural image. (a) Input *Lena* image. (b) Interpolation across phase discontinuities. White arrow points to interpolation artifact. (c) Special interpolation near phase discontinuities. . . 95

4.13 Example of rotation operation in modulation domain. (a) Rotation in real domain. (b) Rotation in modulation domain. (c) Difference image between (a) and (b). 97

4.14 Example of scaling operation in modulation domain. (a) Scaling and Rotation in real domain. (b) Scaling and Rotation in modulation domain. (c) Difference image between (a) and (b). 98

4.15 Example of translation operation in modulation domain. (a) Translation in real domain. (b) Translation in modulation domain. (c) Difference image between (a) and (b). 99

4.16 Amplitude modulation attenuation on *Boat* image. (a) Original image. Filter operations were only applied within the bottom half of the image. (b) Attenuation of amplitude modulation at target frequency characteristics to remove ripples in the water. (c) Modified attenuation of amplitude modulation to leave high frequency ripple content. 101

4.17 Amplitude modulation attenuation on *Wakes* image. (a) Original image. Filter operations were only applied within the bottom half of the image. (b) Attenuation of amplitude modulation at target frequency characteristics to remove one of the boat wakes. 103

4.18 Rotation of specific texture in image. (a) Original *Barbara* image. Image processing is applied to the region within the black rectangle. (b) Region of interest from (a). (c) Previous result from [74]. (d) Rotated texture, using phase growing reconstruction. 104

4.19 (a) Region of interest from original *Barbara* image. (b) Rotation of (a) with spatially rotated windows only. (c) Rotation of (a) with phase growing from windows in (b). 107

4.20 Different scaling factors and orientations for specific texture rotation. (a) Rotation of $\pi/2$ and scaling factor of 2. (b) Rotation of $\pi/2$ and scaling factor of 3. (c) Rotation of $\pi/2$ and scaling factor of 5. (d) Rotation of $\pi/4$ and scaling factor of 3. 109

4.21 Select results from Fig. 4.20 shown in full *Barbara* image. (a) Original *Barbara* image. (b) Rotation of $\pi/2$ and scaling factor of 5 within ROI. (c) Rotation of $\pi/4$ and scaling factor of 3 within ROI. 110

4.22 Rotation of image texture in *Boat* image. (a) Original image. Filter operations were only applied within the bottom half of the image. (b) Modification of target frequency characteristics to rotate ripples in the water. 112

4.23 Removal of bands in *Lena*'s hat via amplitude attenuation on individual components. (a) Previous result from [84]. (b) Reproduced frequency matching AM attenuation. The boundaries of the ROI are visible. (c) Filter operation from (b) excluding the lowest frequency channels. 114

4.24 Removal of bands in *Lena*'s hat via cross-component amplitude attenuation. (a) Input *Lena* image. (b) Cross-Component attenuation over all channels. Filter is also constrained by AM magnitude limits. 116

4.25 Cross-component frequency modulation scaling in *Steps* image. (a) Original image. Filter operations were only applied within the bottom half of the image. (b) Target frequency characteristics scaled by 0.5 to double the number of steps. (c) Target frequency characteristics scaled by 2 to halve the number of steps. 118

Abstract

HIGH-FIDELITY AND PERFECT RECONSTRUCTION TECHNIQUES FOR SYNTHESIZING MODULATION DOMAIN FILTERED IMAGES

Patrick Adrian Campbell, Ph.D.
The University of Oklahoma, 2016

Supervisor: Joseph P. Havlicek

Biomimetic processing inspired by biological vision systems has long been a goal of the image processing research community, both to further understanding of what it means to perceive and interpret image content and to facilitate advancements in applications ranging from processing large volumes of image data to engineering artificial intelligence systems. In recent years, the *AM-FM transform* has emerged as a useful tool that enables processing that is intuitive to human observers but would be difficult or impossible to achieve using traditional linear processing methods.

The transform makes use of the multicomponent AM-FM image model, which represents imagery in terms of amplitude modulations, representative of local image contrast, and frequency modulations, representative of local spacing and orientation of lines and patterns. The model defines image components using an array of narrowband filterbank channels that is designed to be similar to the spatial frequency channel decomposition that occurs in the human

visual system. The AM-FM transform entails the computation of modulation functions for all components of an image and the subsequent exact recovery of the image from those modulation functions. The process of modifying the modulation functions to alter visual information in a predictable way and then recovering the modified image through the AM-FM transform is known as *modulation domain filtering*. Past work in modulation domain filtering has produced dramatic results, but has faced challenges due to phase wrapping inherent in the transform computations and due to unknown integration constants associated with modified frequency content. The approaches developed to overcome these challenges have led to a loss of both stability and intuitive simplicity within the AM-FM model.

In this dissertation, I have made significant advancements in the underlying processes that comprise the AM-FM transform. I have developed a new phase unwrapping method that increases the stability of the AM-FM transform, allowing higher quality modulation domain filtering results. I have designed new reconstruction techniques that allow for successful recovery from modified frequency modulations. These developments have allowed the design of modulation domain filters that, for the first time, do not require any departure from the simple and intuitive nature of the basic AM-FM model. Using the new modulation domain filters, I have produced new and striking results that achieve a variety of image processing tasks which are motivated by biological visual perception. These results represent a significant advancement relative to the state of the art and are a foundation from which future advancements in the field may be attained.

Chapter 1

Introduction

A recurrent theme in image processing and computer vision has been, to the extent possible, to analyze and process images in ways that are similar to and consistent with what is known about the processing of visual information in biological vision systems. It has been a long standing goal of researchers and scientists to create machines and automated processes that can act according to the visually meaningful content within an image, *i.e.* like a human being can. This has been partly out of a desire to further understanding of fundamental theory and satisfy scientific inquiry, but has been perhaps more driven by the potential of such technology to unlock an array of advancements around us with potential to improve our quality of life, from image processing within our smartphones to web-driven computer algorithms that process volumes of image data that humans simply can not.

Critical to meaningful image analysis was the development of Fourier analysis and its application as an image representation to describe frequency content in images. However, the Fourier transform is only descriptive of frequency content at a global scale within an image. Biological vision systems are able to clearly observe and understand underlying image descriptive content such as frequency that is *local* in space. The development of methods by which to mathematically model local image content has, in the last half century or

so, driven image analysis and processing technologies ever further in their capability to work with image information in a way that is intuitive to human beings.

One important image representation model that incorporates local image content has been the *AM-FM* image model. This simple model represents a real-valued image s as [8]

$$s(\mathbf{x}) = a(\mathbf{x}) \cos(\varphi(\mathbf{x})), \tag{1.1}$$

where $a(\mathbf{x})$ is the local amplitude modulation function, which can be thought of as descriptive of changes in local image contrast, and $\nabla\varphi(\mathbf{x})$ is the local frequency modulation function, which describes local spacing and orientation of lines and patterns within the image. Combined with a decomposition of the image into narrowband component channels that allow for smoothly varying amplitude and frequency content, these two modulation functions together provide a meaningful description of the texture within an image as formulated by the multicomponent AM-FM image model [34]

$$z(\mathbf{x}) = \sum_{i=1}^{K} a_i(\mathbf{x}) e^{j\varphi_i(\mathbf{x})}, \tag{1.2}$$

where the complex-valued image z is represented as a sum of K locally coherent AM-FM components. This has led the model to successful application within a wide variety of fields and pursuits. These have included speech analysis [28, 61, 65, 66], image segmentation [9, 30, 42, 108, 122], fingerprint recognition [58], content-based image retrieval [47], infrared target tracking [49, 75, 94], and nonlinear image processing [74, 84, 85] to name a few.

Application of the AM-FM model to the last item mentioned above is a relatively new field of study, which was only made possible once the AM-

FM model had been developed to the point where perfect reconstruction of the original image from its AM and FM functions could be practically carried out. This development was made difficult due to the discrete nature of any practical image processing, which at first necessitated that the AM and FM functions associated with images be approximated, from which close, but not exact reconstructions could be achieved. However, careful advances in filter-bank selection for decomposition [84, 111, 113], fitting of the model to spline based computational methods [97, 98], and addressing the inherent problem of phase wrapping [97–99] pushed the model to a point where the original image *could* be exactly reconstructed from its amplitude and frequency content to within machine error. At this point, representation of an image by its modulation domain functions and the corresponding ability to obtain either modulation domain functions from the image or the image from its modulation domain functions became known as the *AM-FM transform* [84]. This then opened up the possibility of modifying the amplitude and frequency content to achieve image processing goals that are motivated by visual perception and then reconstructing a new image from the modified content, a process known as *modulation domain filtering*.

Previous results have been obtained in the field of modulation domain filtering. These results have shown dramatic image processing tasks that would be difficult or impossible to obtain through traditional filtering methods, however they suffer from two major problems. The results are achieved through the usage of constructs and methods that prevent the development of intuitive modulation domain filters that are driven by visual perception, and the results exhibit low-fidelity reconstruction for filters that modify frequency modulations

due to large estimation errors. The objectives of this dissertation are to address these problems, design new and intuitive modulation domain filtering methods and applications, and develop the AM-FM transform into a tool that corresponds *very* strongly with human visual perception and is capable of driving future research in the field.

My original contributions are in the field of modulation domain filtering and include the following:

- I have completed advancements in the processes underlying perfect reconstruction from the AM-FM model that have enabled the realization of remarkable, new modulation domain filtering results. These results achieve visually meaningful image processing goals that would be difficult or impossible to obtain using traditional linear or nonlinear filtering methods.

- I have developed a new phase unwrapping algorithm which has eliminated nonintuitive phase scaling factors present within earlier algorithms. The algorithm uses new queue-based phase growing techniques to create a non-scaled, unwrapped phase that provides greater computational accuracy than previous methods.

- I have developed representations and methods for the AM-FM transform that identify well with the simple and intuitive nature of the basic AM-FM model. As the intuition of the model corresponds well to biological visual perception, this allows the design of effective and meaningful modulation domain analysis and filtering methods.

- I have explored new reconstruction techniques that are designed for the general reconstruction of images from modified frequency modulation functions. These techniques employ a novel interpolation scheme that eliminates reconstruction artifacts that arise from discontinuities in phase.

- I have created new classes of modulation domain filters that more effectively modify target image texture by employing new filtering techniques.

As a field with many open questions, modulation domain filtering presents many opportunities for future research. The original contributions detailed in this dissertation constitute a basis for the further development of the AM-FM transform and of modulation domain filters capable of realizing new genres of intuitive, visually motivated image processing tasks.

1.1 Dissertation Organization

Chapter 2 of the dissertation gives a discussion of historical work building up to modulation domain filtering, including development of the AM-FM image model and development of perfect reconstruction from the model. Previous results in modulation domain filtering are discussed.

In Chapter 3, I detail the new phase unwrapping algorithm I have developed that produces stable phase images from which meaningful frequency information can be computed. The phase images are conducive to modulation domain filtering.

In Chapter 4, I detail new reconstruction techniques for the recovery of imagery from modified modulation domain functions, and I present a series of new modulation domain filtering results. These results include demonstrations

of scaling, rotation, and translation using the AM-FM transform and the realization of several image processing tasks that are inspired by biological visual perception.

Conclusions about the content of this dissertation and potential future progress are given in Chapter 5.

1.2 Notation and Nomenclature

I will use boldface lettering such as $\mathbf{x}$ for n-dimensional vectors and blackboard bold lettering for vector spaces. Thus, I will denote the n-dimensional spaces of real-, complex-, and integer-valued vectors as $\mathbb{R}^n$, $\mathbb{C}^n$, and $\mathbb{Z}^n$, respectively. Since demodulation of both real- and complex-valued AM-FM models is presented, I have attempted to consistently represent the continuous real-valued model by $s(t)$ or $s(\mathbf{x})$ for the real-valued signal $s : \mathbb{R}^n \to \mathbb{R}$, and the continuous complex-valued model by $z(t)$ or $z(\mathbf{x})$ for the complex-valued signal $z : \mathbb{R}^n \to \mathbb{C}$. The samples of these functions and other discrete signals will be denoted with the vector-valued integer argument $\mathbf{k} \in \mathbb{Z}^n$. Additionally, although a function is a strict mathematical construct and a signal is a manifestation of some physical quantity, I will often use the terms signal, function, and model interchangeably. I will also use the terms *instantaneous* and *modulations* to refer to the amplitude, phase, and frequency functions of the AM-FM model, as in *instantaneous frequency* and *phase modulations*.

I will consistently express Fourier frequency in units of radians and will denote it by Ω or $\mathbf{\Omega}$. Where possible, I will also use the common convention of denoting signals by lower-case letters and their Fourier transforms by the corresponding upper-case letters. Since both Fourier and instantaneous

frequency quantities are pertinent to AM-FM modeling, I will use the term *stationary* to refer to Fourier frequencies and the term *nonstationary* to refer to instantaneous frequencies.

Chapter 2

Background

2.1 Inspiration for AM-FM Analysis

AM-FM image analysis has its roots in more classical signal analysis techniques. This section presents a discussion of the theoretical developments upon which the AM-FM model was built.

2.1.1 Image Composition and Representation

One of the most important developments in signal analysis has been the application of the Fourier expansion concept to signals. The Fourier transform is used to represent signals as a weighted sum of sinusoidal functions. When applied to images, it provides a measure of the orientation and magnitude of frequency content within an image. Critical to human perception of an image, frequency content refers to the periodicity and direction of repeated lines and image patterns as they vary over space. This can be seen in the Fourier transform of a natural image, such as the famous *Lena* image shown in Fig. 2.1. In this image, the transform indicates strong diagonal content that arises from the mirror borders, hat bands, and feathers in the image.

Though instrumental in the development of modern signal processing techniques, the Fourier transform is limited in its ability to describe signal content in a way that matches human perception. Specifically, the transform only


(a)
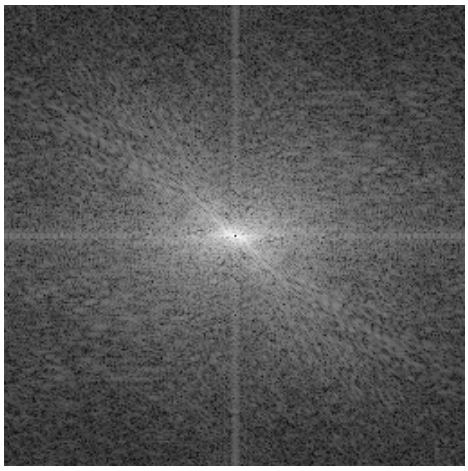
(b)

Figure 2.1: Fourier transform example. (a) Original *Lena* image. (b) Fourier transform of (a).

gives *stationary* information about the signal. Fourier frequency components must exist with a fixed amplitude over *all* time for one-dimensional sound signals or over all space for two-dimensional images. Humans, on the other hand, readily perceive *nonstationary* information from sound or image signals. That is, localized frequency content that is also localized in time or in space.

From the latter half of the 20th century onwards, a great deal of research has been carried out to better match signal analysis to human perception. One of the most important developments in this research was a new representation of signals presented by Gabor in 1946 [27]. Rather than representation by a weighted sum of sinusoidal functions, the Gabor representation consists of a weighted sum of what Gabor referred to as elementary signals. In this new representation, time and frequency are orthogonal coordinates and are furthermore completely symmetrical for all elementary signals. As such, this representation allows for an optimal method of describing signal information that is jointly localized in time and in frequency.

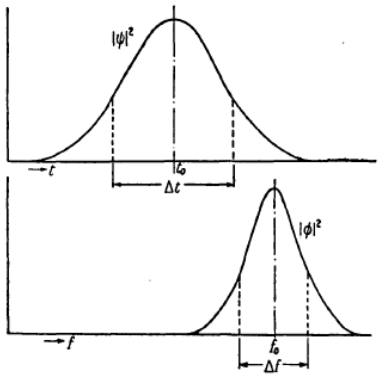

Figure 2.2: From [27], the envelope of the elementary signal in both the time and frequency domains.

The form of the elementary signal is given as [27]

$$\psi(t) = e^{-\alpha^2(t-t_0)^2+j(2\pi f_0 t+\varphi)} \tag{2.1}$$

and its spectrum is [27]

$$\phi(f) = e^{-(\frac{\pi}{\alpha})^2(f-f_0)^2+j(-2\pi t_0(f-f_0)+\varphi)}, \tag{2.2}$$

where α, t_0, f_0, and φ are constants that define the shape of the elementary signal. This shape is illustrated in Fig. 2.2, where Δt and Δf are related to α by

$$\Delta t = \sqrt{\left(\frac{\pi}{2}\right)\frac{1}{\alpha}} \tag{2.3}$$

and

$$\Delta f = \frac{1}{\sqrt{2\pi}}\alpha. \tag{2.4}$$

The principle of localized signal expansion was used in 1977 by Moorer [71] to analyze music. In 1981, Oppenheim discussed the role of phase in signals [86], finding that important signal features tend to only be preserved if phase is maintained, and even that phase information is sufficient to reconstruct signals

to within a scale factor. In 1986, Furukawa [48] presented a Heisenberg-like uncertainty principle applicable to discrete signals.

In [17], Cohen and Lee considered what characteristics would be expected of a 'multicomponent signal', concluding that signals could be locally multicomponent in areas where local bandwidths are well separated. They were motivated by the fact that multicomponent interpretation of signals had often been used for meaningful signal analysis.

Instantaneous Frequency

By 1990, the movement towards measurement of localized frequency information had led to the concept of instantaneous frequency as a description of content. An idealized concept, instantaneous frequency is the frequency content of a signal at a single point in time or space. Cohen and Lee derived expressions for standard deviation of instantaneous frequency in [17] and discussed instantaneous bandwidth as an indication of frequency spread at a point in time. In [18], they furthermore stated that global bandwidth has an amplitude modulation and frequency modulation contribution, where the amplitude modulation contribution is the average of instantaneous bandwidth. Boashash then discussed definitions of and mathematical models for instantaneous frequency in [6]. He noted methods of estimating instantaneous frequency, including differentiation of phase followed by application of smoothing filters.

2.1.2 Analysis Informed by Study of Biological Vision

Use of Gabor filters for image analysis purposes gained support in the 1980's as research found the Gabor representation to be similar to biological vision

systems.

In [53, 54] it was proposed that texture could be considered to be made of fundamental features called "textons", and it was found that only the first-order statistics of these textons have perceptual significance. These textons are consistent with feature detection found in early biological visual processing [8]. Gabor filters have been found to efficiently extract these textons and compute their statistics.

It was observed in [59,69] that the organization of cortical simple cells in the primary visual cortex resembles the profile of Gabor Elementary Functions. A model for simple cells using 2-D extensions of Gabor Elementary Functions was developed in [19]. Quantitative data from recordings from many cells was presented in [50–52], supporting the theory that the behavior of cortical simple cells can be well modeled using Gabor filters.

An image representation using decomposition into two-dimensional Gabor Elementary Functions was developed in [88,89]. Citing the wavelet form of cortical simple cells that is similar to Gabor Elementary Functions, a Gaborian pyramid filterbank was derived for image analysis. A single channel decomposition from this filterbank can be seen in Fig. 2.3, illustrating the type of texture information isolated by a Gabor Elementary Function. The reconstruction of an image from the filterbank can also be seen in [88].

The discovery of a model that closely approximates biological vision systems allowed for new advances in the analysis of image content. New analysis algorithms inspired by the model were presented in [93], and texture discrimination and segmentation methods were developed in [16] and [89]. New work


(a)
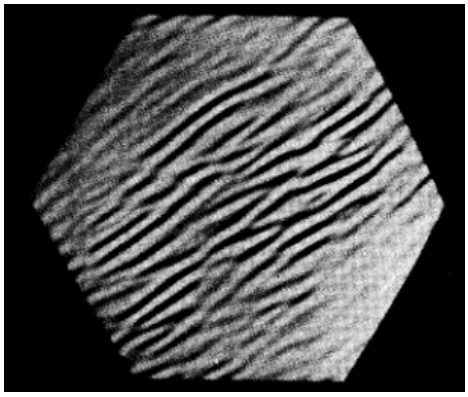
(b)

Figure 2.3: From [89] (a) Input *Lena* image. (b) A Gabor decomposition channel at 30 degrees.

on detecting textons and texture boundaries was presented in [120]. The model also profoundly influenced the development of the AM-FM model described in upcoming sections.

2.1.3 Signal Reconstruction

In addition to the analysis of image content using the signal representations discussed above, efforts were made to reconstruct speech signals from the representations. Quatieri performed reconstruction of speech using only magnitudes of component sine waves in [70]. In [86], Oppenheim illustrated the intelligibility of phase reconstruction for images and speech. In [88], Porat and Zeevi gave a comparison of reconstructed signals from Gabor decompositions, an example of which can be seen in Fig. 2.4.

2.2 AM-FM Demodulation and Analysis

An important advancement in the ability to mathematically describe and analyze signal content that relates to human perception was the development of



Figure 2.4: From [88] (a) Input *Lena* image. (b) Reconstruction from 12626 5x5 Gabor elementary functions. (c) Reconstruction from 4062 GEF's.

the AM-FM model as a signal representation.

The AM-FM representation, first described in 1990 in [8], is given for a real-valued signal $s : \mathbb{R}^n \to \mathbb{R}$, $\mathbf{x} \in \mathbb{R}^n$ as

$$s(\mathbf{x}) = a(\mathbf{x}) \cos(\varphi(\mathbf{x})), \tag{2.5}$$

where $a : \mathbb{R}^n \to [0, \infty)$ is the amplitude modulation of the signal and $\varphi : \mathbb{R}^n \to \mathbb{R}$ is the phase modulation of the signal. Within this model, the values of the amplitude modulation can be thought of as local contrast, and the gradient of the phase $\nabla\varphi(\mathbf{x})$ gives the instantaneous frequency of the signal which describes the local magnitude and orientation of repeated lines and patterns in the image.

For a complex-valued signal $z : \mathbb{R}^n \to \mathbb{C}$, the AM-FM signal model is given by

$$z(\mathbf{x}) = a(\mathbf{x}) e^{j\varphi(\mathbf{x})}. \tag{2.6}$$

In [8], Bovik *et al.* used the model to differentiate image texture regions, which were first decomposed into frequency and orientation channels using two-dimensional Gabor functions. While application of the AM-FM model to a general image can result in highly variable modulation functions, the preliminary Gabor decomposition allowed slowly varying amplitude and phase channel envelopes for analysis. Examples were given of texture segmentation via comparison of AM functions and via comparison of FM functions to detect textures of differing phase.

Justification for the use of Gabor filter decomposition in AM-FM models was given in [7], which stated that the trade-off between spectral selectivity and accuracy in boundary localization is optimized by Gabor filters. The first

discussion of a multicomponent version of the AM-FM model based upon a Gabor filterbank was presented in [34]. This led to the formalization of the complex-valued multicomponent AM-FM image model

$$z(\mathbf{x}) = \sum_{i=1}^{K} a_i(\mathbf{x}) e^{j\varphi_i(\mathbf{x})}. \tag{2.7}$$

In this model, the image is represented as a sum of K locally coherent AM-FM components, where a_i and φ_i are the amplitude and frequency modulations of the ith component.

An approach measuring emergent image frequencies was then given in [9], where the emergent frequencies are taken across channels at all points in an image. It was analytically shown that Gabor channel filters maximally localize the potential emergent frequency solutions in space and frequency, and an iterative algorithm was given to compute frequencies and amplitudes. Local spatio-frequency moments were used in [105] to compute the shape of curved surfaces from texture.

Lu and Doerschuk presented in [61] a Model-Based Demodulation Algorithm (MBDA) which considers sums of AM-FM models in the presence of noise to statistically determine amplitude and frequency modulations. They stated that the parameters of the algorithm correspond to decomposition by Gabor filters and they applied the MBDA to speech processing examples.

These first approaches to the AM-FM model presented results that highly correspond to human visual perception, as demonstrated in Fig. 2.5.

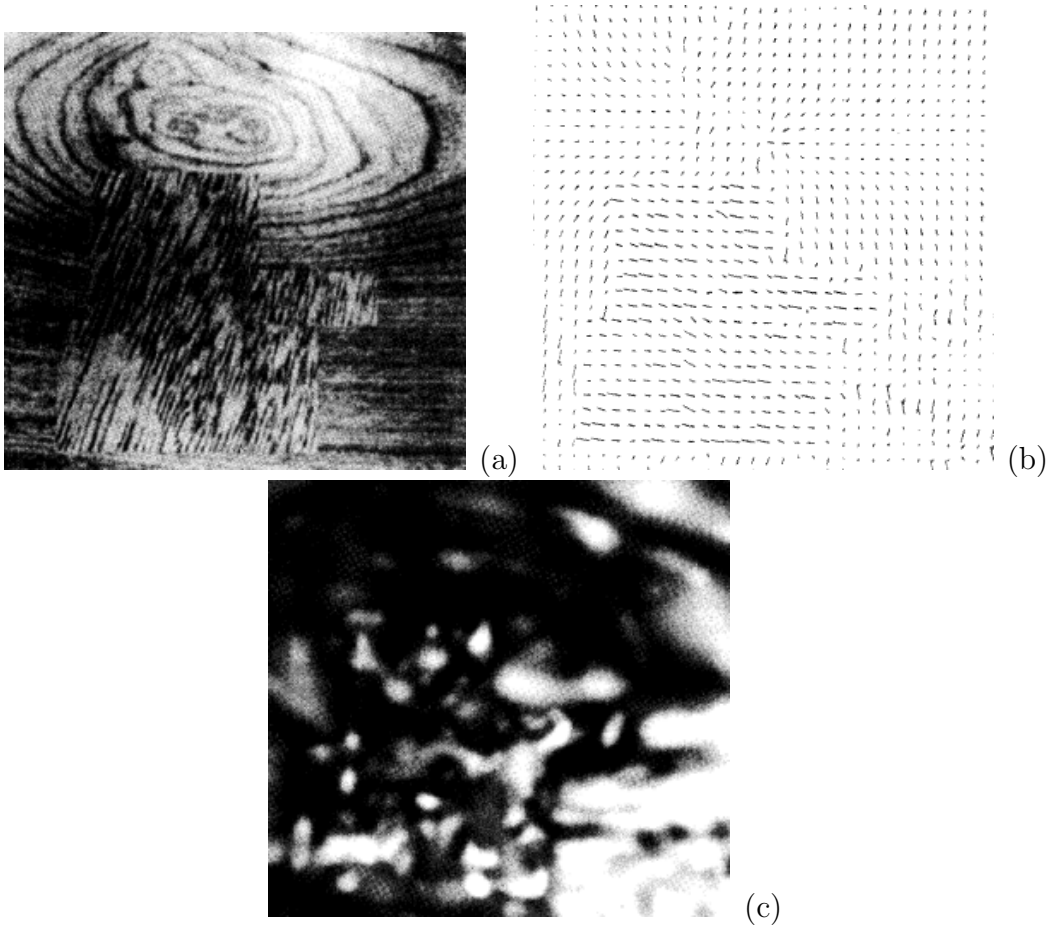

Figure 2.5: From [9] (a) Input image of juxtaposed textures. (b) Needle diagram showing computed emergent image frequencies. (c) Computed amplitude modulation.

2.2.1 Gabor Filterbank Decomposition

The filterbanks used for decomposition of images into narrowband channels of smoothly varying local amplitude and frequency modulation functions are composed of a designed array of Gabor filters. In the space domain, the Gabor filter impulse response is given by [40]

$$g_m(\mathbf{x}) = \frac{1}{\sigma_m\sqrt{2\pi}} \exp[-\frac{1}{4\sigma_m^2}\mathbf{x}^T\mathbf{x}] \exp[j2\pi\mathbf{\Omega}_m^T\mathbf{x}], \tag{2.8}$$

where the radial center frequency is $|\mathbf{\Omega}_m|$ and the orientation is $\arg[\mathbf{\Omega}_m]$.

In the frequency domain, the Gabor filter is given by [35]

$$G_m(\mathbf{\Omega}) = \sqrt[4]{8\pi\sigma_m^2} \exp[-4\pi^2\sigma_m^2(\mathbf{\Omega} - \mathbf{\Omega}_m)^2]. \tag{2.9}$$

The shape of each filter is determined by the parameter η, which is the desired fraction of the filter peak at which the filters in the filterbank will intersect, and by the parameter

$$\sigma_m = \frac{\sqrt{-\ln\eta}}{2\pi\gamma\mathbf{\Omega}_m} \tag{2.10}$$

where

$$\gamma = \frac{(2^B - 1)^2}{(2^B + 1)^2} \tag{2.11}$$

and B is the bandwidth of the filter in octaves.

The positioning of the filters in the Gabor filterbank is determined by the ratio R between filter center frequencies along each ray of the filterbank and by the angular spacing A between adjacent rays. For a chosen ratio R, the angular spacing is given by [40]

$$A = 2\arcsin\left[(4R)^{-\frac{1}{2}}\{(R^2+1)(\gamma-1) + 2R(\gamma+1)\}^{\frac{1}{2}}\right], \tag{2.12}$$

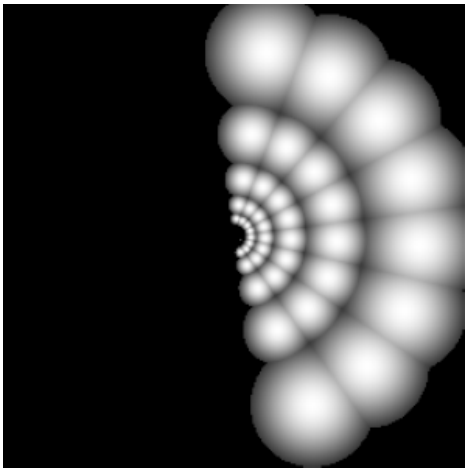

Figure 2.6: From [30], 40 channel Gabor filterbank used for decomposition.

where γ is computed from (2.11) for a chosen bandwith B.

In much of the work discussed here, parameters to the Gabor filterbank were taken as or similar to $\eta = \frac{1}{2}$ and $B = 1$ [40], so that each octave filter would intersect with its neighbors at half-peak amplitude. Choosing then $R = 1.8$ would produce a filterbank such as the 40 channel filterbank shown in Fig. 2.6.

2.2.2 Demodulation Via Energy Separation

Several approaches to computing or estimating the amplitude modulation $a(\mathbf{x})$ and the frequency modulation $\nabla\varphi(\mathbf{x})$ from (2.5), a process known as *demodulation*, have been proposed. Among the most researched approaches are demodulation using the Teager-Kaiser energy separation and demodulation based upon the complex analytic image.

Kaiser presented an algorithm that can compute the 'energy' required to generate a signal f in [55], and he discussed its application to speech processing.

Called the Teager-Kaiser Energy Operator, its 1-D discrete equation is [55, 63]

$$\Psi_d[f(k)] = f^2(k) - f(k-1)f(k+1) \tag{2.13}$$

and its 1-D continuous form is [56, 63]

$$\Psi[f(t)] = [f'(t)]^2 - f(t)f''(t). \tag{2.14}$$

The energy operator was applied to the estimation of amplitude and frequency modulation functions in [56,68]. It was shown that the operator gives approximate equality to the product of the squared AM and FM functions, as in

$$\Psi\{a(t)\cos[\phi(t)]\} \approx a^2(t)\omega^2(t) \tag{2.15}$$

where $\omega(t) = \frac{d\phi(t)}{dt}$, and the operator was used to track the AM envelope of a signal, the instantaneous frequency, and product of the two. Then, in [65–67], the energy operator was used to separate the energy product into AM and FM components using combinations of instantaneous outputs from the operator applied to the signal and its one-step shifts. This resulted in an approximation called the *Energy Separation Algorithm* (ESA). The continuous 1-D Energy Separation Algorithm (CESA) [65] is given by

$$|a| \approx \frac{\Psi[f(t)]}{[\Psi(\frac{df}{dt})]^{\frac{1}{2}}}, \tag{2.16}$$

$$|\omega| \approx \left[\frac{\Psi(\frac{df}{dt})}{\Psi[f(t)]}\right]^{\frac{1}{2}} \tag{2.17}$$

and the discrete 1-D Energy Separation Algorithm (DESA) [65] is given by

$$\left|a(k)\right| \approx \frac{2\Psi_d[f(k)]}{\{\Psi_d[f(k+1)-f(k-1)]\}^{\frac{1}{2}}}, \tag{2.18}$$

$$\Omega(k) \approx \arcsin\left(\left\{\frac{\Psi_d[f(k+1)-f(k-1)]}{4\Psi_d[f(k)]}\right\}^{\frac{1}{2}}\right). \tag{2.19}$$

Bounds for the approximation error were derived in [67].

In [12], the performance of the ESA was shown to improve when the signal is first separated into channels and a multidimensional energy operator and multidimensional energy separation algorithm to estimate the AM and FM functions of images was presented in [63, 64]. The continuous 2-D energy operator given was

$$\Phi(f) = \left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2 - f\left(\frac{\partial^2 f}{\partial x} + \frac{\partial^2 f}{\partial y}\right). \tag{2.20}$$

This operator was used to produce the continuous 2-D Energy Separation Algorithm

$$|a| \approx \frac{\Phi(f)}{[\Phi(\frac{\partial f}{\partial x}) + \Phi(\frac{\partial f}{\partial y})]^{\frac{1}{2}}}, \tag{2.21}$$

$$|\omega_x| \approx \left[\frac{\Phi(\frac{\partial f}{\partial x})}{\Phi(f)}\right]^{\frac{1}{2}}, \tag{2.22}$$

$$\left|\omega_y\right| \approx \left[\frac{\Phi(\frac{\partial f}{\partial y})}{\Phi(f)}\right]^{\frac{1}{2}}. \tag{2.23}$$

The discrete version of the 2-D energy operator was also presented in [63] as

$$\Phi_d[f(m,n)] = 2f^2(m,n) - f(m-1,n)f(m+1,n) - f(m,n-1)f(m,n+1) \tag{2.24}$$

along with the resulting discrete 2D Energy Separation Algorithm

$$\left|a(m,n)\right| \approx \frac{2\Psi_d[f(m,n)]}{\{\Psi_d[f(m+1,n)-f(m-1,n)]+\Psi_d[f(m,n+1)-f(m,n-1)]\}^{\frac{1}{2}}}, \tag{2.25}$$

$$\left|\Omega_1(m,n)\right| \approx \arcsin\left(\left\{\frac{\Psi_d[f(m+1,n)-f(m-1,n)]}{4\Psi_d[f(m,n)]}\right\}^{\frac{1}{2}}\right), \tag{2.26}$$

$$\left|\Omega_2(m,n)\right| \approx \arcsin\left(\left\{\frac{\Psi_d[f(m,n+1)-f(m,n-1)]}{4\Psi_d[f(m,n)]}\right\}^{\frac{1}{2}}\right). \tag{2.27}$$

In [63], the multidimensional DESA algorithms were applied to image channels obtained via Gabor decomposition.

Application To Speech and Image Analysis

The Energy Separation Algorithm was applied to speech modeling in [12, 65, 66], with bandpass filtered speech shown to give corresponding AM and FM functions in [66], and with the modeling of noisy signals considered in [12]. An example of ESA demodulation of a speech signal can be seen in Fig. 2.7. The multidimensional DESA was applied to image texures in [63], with an example of ESA demodulation of a texture image shown in Fig. 2.8.

2.2.3 Analytic Image Based Demodulation

Another set of highly researched demodulation techniques is based on computation of a complex *analytic image* that is associated with a real image. These

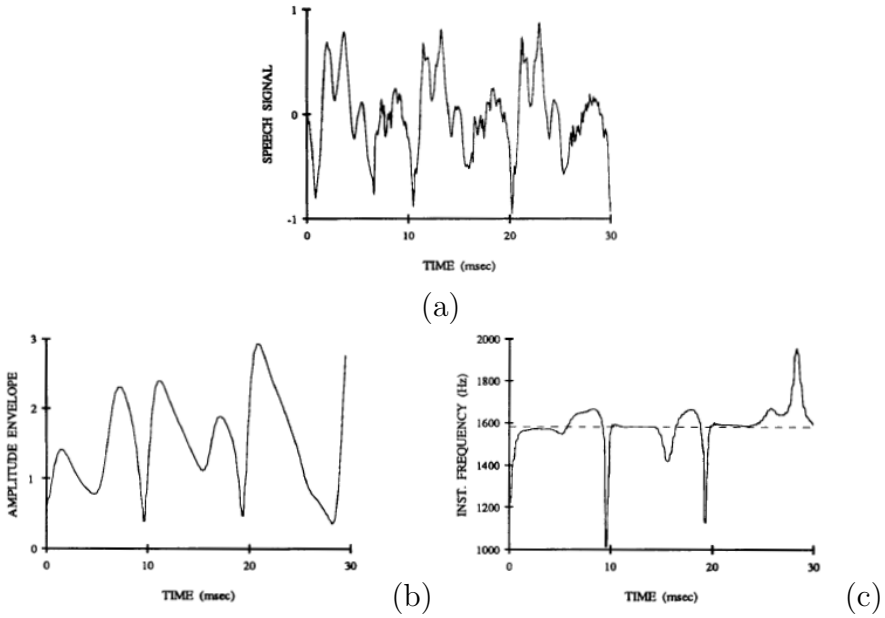

Figure 2.7: From [66], DESA applied to speech. (a) Signal s(n) from speech vowel /E/. (b) Estimated amplitude envelope using DESA, from (a) filtered by Gabor filter centered around 1580 Hz. (c) Estimated instantaneous frequency, smoothed by an 11-point median filter. Dotted line is Gabor filter center.

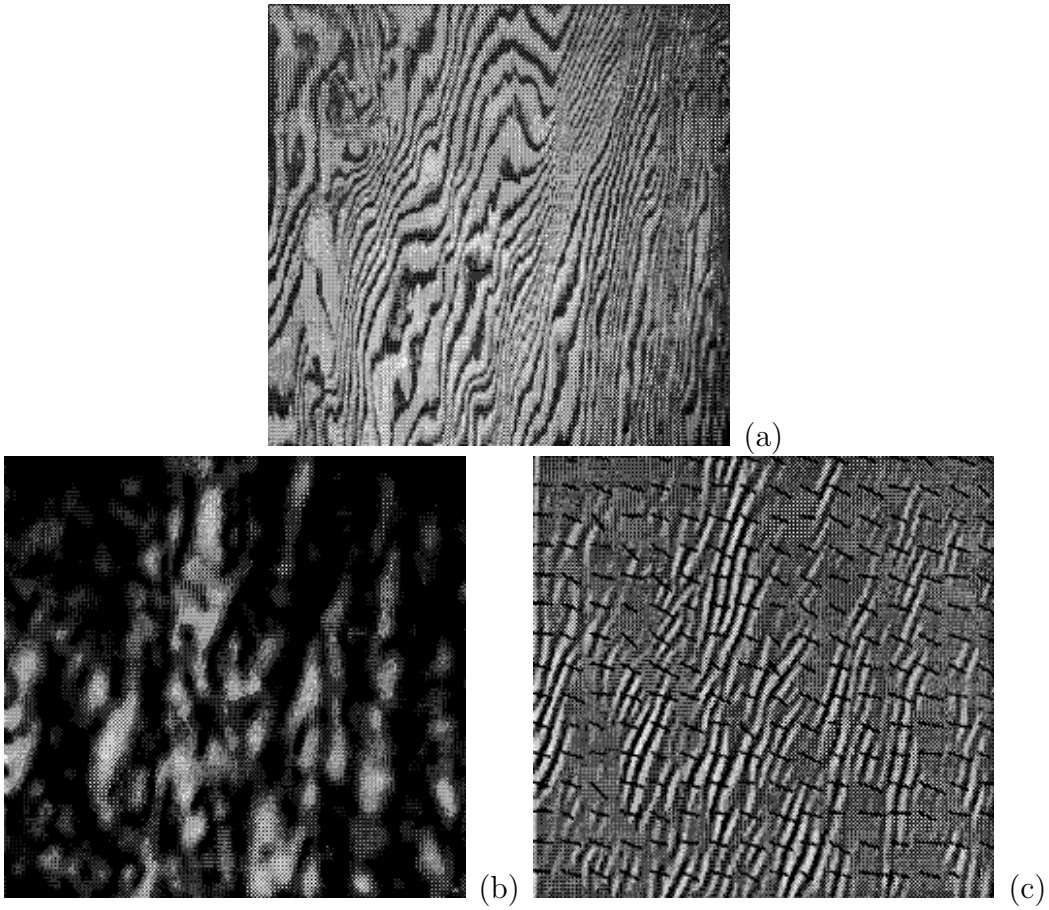

Figure 2.8: From [63], DESA applied to a texture image. (a) Original *wood* texture image. (b) Estimated amplitude envelope using DESA, from (a) filtered by Gabor filter centered around horizontal and vertical frequencies of 35.5 and 14.7 cycles per image. (c) Estimated instantaneous frequency vectors, superimposed upon the bandpass image of (a).

techniques have found wide usage in analysis and applications involving the AM-FM model.

Demodulation based upon the concept of the analytic image was presented in [8]. Given a complex image $z(\mathbf{x})$, computation of the amplitude and of the phase in (2.6) is given by [8, 29, 34, 40]

$$a(\mathbf{x}) = |z(\mathbf{x})| \tag{2.28}$$

for the amplitude modulation, and by [8, 97]

$$\varphi(\mathbf{x}) = \arctan\left(\frac{\mathrm{Im}\left[z(\mathbf{x})\right]}{\mathrm{Re}\left[z(\mathbf{x})\right]}\right) \tag{2.29}$$

for the phase modulation.

The frequency modulation $\nabla\varphi(\mathbf{x})$ can then be taken as the gradient of the phase modulation in (2.29), however since the arctangent function in that equation has a range of $(-\pi, \pi]$, phase discontinuities are introduced that have no relation to visual image content. As such, alternative methods for the estimating frequency content have been developed, such as the frequency demodulation developed in [34] using a quasi-eigenfunction approximation, given as

$$\nabla\varphi(\mathbf{x}) = \mathrm{Re}\left[\frac{\nabla z(\mathbf{x})}{jz(\mathbf{x})}\right]. \tag{2.30}$$

The Analytic Image

Introduced by Gabor in [27], the analytic representation of a real signal, called the *analytic signal* removes redundant frequency information from the real-valued signal and facilitates the demodulation of the signal. Gabor gave the analytic signal $z(t)$ corresponding to a one dimensional continuous real signal

$s(t)$ as

$$z(t) = s(t) + jq(t), \tag{2.31}$$

where $q(t)$ is the signal in quadrature to $s(t)$. He then noted that the analytic signal could be computed effectively in the frequency domain, as the Fourier transform of $z(t)$, $Z(\Omega) = \mathcal{F}[z(t)]$, can be found directly from the Fourier transform of $s(t)$, $S(\Omega) = \mathcal{F}[s(t)]$, according to

$$Z(\Omega) = \begin{cases} 2S(\Omega), & \Omega > 0 \\ 0, & \text{otherwise.} \end{cases} \tag{2.32}$$

The quadrature signal $q(t)$ is related to $s(t)$ in that it is the Hilbert transform of $s(t)$, defined by

$$\mathcal{H}\left[s(t)\right] = s(t) * \frac{1}{\pi t} = \frac{1}{\pi} \int_{\mathbb{R}} \frac{s(\tau)}{t - \tau} d\tau, \tag{2.33}$$

which was noted in [29] to also be readily computed in the frequency domain from

$$\mathcal{F}\Big\{\mathcal{H}\left[s(t)\right]\Big\} = -j \operatorname{sgn}(\Omega) S(\Omega), \tag{2.34}$$

where $S(\Omega) = \mathcal{F}\Big\{s(t)\Big\}$. This leads to a direct computation of the analytic signal in the frequency domain similar to that given by Gabor:

$$Z(\Omega) = \begin{cases} 2S(\Omega), & \Omega > 0 \\ S(\Omega), & \Omega = 0 \\ 0, & \Omega < 0, \end{cases} \tag{2.35}$$

The definition of the analytic signal was used in [8] to accomplish demodulation by first decomposing the image using Gabor filters, and then noting that the real and imaginary parts of the filters closely approximate a Hilbert transform pair. This allowed the Gabor filter to be interpreted as the analytic signal z of the Gabor channel, from which the amplitude modulation function

could be taken as that from (2.28), and the frequency modulation functions could be computed from

$$f_x(x,y) = \frac{\mathrm{Re}\left[z(x,y)\right] \cdot D_x \mathrm{Im}\left[z(x,y)\right] - \mathrm{Im}\left[z(x,y)\right] \cdot D_x \mathrm{Re}\left[z(x,y)\right]}{a^2(x,y)} \quad (2.36)$$

and

$$f_y(x,y) = \frac{\mathrm{Re}\left[z(x,y)\right] \cdot D_y \mathrm{Im}\left[z(x,y)\right] - \mathrm{Im}\left[z(x,y)\right] \cdot D_y \mathrm{Re}\left[z(x,y)\right]}{a^2(x,y)}, \quad (2.37)$$

where $D_x = \frac{\partial}{\partial x}$ and $D_y = \frac{\partial}{\partial y}$.

A more direct computation was described in [32] by presenting the two dimensional analytic image associated with a real image. The analytic image is based on a multidimensional directional Hilbert transform, an adjusted, continuous version of which is given in [29, 44] as

$$\mathcal{F}\Big\{\mathcal{H}_i[s(\mathbf{x})]\Big\} = -j\,\mathrm{sgn}_{adj}(\mathbf{\Omega}^T\mathbf{e}_i)S(\mathbf{\Omega}), \quad (2.38)$$

where $\mathbf{e}_i$ is the unit vector in the x_i direction, $\mathbf{\Omega} = [\Omega_1\ \Omega_2\ \ldots\ \Omega_n]^T$ is the argument of the n-dimensional Fourier transform, and

$$\mathrm{sgn}_{adj}\,\mathbf{\Omega} = \sum_{i=1}^{n} \mathrm{sgn}_*\,\Omega_i \prod_{k=1}^{i-1}\Big(1 - |\mathrm{sgn}_*\,\Omega_k|\Big), \quad (2.39)$$

where

$$\mathrm{sgn}_*\,x = \begin{cases} 1, & x > 0 \\ 0, & x = 0 \\ -1, & x < 0, \end{cases} \quad (2.40)$$

or, as noted in [97], where

$$\mathrm{sgn}_{adj}\,\mathbf{\Omega} = \begin{cases} \mathrm{sgn}_*\,\Omega_\nu, & \mathbf{\Omega} \neq \mathbf{0} \\ 0, & \mathbf{\Omega} = \mathbf{0}, \end{cases} \quad (2.41)$$

where $\nu = \min\big\{i \in [1,n] : \Omega_i \neq 0\big\}$.

The discrete version of the multidimensional, adjusted, directional Hilbert transform is given in [29, 97] as

$$\mathcal{F}\Big\{\mathcal{H}_i[s(\mathbf{k})]\Big\} = -j\,\mathrm{sgn}_{per}(\mathbf{\Omega}^T e_i)S(\mathbf{\Omega}), \tag{2.42}$$

where $k \in \mathbb{Z}^2$, and $\mathrm{sgn}_{per}(\mathbf{\Omega})$ is a periodic sign function

$$\mathrm{sgn}_{per}\,\mathbf{\Omega} = \sum_{i=1}^{n} \mathrm{sgn}_{\bigstar}\,\Omega_i \prod_{k=1}^{i-1}\Big(1 - \Big|\mathrm{sgn}_{\bigstar}\,\Omega_k\Big|\Big), \tag{2.43}$$

where

$$\mathrm{sgn}_{\bigstar}\,x = \begin{cases} 1, & 0 < W\{x\} < \pi \\ 0, & W\{x\} \in \{0, \pi\} \\ -1, & -\pi < W\{x\} < 0, \end{cases} \tag{2.44}$$

and where $W\{\cdot\}$ is a wrapping operator given by

$$W\{\omega\} = \begin{cases} \pi, & \omega \in \{x \in \mathbb{R} : x = \pi(2m+1),\ \forall m \in \mathbb{Z}\} \\ \omega \bmod \pi, & \text{otherwise.} \end{cases} \tag{2.45}$$

The discrete adjusted Hilbert transform is then used to obtain a discrete analytic image $z : \mathbb{Z}^2 \to \mathbb{C}$ according to

$$z(\mathbf{k}) = s(\mathbf{k}) + jq(\mathbf{k}), \tag{2.46}$$

where $q(\mathbf{k})$ is obtained by applying the discrete partial Hilbert transform to $s(\mathbf{k})$. For an for $N \times L$ image, direct calculation of the spectrum of the analytic image from the spectrum of the input image is then given by [29, 97]

$$Z(u,v) = \begin{cases} 2S(u,v), & u = 1, 2, \ldots, \frac{N}{2} - 1, \\ 0, & u = \frac{N}{2} + 1, \frac{N}{2} + 2, \ldots, N - 1, \\ 2S(u,v), & u = 0,\ v = 1, 2, \ldots, \frac{L}{2} - 1, \\ 2S(u,v), & u = \frac{N}{2},\ v = 1, 2, \ldots, \frac{L}{2} - 1, \\ 0, & u = 0,\ v = \frac{L}{2} + 1, \frac{L}{2} + 2, \ldots, L - 1, \\ 0, & u = \frac{N}{2},\ v = \frac{L}{2} + 1, \frac{L}{2} + 2, \ldots, L - 1, \\ S(u,v), & \text{otherwise.} \end{cases} \tag{2.47}$$

The analytic image was used in [32] to derive corresponding demodulation algorithms that operate on the complex image. Once a frequency modulation function had been computed, the new AM equation was given as

$$a(\mathbf{x}) \approx \left| \frac{z_i(\mathbf{x})}{G_i[\nabla\varphi(\mathbf{x})]} \right|, \tag{2.48}$$

where G_i is the Gabor function for the ith component in (2.7).

Further consideration of the directional Hilbert transform was also presented in [45]. There, it was stated that the directional Hilbert transform acting in horizontal and vertical directions gives different AM-FM models. Relationships between the models were investigated and it was shown that they are 'essentially equivalent'.

The Quasi-EigenFunction Approximation

Analytic image based demodulation algorithms were developed through the use of a quasi-eigenfunction approximation, or QEA, that was introduced in [34]. The approximation, which is a 'quasi-extension of the eigenfunction concept of linear system theory', posits that the response of a linear system g with Fourier transform G given input signals of the form $Ae^{j\varphi}$ can be approximated as

$$g\{Ae^{j\varphi}\} = G(\nabla\varphi)Ae^{j\varphi}. \tag{2.49}$$

This is similar to the concept of the eigenfunction/eigenvalue pair in linear system theory, where $G(\nabla\varphi)$ is the eigenvalue associated with the eigenfunction $Ae^{j\varphi}$ for the system g, and this similarity was shown in [29,33,34] to be close to equality for signals $Ae^{j\varphi}$ with smoothly varying amplitude and frequency modulations. The quasi-eigenfunction approximation was then used [11, 29,

33, 34] to develop the frequency demodulation estimate in (2.30) and to give bounds for the estimation error. The QEA was also applied to the discrete Teager-Kaiser operator in [10, 11] to derive bounds for ESA demodulation approximations.

Bounds on approximation errors for a multidimensional QEA were derived in [29, 32, 42], and the variation of approximation error bounds was discussed in [33]. The multi-dimensional, discrete QEA was used to derive new discrete frequency demodulation algorithms [29, 38, 42]

$$\mathbf{e}_i^T \nabla\varphi(\mathbf{k}) \approx \mathbf{e}_i^T \nabla\widehat{\varphi}(\mathbf{k}) = \arcsin\left[\frac{z(\mathbf{k}+\mathbf{e}_i) + z(\mathbf{k}-\mathbf{e}_i)}{2z(\mathbf{x})}\right], \tag{2.50}$$

known as the multidimensional sine algorithm, where $\mathbf{e}_i$ is the unit vector in the k_i direction, and

$$\mathbf{e}_i^T \nabla\varphi(\mathbf{k}) \approx \mathbf{e}_i^T \nabla\widehat{\varphi}(\mathbf{k}) = \arccos\left[\frac{z(\mathbf{k}+\mathbf{e}_i) + z(\mathbf{k}-\mathbf{e}_i)}{2z(\mathbf{x})}\right], \tag{2.51}$$

known as the multidimensional cosine algorithm. It was noted that either of these algorithms could be used independently to estimate the frequency modulation function to within π radians, or that they could be used together to estimate the frequency modulation function to within 2π radians.

Comparison Against ESA Based Analysis

Potamiamos and Maragos performed a comparison between Energy Separation Algorithm and Analytic Signal based demodulation in [90], and found the results of the two approaches to be similar, with those from the ESA to be more suitable in some cases.

Vakman performed a separate comparison in [118], and concluded that only the analytic signal based approach meets the following physical conditions:

(A) Amplitude continuity and differentiability: for a small, additive change in the real signal $s(t)$, the quadrature signal $q(t)$ and, thus, the complex signal amplitude $\vartheta(t)$ should also only change by a small amount.

(B) Phase independence of scaling and homogeneity: the phase $\psi(t)$ associated with the real signal $s(t)$ should be the same as that associated with the scaled real signal $\alpha s(t)$, $\alpha \in \mathbb{R}^+$.

(C) Harmonic correspondence: when the real signal $s(t)$ has a constant amplitude and a linear phase, the complex signal should preserve these quantities as seen in the Euler identity

$$z(t) = Ae^{j(\omega t+\omega_0)} = A\cos(\omega t+\omega_0)+jA\sin(\omega t+\omega_0) = s(t)+jq(t). \quad (2.52)$$

He argued that the analytic signal based approach also has an advantage in accuracy and simplicity. While research has continued using extensions to the Teager-Kaiser energy operator [22, 60], the remainder of this dissertation will consider work done using analytic signal based demodulation as a foundation.

2.2.4 Component Analysis Paradigms

The analytic signal based demodulation techniques described above have been used to formulate several different approaches to texture analysis, which are briefly summarized in this section.

Tracked Components Analysis

Tracked Multicomponent Analysis was presented and described in [31, 37, 40]. This approach to analysis considers an image texture to be composed of distinct components which are tracked across channels using a Bayesian estimation

scheme. The components tracking is carried out with one-dimensional Kalman filters and the components are then isolated for analysis.

A demonstration of Tracked Components Analysis was given in [41]. 1-D signal reconstructions from computed components were shown for the first time in [35]. The first image reconstruction results were given in [36, 37]. Notable reconstruction results were also given in [38, 40]. Reconstructions including postfilter smoothing processes were introduced in [39]. Tracked Components Analysis provided reconstructions from relatively few isolated AM-FM components, as seen in Fig. 2.9. Tracked Components Analysis is significant because is was the first approach to successfully reconstruct a 2D image from computed modulations, albeit with low fidelity.

Dominant Component Analysis

Introduced in [46] and described in [42], Dominant Component Analysis (DCA) computes a single representative component for an image texture which is then analyzed. This more compact representation still leads to effective analysis and reconstruction results, such as those given in [38, 42]. As seen in Fig. 2.10, a single computed dominant component is capable of retaining significant image structure.

Channelized Components Analysis

Channelized Components Analysis (CCA) computes and analyzes one AM-FM image component per filterbank channel. As described in [39, 42, 43], this approach leads to many more components than the other analysis approaches. However, these components may be used for higher-quality analysis and recon-

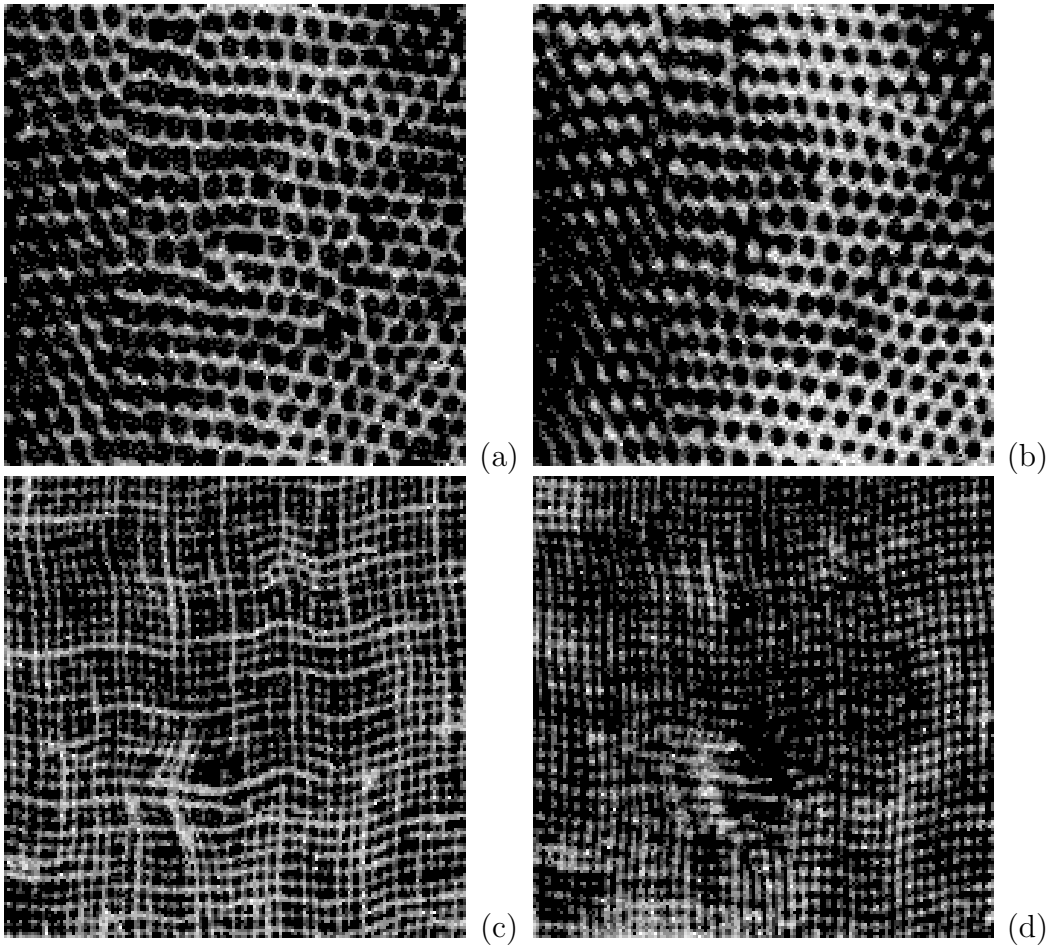

Figure 2.9: From [39], reconstruction from Tracked Components Analysis. (a) Original *Reptile* image. (b) Reconstruction of (a) from six tracked AM-FM components. (c) Original *Raffia* image. (d) Reconstruction of (c) from nine tracked AM-FM components.


(a)
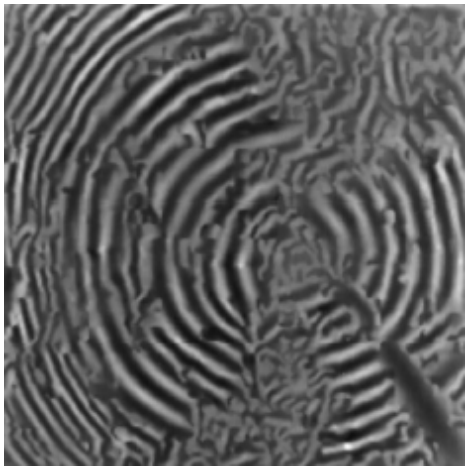
(b)

Figure 2.10: From [42], reconstruction from Dominant Component Analysis. (a) Original *Tree* image. (b) Reconstruction of (a) from a single dominant AM-FM component.

struction results, such as those seen in [39, 42]. Examples of reconstructions from Channelized Components Analysis are shown in Fig. 2.11, where it can be seen that the reconstructions capture a large amount of image structure, though they are not exact. While CCA represented a significant leap forward in terms of achieving high fidelity AM-FM image reconstructions, it cannot provide perfect reconstruction due to the inherently non-orthogonal Gabor filters used for the CCA analysis filterbank.

2.2.5 Practical Analysis Applications

The work described above to develop the AM-FM model has found many uses in varied applications. The model was applied to the computation of flow lines in [46]. In [92], Dominant Component Analysis was used to inform the computation of snakes, which are 1D structures that can track visually meaningful image structures by encircling them. These snakes, or active contours, were used to perform image segmentation.

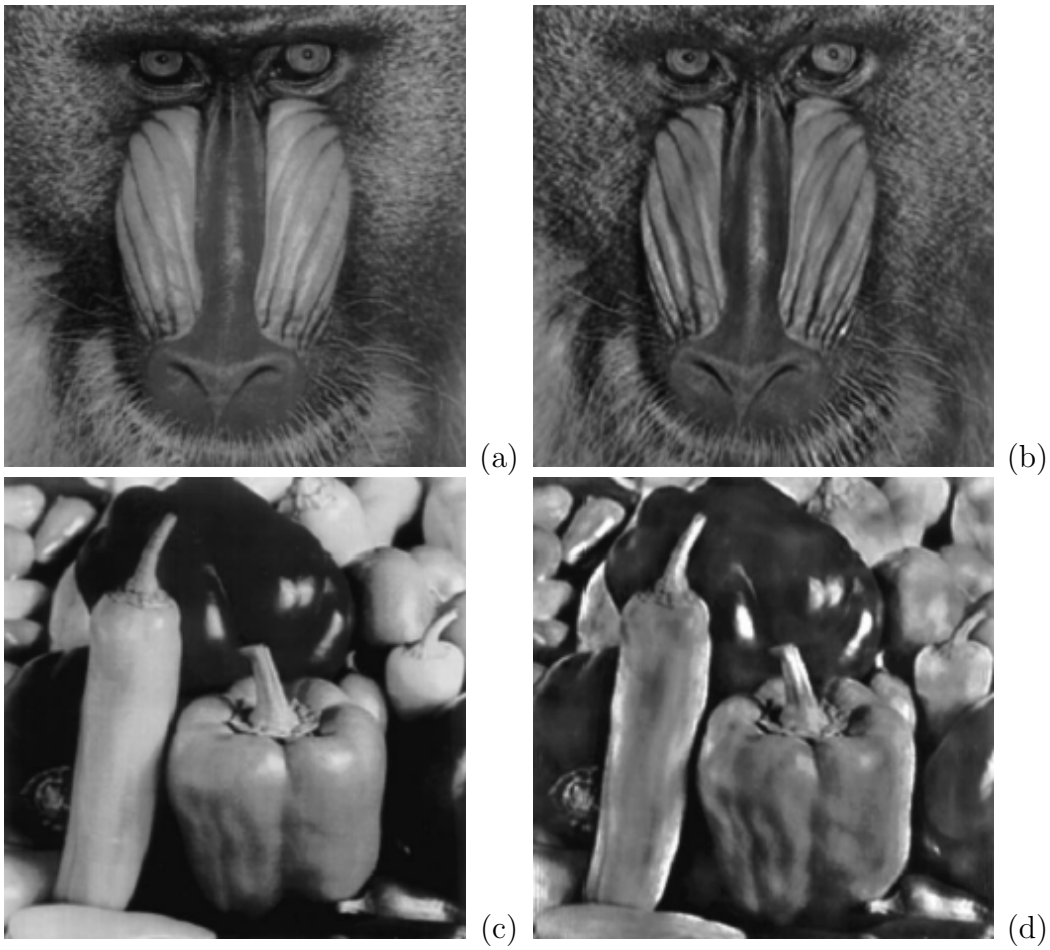

Figure 2.11: From [42], reconstruction from Channelized Components Analysis. (a) Original *Mandrill* image. (b) Reconstruction of (a) from a 43 channelized AM-FM components. (c) Original *Peppers* image. (d) Channelized Components Analysis Reconstruction of (c).

Developmental work on the AM-FM model also helped to expand the range of the practical application of the model. General n-dimensional analysis for the model was presented in [42]. A definition for multidimensional instantaneous bandwidth was given in [41], and histograms were employed to test the validity of the multicomponent interpretation. In [43], discontinuities in the empirically computed image phase were shown to lead to 'wideband frequency excursions' that degrade model quality. Postfilters were presented in [43] to ameliorate the excursions.

A new reconstruction algorithm for analysis-only AM-FM models was presented in [106] that eliminated blocking artifacts found in earlier reconstruction work, particularly along borders between the spatial support of components extracted by Tracked Components Analysis and between phase reconstruction blocks in the CCA approach (where non-perfect phase reconstruction is typically carried out independently in fixed size image blocks using independent phase initial values for each block). The new algorithm proposed in [106] operated by making use of phase reconstruction boundary conditions in a multipath interpolative scheme.

Image Segmentation

A major area of application of the AM-FM model has been image segmentation. Promising texture segmentation results using dominant modulations were shown in [30] and results for both Dominant Component Analysis based segmentation and Channelized Components Analysis based segmentation were given in [42].

A statistical clustering algorithm based on a similarity measure was

combined with Dominant Component Analysis in [108], with results refined by morphological filtering and connected components labeling for postprocessing. This algorithm was expanded to be partially unsupervised in [109], requiring only that the number of regions be specified *a priori*. In [123], the maximum number of textured regions was estimated via modulation domain density clustering to achieve fully unsupervised modulation domain image segmentation. Furthermore, horizontal and vertical spatial information was also considered as a part of the segmentation algorithm.

Fully unsupervised image Bayesian segmentation based on Dominant Component Analysis was presented in [122], employing a feature space modeled as mixture of Gaussians. An expectation-maximization algorithm was used to estimate the parameters of the mixture and estimates of pixel class labels were then obtained.

Other Applications

The AM-FM model was used to synthesize missing texture in occluded regions in [1], where texture inpainting was achieved by a reaction-diffusion process. In this approach, diffusion is analogous to a smoothing of the image and reaction then synthesizes the missing texture. The AM-FM model allows localization of the reaction filters and the new method was shown to result in seamless transitions with the original texture boundaries as seen in Fig. 2.12.

AM-FM analysis was used to create an effective content-based image retrieval system from libraries of images in [47]. The performance of this system was shown to have an advantage against other Gabor magnitude response methods.

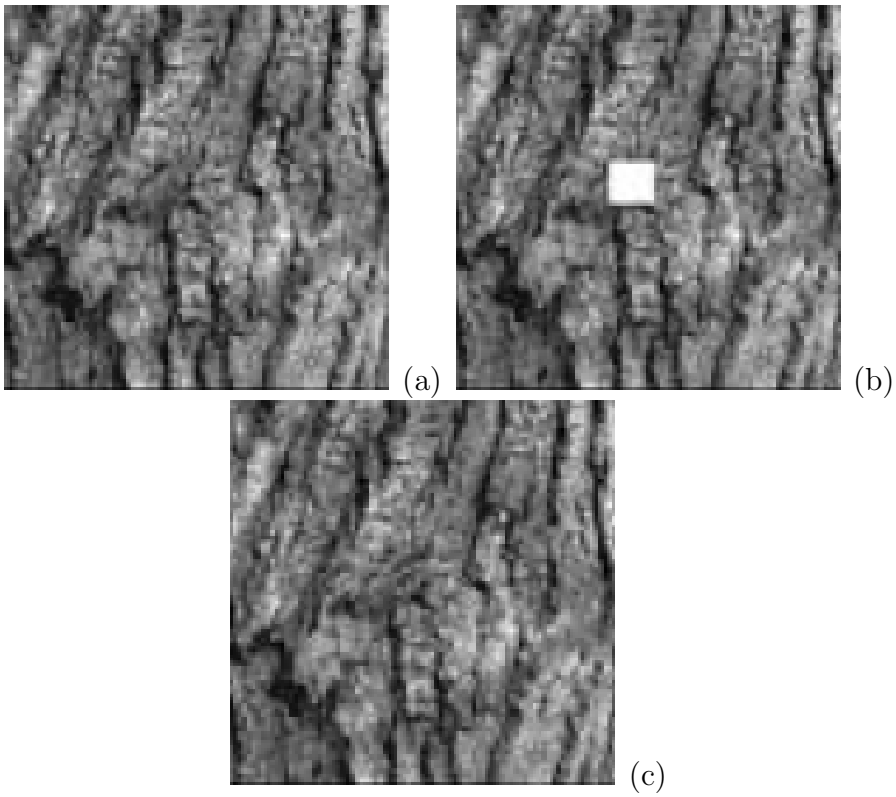

Figure 2.12: From [1], texture inpainting applied to a texture image. (a) Original *Bark* texture image. (b) Occluded image from (a). (c) Image from (b) after texture inpainting by AM-FM reaction-diffusion.

Fingerprint analysis also saw promising work resulting from the employment of the AM-FM model. In [58], an algorithm was presented to detect reference points in fingerprint images that had been analyzed using the AM-FM model. Performance advantages of this approach were demonstrated against competing methods.

2.3 Orthogonal Filterbank Decomposition

Although the AM-FM models described above found many uses in analysis applications and resulted in high quality reconstruction of images, perfect reconstruction efforts were hampered due to the inherently non-orthogonal nature of the Gabor filterbank. Gabor filters are very desirable for AM-FM analysis because they provide optimal joint spatiospectral localization which leads to locally coherent (*i.e.* locally smooth) image components. However, Gabor filters cannot be made orthogonal, which precludes the possibility of perfect reconstruction from filterbank decomposition unless an undesirable non-smooth family of biorthogonal functions is introduced for reconstruction (as is done in the Gabor transform [4]). Accordingly, several efforts were made to address the trade-off between using Gabor filters that provide smooth image components which correspond well to human visual perception but cannot provide perfect reconstruction and seeking for some alternative family of orthogonal filters that might provide perfect reconstruction while still retaining the many desirable properties of Gabor filters.

2.3.1 Orthogonal Filterbanks for Reconstruction

An orthogonal multiresolution representation was presented by Mallat in 1989 [62], where he defined the wavelet representation. In that work, he used the representation to study applications of image data compression and texture discrimination. Herley and Vetterli then used similar wavelets to construct perfect reconstruction wavelet filterbanks in [119]. Giving a discussion of the wavelet transform versus the classical short-time Fourier transform, they also described the conditions needed for perfect reconstruction. Their proposed filterbanks were biorthogonal with symmetries, with complementary high pass filters used in conjunction with the given low pass filters to achieve perfect reconstruction.

In 1999, to be able to quantify the joint localization of the analysis filterbank channels in both time and frequency, Przebinda *et al.* introduced a measure related to the discrete uncertainty principle in [20] for the phase plane compactness of general discrete-time signals. They presented a lower limit for compactness and showed that Gaussians may not be the most compact basis functions when discretized. In [21], they termed this measure the 'discrete Hirschman (1957) uncertainty principle', and showed that the discrete Hirschman measure does not correspond to the continuous Hirschman uncertainty principle (which generalizes Heisenberg uncertainty) no matter how densely it is sampled.

A similar discrete uncertainty measure was developed by Tay *et al.* in [111] to design a wavelet filter with optimal localization, and the filter was used to construct a separable 2-D wavelet transform, which was applied to test images. In [110], the new wavelet transform was used to obtain an estimate of the number of textured regions in an image. This estimate was obtained from

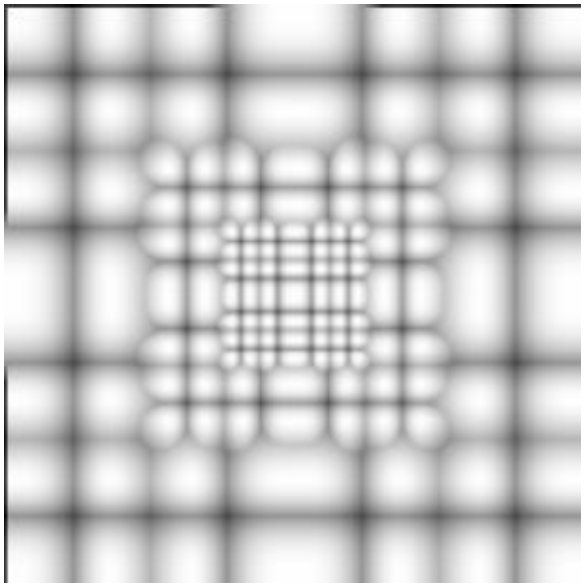

Figure 2.13: From [97], the discrete wavelet transform filterbank developed by Tay. The gray levels in this representation have been scaled to improve display.

the clustering of wavelet coefficients computed over disjoint blocks. In [112], an uncertainty measure was presented for entire wavelet filterbanks, taken as the geometric mean of per-filter uncertainty measures. Here, comparative test results were given for different channel count filter banks. Then, in [113], a new new non-separable wavelet filter bank was constructed to allow for a wavelet transform with orientation selection. Starting with the previously developed separable filterbank, the channels were decomposed in the frequency domain into two non-separable filters that provide orientation information. A representation of the developed separable filterbank can be seen in Fig. 2.13.

2.3.2 The Steerable Pyramid Filterbank

A separately developed orthogonal filterbank was the *Steerable Pyramid* filterbank presented in 1991 by Freeman [26]. The filters in this filterbank are linear

combinations of basis filters that can be adaptively steered to any orientation. Freeman showed that the filter output is analytically determined as a function of orientation. In [3], Daubechies *et al.* used the wavelet transform to decompose images in a pyramidal structure. Using this structure, they presented a progressive image transmission scheme that quickly transmits recognizable image structure. In [96], Simoncelli and Freeman stated that a drawback of orthogonal wavelet transforms is that they are not 'shiftable', meaning that they exhibit instability under translation, dilation, or rotation of the input signal. They proceeded to present one and two-dimensional shiftable transforms.

Then, in 1995, Simoncelli presented a description of Steerable Pyramid decomposition and of the implementation of the Steerable Pyramid transform [95]. While not used for perfect reconstruction results at the time, Steerable Pyramid decomposition was effectively used in a variety of analysis applications. An example of a Steerable Pyramid filterbank used later in this dissertation is shown in Fig. 2.14 and examples of image components obtained from decomposition using the filterbank can be seen in Fig. 2.15.

2.4 Spline Representation and Computation

A separate problem preventing perfect reconstruction from the AM-FM model was the inability to differentiate discrete phase or integrate discrete frequency modulation functions without error. While these calculations could be carried out exactly in the continuous domain, discrete versions of the calculations have always necessitated estimates of some sort. To greatly reduce the error associated with the discrete calculations, work was put towards fitting splines to the AM-FM functions.

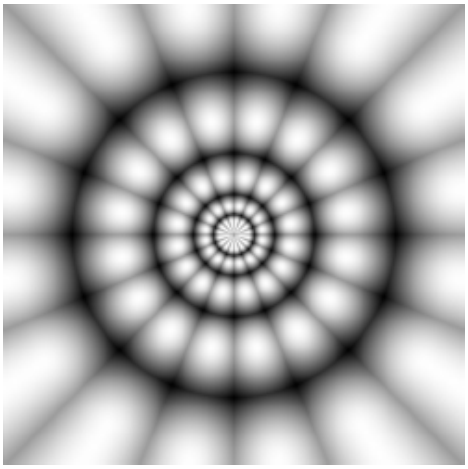

Figure 2.14: From [73], frequency domain representation of example Steerable Pyramid filterbank.

2.4.1 Spline Representation

The space of polynomial splines is given by [2, 114–116]

$$\mathcal{S}^p = \left\{ g^p(x) = \sum_{k\in\mathbb{Z}} y(k)\beta^p(x-k),\ x \in \mathbb{R},\ y \in \ell^2(\mathbb{Z}) \right\}, \tag{2.53}$$

where $\beta^p(x)$ is the B-spline of order p given by

$$\beta^p(x) = \sum_{i=0}^{p+1} \frac{(-1)^i}{p!} \binom{p+1}{i} \left(x + \frac{p+1}{2} - i \right)^p \mu\left(x + \frac{p+1}{2} - i \right) \tag{2.54}$$

and where $\mu(x)$ is the unit step function

$$\mu(x) = \begin{cases} 0, & x < 0 \\ 1, & x \geq 0. \end{cases} \tag{2.55}$$

The set of shifted B-splines $\{\beta^p(x-k),\ k \in \mathbb{Z}\}$ is a basis for the spline space $\mathcal{S}^p$ [116]. This allows a signal $g^p(x)$ to be represented in the spline space as a weighted sum of the basis functions, for which interpolation, differentiation, and integration operations can be accurately carried out. A theoretical

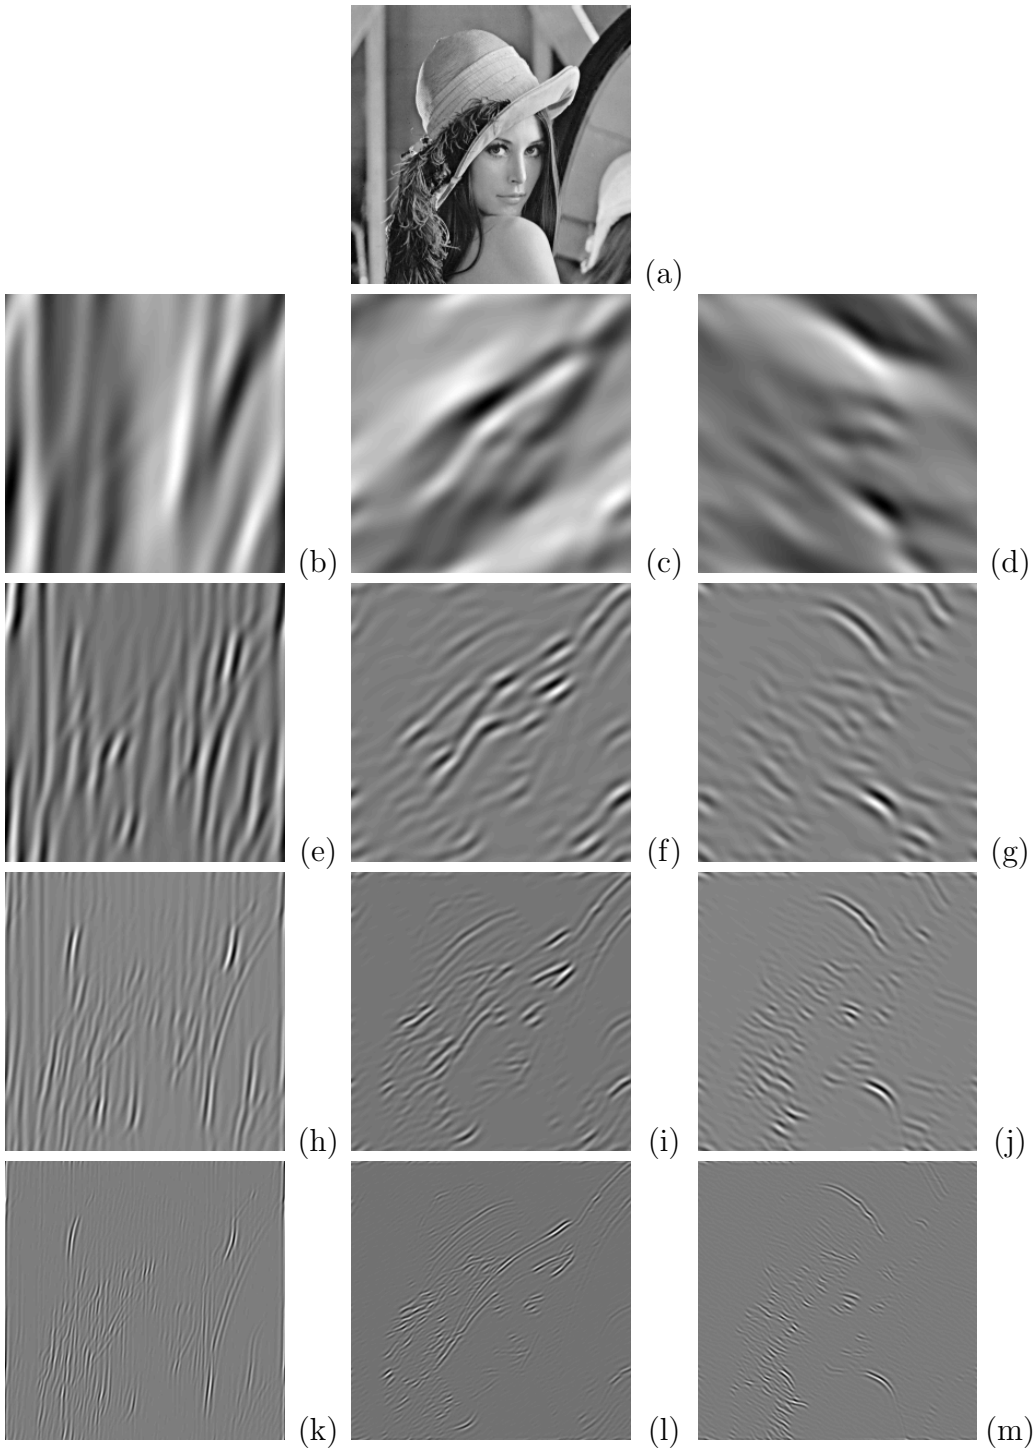

Figure 2.15: Example image components obtained via Steerable Pyramid decomposition. (a) Original *Lena* image. (b-m) Components at orientations: 7.5° (column 1), 60° (column 2), 120° (column 3).

framework supporting the general use of splines in discrete signal processing was developed by Unser, Aldroubi, and Eden in [115], where they presented algorithms for the continuous representation of discrete signals using B-splines and gave signal reconstruction with interpolation. In [116, 117], they described efficient filtering techniques to represent and process discrete signals with B-splines, while in [114] they discussed spline image processing and made further justifications for the usage of splines. They noted that useful examples of image processing operations that could be made theoretically rigorous by spline representation include edge detection via gradients and rotation and scaling via interpolation.

The weights $y(k)$ in (2.53) that are assigned to each spline basis function are the *spline coefficients*. By noting that all discrete B-splines $b^p(k)$ are symmetric FIR filters, Unser *et al.* developed methods by which to efficiently compute the spline coefficients using inverse filters. Spline coefficients may be computed from signal samples $g(k)$ recursively according to [97, 114, 117]

$$c^+(k) = g(k) + z_i c^+(k-1), \tag{2.56}$$

$$c^-(k) = z_i \left(c^-(k+1) - c^+(k) \right), \tag{2.57}$$

where z_i is the ith root inside the unit circle of the inverse filter for the B-spline. For B-splines of order $p = 3$, which is the order of the B-splines used in this dissertation, z_i is given in [114, 117] as $z_1 = 2 + \sqrt{3}$.

The equations (2.56) and (2.57) require boundary conditions before they can be applied. For an input signal g of length K the boundary conditions

developed by Unser *et al.* [114, 117] are given as

$$c^+(0) = \frac{1}{1 - z_i^{2K-2}} \sum_{m=0}^{2K-3} g(m) z_i^m, \tag{2.58}$$

$$c^-(K-1) = \frac{-z_i}{1 - z_i^2} \left(c^+(K-1) + z_i c^+(K-2) \right), \tag{2.59}$$

where the signal g in (2.58) is extended with mirror symmetric boundary conditions. Thus, to apply a filter $b^p(k)$ to an input signal $g(k)$ and obtain its spline coefficients $y(k)$ via the recursive implementation of the spline interpolation, the initial condition of (2.58) must first be calculated, allowing the application of the causal filter given by (2.56). Then from the resulting intermediate signal, the final condition of (2.59) must be calculated, allowing the application of the anticausal filter given by (2.57). The resulting signal then contains the desired spline coefficients $c^-(k) = y(k)$.

The representation of a multidimensional signal $g^p(\mathbf{x})$ in the space of polynomial splines can be accomplished by calculating the spline coefficients $y(\mathbf{k})$ present in the tensor product spline model [91, 114–117]

$$\begin{aligned} g^p(\mathbf{x}) &= \sum_{\mathbf{k} \in \mathbb{Z}^n} y(\mathbf{k}) \prod_{i=1}^{n} \beta^{p_i}(x_i - k_i) \\ &= y(\mathbf{k}) * \beta^{p_1}(x_1) * \beta^{p_2}(x_2) * \cdots * \beta^{p_n}(x_n), \end{aligned} \tag{2.60}$$

where $\mathbf{x} = [x_1 \ x_2 \ \ldots \ x_n]^T$ and $\mathbf{k} = [k_1 \ k_2 \ \ldots \ k_n]^T$. The B-spline basis functions in (2.60) are separable [116], and so the multidimensional spline interpolation is accomplished by applying the recursive filter implementation (2.56)–(2.59) independently in each dimension.

2.4.2 Spline Differentiation and Integration

The representation of a discrete signal in the space of polynomial splines can be thought of as the assignment of a weighted B-spline to each point of the

discrete signal. Once the spline coefficients of the representation are known, the following very useful properties can be observed [116]:

(A) The sum of the piecewise splines exactly constitutes the original signal.

(B) The piecewise splines are continuous where they join together, making the entire representation continuous and differentiable up to order $n-1$ for splines of order n.

The derivative of the order-p B-spline is given by [114, 116, 117]

$$\frac{d}{dx}\beta^p(x) = \beta^{p-1}\left(x+\frac{1}{2}\right) - \beta^{p-1}\left(x-\frac{1}{2}\right), \tag{2.61}$$

which leads to the derivative of the spline interpolant $g^p(x)$ given by

$$\begin{aligned}\frac{d}{dx}g^p(x) &= \sum_{k\in\mathbb{Z}} y(k)\frac{d}{dx}\beta^p(x-k) \\ &= \sum_{k\in\mathbb{Z}} y(k)\left(\beta^{p-1}\left(x-k+\frac{1}{2}\right) - \beta^{p-1}\left(x-k-\frac{1}{2}\right)\right) \\ &= \sum_{k\in\mathbb{Z}} \big(y(k)-y(k-1)\big)\,\beta^{p-1}\left(x-k+\frac{1}{2}\right),\end{aligned} \tag{2.62}$$

as well as the second derivative of $g^p(x)$ [114, 116, 117] which is given by

$$\frac{d^2}{dx^2}g^p(x) = \sum_{k\in\mathbb{Z}} \big(y(k+1)-2y(k)+y(k-1)\big)\,\beta^{p-2}(x-k), \tag{2.63}$$

where $x \in \mathbb{R}$ and $k \in \mathbb{Z}$.

The derivatives can be efficiently computed by calculating the first- and second-order differences according to [97]

$$dy(k) = y(k) - y(k-1) \tag{2.64}$$

$$d^2y(k) = y(k+1) - 2y(k) + y(k-1) \tag{2.65}$$

and then applying the shifted filter $c^{p-1}(k)$ to $dy(k)$ to compute the first derivative samples $\dot{g}(k)$ and the centered filter $b^{p-2}(k)$ to $d^2y(k)$ to compute the second derivative samples $\ddot{g}(k)$. The shifted filter $c^{p-1}(k)$ is the inverse of the discrete shifted B-spline $c^p(k) = \beta^p(x + \frac{1}{2})\big|_{x=k}$ developed by Unser *et al.* [2, 115, 116].

Integration of the derivative samples for recovery of the original discrete signal $g(k)$ can be carried out provided that initial conditions, analogous to the integration constant, are supplied. These initial conditions can be determined from the transfer functions of the filters used to compute the derivatives. In the case where a B-spline of order 3 is used to compute the spline interpolant, also referred to as a cubic spline interpolant, the z-transforms of the derivative filters are given by [114–117]

$$\mathcal{D}_1(z) = \left(\frac{6}{z + 4 + z^{-1}}\right)\left(1 - z^{-1}\right)\left(\frac{z+1}{2}\right), \tag{2.66}$$

$$\mathcal{D}_2(z) = \left(\frac{6}{z + 4 + z^{-1}}\right)\left(z - 2 + z^{-1}\right)(1), \tag{2.67}$$

where $\mathcal{D}_1(z)$ is the z-transform of the first derivative filter and $\mathcal{D}_2(z)$ is the z-transform of the second derivative filter. Upon analysis of these transfer functions, it can be seen that zeros exist at $z = 1$ and $z = -1$ for the first derivative filter, and at only $z = 1$ for the second derivative filter. This implies that two initial conditions are needed to recover $g(k)$ from its first derivative, whereas only one initial condition is needed to recover $g(k)$ from its second derivative.

Using (2.66), the signal samples $g(k)$ can be computed from the first derivative samples $\dot{g}(k)$ according to the integration filter [97]

$$g(k) = \dot{g}(k) * \left(c^{p-1}\right)^{-1}(k) * \mu(k) * b^p(k). \tag{2.68}$$

For the case of a cubic spline interpolant, the z-transform of the above integration filter is given by [97]

$$\mathcal{I}(z) = \left(\frac{2}{z+1}\right)\left(\frac{1}{1-z^{-1}}\right)\left(\frac{z+4+z^{-1}}{6}\right). \tag{2.69}$$

This transfer function shows that the spline integration filter is equivalent to the inverse of the spline derivative filter in (2.66) , meaning that $\mathcal{I}(z) = (\mathcal{D}_1)^{-1}(z)$. Similar results can also be obtained for the case of recovering $g(k)$ from its second derivative samples [97]. Thus, the original discrete signal $g(k)$ can be recovered from its first derivative samples by using the inverse filter $(\mathcal{D}_1)^{-1}(z)$ and the samples $g(0)$ and $g(1)$ as initial conditions, and $g(k)$ can be recovered from its second derivative samples by using the inverse filter $(\mathcal{D}_2)^{-1}(z)$ and the sample $g(0)$ as the initial condition.

2.5 Phase Unwrapping

By adding the spline representation and computation techniques discussed in Section 2.4 into the discrete AM-FM model, it becomes possible in theory to replace the earlier frequency estimation algorithms with a direct and more accurate calculation of frequency information as the gradient of the phase modulation from (2.29), but for one remaining obstacle. The arctangent function in (2.29) is multivalued and its principal branch has a range of $(-\pi, \pi]$, which introduces undesirable non-smooth jumps in the computed phase where the phase modulation function crosses branch cuts. To make this clear, consider that (2.29) can also be written as

$$W\{\varphi(\mathbf{k})\} = \arctan\left(\frac{\mathrm{Im}\left[z(\mathbf{k})\right]}{\mathrm{Re}\left[z(\mathbf{k})\right]}\right), \tag{2.70}$$

where $\mathbf{k} \in \mathbb{Z}^n$, $z(\mathbf{k})$ is the discrete analytic image, and $W\{\cdot\}$ is the wrapping operator that maps each phase value to its modulo-2π congruent value in $(-\pi, \pi]$. The phase samples directly calculated from (2.70) are known as the *principal phase values* (ppv's), and the resulting phase function is referred to as a *wrapped phase*. Because of the branch cuts in the arctangent function, the ppv's computed by (2.70) contain jumps that have no relation to visual image content, which renders any frequency information computed from them ineffective for analysis. In fact, the AM-FM frequency estimation algorithms presented thus far were developed to avoid direct use of the computed phase [8].

The wrapped phase can be made much more useful if the principal phase values are modified so that they are allowed to fall outside of the range $(-\pi, \pi]$, producing an equivalent phase function that is locally smooth and corresponds well to human visual perception of the visual information in an image. The process of modifying the principal phase values in this way is referred to as *phase unwrapping*, a problem which has been extensively studied. It is encountered in, *e.g.*, synthetic aperture radar (SAR) [23, 25, 121], magnetic resonance imaging [15, 101], texture discrimination and image segmentation [8], acoustic imaging, projection and diffraction tomography, and adaptive optics [124].

In one dimension, phase unwrapping is a relatively simple problem to solve. Because the phase values in the AM-FM model from (2.5) or (2.6) are arguments of sinusoidal functions, one may add any integer multiple of 2π to them without changing the signal whatsoever. As such, phase unwrapping is the process of selecting the best multiple of 2π, or *branch* of the sinusoidal function, to maximize local smoothness of the phase function. The unwrapped

phase can be expressed as [97]

$$\varphi(\mathbf{k}) = W\{\varphi(\mathbf{k})\} + 2\pi b(\mathbf{k}), \tag{2.71}$$

where $b : \mathbb{Z}^n \to \mathbb{Z}$ is the branch selection function for each principal phase value. In one dimension, the most common method of unwrapping phase is to assume that the difference between any two neighboring phase samples should be no more than π radians [14,57,97,102,103,107] and to select branches accordingly. This results in a smooth unwrapped phase in one dimension.

Unfortunately, in two dimensions phase unwrapping becomes more complex as each principal phase value has multiple neighboring phase values from which to select a branch. Thus, for each principal phase value, one neighbor must be chosen to conform to, and in general it is impossible to select branches such that the difference between *all* neighboring phase values will be less than π [97].

Many multidimensional phase unwrapping algorithms have been proposed, out of which a few of are noted here. In [103], Spagnolini presented a 2-D phase unwrapping algorithm given separate instantaneous frequency estimates that seeks to minimize discontinuity in the unwrapped phase. He noted that the presence of noise in signals results in the inability to distinguish between aliasing due to noise or due to steep phase slopes. Jain *et al.* presented a block least-squares method for 2-D phase unwrapping in [104]. This method breaks an image into small blocks of one phase wrap, unwraps the blocks, and then merges the blocks back together. It was shown to successfully unwrap real and synthetic images. Valadao gave an energy minimization framework for phase unwrapping in [5]. The algorithms given for the framework were termed 'phase

unwrapping max-flow/min-cut', and solve optimization problems using graph cut techniques.

In [97–100], Sivley and Havlicek developed a new multidimensional phase unwrapping algorithm for the AM-FM model that was based on the spline representation and computation techniques described in the previous section. This algorithm made use of a least-squares phase unwrapping method in which an estimated unwrapped phase with a second derivative closest in the least-squares sense to the second derivative of the principal phase values is computed and then made congruent with the principal phase values to arrive at a 'consistent' estimation of amplitude, phase, and frequency information.

The estimated unwrapped phase modulation, denoted by $\psi(\mathbf{k})$, is arrived at using the common least-squares minimization procedure, where the error function is defined as [99]

$$\varepsilon = \sum_{\mathbf{k}} \left| \partial^2 \psi(\mathbf{k}) - \partial^2 \widehat{\varphi}(\mathbf{k}) \right|^2. \tag{2.72}$$

In (2.72), $\partial^2 \psi(\mathbf{k})$ is the second derivative of the desired estimated phase modulation and $\partial^2 \widehat{\varphi}(\mathbf{k})$ is the second derivative of the principal phase values, estimated as [99]

$$\partial^2 \widehat{\varphi}(\mathbf{k}) = \mathrm{Re}\left[\frac{\partial^2 z(\mathbf{k})}{j z(\mathbf{k})} \right] + \mathrm{Im}\left[\left(\frac{\nabla z(\mathbf{k})}{j z(\mathbf{k})} \right)^2 \right]. \tag{2.73}$$

An efficient method for direct calculation of the estimated phase modulation $\psi(\mathbf{k})$ using DFT techniques was developed in [97]. Since the spline derivative operation is separable, the direct calculation makes use of a 1-D second derivative filter $h(k)$, where $k \in \mathbb{Z}$, such that (2.72) is now expressed

as [97]

$$\varepsilon = \sum_{i=1}^{n} \sum_{\mathbf{k}} \Big(\psi(\mathbf{k}) * h(\mathbf{e}_i^T \mathbf{k}) - \mathbf{e}_i^T \nabla \widehat{\varphi}(\mathbf{k}) \Big)^2, \tag{2.74}$$

where $\mathbf{e}_i$ is the unit vector in the k_i direction. The $z-$transform of the second derivative filter $h(k)$ is given as [99]

$$H(z) = \frac{6z - 12 + 6z^{-1}}{z + 4 + z^{-1}}. \tag{2.75}$$

Differentiation of the estimated unwrapped phase modulation gives a frequency modulation that corresponds well to visual perception of the image content. However, the estimated phase is not *congruent* with the principal phase values, meaning that each estimated unwrapped phase value $\psi(\mathbf{k})$ is not a multiple of 2π added to the principal phase value $\widehat{\varphi}(\mathbf{k})$ as per (2.71) and cannot be used to perfectly reconstruct the original image.

An unwrapped phase modulation that retains the structure of the estimated unwrapped phase modulation but is congruent with the principal phase values is obtained by defining the desired unwrapped phase modulation from (2.71) to be a scaled version of the least-squares phase $\psi(\mathbf{k})$ according to [97]

$$\varphi(\mathbf{k}) = W\{\varphi(\mathbf{k})\} + 2\pi b(\mathbf{k}) \approx \gamma \psi(\mathbf{k}), \tag{2.76}$$

where $W\{\varphi(\mathbf{k})\}$ is the wrapped phase modulation containing the principal phase values and the branch function $b(\mathbf{k})$ is defined as [97]

$$b(\mathbf{k}) = \text{round}\left(\frac{\gamma \psi(\mathbf{k}) - W\{\varphi(\mathbf{k})\}}{2\pi} \right), \tag{2.77}$$

where γ is a large, real-valued scaling factor.

Selection of a large value for γ controls discontinuities in the original phase structure by making them insignificant compared to the magnitude of

the scaled phase. For practical computation, a value of $\gamma = 300$ was advocated by [97] to ensure that the computed scaled phase modulation $\varphi(\mathbf{k})$ embodies the same nonstationary signal structure as the estimated least-squares phase modulation $\psi(\mathbf{k})$.

The frequency modulation $\nabla\varphi(\mathbf{k})$ was then defined as the spline gradient of the scaled phase modulation $\varphi(\mathbf{k})$. By this definition, $\nabla\varphi(\mathbf{k})$ is approximately γ times larger than the frequency modulation that would be obtained directly from the principal phase values and thus is a scaled representation of the visually significant image structure [97]. The advantage to this framework is that the scaled phase modulation $\varphi(\mathbf{k})$ could be exactly recovered from the scaled frequency modulation $\nabla\varphi(\mathbf{k})$, and once divided by the scaling factor γ, could be used to exactly recover the original image, while still providing visually meaningful frequency modulations such as those seen in Fig. 2.16.

2.6 Perfect AM-FM Reconstruction

The development of a suitable orthogonal filterbank, spline-based representation, differentiation, and integration, and successful phase unwrapping algorithms culminated in the long-sought goal of perfect reconstruction from the multicomponent AM-FM representation of an image. A usable orthogonal filterbank allowed components to be independent of each other and thus accurately summed back together into the original image. Spline-based computational methods removed error from the operations involved in demodulation and reconstruction. Finally, successful phase unwrapping provided a usable and visually meaningful frequency modulation function from which the original phase could be recovered, allowing recovery of the original image components

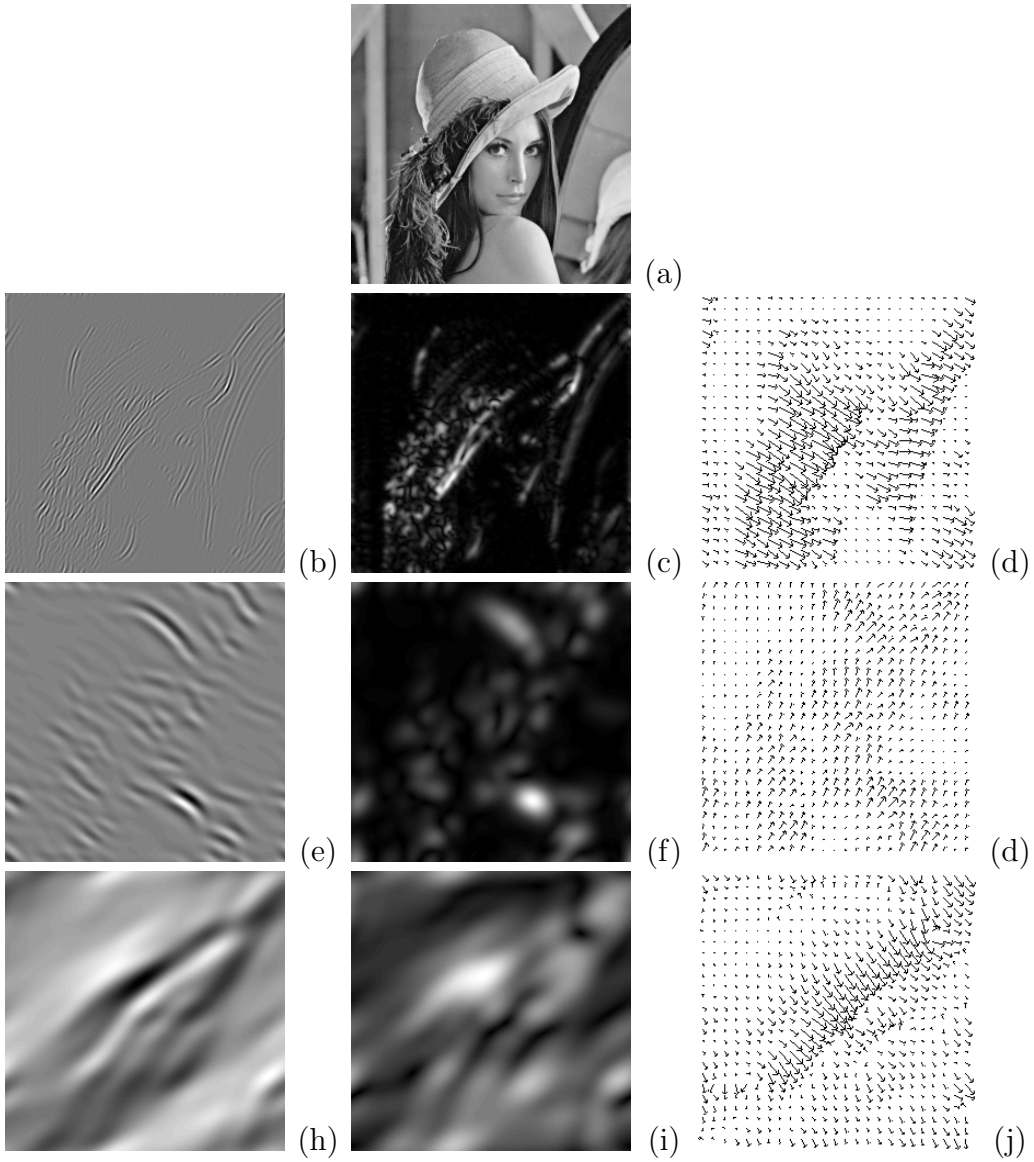

Figure 2.16: Example modulation functions for image components. (a) Original *Lena* image. (b-j) Example components (column 1) and corresponding amplitude modulations (column 2) and frequency modulations (column 3).

and then the original image.

This perfect reconstruction AM-FM model was presented in [97,99,100] using a separable wavelet filterbank that is then decomposed into non-separable orientation selective channels as described in Section 2.3. Each of the image components were then fit with splines and the phase of each component was unwrapped. Then the components were demodulated to obtain a perfect reconstruction representation. At this point, reconstructed images from the representation were visually indistinguishable from the original image.

Later work with the perfect reconstruction AM-FM image model, now also referred to as the *AM-FM transform*, replaced the decomposition filterbank with the orthogonal Steerable Pyramid filterbank described in Section 2.3.2. First employed together in [84], the Steerable Pyramid and AM-FM transform enabled much of the work described in the next section.

2.7 Recent Application of the AM-FM Model

Work was presented on a new multicomponent AM-FM model that allows asymptotically exact reconstruction of speech signals in [28]. This new model is informed by both the ESA and Analytic Image based demodulation approaches.

Computation of frequency modulation functions from the monogenic signal was demonstrated in [73, 79]. In this work, it was shown that AM-FM functions computed from partial Hilbert transform are similar to those computed from the monogenic signal although it was argued that the partial Hilbert approach is a more efficient representation.

Modulation domain filters for removing 'beat type' noise in corrupted

video were presented in [87]. This noise removal technique can be generalized for other fixed pattern noise types.

An algorithm to fuse infrared images was given in [81]. Infrared images were fused at both lowpass and highpass frequency ranges, and results were shown to be superior to other competing methods.

An AM-FM model based color to grayscale conversion algorithm was presented in [83]. Modulation domain features of color images were used to inform an algorithm that computed the mixing coefficients of the color channels. Favorable performance was demonstrated against other methods.

A new algorithm to compute modulation functions was also presented in [82]. It was shown that previous algorithms had a wrapped orientation problem in the phase gradient, and the problem was addressed by imposing local smoothness constraints on the phase. This algorithm is not used for image reconstruction, but leads to improved analysis capabilities.

2.7.1 Infrared Target Tracking

The AM-FM model has found considerable application for tracking moving targets in infrared video. The identification of infrared targets and backgrounds using modulation domain features was presented in [75]. The features considered included dominant modulations and it was shown that IR imagery possesses rich texture structure.

A dual-domain infrared target tracker was developed in [72], using feature information from both the image domain and the modulation domain. The new tracker made use of an SIR particle filter track processor and correlation based detection processes. The tracker was shown to have superior

performance on infrared missile closure sequences. The dual-domain target tracker was updated in [49] to run fully autonomously after the first frame of video. This allowed new consistency checks that can detect when the target appearance model needs to be updated, thereby solving a notoriously difficult problem in general target tracking for the infrared case.

A different modulation-domain-only target tracker was described in [94]. This tracker employed a new target appearance model update method that used a new measure of target signature evolution. The modulation-domain-only tracker was noted to have a reduced computational cost compared to the dual-domain trackers.

2.7.2 Texture Component Decomposition

Decomposition of texture into new texture components was explored in [73, 77, 78]. An algorithm was given to iteratively extract texture components based on a coherency measure. Examples of these texture components were given for Brodatz textures. Perfect reconstruction of the original image from computed texture components was also demonstrated.

Further decomposition of images into both textured and non-textured, or 'cartoon', components was demonstrated in [73, 80].

2.8 Modulation Domain Filtering

One of the principal aspirations behind the development of the AM-FM transform has been to gain the ability to perform signal processing operations directly on the computed modulation functions of an image. Because modulation functions correspond highly with human visual perception, filtering operations


(a)

(b)

Figure 2.17: From [84], example of modulation domain filtering on amplitude modulation functions. (a) Original *Lena* texture image. (b) Amplitude modulation attenuation within the black rectangle of (a) to remove the bands of Lena's hat.

performed in the modulation domain are capable of biological vision inspired processing tasks that are not easily obtainable using traditional linear filtering methods. The process of applying these operations in the modulation domain and successfully reconstructing from the filtered modulation functions is referred to as *modulation domain filtering* and is the primary topic discussed in the remainder of this dissertation.

2.8.1 AM Filters

Initial modulation domain filtering results were focused on performing modifications to the amplitude modulation only. Results of this type of filtering were first presented in [84], where AM functions were attenuated at locations matching specified FM characteristics. Among other results, this filtering was applied to the *Lena* image to remove the bands in Lena's hat as seen in Fig. 2.17.

AM based modulation domain filtering was also shown in [76]. Filtering

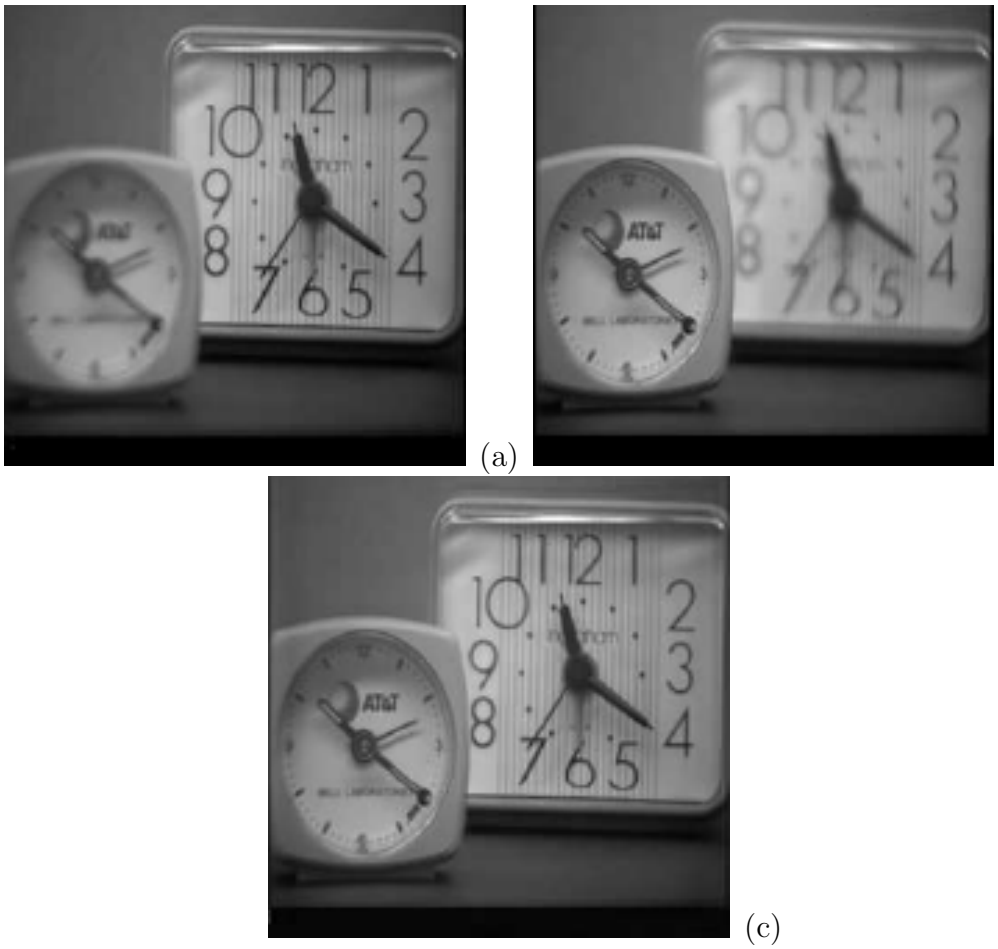

Figure 2.18: From [76], example of modulation domain filtering on amplitude modulation functions. (a),(b) Input clock images. (c) AM-FM fusion of (a) and (b).

was applied to obtain examples of frequency selective filtering, image restoration, and image fusion. The image fusion technique was used to combine two differently focused images of clocks into a single image that is focused everywhere, as shown in Fig. 2.18.

2.8.2 FM Filters

Modulation domain filters that modify the frequency modulation functions have turned out to be an inherently more difficult problem than modification of AM functions alone. This is primarily because reconstruction of the phase, and in turn the image, from the frequency modulation functions is achieved by integrating the FM functions. This integration is dependent upon an integration constant, which in the context of AM-FM image reconstruction considered here is known as an *initial phase condition*. For an unmodified frequency modulation, these initial phase conditions can be taken from the principal phase of the original image. However, once the FM function has been modified the initial phase conditions are generally unknown and can be very difficult to deduce.

Early approaches for estimating the initial phase conditions needed in order to reconstruct by integrating the modified frequency modulations were presented in [13]. A practical algorithm was given and incorporated into an integration scheme on different subwindows within the image. This algorithm, as well as another with estimated initial phase conditions based on a least squares solution of Poisson equations, was presented in [74]. The least squares based algorithm was used to perform pointwise rotation and isotropic zooming to a frequency vector field, an example of which is shown in Fig. 2.19. The subwindow integration based algorithm was also used to perform pointwise frequency vector field operations, as seen in Fig. 2.19 and Fig. 2.20. The 'flattening' filter in Fig. 2.19 (c) is achieved by setting the horizontal components of all frequency vectors to zero, while setting the vertical components to the magnitude of the original vectors. In Fig. 2.20, a pointwise rotation of $\frac{\pi}{2}$ radians is only applied to locations matching frequency orientation and magnitude

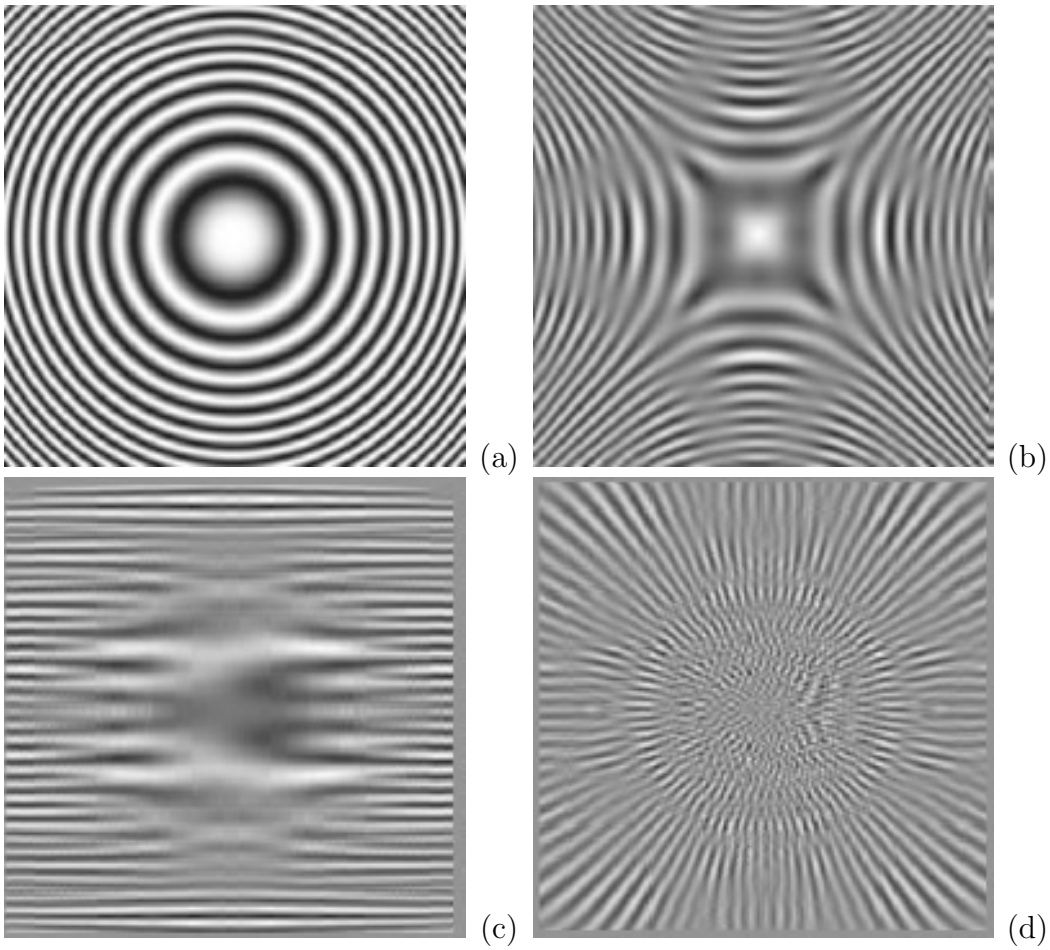

Figure 2.19: From [74], modulation domain filtering on frequency modulation functions. (a) Original *Chirp* image. (b) Rotation of frequency vectors in (a) using least squares based reconstruction. (c) Flattening of frequency vectors in (a) using integration based reconstruction. (d) Rotation and inversion of frequency vectors in (a) using integration based reconstruction.

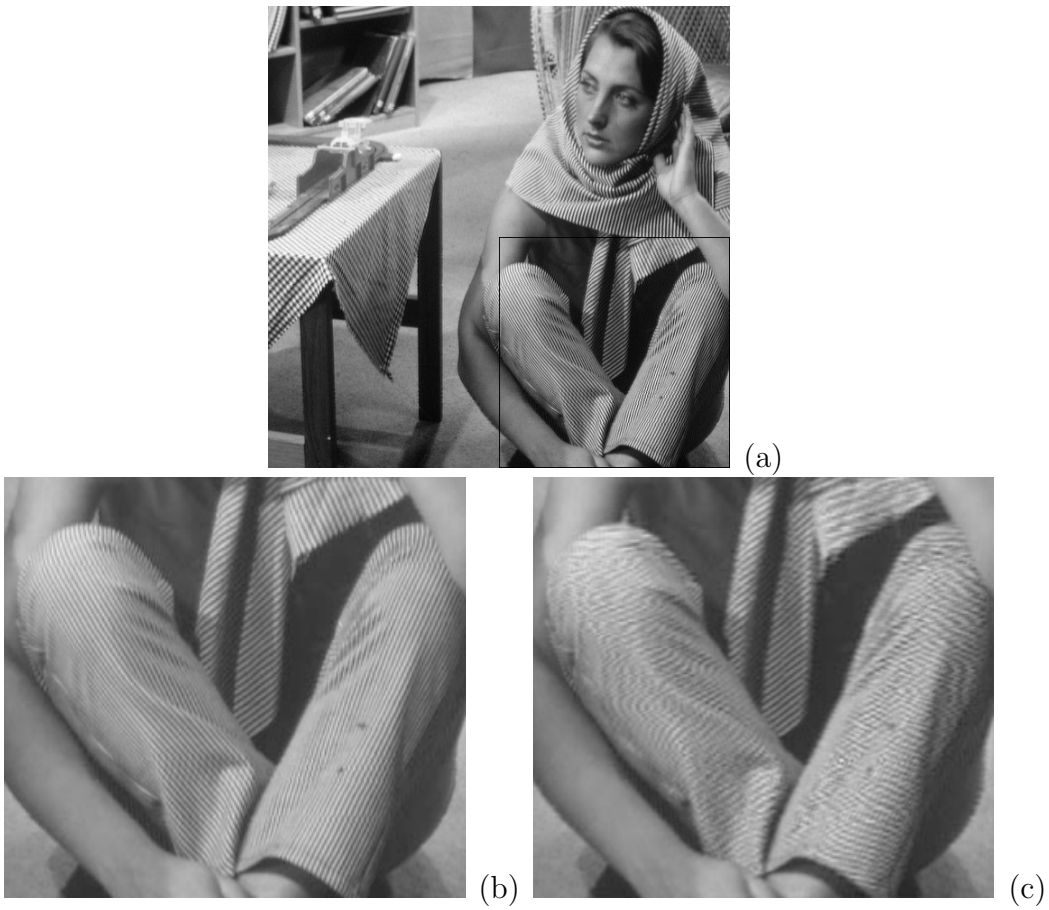

Figure 2.20: From [74], modulation domain filtering on frequency modulation functions in a natural image. (a) Original *Barbara* image. Processing was performed within the black rectangle. (b) Detail view of original image. (c) Result of FM processing to rotate the stripes on the pants.

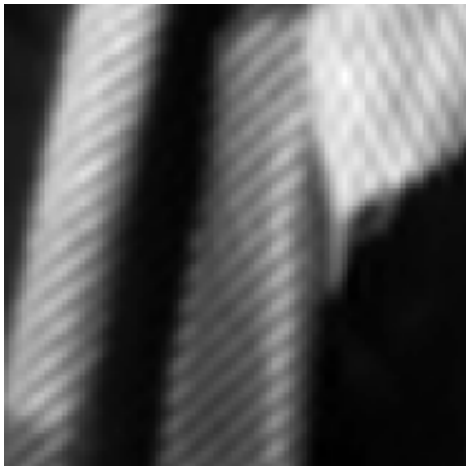
(a)

(b)

Figure 2.21: From [85], modulation domain scaling and rotation. (a) *Barbara* image zoomed in by 2.0. (d) *Boat* image rotated by 45 degrees.

characteristics that describe the lines in *Barbara*'s pants, effectively rotating them while leaving untouched other aspects of the image such as the folds in the pants. While crude in nature, these early results illustrate some of the types of remarkably powerful and visually meaningful signal processing results that can be obtained by modulation domain filtering and by FM processing in particular. Note that results of this type cannot be readily achieved by any conventional form of linear or nonlinear filtering.

New generalized AM and FM functions for modulation domain filtering with artifact free reconstruction were presented in [85]. A new version of the AM-FM model was developed that incorporates a generalized AM function $A_i(\mathbf{k})$ for each component i according to

$$\begin{aligned} s_i(\mathbf{k}) &= a_i(\mathbf{k}) \cos[\psi_i(\mathbf{k})] \\ &= A_{1i}(\mathbf{k}) \cos[\varphi_i(\mathbf{k})] + A_{2i}(\mathbf{k}) \sin[\varphi_i(\mathbf{k})], \end{aligned} \tag{2.78}$$

where $\varphi_i(\mathbf{k})$ is a smoothed least-squares unwrapped phase and where

$$A_{1i}(\mathbf{k}) = a_i(\mathbf{k}) \cos[p_i(\mathbf{k})] \tag{2.79}$$

and

$$A_{2i}(\mathbf{k}) = -a_i(\mathbf{k}) \sin[p_i(\mathbf{k})], \tag{2.80}$$

where $p_i(\mathbf{k})$ is a phase congruence term chosen to enforce the relation

$$\psi_i(\mathbf{k}) = \varphi_i(\mathbf{k}) + p_i(\mathbf{k}). \tag{2.81}$$

The generalized function was used to achieve high quality modulation domain scaling, rotation, and translation results, as seen in Fig. 2.21.

2.9 Summary

In this chapter, I presented the historical development and application of the AM-FM transform and modulation domain filtering. This included a discussion of work in the 20th century that established that biological vision systems involve mechanisms that are very similar to expansion by 2D Gabor elementary functions, which exhibit optimal joint localization in space and frequency. The development of the AM-FM model was then presented, with the usage of the Gabor decomposition contributing heavily to the suitability of the model to image analysis based on biological visual perception. Different methods of estimating the frequency modulations of the model were reviewed and multiple methods of AM-FM based analysis were described. Early results from the model were covered, with a focus on efforts to estimate image reconstructions from the model.

Perfect reconstruction from the model required the development of solutions to the challenges of orthogonal filterbank design, spline based computational methods in the model, and phase unwrapping. Each of these topics was covered in turn and the resulting emergence of the AM-FM transform with

perfect reconstruction from the model was discussed. The remainder of the chapter then focused on the usage of the AM-FM transform in relatively recent work, with a strong focus on the usage of the transform to reconstruct images from modified modulation functions, *i.e.* perform modulation domain filtering. Results were presented from past modulation domain filters which modify both amplitude and frequency modulations and the difficulties inherent in reconstruction from modified frequency modulations were noted. Modifications made to the AM-FM model that overcome these difficulties were presented.

Chapter 3

Phase Unwrapping for Modulation Domain Filtering

The phase unwrapping algorithm described in Section 2.5 was a breakthrough that enabled perfect reconstruction from the multicomponent AM-FM model (2.7). It also made possible modulation domain filters which modify AM functions. Unfortunately, the algorithm results in difficulty when reconstructing from modulation domain filters which modify frequency information due to the frequency and phase *scaling factor* that is a core part of the algorithm.

In this chapter, I will describe a new phase unwrapping algorithm that makes it possible to obtain high fidelity image reconstructions from computed AM-FM models without the usage of a phase scaling factor. When the frequency modulation functions of the AM-FM model are unmodified, the new phase unwrapping algorithm still allows for perfect reconstruction. The unwrapped phase produced by the algorithm also still provides accurate frequency modulation functions that are in good agreement with visual perception of the image.

In the new approach, frequency information from an initial least squares estimate of the unwrapped phase is used to guide selection of refined phase values that are congruent with the principal phase of the image. The selection process applies a queue-based region growing strategy to compute the final

unwrapped phase solution with sparse branch cuts that tend to be placed only in areas with low visual impact.

3.1 Challenges Presented by Modulation Domain Filtering

As described in Section 2.2.3, computation of $\varphi_i(\mathbf{x})$ in the multcomponent AM-FM model

$$z(\mathbf{x}) = \sum_{i=1}^{K} z_i(\mathbf{x}) = \sum_{i=1}^{K} a_i(\mathbf{x}) e^{j\varphi_i(\mathbf{x})} \tag{3.1}$$

by the inverse trigonometric function in

$$\varphi_i(\mathbf{x}) = \arctan\left(\frac{\mathrm{Im}\left[z_i(\mathbf{x})\right]}{\mathrm{Re}\left[z_i(\mathbf{x})\right]}\right) \tag{3.2}$$

results in a *wrapped* phase, also referred to as the *principal phase*, that is mathematically correct in the sense that it provides perfect reconstruction, but contains sharp discontinuities along branch cuts of the multi-valued inverse trigonometric function. These phase discontinuities are undesirable because they arise solely from mathematics and are both unrelated to and inconsistent with visual perception of the image. In order to obtain a representation that provides a visually meaningful interpretation of the FM functions as phase gradient, $\varphi_i(\mathbf{x})$ must be *unwrapped* to remove these branch cuts to the extent possible. However, in order to maintain perfect reconstruction, the unwrapped phase at every pixel must be be congruent with the principal phase values in (3.1).

There is an additional constraint for phase information that is unwrapped within the context of the AM-FM image model in order to allow for perfect reconstruction of the original image: each value $\widehat{\varphi}_i(\mathbf{x})$ of the unwrapped

phase must be congruent with the principal phase $\varphi_i(\mathbf{x})$ modulo 2π:

$$\widehat{\varphi}_i(\mathbf{x}) = \varphi_i(\mathbf{x}) + 2\pi b, b \in \mathbb{Z}. \tag{3.3}$$

In Section 2.5, an LMS technique was described that produces an unwrapped phase estimate $\psi_i(\mathbf{x})$ for the AM-FM model having a second derivative as close as possible to the second derivative of the principal phase in a least-squares sense. However, $\psi_i(\mathbf{x})$ is not itself congruent with the principal phase. A congruent unwrapped phase is computed from $\psi_i(\mathbf{x})$ through the use of a scaling factor. Upon multiplying $\psi_i(\mathbf{x})$ by a sufficiently large scale factor and then selecting the closest congruent phase values, a *congruent* unwrapped phase is created that retains the *structure* of the LMS estimate in a scaled form.

The scaling factor, which is typically chosen to be 300, causes the difference between neighboring values in both the phase and frequency modulation functions to become very large, but the scaled information remains completely accurate as long as the scaling factor is included in analysis operations and reconstruction from unmodified frequency information. However, if the scaled frequency vector field is modified in any way that requires the reconstructed phase modulation to be estimated from the modified frequency modulation, then the resulting reconstructed phase values will differ significantly from their ideal correct values due to an estimation error that is multiplied by the scaling factor. Once these reconstructed phase values are applied to the AM-FM image model (3.1), the output of the model suffers a substantial loss of fidelity.

3.2 Frequency Guided Phase Unwrapping

Phase unwrapping techniques may be generally classified as either path-following algorithms or minimum-norm algorithms [98, 124]. The phase unwrapping method introduced here uses elements of both approaches. It employs the LMS unwrapped estimates described in Section 2.5, but uses a region growing technique similar to those found in [24, 25, 121, 124] to obtain a congruent unwrapped phase that is also unscaled, has minimal local variation, and closely agrees with the smooth (but incongruent) LMS phase estimate. As described in Section 2.5, the LMS unwrapped phase estimates are computed on the pixel grid using a DCT-based technique. However, unlike in Section 2.5, where phase scaling is applied to impose congruency, here the LMS solution is directly interpolated with cubic tensor product splines and differentiated analytically to obtain an initial estimate of the FM field. This initial estimate is used to guide a queue-based phase growing algorithm that produces a final unwrapped and non-scaled phase modulation.

The initial FM estimate is in good agreement with visual perception but cannot provide perfect reconstruction because the LMS phase estimates are incongruent with the principal phase values. It is instead used to guide a queue-based region growing algorithm in order to obtain an improved unwrapped phase solution $\widehat{\varphi}_i$ that is congruent with the principal phase but has a gradient that is close to the initial FM estimate in the sense that it is also in good agreement with visual perception of the image. The region growing algorithm seeks to unwrap phase values smoothly starting from dominant/salient seed points characterized by high amplitude. For each component t_i in (3.1), the AM function a_i is thresholded at 20% of peak. Connected components labeling

is performed on the set of pixels where the AM exceeds this threshold. For each resulting connected component, the pixel with maximum AM is placed in a queue where it will seed a region in the final unwrapped phase solution. For each of these seed pixels, the final phase solution is computed by (3.3) with

$$b = \arg\min_{b \in \mathbb{Z}} \left| (\varphi_i(\mathbf{x}) + 2\pi b) - \psi_i(\mathbf{x}) \right|, \tag{3.4}$$

which is the congruent phase value that is closest in absolute value to the LMS unwrapped phase. The seed pixel is then removed from the head of the queue and its 4-neighbors are added to the tail of the queue. When all seed pixels have been processed in this way, the queue is initialized.

Phase unwrapping is then performed for all remaining pixels in t_i using the region growing algorithm given in Table 3.1. In each iteration of the main Repeat loop, an improved unwrapped phase value is assigned to the pixel $t_i(\mathbf{n})$ at the head of the queue. This is accomplished by looping over the N, S, E, and W neighbor pixels. For each neighbor $t_i(\mathbf{m})$ where an unwrapped phase has already been assigned, lines `13-14` select a candidate congruent phase value $\widehat{\varphi}^{\text{dir}}$ with minimum absolute difference from the sum (for a N or E neighbor) or difference (for a S or W neighbor) of the appropriate (horizontal or vertical) component of the LMS frequency estimate at $\mathbf{m}$ and the unwrapped phase at $\mathbf{m}$. These candidates will all be congruent, differing by integer multiples of 2π as a consequence of lying on different branches of arctan in (3.2). Lines `17` and `25` assign the unwrapped phase at $\mathbf{n}$ by selecting the candidate with largest AM, which tends to suppress branch cuts from occurring in neighborhoods where the amplitude is large. Any 4-neighbors of pixel $t_i(\mathbf{n})$ where an unwrapped phase has not been assigned are added to the tail of the queue for subsequent processing (line `11`).

Table 3.1: Queue-based phase region growing algorithm

```
01 Repeat
02   n ← spatial coordinate of pixel at head of queue
03   remove pixel t_i(n) from head of queue
04   a_max = 0
05   for dir ∈ {N, S, E, W} {
06     m ← spatial coordinate of neighbor dir
07     if (dir ∈ {N, E}) op = " + "
08     else op = " − "
09     if (dir ∈ {E, W}) e = [1 0]^T
10     else e = [0 1]^T
11     if (φ̂_i(m) not assigned) place pixel t_i(m) at tail of queue
12     else {
13       b^dir = arg min_{b∈Z} |(φ_i(x) + 2πb)
                              − (φ̂_i(m) op e^T ∇ψ_i(m))|
14       φ̂^dir = φ_i(x) + 2πb^dir
15       if (a_i(m) > a_max){
16         a_max = a_i(m)
17         φ̂^s = φ̂^dir
18   } } }
19   if ((φ̂_i(n) assigned ) and (φ̂_i(n) ≠ φ̂^s)){
20     for dir ∈ {N, S, E, W} {
21       m ← spatial coordinate of neighbor dir
22       if ((φ̂_i(m) assigned) and ((φ̂^s ≠ φ̂^dir)){
23         place pixel t_i(m) at head of queue for reprocessing
24   } } }
25   φ̂_i(n) = φ̂^s
26 Until queue is empty
```

The conditional loop in lines `19-24` identifies 4-neighbors of pixel $t_i(\mathbf{n})$ where a suboptimal unwrapped phase was previously assigned by using the wrong neighbor $\widehat{\varphi}_i(\mathbf{m})$ in line `13`. This occurs because the unwrapped phase at the best neighbor was not yet available due to the pixel processing order. To improve local smoothness and redirect branch cuts towards local amplitude minima, such 4-neighbors are returned to the *head* of the queue for immediate re-processing (line `23`). The test in line `19` avoids the possibility of an infinite recursion by ensuring that the *neighbors* of any such re-queued pixel will not themselves be returned to the queue unless a previously assigned phase value is actually changed.

When the algorithm terminates, the unwrapped phase is congruent with the principal values and agrees closely with the LMS estimate while avoiding the use of phase scaling. The real part of component t_{20} of the *Chirp* image is given in Fig. 3.1(a). The unwrapped phase computed by frequency guided phase unwrapping is given in Fig. 3.1(b), where because of nonconservative residues in the phase [24], undesirable branch cuts still arise. However, the frequency guided phase unwrapping algorithm tends to place them in low-amplitude, low-saliency regions where their contribution to the overall image reconstruction is small. The FM field $\nabla\varphi_{20}(\mathbf{x})$ obtained by differentiating the unwrapped phase is given in Fig. 3.1(c). For each image component t_i in (3.1), the FM function $\nabla\varphi_i(\mathbf{x})$ is then obtained by fitting the unwrapped phase samples with cubic tensor product splines and differentiating analytically.

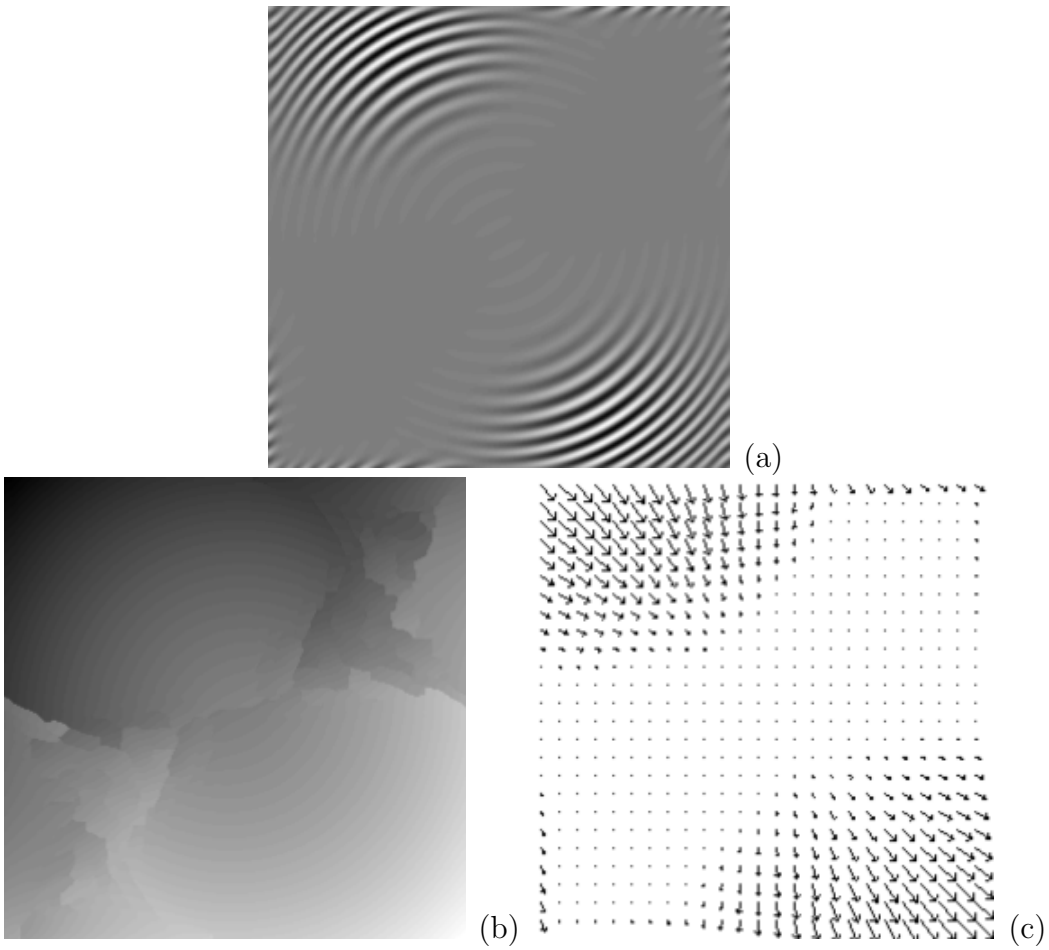

Figure 3.1: Phase unwrapping example. (a) Real part of component t_{20} of *Chirp* image. (b) Unwrapped phase from (a) using frequency guided phase unwrapping. (c) FM function $\nabla\varphi_{20}$ obtained by differentiating (b).

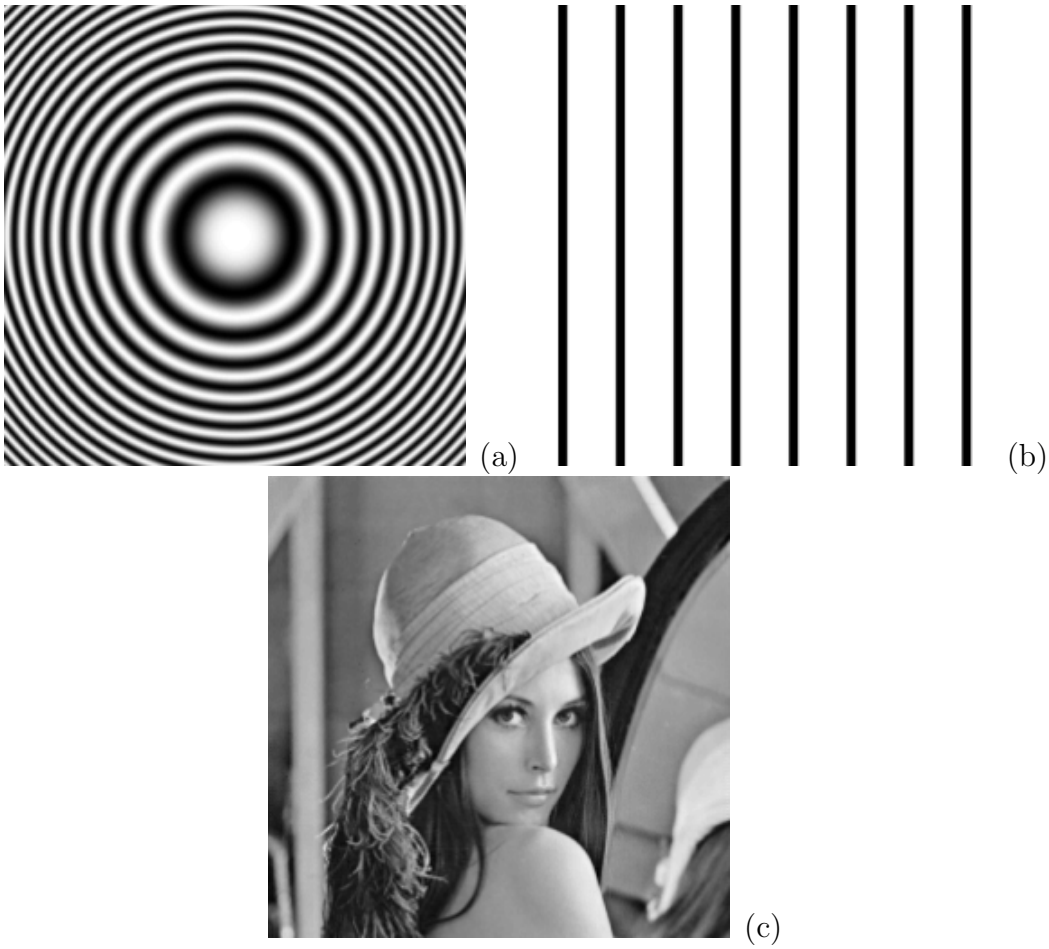

Figure 3.2: Perfect reconstruction examples. (a) Reconstructed *Chirp* image. (b) Reconstructed *Lines* image. (c) Reconstructed *Lena* image.

3.3 Improved Stability in the AM-FM Transform

The motivation for frequency guided phase unwrapping was to develop a suitable means for AM-FM reconstruction of a processed image after modification of the FM functions by modulation domain filtering. However, in this section I present reconstruction from the *unfiltered* modulation functions of the three images given in Fig. 3.2, since this provides known ground truth thereby enabling quantitative evaluation of the reconstruction errors. For each image,

Table 3.2: Mean squared reconstruction error for *Chirp*, *Lines*, and *Lena* images using different phase unwrapping algorithms for AM-FM reconstruction.

Algorithm	Initial Phase Vals	MS Reconstruction Error		
		Chirp	*Lines*	*Lena*
Phase Scaling [98]	stored	10^{-8}	10^{-8}	10^{-8}
Frequency Guided	stored	10^{-11}	10^{-12}	10^{-12}
Frequency Guided	estimated	10^{-6}	10^{-6}	10^{-5}

I computed an AM-FM model with $K = 40$ in (3.1). I reconstructed using the phase scaling method given in [98] and described in Section 2.5, the frequency guided phase unwrapping method described in Section 3.2 with multiple initial phase values stored, and the frequency guided phase unwrapping method with only a single initial phase value stored and the rest computed by a technique that will be given in Section 4.1.1. Mean squared reconstruction errors for each case are given in Table 3.2 for floating point pixels in the range [-1,1] (the pixels were scaled for modeling against (3.1)).

For all nine cases given in Table 3.2, the reconstructed images were visually indistinguishable from the originals. Fig 3.2(a) shows the reconstruction of *Chirp* obtained by adding all 40 components, each reconstructed by applying the unwrapped phase for the component to equation (3.1). Similar high fidelity reconstructions for *Lines* and *Lena* are given in Fig. 3.2(b) and (c). The main significance of these results is twofold. Line two of Table 3.2 demonstrates that the proposed phase unwrapping technique *substantially* improves the quality of the computed FM field. Line three shows that the proposed initial phase value estimation technique is capable of delivering high fidelity image reconstructions without the need for saving multiple phase initial conditions, which

is important for reconstructing from filtered AM and FM functions where the true initial phase values are unknown.

3.4 Summary

The phase unwrapping solution discussed in Section 2.5 made use of a phase scaling factor to achieve smoothly unwrapped phase that delivers frequency modulations that correspond well to visually meaningful image content. In this chapter, I discussed the impact of this phase scaling factor on reconstruction from modified frequency modulations and described how it can amplify reconstruction errors. To mitigate this problem, I presented a new phase unwrapping algorithm that does not employ a phase scaling factor. The new algorithm makes use of a smooth unwrapped phase estimate that has a second derivative as close as possible in a least-squares sense to the second derivative of the principal phase, which is obtained by direct calculation from the phase demodulation in (3.1). This smooth unwrapped phase cannot be used for reconstruction, but is used to guide the unwrapping of the principal phase using a queue-based phase growing methodology. The queue-based phase growing algorithm is designed to direct phase discontinuities in the unwrapped phase towards locations of minimal amplitude response where they contribute the least to potential reconstruction errors. I then demonstrated that the new phase unwrapping algorithm results in reconstructions that are more accurate than those obtained from phase unwrapping with phase scaling.

Chapter 4

Image Reconstruction From Processed Modulations

The new phase unwrapping technique described in Chapter 3 has allowed the development of new modulation domain filtering methods that do not rely upon phase scaling factors and thus exhibit greater stability in filtering computations. Furthermore, these new modulation domain filtering methods have been developed to operate directly from the unaltered AM-FM model, allowing the intuitive nature of the model to be a part of modulation domain filtering operations.

This chapter first describes the techniques investigated to enable reconstruction after modulation domain filtering. These techniques focus on the problem of reconstructing phase modulations from modified frequency modulation functions. Results are then presented for new modulation domain filtering methods. Examples are given for standard image rotation, scaling, and translation operations in the modulation domain and for modulation domain filters which attenuate amplitude modulation. The rotation and scaling of specific image texture is then presented, followed by results from new cross-component modulation domain filters.

4.1 Reconstruction Techniques

This section describes the investigation of multiple techniques for recovering phase from frequency modulation information. The techniques are presented in terms of their ability to successfully reconstruct from unmodified modulation functions and their applicability to reconstruction from modified frequency information.

4.1.1 Image-Wide Integration

The spline-based integration technique employed in Section 2.4 is one that recovers the phase from the frequency modulation function by integrating the entire frequency modulation image given initial phase values from the top left of the image. Reconstruction from the AM-FM model is performed by reconstructing the phase of each component and then applying (2.7) to arrive at a final reconstructed image.

As noted in Section 2.4, spline based integration from the first derivative requires two initial phase values. This requirement is observed in the equation used for spline integration. Once the first derivative samples of the phase have been found, the recovery of the phase at any one pixel from the derivative samples is dependent upon the phase values at the two prior pixels on the same row or column according to

$$x(k) = x(k-2) + \frac{x'(k-2) + 4x'(k-1) + x'(k)}{3}. \tag{4.1}$$

Starting with two original phase values, this equation can be used to iteratively integrate an entire row or column.

Since the 2D cubic tensor product splines are separable [116], the inte-



Figure 4.1: Perfect reconstruction from unmodified modulation functions using standard image-wide spline integration.

gration can be carried out separably as a 1D operation row by row or column by column as indicated in (4.1). Thus, the integration method used in Section 2.4 uses the original phase values from the first two pixels at the top of the first column to integrate the entire first column, and then uses the original phase values from the first two pixels at the top of the second column to integrate the entire second column. With phase values recovered for the first two columns, integration proceeds row by row, with each row integrated using the phase values from the first two pixels at the left of the row.

This means that the entire phase image for each component, and thus the entire original image, is recovered from each component's AM function, FM function, and the original values of the 4 pixels at the upper left of each component's phase image. These spline computational techniques result in perfect reconstruction, as seen in Fig. 4.1, which is a reconstructed image that has a mean squared error of 6.0288×10^{-12} against the original pixel values.



Figure 4.2: Rotated *Lena* image reconstructed directly from rotated AM and PM functions.

Integration via Subwindows

The spline integration technique described above works extremely well for reconstruction from modulation functions when the frequency modulation functions have not been modified. However, when the frequency modulation functions have been modified by AM-FM filtering, the requirement of initial phase values from the top left corner of the image can quickly become untenable. As an example, when the frequency modulation has been rotated, the original phase values in the top left corner no longer represent the initial integration constants needed for recovery of the rotated phase. This can be seen in Fig. 4.2, in which the amplitude, phase, and frequency modulations have been rotated using bicubic interpolation and the image was reconstructed directly from the amplitude and phase modulations.

An approach to overcoming this problem is to start the spline integration computations from initial phase conditions at a different point in the image. Starting from four neighboring initial phase values at some central point in

the image, spline integration can be applied to recover the entire phase by integrating in different directions within the four quadrants (subwindows) of the image. In practice, the four initial phase values are chosen to be those near the coordinates (r_1, c_1) of the global maximum of the amplitude modulation function, as $\varphi(r_1, c_1)$, $\varphi(r_1, c_1 - 1)$, $\varphi(r_1 - 1, c_1)$, and $\varphi(r_1 - 1, c_1 - 1)$. The region lying above and to the left of the initial phase values is then integrated by first integrating up the columns $c_1 - 1$ and c_1 and then integrating from right to left along rows from r_1 to the top of the image. To reconstruct the phase in the region lying below and to the left of the initial phase values, integration is carried out down the columns $c_1 - 1$ and c_1 and then from right to left across rows from $r_1 + 1$ to the bottom of the image. Analogous procedures are used to reconstruct the phase in the regions lying above and to the right and below and to the right of the initial phase values.

Similar to the integration technique that starts from the upper left of the image, this technique works very well for perfect reconstruction from unmodified frequency modulations, as it simply starts integration from four initial phase values at some central location in the image. This can been seen in the reconstructed image in Fig. 4.3, which has a mean squared error of 5.8676×10^{-12} relative to the original pixel values.

When the frequency field has been rotated, the four initial phase values at the central point of integration still do not represent the initial integration constants needed for recovery of the rotated phase. However, high quality reconstruction results can be obtained by starting from a single initial phase value, and estimating the other 3 initial phase values needed for reconstruction. Starting from the single initial phase value, the neighboring phase values can



Figure 4.3: Perfect reconstruction from unmodified modulation functions using spline integration split into four subwindows.

be simply estimated according to

$$\varphi(r_1, c_1 - 1) = \varphi(r_1, c_1) + \min\left(\nabla\varphi_x(r_1, c_1 - 1), \nabla\varphi_x(r_1, c_1)\right) + \left|\nabla\varphi_x(r_1, c_1 - 1) - \nabla\varphi_x(r_1, c_1)\right|/2, \tag{4.2}$$

$$\varphi(r_1 - 1, c_1) = \varphi(r_1, c_1) + \min\left(\nabla\varphi_y(r_1 - 1, c_1), \nabla\varphi_y(r_1, c_1)\right) + \left|\nabla\varphi_y(r_1 - 1, c_1) - \nabla\varphi_y(r_1, c_1)\right|/2, \tag{4.3}$$

$$\varphi(r_1 - 1, c_1 - 1) = \varphi(r_1 - 1, c_1) + \min\left(\nabla\varphi_x(r_1 - 1, c_1 - 1), \nabla\varphi_x(r_1 - 1, c_1)\right) + \left|\nabla\varphi_x(r_1 - 1, c_1 - 1) - \nabla\varphi_x(r_1 - 1, c_1)\right|/2. \tag{4.4}$$

It may be observed that this estimation will give more accurate estimates in areas where the change in phase is more linear, or in areas that are closer to constant frequency. The likelihood of this condition being true is assisted by selecting integration start points at areas of maximum amplitude, where the response to the channel filter is strong.



Figure 4.4: Estimated reconstruction from unmodified modulation functions using spline integration split into four subwindows. Integration is computed from a single initial phase value, and other required initial phase values are estimated.

Although the initial phase conditions used for integration are now estimated, this technique is still capable of providing high quality reconstruction. This can be seen in the reconstruction from unmodified modulation functions shown in Fig. 4.4. The reconstructed image is obtained from a single initial phase value from which the other initial phase values are estimated according to (4.2)–(4.4). The MSE of the reconstruction is 5.7322×10^{-5} relative to the original image pixel values.

Estimation from a single initial phase value, however, is not sufficient for quality reconstruction from modified frequency content. This is illustrated in the reconstruction result shown in Fig. 4.5, in which the frequency modulation function has been rotated and the phase has been recovered from a single initial phase value. The image could possibly be described as being 'out of phase with itself,' as in the rows and columns of the image don't quite seem to match up with each other. This is accurate, and can be better seen in the


(a)
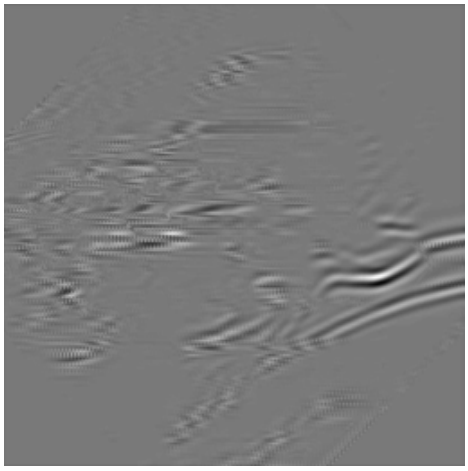
(b)

Figure 4.5: Estimated reconstruction from rotated modulation functions using spline integration split into four subwindows. (a) Reconstruction of entire rotated image. (b) Reconstruction of single component of image.

single component reconstruction shown in Fig. 4.5 (b). In each component, the reconstruction is accurate around the starting point, but quickly degrades as the integration computation of each row falls out of synchronization with its neighboring rows.

4.1.2 Subwindow Integration from Multiple Points

Since the above results for reconstruction from modified frequency content only give quality reconstruction near the point from which the integration is started, the next approach is to start integration from multiple locations in order to try to achieve quality reconstruction results for the most significant portions of the image content. A logical extension of the above approach is to define multiple windows within the image for integration, and then to start integration for each from a central point of smooth frequency using the subwindow integration method described in Section 4.1.1.



Figure 4.6: Reconstruction using windowed integration from multiple points without phase growing step.

For this approach, I have developed an algorithm which assigns windows for integration by first creating a binary mask from the amplitude modulation function using a threshold equal to a percentage of the max amplitude response in the image. Connected components labeling is then applied to the mask to find separate regions within the image and integration windows are defined as the spatial extents of these regions. Subwindow integration as described in Section 4.1.1 is then applied to each of these regional windows, starting from the point of greatest amplitude response within each window.

This technique is able to achieve quality reconstruction within the integration windows, but artifacts and distortion generally occur along the boundaries between windows. This can be seen in Fig. 4.6, which is a reconstruction from regional integration windows only, using a binary mask threshold of 40 percent of the maximum amplitude response per channel. The resultant image appears very 'blocky', with missing image content between the integration windows.

Image content in between the integration windows is filled in using a queue based phase growing method similar to the queue based method used for phase unwrapping in Section 3.2. In this application, the seeds for the phase growing queue are all border pixels around the integration windows. Phase values are then grown out from each of the windows in turn, with phase values from different regions meeting at midway points between the integration windows.

As the locations where phase values grown from different regions will usually have large differences between neighboring phase values, minimizing the reconstruction error demands that these locations be restricted to occur only in areas of low amplitude response where they will not contribute significantly to visual perception of the final reconstructed image. As the midpoints between integration windows may not fall at locations of minimum amplitude response, a requeueing mechanism was also developed for the queue based phase growing technique. In this mechanism, as the phase for a pixel is computed, if its neighbors are found to have already been processed then those neighboring pixels are prepended to the queue. In this way, pixels are checked to determine whether or not they have computed a phase based on a suboptimal neighboring phase and if so they are immediately recomputed from the optimal neighboring phase. Any potential suboptimal neighbors of such pixels are also requeued, causing recomputations to occur until lines of minimum amplitude are encountered.

As the reader may have already conjectured, the potential for many pixels to be requeued and recomputed many times causes the above algorithm to run much slower with the requeuing functionality enabled. The requeuing functionality results in optimal placement of locations where phase values


(a)

(b)

Figure 4.7: Estimated reconstruction using windowed integration from multiple points with phase growing step. (a) Result of phase growing algorithm run with requeuing. (b) Result of phase growing algorithm run without requeuing.

grown from different integration windows meet and the additional processing results in a slight improvement in the observed quality of the reconstructions. This can be seen by comparing the two images in Fig. 4.7. In Fig. 4.7 (a), the reconstruction algorithm was run with requeing enabled. The resulting MSE is 0.0030 compared to the original image pixel values. In Fig. 4.7 (b), the reconstruction algorithm was run without requeing enabled; the resulting MSE is 0.0032 compared to the original image pixel values.

The integration of multiple windows and subsequent phase growing algorithm described above was used for results seen in Section 2.8.2.

4.1.3 Phase Growing from Multiple Points

Although the integration of separate windows followed by phase growing can produce a reconstructed phase with phase values grown from different regions meeting at lines of minimum amplitude, it does so very inefficiently. An al-

ternative approach which bypasses integration windows and grows phase for *all* pixels from neighboring pixels of higher amplitude is able to efficiently and successfully grow the phase such that values grown from different sections meet at lines of minimum amplitude.

This approach makes use of a dependency image, which is simply an image in which each pixel contains a value that indicates which of its neighbors has the highest amplitude response that is greater than its own. For the purposes of phase growing, this acts as a guide for each pixel that specifies which neighbor to use when computing its own phase (that is, which neighbor the pixel is dependent on). For the reconstruction of phase from a frequency modulation image, a phase growing queue is seeded with all pixels in the dependency image which have no dependency, *i.e*, those which lie at local maxima of the amplitude function. Pixels are then processed according to the following rules. If a pixel has no dependency, then the phase value for that pixel is the value from the original phase image. Otherwise, the phase value for the pixel is estimated from the neighbor of dependency according to

$$x(k) = x(k-1) + x'(k-0.5), \tag{4.5}$$

where $x(k-1)$ is the phase of the neighboring pixel and $x'(k-0.5)$ is the gradient at the midpoint between them. The value of $x'(k-0.5)$ is acquired as a linear interpolation between the gradients $x'(k-1)$ and $x'(k)$. I also completed tests where $x'(k-0.5)$ was instead obtained by computing cubic interpolations for midpoints in the frequency modulation image and resulting reconstructions were not noticeably different, leading to an inference that the phase estimate is often closely approximated by linear estimation.



Figure 4.8: Estimated reconstruction from unmodified modulation functions using phase growing only algorithm.

The relative simplicity of this algorithm leads to much faster processing than the windowed integration algorithm given in Section 4.1.2. This algorithm produces high fidelity reconstruction results, as seen in Fig. 4.8 where the MSE of the reconstruction is MSE 3.0264×10^{-4} compared to the pixel values of the original image. This is somewhat less than the reconstruction error observed from the windowed integration algorithm and is likely due in large part to a much higher number of starting phase reconstruction points which take accurate values from the original phase.

Even so, because the phase growing algorithm described in this section starts from single points, it is suitable for reconstructing phase from modified frequency modulations. An example of the algorithm used to compute a reconstruction from a rotated frequency function is shown in Fig. 4.9 (a). This reconstruction is visually very similar to the image reconstructed directly from rotated phase without reconstruction from a frequency function in Fig. 4.2.

In both of the reconstructed images shown in Figs. 4.2 and 4.9, slight


(a)
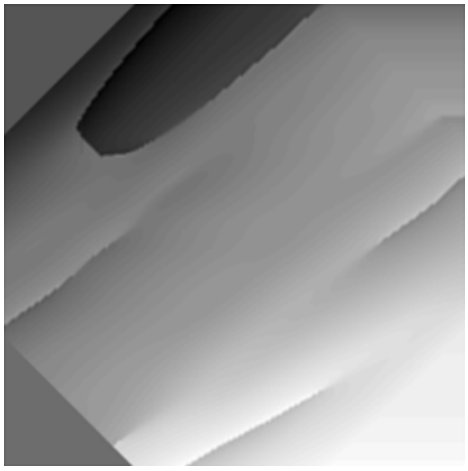
(b)

Figure 4.9: Estimated reconstruction from rotated modulation functions using phase growing only algorithm. (a) Reconstruction of entire rotated image. White arrow points to interpolation artifact. (b) Reconstruction of phase from rotated frequency modulation for a single component. The component displays a phase discontinuity that is visible as artifacts in (a).

artifacts are visible that are due to discontinuities in the unwrapped phase. One major example of such a discontinuity can be seen in the upper left part of the single component phase shown in Fig. 4.9 (b). These discontinuities do not appear in reconstructions from unmodified frequency content, but do appear in the reconstruction of rotated content due to the sub-pixel interpolation necessary when rotating the phase and frequency. A method to resolve this problem is given in Section 4.2.

4.2 Avoidance of Phase Discontinuities During Interpolation

As noted in Section 4.1.3 and shown in Fig. 4.9, discontinuities in the unwrapped phase can cause artifacts when subpixel interpolation is applied during filtering operations. These artifacts occur when the subpixel value is in-

terpolated from two neighboring phase values that are on opposite sides of a discontinuity, causing the subpixel value to fall somewhere in between the very different phase values and to be unrelated to image content.

This problem can be addressed by interpolating values for pixels near phase discontinuities in such a way that the interpolation operation does not consider values from across the phase discontinuity. The algorithm I developed to achieve this first identifies all phase discontinuities in the unwrapped phase as being locations where neighboring phase values differ by more than π. Then, because bicubic interpolation considers values from three pixels, any requested subpixel values within three pixels of a phase discontinuity are specially interpolated in the following way. If three pixel values are available in the direction opposite of the phase discontinuity, the new subpixel value is interpolated or extrapolated from those three pixels using bicubic interpolation. Otherwise, if two pixel values are available in the direction opposite of the phase discontinuity, the new subpixel value is interpolated or extrapolated from those two pixels using linear interpolation. If only one pixel value is available in the direction opposite of the phase discontinuity, the new subpixel value is copied from that pixel value.

This special handling of subpixel interpolation near phase discontinuities results in interpolated values that match the original image content much more closely, as demonstrated in Fig. 4.10. Fig. 4.10 (a) shows a zoomed in section of the original unrotated channel 19 component of the *Lena* image. Fig. 4.10 (b) shows the reconstructed rotated version of (a), with rotated phase subpixel values interpolated using standard bicubic interpolation across phase discontinuities. Fig. 4.10 (c) shows the reconstructed rotated version of (a),

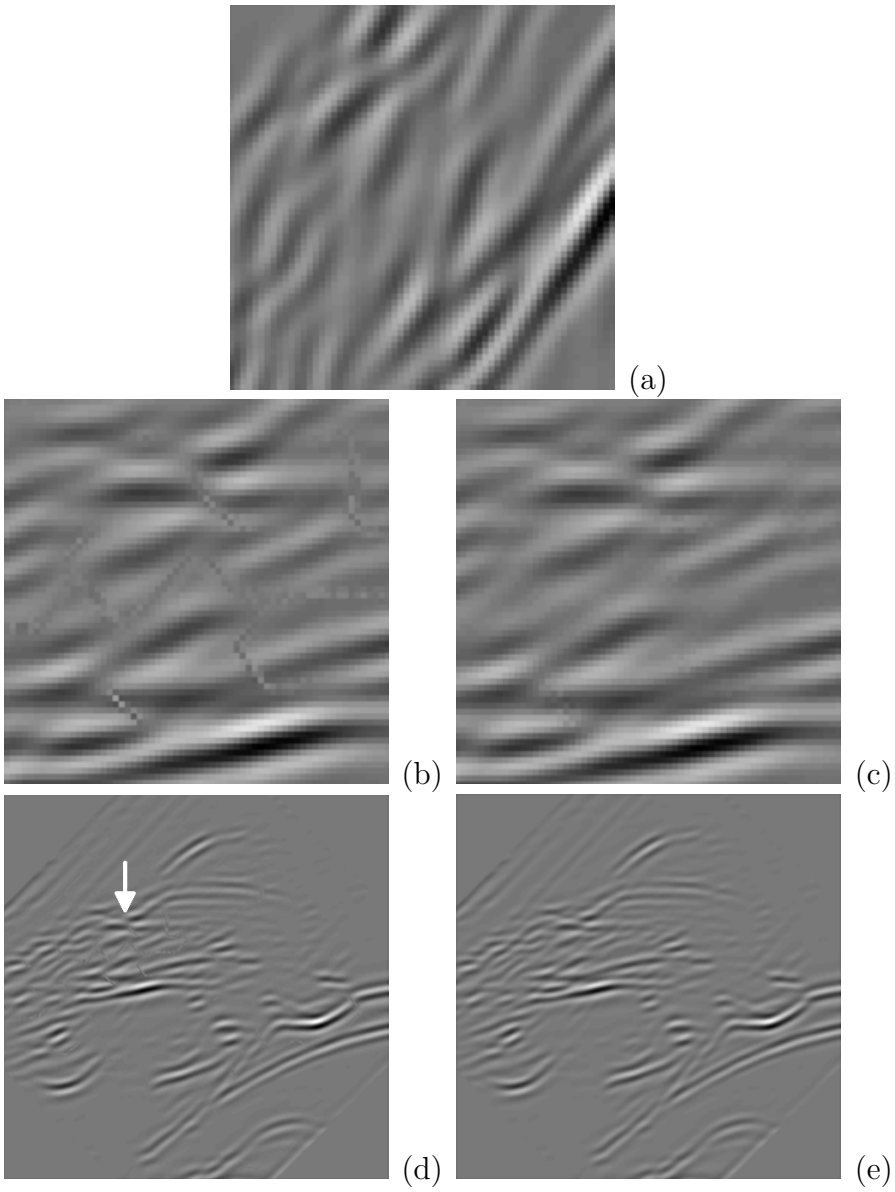

Figure 4.10: Comparison of interpolation methods for reconstruction of rotated mid-frequency channel image. (a) Unrotated region from zoomed in section of component image. (b) Reconstruction of rotated region from (a) with interpolation across phase discontinuities. (c) Reconstruction of rotated region from (a) with special interpolation algorithm used near phase discontinuities. (d) Reconstruction of entire channel image with interpolation across phase discontinuities. White arrow points to interpolation artifacts. (e) Reconstruction of entire channel image with special interpolation near phase discontinuities.

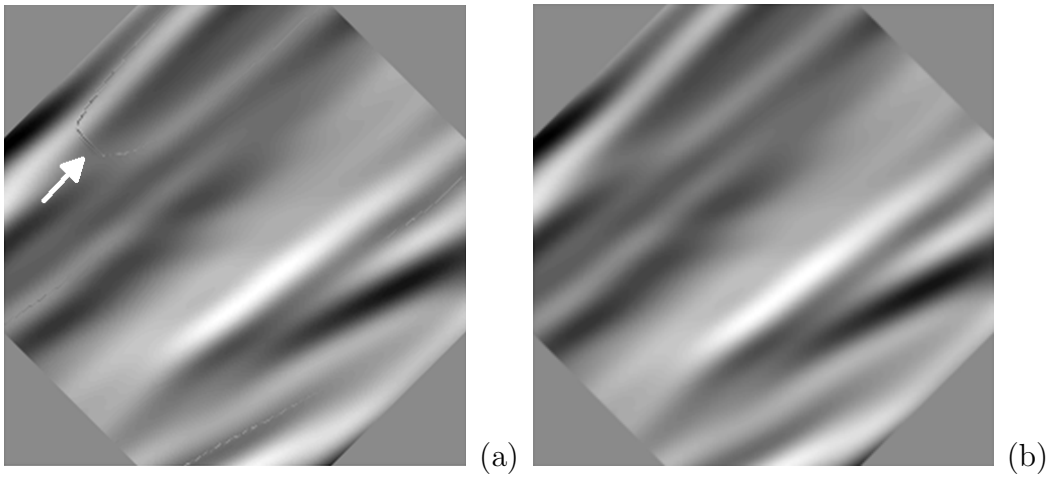

Figure 4.11: Comparison of interpolation methods for reconstruction of rotated low-frequency channel image. (a) Interpolation across phase discontinuities. White arrow points to interpolation artifact. (b) Special interpolation near phase discontinuities.

with rotated phase subpixel values interpolated using the special interpolation algorithm near phase discontinuities. It can be seen that artifacts visible in Fig. 4.10 (b) are no longer visible in Fig. 4.10 (c). The same results can be seen in the reconstructions of the entire channel image in Fig. 4.10 (d) and (e).

Similar results for a complete channel rotation are shown in Fig. 4.11. Artifacts in reconstruction from standard interpolation across phase discontinuities are no longer visible in reconstruction from special interpolation that avoids phase discontinuities. The phase for the channel component in Fig. 4.11 was shown in Fig. 4.9 (b).

The result for a complete image rotation is shown in Fig. 4.12. Again, reconstruction from standard interpolation across phase discontinuities results in visible artifacts that are unrelated to image content, while such artifacts are no longer visible in reconstruction using special interpolation near phase discontinuities.

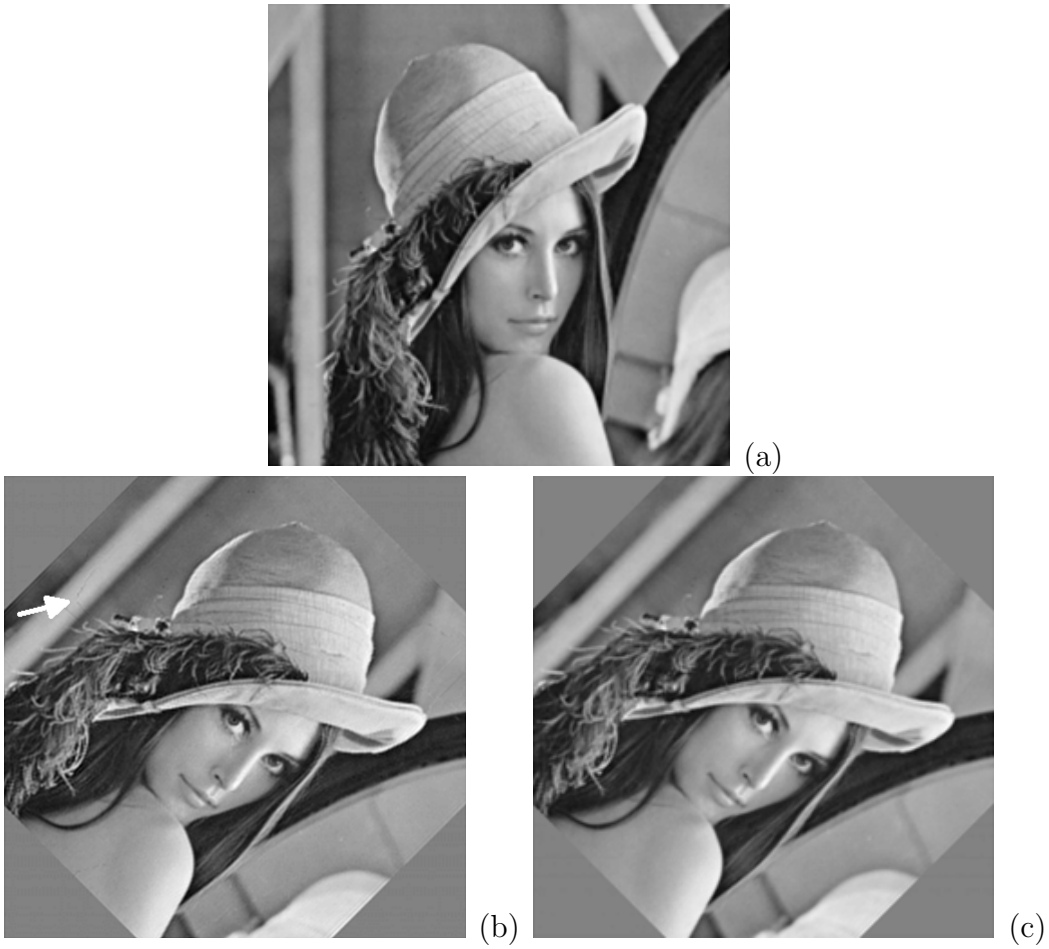

Figure 4.12: Comparison of interpolation methods for reconstruction of rotated natural image. (a) Input *Lena* image. (b) Interpolation across phase discontinuities. White arrow points to interpolation artifact. (c) Special interpolation near phase discontinuities.

4.3 Modulation Domain Scaling, Rotation, and Translation

With the development of special phase interpolation techniques near phase discontinuities, high quality modulation domain scaling, rotation, and translation operations became possible for the AM-FM transform without the need for counterintuitive modifications to the basic AM-FM model as were proposed and used in [84, 85, 97–100]. While such modifications as described in Section 2.8 were successful in achieving high fidelity AM-FM image reconstructions, they are undesirable in the sense that they compromise the simplicity, elegance, and intuitive nature of the basic model (2.5) and they also weaken the link to the processing that is known to occur in biological vision systems. The results presented here are the first to achieve modulation domain scaling, rotation, and translation from the AM-FM transform without the use of phase scaling factors or modified AM functions.

These scaling, rotation, and translation results were obtained by performing the filtering operations directly on the unwrapped phase, computing a corresponding frequency modulation from the filtered phase, and then reconstructing the final modified image. For each of the results shown in this section, a corresponding operation was performed in the pixel domain using standard bicubic interpolation for comparison purposes.

The pixel domain versus modulation domain rotation of *Lena* is shown in Fig. 4.13, with the pixel domain rotation shown in Fig. 4.13 (a), the modulation domain rotation shown in Fig. 4.13 (b), and the difference image between the two shown in Fig. 4.13 (c). The difference image in Fig. 4.13 (c) has been independently scaled for full contrast to visully amplify the details. The images

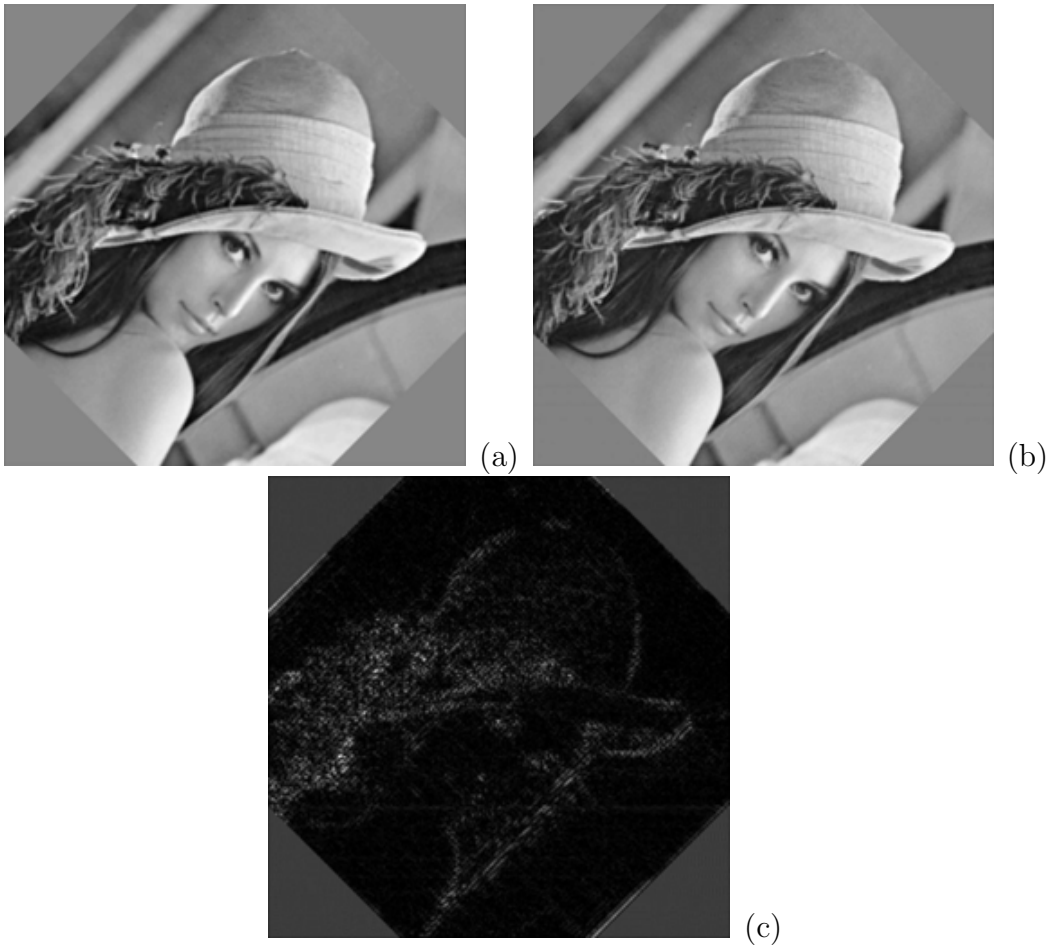

Figure 4.13: Example of rotation operation in modulation domain. (a) Rotation in real domain. (b) Rotation in modulation domain. (c) Difference image between (a) and (b).

have been rotated from the original by $\pi/4$ radians. The mean squared error between the pixel values of the two rotation images is 8.094×10^{-4} and the maximum single pixel value difference between the two images is 0.21533. The largest differences occur near the edges of the original image.

Fig. 4.14 shows similar results for the scaling and rotation of *Lena*. In this figure, the images have been rotated from the original by $\pi/4$ radians and scaled by a factor of 2. The mean squared error between the pixel domain and

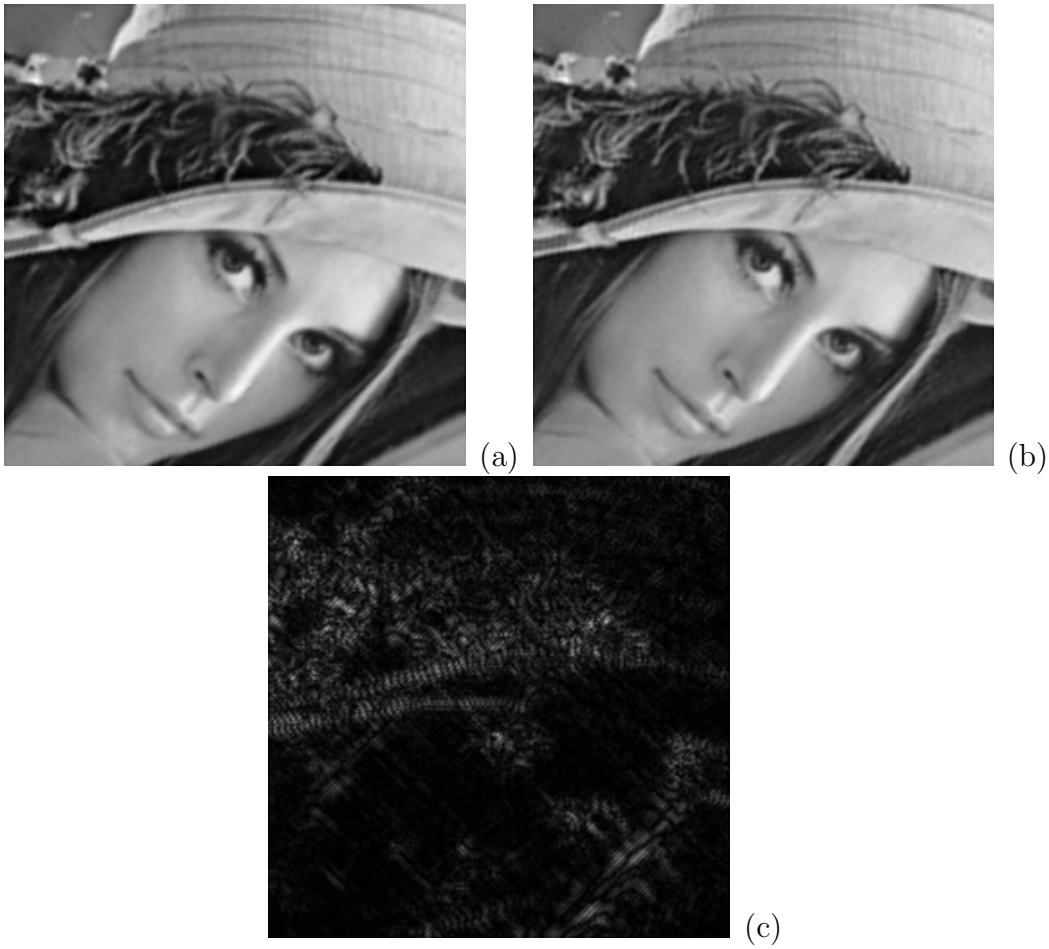

Figure 4.14: Example of scaling operation in modulation domain. (a) Scaling and Rotation in real domain. (b) Scaling and Rotation in modulation domain. (c) Difference image between (a) and (b).

modulation domain result images is 5.006×10^{-4} and the maximum single pixel difference is 0.20546. In this result, there is no difference at the edge of the original image, but the MSE and maximum pixel difference are comparable to the rotation only result.

Translation results are shown in Fig. 4.15. In this figure, the *Lena* image has been translated by 42.3 pixels in the horizontal direction, and 57.8 pixels in the vertical direction. The mean squared error between the pixel domain and

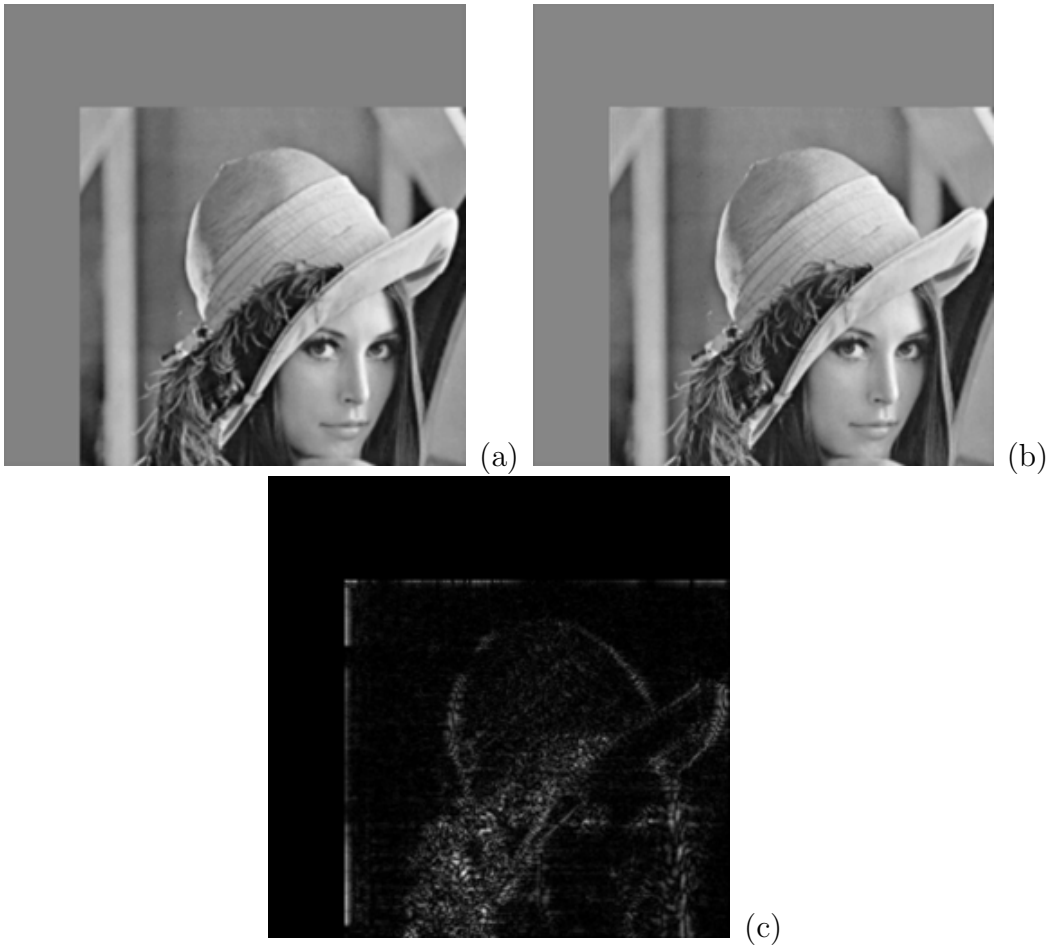

Figure 4.15: Example of translation operation in modulation domain. (a) Translation in real domain. (b) Translation in modulation domain. (c) Difference image between (a) and (b).

modulation domain result images is 2.7125×10^{-4} and the maximum single pixel difference is 0.18129. Here again, the largest differences occur near the edges of the original image. But both the MSE and the maximum pixel difference are slightly less than those of the other results.

4.4 Filtering via Attenuation of Amplitude Modulation

In this section, the phase unwrapping technique developed in Chapter 3 is applied to new filters that attenuate the amplitude modulations of images in order to obtain intuitive and visually meaningful image processing results. This class of modulation domain filters was discussed in Section 2.8.1, with the first filter results of this type presented in [84]. The new filters designed here are the first to achieve amplitude modulation attenuation without the use of phase scaling factors in the phase and frequency modulations. Here new filters are designed that dramatically modify images in ways that would be difficult or impossible using traditional linear or nonlinear filtering methods.

One of the new modulation domain filters presented here was designed to modify the water ripples present in the *Boat* image shown in Fig. 4.16 (a). In this image, the ripples of vertical frequency can be isolated and removed by applying a filter that sets the amplitude modulation to zero at locations where the frequency modulation characteristics match the ripples in the water. Denoting the area of the image to which the filter is applied, or the *region of interest* (ROI), as $\mathcal{X}$, the new attenuated amplitude modulations $\widehat{a}_i$ are given as

$$\widehat{a}_i(\mathbf{k}) = \begin{cases} 0, & \left(\left|\arg[\nabla\varphi_i(\mathbf{k})]\right| - \pi/2\right) < \pi/4 \\ & \text{and } \left|\nabla\varphi_i(\mathbf{k})\right| \in [\pi/8, 2\pi] \\ & \text{and } \mathbf{k} \in \mathcal{X}, \\ a_i(\mathbf{k}), & \text{otherwise,} \end{cases} \tag{4.6}$$

where i denotes the i_{th} component of the image, a_i and $\nabla\varphi_i$ are the respective unfiltered amplitude and frequency modulations, $\arg[\nabla\varphi_i(\mathbf{k})]$ is the orientation of the frequency modulation at $\mathbf{k}$ in rad, and $\left|\nabla\varphi_i(\mathbf{k})\right|$ is the magnitude of the frequency modulation at $\mathbf{k}$ given in units of rad/pixel.


(a)


(b)


(c)

Figure 4.16: Amplitude modulation attenuation on *Boat* image. (a) Original image. Filter operations were only applied within the bottom half of the image. (b) Attenuation of amplitude modulation at target frequency characteristics to remove ripples in the water. (c) Modified attenuation of amplitude modulation to leave high frequency ripple content.

The result of applying (4.6) to the *Boat* image can be seen in Fig. 4.16 (b), where the region of interest is the area of the image beneath the boat. It can be seen that the water takes on a very smooth appearance within the ROI, but that the horizontal frequency structure of the reflected boat remains mostly intact.

A slightly modified version of (4.6) can be applied to the *Boat* image to leave some high frequency vertical content on the water surface. This filter operates with a lower maximum frequency magnitude constraint and is given by

$$\widehat{a}_i(\mathbf{k}) = \begin{cases} 0, & (\big|\arg[\nabla\varphi_i(\mathbf{k})]\big| - \pi/2) < \pi/4 \\ & \text{and } \big|\nabla\varphi_i(\mathbf{k})\big| \in [\pi/8, \pi/2] \\ & \text{and } \mathbf{k} \in \mathcal{X}, \\ a_i(\mathbf{k}), & \text{otherwise.} \end{cases} \tag{4.7}$$

The application of (4.7) to the *Boat* image can be seen in Fig. 4.16 (c), where the water surface still appears less turbulent relative to Fig. 4.16 (a). However, the small amount of surface ripple in Fig. 4.16 (c) causes the filtered region of interest to blend in more naturally with the area outside of the ROI.

A similar filter was designed to operate on the *Wakes* image shown in in Fig. 4.17 (a), which displays an interference pattern created by two crossing boat wakes in an ocean. This filter is given by

$$\widehat{a}_i(\mathbf{k}) = \begin{cases} 0, & (\big|\arg[\nabla\varphi_i(\mathbf{k})]\big| - \pi/2) < \pi/8 \\ & \text{and } \big|\nabla\varphi_i(\mathbf{k})\big| \in [\pi/8, 2\pi] \\ & \text{and } \mathbf{k} \in \mathcal{X}, \\ a_i(\mathbf{k}), & \text{otherwise} \end{cases} \tag{4.8}$$

and when applied to the *Wakes* image, the filter successfully removes most of one of the boat wakes as seen in Fig. 4.17 (b).


(a)

(b)

Figure 4.17: Amplitude modulation attenuation on *Wakes* image. (a) Original image. Filter operations were only applied within the bottom half of the image. (b) Attenuation of amplitude modulation at target frequency characteristics to remove one of the boat wakes.

4.5 Rotation and Scaling of Target Image Texture

While global rotation of the entire amplitude and phase modulation images works well for whole-image rotations as shown in Section 4.3, different techniques must be applied when rotating a specific region or object. For the task of rotating the stripes in the pants of the *Barbara* image, a task which is shown in Fig. 4.18, the texture is rotated by performing a pointwise rotation of all frequency modulation vectors that match the target frequency characteristics and then reconstructing. In this section, a combination of previously discussed reconstruction methods will be employed to achieve new results that meet this image processing goal.

Note that rotation is applied only to the region of interest $\mathcal{X}$ shown in Fig. 4.18 (a). Within that region, all frequency vectors are pointwise rotated



Figure 4.18: Rotation of specific texture in image. (a) Original *Barbara* image. Image processing is applied to the region within the black rectangle. (b) Region of interest from (a). (c) Previous result from [74]. (d) Rotated texture, using phase growing reconstruction.

according to

$$\arg[\nabla\widehat{\varphi}_i(\mathbf{k})] = \arg[\nabla\varphi_i(\mathbf{k})] + \begin{cases} \pi/2, & \arg[\nabla\varphi_i(\mathbf{k})] \in [-\pi/2, 3\pi/16] \\ & \text{and } \left|\nabla\varphi_i(\mathbf{k})\right| \in [\pi/2, 2\pi] \\ & \text{and } \mathbf{k} \in \mathcal{X}, \\ 0, & \text{otherwise,} \end{cases} \tag{4.9}$$

where $\nabla\widehat{\varphi}_i$ is the filtered frequency modulation of the ith component of the image. Fig. 4.18 (c) shows the previous results for this task as described in Section 2.8.2. Those results applied the pointwise rotation filter (4.9) and then reconstructed using an earlier version of the windowed integration algorithm described in Section 4.1.2.

The results in Fig. 4.18 (d) are obtained by applying the pointwise rotation filter (4.9) and then reconstructing with the phase growing algorithm described in Section 4.1.3. While the two results are similar, it can be seen that the phase growing based reconstruction gives a somewhat cleaner and more natural looking result, which is likely due to the algorithm's use of more starting points for phase reconstruction.

A separate method for rotating a single, specific texture is to combine the spatial rotation method from Section 4.3 and the phase growing method from Section 4.1.3. Spatial rotation is not possible for texture in general because the rotation must be defined around a single point in space. As a texture region can be arbitrarily shaped, this leaves many points in the texture with no corresponding point to take a rotated value from. However, by combining spatial rotation for pixels where it is available and phase growing for other pixels, one can arrive at a reconstruction that is very clean for spatially rotated parts of the texture and provides coherent reconstruction from phase growing

for the other parts.

Ideally, the pixels for which quality spatial rotation is available would be automatically detected, however in this work windows for spatial rotation were manually specified. For each spatial rotation window, spatial rotation was carried out on a per-component basis around the center of the window in the following way. First, for each component, the pointwise frequency vector rotation of (4.9) is applied to all pixels in the image that match the target texture characteristics. Then, in each rotation window, for each pixel that matches the target texture characteristics, if the corresponding pre-rotation point from which it would take a new value also matches the target texture characteristics, then new interpolated phase and frequency values are computed for the pixel from the pre-rotation point. Once all possible pixels within the rotation window have been spatially rotated, the target pixels that neighbor the spatially rotated pixels are added to the queue that will be used for phase growing as seeds. The queue-based phase growing algorithm is then run to compute a new phase for all non-spatially rotated pixels, and the final image reconstruction is computed from the completed phase.

It may be noted that, while the spatial rotation described above is similar to the image-wide spatial rotation described in Section 4.3, it differs in that application of the spatial rotation is driven by information from the modulation domain functions.

For the rotation of *Barbara*'s pants, the results of the selected spatial rotation windows only, prior to the phase growing step, can be seen in Fig. 4.19 (b). The completed result after the phase growing step is shown in Fig. 4.19 (c). The final result exhibits relatively sharp and clear lines compared to previous

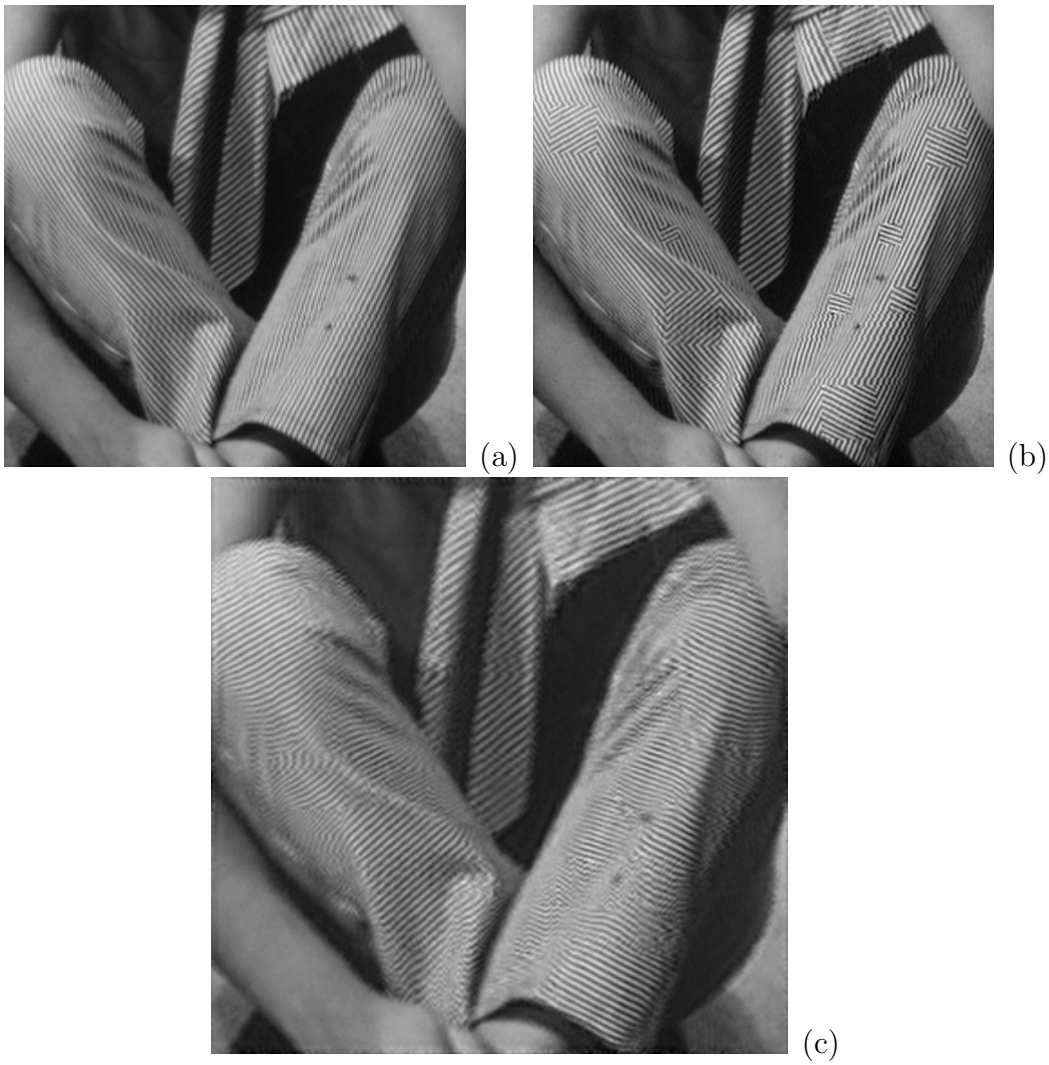

Figure 4.19: (a) Region of interest from original *Barbara* image. (b) Rotation of (a) with spatially rotated windows only. (c) Rotation of (a) with phase growing from windows in (b).

results, with texture reconstruction still showing some distortion as distance from the spatially rotated windows increases.

Spatial rotation combined with phase growing also gives good results for scaling of texture, as shown in Figs. 4.20 and 4.21. Each part of the figure presents a result in which the lines in *Barbara*'s pants are both rotated and scaled using the pointwise rotation and scaling filters

$$\arg[\nabla\widehat{\varphi}_i(\mathbf{k})] = \arg[\nabla\varphi_i(\mathbf{k})] + \begin{cases} \theta_f, & \arg[\nabla\varphi_i(\mathbf{k})] \in [-\pi/2, 3\pi/16] \\ & \text{and } \left|\nabla\varphi_i(\mathbf{k})\right| \in [\pi/2, 2\pi] \\ & \text{and } \mathbf{k} \in \mathcal{X}, \\ 0, & \text{otherwise,} \end{cases} \tag{4.10}$$

$$\left|\nabla\widehat{\varphi}_i(\mathbf{k})\right| = \left|\nabla\varphi_i(\mathbf{k})\right| \times \begin{cases} \lambda_s, & \arg[\nabla\varphi_i(\mathbf{k})] \in [-\pi/2, 3\pi/16] \\ & \text{and } \left|\nabla\varphi_i(\mathbf{k})\right| \in [\pi/2, 2\pi] \\ & \text{and } \mathbf{k} \in \mathcal{X}, \\ 1, & \text{otherwise,} \end{cases} \tag{4.11}$$

where θ_f is the desired angle of rotation for the filter and λ_s is the desired scaling factor. The pointwise filters are followed by spatial rotation and scaling operations in the same manner as the spatial rotation described for the results in Fig. 4.19. Scaling results for parts (a), (b), and (d) remain relatively clear, while it can be seen that for the higher scaling factor in (c), the texture starts to appear blurry. The reconstructions in Fig. 4.20 (b), (c), and (d) use one less spatially rotated window than those shown in Fig. 4.19 (b). This was done in order to avoid a noticeable phase mismatch that occurs in the middle right pants leg of (a) where phases from different starting points grow into each other. However, this comes with the trade-off that the area becomes noisy instead from the phase growing algorithm. The optimal selection of starting spatial rotation windows remains an open question and, preferably, a method

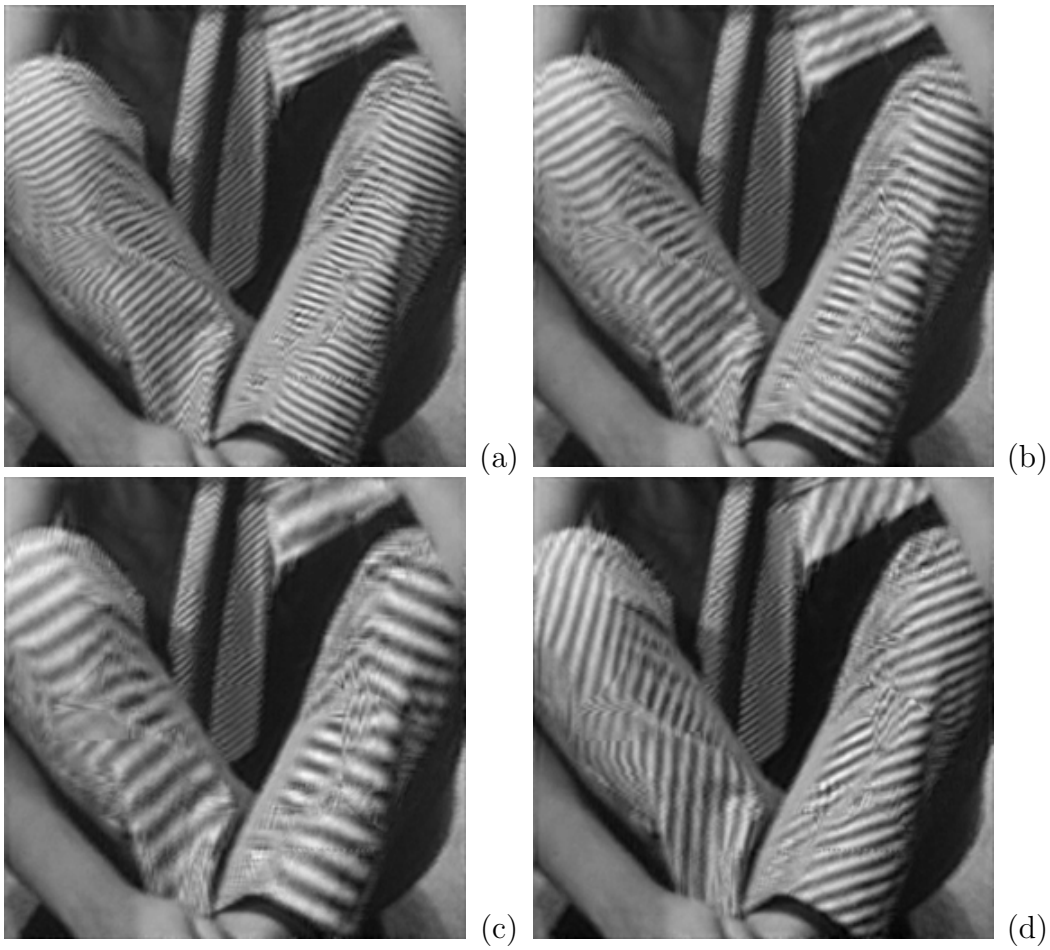

Figure 4.20: Different scaling factors and orientations for specific texture rotation. (a) Rotation of $\pi/2$ and scaling factor of 2. (b) Rotation of $\pi/2$ and scaling factor of 3. (c) Rotation of $\pi/2$ and scaling factor of 5. (d) Rotation of $\pi/4$ and scaling factor of 3.



Figure 4.21: Select results from Fig. 4.20 shown in full *Barbara* image. (a) Original *Barbara* image. (b) Rotation of $\pi/2$ and scaling factor of 5 within ROI. (c) Rotation of $\pi/4$ and scaling factor of 3 within ROI.

for automatically selecting the windows should be developed. Reconstruction results might also be improved by devising a way to make phases from different regions grow into each other more smoothly.

A new filter using the pointwise and spatial rotation methods described above was also designed for the *Boat* image shown in Fig. 4.22 (a). This filter carries out the pointwise rotation

$$\arg[\nabla\widehat{\varphi}_i(\mathbf{k})] = \arg[\nabla\varphi_i(\mathbf{k})] + \begin{cases} -\pi/8, & (\big|\arg[\nabla\varphi_i(\mathbf{k})]\big| - \pi/2) < \pi/4 \\ & \text{and } \big|\nabla\varphi_i(\mathbf{k})\big| \in [\pi/8, 2\pi] \\ & \text{and } \mathbf{k} \in \mathcal{X}, \\ 0, & \text{otherwise} \end{cases} \tag{4.12}$$

with a region of interest $\mathcal{X}$ that is the lower half of the image. The pointwise rotation filter is followed by analagous spatial rotation in manually selected windows. The resulting filtered image, shown in Fig. 4.22 (b), displays rotated ripples on the surface of the water within the ROI, while the reflection of the boat in the water remains clearly observable.

4.6 Cross-Component Amplitude Attenuation

In this section, a new modulation domain filtering technique is presented that is an extension of the previous amplitude modulation attenuation technique discussed in Section 4.4. While previous results applied amplitude modulation attenuation individually for each component, the approach described here makes use of information from all components when applying amplitude modulation attenuation. This is done in order to achieve the image processing goal on targeted texture only while preserving the non-targeted texture in its original state.

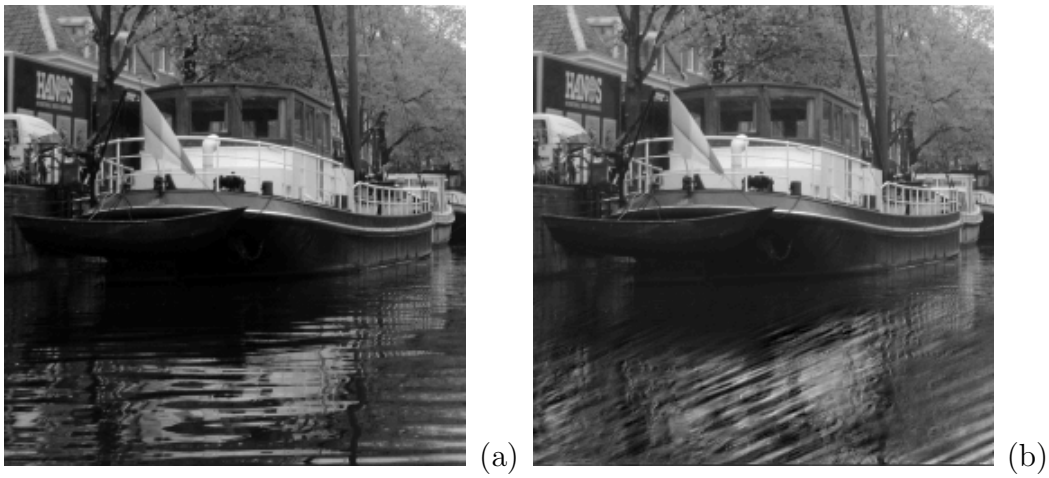

Figure 4.22: Rotation of image texture in *Boat* image. (a) Original image. Filter operations were only applied within the bottom half of the image. (b) Modification of target frequency characteristics to rotate ripples in the water.

After first specifying a region of interest within the image, all pixels within the ROI are analyzed across all components to detect which pixels are candidates for attenuation. Any pixel which matches the frequency characteristics of the target texture to be attenuated becomes a candidate for attenuation. However, any pixel which matches the frequency characteristics of other, non-target textures within the image is set to never be marked as a candidate for attenuation. If such a non-target texture pixel had already been marked as a candidate for attenuation, that candidacy is removed. Once candidate pixels have been identified, the amplitude modulation functions are attenuated on a per-component basis for any candidate pixels which match the frequency characteristics of the target texture to be attenuated.

In this section, the above algorithm will be applied to the visually intuitive task of removing the bands from *Lena*'s hat that is shown in Fig. 4.23. Part (a) of Fig. 4.23 is the previously reported result from [84] for this task,

which was described in Section 2.8.1. Parts (b) and (c) of Fig. 4.23 are reproductions of the task, applying amplitude modulation attenuation without the consideration of information from multiple components. For this filter, the amplitude modulation of each image component was attenuated within a region of interest $\mathcal{X}$ according to [84]

$$\widehat{a}_i(\mathbf{k}) = \begin{cases} 0, & (\arg[\nabla\varphi_i(\mathbf{k})] - \pi/4) < \pi/8 \\ & \text{and } \mathbf{k} \in \mathcal{X}, \\ a_i(\mathbf{k}), & \text{otherwise.} \end{cases} \tag{4.13}$$

The final reconstructed result after attenuation of all components can be seen in Fig. 4.23 (b), where the lower and right boundaries of the region of interest are clearly visible. Fig. 4.23 (c) shows the same attenuation filter with an added frequency magnitude constraint according to

$$\widehat{a}_i(\mathbf{k}) = \begin{cases} 0, & (\arg[\nabla\varphi_i(\mathbf{k})] - \pi/4) < \pi/8 \\ & \text{and } \left|\nabla\varphi_i(\mathbf{k})\right| \in [\pi/8, 2\pi] \\ & \text{and } \mathbf{k} \in \mathcal{X}, \\ a_i(\mathbf{k}), & \text{otherwise.} \end{cases} \tag{4.14}$$

In parts (a) and (b) of Fig. 4.23, the modulation domain filter also has the unintended affect of severely washing out some of the feathers in *Lena*'s hat. In part (c), most high frequency content has been removed from the feathers but their shape is still recognizable due to the unaffected low frequency content. In all parts of Fig. 4.23, the brim of Lena's hat within the ROI has also been removed.

Application of the cross-component amplitude attenuation filter described above is shown in Fig. 4.24. The cross-component amplitude attenuation filter is implemented by first analyzing all components to isolate non-target

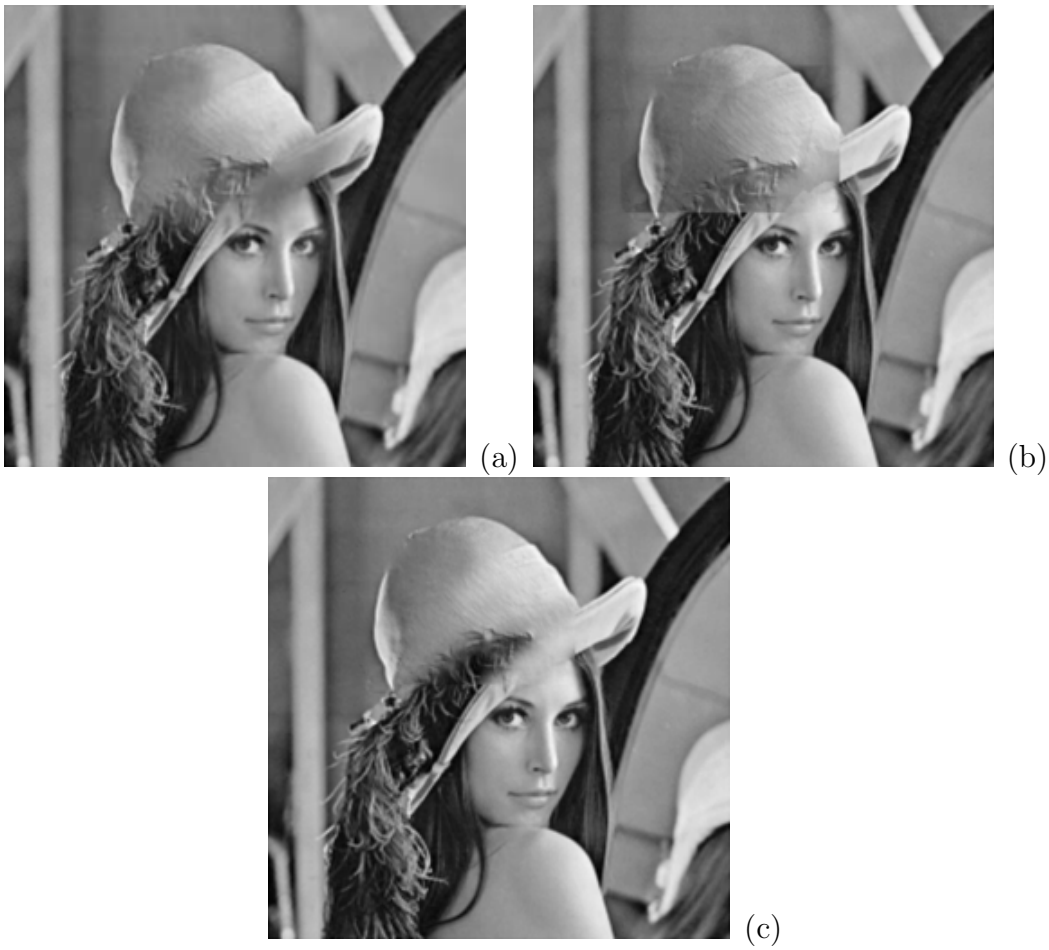

Figure 4.23: Removal of bands in *Lena*'s hat via amplitude attenuation on individual components. (a) Previous result from [84]. (b) Reproduced frequency matching AM attenuation. The boundaries of the ROI are visible. (c) Filter operation from (b) excluding the lowest frequency channels.

texture pixels $\mathcal{X}_{nt}$ according to

$$\mathcal{X}_{nt}(\mathbf{k}) = \begin{cases} 1, & \arg[\nabla\varphi_i(\mathbf{k})] \in [\pi/2, \pi] \\ & \text{or } \mathbf{k} \notin \mathcal{X} \\ & \text{or } \mathcal{X}_{nt}(\mathbf{k}) = 1, \\ 0, & \text{otherwise} \end{cases} \tag{4.15}$$

for all $i \in [1, K]$, where K is the number of image components. Then for each component, the amplitude modulation is attenuated according to

$$\widehat{a}_i(\mathbf{k}) = \begin{cases} 0, & \arg[\nabla\varphi_i(\mathbf{k})] \in [\pi/8, \pi/2] \\ & \text{and } a_i(\mathbf{k}) < 0.1 \\ & \text{and } \mathcal{X}_{nt}(\mathbf{k}) = 0, \\ a_i(\mathbf{k}), & \text{otherwise.} \end{cases} \tag{4.16}$$

In Fig. 4.24 (b) it can be seen that the bands of *Lena*'s hat have been removed with minimal effect on the feathers within the ROI. In this filter, the range of the target frequency characteristics was expanded slightly to $[\pi/8, \pi/2]$ in order to successfully remove more of the bands from the hat. The range of the non-target frequency characteristics was set to $[\pi/2, \pi]$ for the filter.

Parts of the brim of *Lena*'s hat were also found to have a higher amplitude response than the bands of the hat. As a result, parts of the brim of the hat could remain unaffected by the filter by including an amplitude magnitude constraint as a part of the target texture characteristics. For this filter, the target texture was defined as the above noted frequency range and an amplitude magnitude of less than 0.1. The right section of the brim of the hat is still affected by the filter, but looking at the original image in Fig. 4.24 (a) reveals that, to a human observer, the right section of the hat brim in fact looks identical to the bands on the hat. Thus, a higher level of sophistication


(a)

(b)

Figure 4.24: Removal of bands in *Lena*'s hat via cross-component amplitude attenuation. (a) Input *Lena* image. (b) Cross-Component attenuation over all channels. Filter is also constrained by AM magnitude limits.

and complexity would be required in the filter design to retain the hat brim while attenuating the hat bands. Developing such more sophisticated filters is an important topic for future research.

4.7 Cross-Component Scaling of Target Image Texture

The frequency scaling method from Section 4.5 was expanded to include the cross-component non-target texture identification described in Section 4.6 in order to design a modulation domain filter that can effectively modify frequency content without affecting non-target texture. In this section, such an approach is used to apply frequency scaling to the *Steps* image shown in Fig. 4.25 (a). In that image, the people walking up the steps share some frequency characteristics with the steps. However, cross-component scaling can be used to apply scaling to the steps without altering the people in the image.

Similar to the cross-component amplitude attenuation filter described in

Section 4.6, the cross-component texture scaling filter is implemented by first analyzing all components to isolate non-target texture pixels $\mathcal{X}_{nt}$ according to

$$\mathcal{X}_{nt}(\mathbf{k}) = \begin{cases} 1, & ((a_i(\mathbf{k})/\max a_i) > 0.3 \\ & \text{and } \left|\nabla\varphi_i(\mathbf{k})\right| \in [5\pi/32, 2\pi]) \\ & \text{or } \mathbf{k} \notin \mathcal{X} \\ & \text{or } \mathcal{X}_{nt}(\mathbf{k}) = 1, \\ 0, & \text{otherwise} \end{cases} \tag{4.17}$$

for all $i \in [1, K]$, where K is the number of image components. In this filter, the non-target texture signature is specified to be pixels of high amplitude response within a range of high frequency magnitudes. This signature corresponds to the sharp edges of the people relative to the steps and provides a method by which to prevent the people from being scaled by the filter.

Once the non-target texture pixels have been identified, the frequency modulation is pointwise scaled for each component according to

$$\left|\nabla\widehat{\varphi}_i(\mathbf{k})\right| = \left|\nabla\varphi_i(\mathbf{k})\right| \times \begin{cases} \lambda_s, & (\left|\arg[\nabla\varphi_i(\mathbf{k})]\right| - \pi/2) < \pi/4 \\ & \text{and } \left|\nabla\varphi_i(\mathbf{k})\right| \in [\pi/8, 2\pi] \\ & \text{and } \mathcal{X}_{nt}(\mathbf{k}) = 0, \\ 1, & \text{otherwise,} \end{cases} \tag{4.18}$$

where λ_s is the desired scaling factor. As described in the scaling method of Section 4.5, the pointwise scaling operation is followed by a spatial scaling operation within manually selected windows.

The above procedure was applied to the *Steps* image with a region of interest equivalent to the lower half of the image, which contains the darker colored steps. The results for scaling the steps by factors of 0.5 and 2 are shown in Fig. 4.25 (b) and Fig. 4.25 (c) respectively, where it can be seen that by using information from other components the filter is able to modify the



(a)



(b)



(c)

Figure 4.25: Cross-component frequency modulation scaling in *Steps* image. (a) Original image. Filter operations were only applied within the bottom half of the image. (b) Target frequency characteristics scaled by 0.5 to double the number of steps. (c) Target frequency characteristics scaled by 2 to halve the number of steps.

texture of the steps and not the people. The targeted frequency orientation range of the filter also keeps in place edge structure of horizontal frequency between the steps that helps human observers to recognize the scaled texture as steps. While these results do achieve the stated image processing goal to a large degree, the development of methods to increase the sharpness of the results and improve the identification of non-target texture areas are important directions for future work with this class of modulation domain filters.

4.8 Summary

In this chapter, I investigated a variety of phase reconstruction and modulation domain filtering methods. I first presented multiple techniques for reconstruction of phase from frequency modulations. These techniques were presented in terms of their ability to perform perfect or estimated reconstruction in rectangular or arbitrarily shaped regions. I identified a source of reconstruction error that arises from discontinuities in the unwrapped phase and I presented a novel interpolation technique which mitigates the reconstruction error.

I then used the new reconstruction techniques to demonstrate several new modulation domain filters. The ability of the new techniques to provide high-fidelity whole-image rotation, scaling, and translation was demonstrated. New results from modulation domain filters which modify amplitude modulation using non-scaled unwrapped phase were presented. A class of modulation domain filters employing a combination of spatial and pointwise operations was shown to produce effective rotation and scaling of target texture regions in images. The results obtained by this class of filters represent both significant advancement over past approaches to similar image processing goals and

results for newly defined image processing tasks.

New cross-component modulation domain filters were presented for the modification of both amplitude and frequency modulations. These filters isolate both target and non-target texture regions in images, with the non-target texture regions identified by examining modulation information from multiple image components. This allows these new filters to perform more accurate implementation of visually meaningful image processing goals. Results were shown for both cross-component amplitude modulation attenuation and cross-component frequency modulation scaling filters.

Chapter 5

Conclusion and Future Work

Recent years have seen the introduction of the AM-FM transform as a tool to carry out biologically motivated image processing tasks via filtering in the modulation domain. These tasks would be difficult or impossible to achieve by traditional liner filtering. Results in modulation domain filtering have been dramatic, but the development of a general theory for image processing in the modulation domain has been slowed by the difficulties inherent in reconstructing an image from modified frequency modulation functions. Early solutions to the reconstruction problem such those given in [76, 82, 84, 85, 97, 99] required mathematical constructs that led to a loss of the intuitive nature of the AM-FM model. In this dissertation I have developed new techniques for use in the AM-FM transform that focus on enabling the design of intuitive modulation domain filters. I have employed these techniques to develop several new modulation domain filters that are inspired by biological visual perception and that have given new and remarkable results.

5.1 Dissertation Summary

Following the brief introduction in Chapter 1, I presented in Chapter 2 a detailed overview of the theoretical advancements that led to the AM-FM transform and its usage in modulation domain filtering. The development of the

AM-FM model was presented, and its strong link to biological vision systems was discussed. The developments that led to perfect reconstruction from the model and the corresponding AM-FM transform were described in detail. Past results from the usage of the transform to perform modulation domain filtering were presented for filters which modify both amplitude and frequency modulations.

In Chapter 3, I discussed the disadvantages of using a scaling factor during phase unwrapping as related to reconstruction from modified frequency modulations, and I presented a new phase unwrapping algorithm that avoids any scaling of phase. The new algorithm makes use of a smoothly unwrapped phase to guide the unwrapping of non-scaled phase using a queue-based phase growing methodology. I demonstrated that the new phase unwrapping algorithm results in reconstructions that are more accurate than those obtained from phase unwrapping with phase scaling.

Then in Chapter 4, I investigated several new modulation domain filtering techniques. I presented multiple techniques for the reconstruction of phase from frequency modulations, and I presented a novel interpolation technique to improve the fidelity of reconstruction from modified frequency content. I then demonstrated a variety of new modulation domain filters that make use of the new reconstruction techniques and the new phase unwrapping algorithm from Chapter 3. New results were presented for modulation domain filters which modify amplitude modulation and for a new class of modulation domain filters which employs a combination of spatial and pointwise operations to rotate and scale target texture regions. Results were also presented for new cross-component modulation domain filters that operate based on target and

non-target texture regions to modify both amplitude and frequency modulations. The results shown in Chapter 4 represent both improved filtering results relative to past modulation domain filtering efforts and dramatic new modulation domain filtered images.

5.2 Contributions to the Field

The original contributions of this dissertation include the following:

- I developed a new phase unwrapping algorithm that improves the stability of the calculations in the AM-FM transform by eliminating the need for a phase scaling factor in the algorithm. Unwrapped phases given by the algorithm still provide meaningful frequency information and perfect reconstruction in the AM-FM transform.

- I investigated multiple reconstruction techniques that can be used for perfect reconstruction from unmodified frequency modulation functions and estimated reconstruction from modified frequency modulation functions.

- I expanded the functionality of modulation domain filters that modify both amplitude and frequency modulations to identify target and non-target texture regions. Non-target texture is isolated from modulation information of all channel components and improves the ability of the filter to accurately modify only target texture.

- I demonstrated new results for image processing tasks that are inspired by biological visual perception and achieved through modulation domain filtering. The results presented represent both a significant advancement

over past results in the field and new directions in modulation domain image processing.

- I designed and demonstrated practical modulation domain filters that are elegant in that, for the first time, they do not require any departure from the simple and intuitive nature of the basic AM-FM model.

These advances in the state-of-the-art set the stage for the emergence of the AM-FM transform as a highly useful tool for performing intuitive, visually motivated image processing and have laid a foundation for the future development of both the transform and new modulation domain filtering techniques.

5.3 Future Work

Modulation domain filtering is a burgeoning but still nascent field of study with a great deal of potential for future development. Here I will describe a few of the directions that research in this field might take.

Although several different modulation domain filters have now been defined and demonstrated, strict definitions for classes of modulation domain filters are still unknown. A categorical investigation into the fundamental processes that can be achieved through modulation domain filtering and a description of the ways they can be put together would be beneficial to the development and understanding of future modulation domain filters.

The development of new types of filters opens up a related field of study that investigates what visually motivated image processing objectives might be possible or desired. While we can look to traditional image processing tasks for examples of filters and modifications that have been beneficial in the past,

the ability to modify the fundamental structure of an image as a function of its *perceived meaning* is a new capability that is not yet fully understood. The identification of potential ways in which human observers might wish for their perception of an image to change and what ways of changing that perception are possible are open questions that will require consideration. Merely as introductory examples, such tasks might include the replacement of textures on surfaces in images with entirely new complex textures or the enhancement of very specific complex textures, or parts of textures, within an image. It is likely that some potential, desired ways of changing the visual meaning of an image have not yet been imagined.

Several technical challenges remain in modulation domain filtering. The reconstruction algorithms presented in Section 4.1 provide methods by which to perfectly reconstruct phase from rectangular regions and reconstruct estimated phase from arbitrary regions. However, a method by which to perfectly reconstruct phase from an arbitrary region has yet to be developed. The spline based integration methods which are required for perfect reconstruction are not immediately amenable to the path following phase growing algorithm described because the separable nature of the spline representation precludes integration along arbitrary paths. Instead, a new method would need to be developed that is capable of performing row and column integration in such a way that integration is completely carried out over a region of arbitrary shape. Such a method would potentially allow perfect reconstruction of targeted texture regions, which are generally arbitrary in shape.

In Section 4.5, the filtering algorithm described made use of both pointwise and spatial filtering operations. The spatial operations are carried out

within manually selected windows and greatly contribute to the quality of the final reconstruction. While manual placement of the windows has been very useful as a proof of concept, modulation domain filters ideally alter image texture automatically according to a design without any human in the loop. The optimal and automatic placement of these spatial windows is an important topic for future research with several open questions. In the results obtained so far, it has already been noted that there exists a trade-off between the placement of windows which can lead to noticeable differences in texture phase as different reconstruction regions grow together and the non-placement of windows in which those same areas suffer from noisy reconstruction. These obstacles may be ameliorated by careful placement of spatial windows to avoid phase differences or improved phase growing at regional boundaries.

The further investigation of cross-component filtering techniques may lead to new functionality in or expanded classes of modulation domain filters. As demonstrated in Sections 4.6 and 4.7, texture information is often best described in information from multiple AM-FM image components. The relationships between these components and ways to use information from multiple components to better inform modulation domain filters demand further investigation.

In discovering how the visually perceived elements of imagery are constructed well enough to take them apart and put them back together in new ways, we better understand the fundamental nature of how we visually perceive image content. I believe that further investigation into this field of study will yield discoveries which will enhance current image processing technologies used in everyday life and lead to new technologies and applications not yet imagined.

Bibliography

[1] S. T. Acton, D. P. Mukherjee, J. P. Havlicek, and A. C. Bovik, "Oriented texture completion by AM-FM reaction-diffusion," *IEEE Trans. Image Proc.*, vol. 10, no. 6, pp. 885–896, June 2001.

[2] A. Aldroubi, M. Unser, and M. Eden, "Cardinal spline filters: Stability and convergence to the ideal sinc interpolator," *Signal Processing*, vol. 28, no. 2, pp. 127–138, Aug. 1992.

[3] M. Antonini, M. Barlaud, P. Mathieu, and I. Daubechies, "Image coding using wavelet transform," *IEEE Trans. Image Proc.*, vol. 1, no. 2, pp. 205–220, Apr. 1992.

[4] M. J. Bastiaans, "Gabor's expansion of a signal into Gaussian elementary signals," *Proceedings of the IEEE*, vol. 68, no. 4, pp. 538–539, Apr. 1980.

[5] J. M. Bioucas-Dias and G. Valadao, "Phase unwrapping via graph cuts," *IEEE Trans. Image Proc.*, vol. 16, no. 3, pp. 698–709, Mar. 2007.

[6] B. Boashash, "Estimating and interpreting the instantaneous frequency of a signal," *Proceedings of the IEEE*, vol. 80, no. 4, p. B. Boashash, Apr. 1992.

[7] A. C. Bovik, "Analysis of multichannel narrow-band filters for image texture segmentation," *IEEE Trans. Signal Proc.*, vol. 39, no. 9, pp. 2025–2043, Sep. 1991.

[8] A. C. Bovik, M. Clark, and W. S. Geisler, "Multichannel texture analysis using localized spatial filters," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 12, no. 1, pp. 55–73, Jan. 1990.

[9] A. C. Bovik, N. Gopal, T. Emmoth, and A. Restrepo, "Localized measurement of emergent image frequencies by Gabor wavelets," *IEEE Trans. Info. Theory*, vol. 38, no. 2, pp. 691–712, Mar. 1992.

[10] A. C. Bovik, J. P. Havlicek, and M. D. Desai, "Theorems for discrete filtered modulated signals," in *Proc. IEEE Int'l. Conf. Acoust., Speech, Signal Proc.*, vol. 3, Minneapolis, MN, Apr. 27-30, 1993, pp. 153–156.

[11] A. C. Bovik, J. P. Havlicek, D. S. Harding, and M. D. Desai, "Limits on discrete modulated signals," *IEEE Trans. Signal Proc.*, vol. 45, no. 4, pp. 867–879, Apr. 1997.

[12] A. C. Bovik, P. Maragos, and T. F. Quatieri, "AM-FM energy detection and separation in noise using multi-band energy operators," *IEEE Trans. Signal Proc.*, vol. 41, no. 12, pp. 3245–3265, Dec. 1993.

[13] P. A. Campbell, "AM-FM image processing toolbox," Master's thesis, The University of Oklahoma, 2008.

[14] S. Chavez, Q. S. Xiang, and L. An, "Understanding phase maps in MRI: A new cutline phase unwrapping method," *IEEE Trans. Med. Imag.*, vol. 21, no. 8, pp. 966–977, Aug. 2002.

[15] N. H. Ching, D. Rosenfeld, and M. Braun, "Two-dimensional phase unwrapping using a minimum spanning tree algorithm," *IEEE Trans. Image Proc.*, vol. 1, no. 3, pp. 355–365, Jul. 1992.

[16] M. Clark and A. C. Bovik, "Texture discrimination using a model of visual cortex," in *Proc. IEEE Int'l. Conf. Syst., Man, Cybern.*, Atlanta, GA, 1986.

[17] L. Cohen and C. Lee, "Standard deviation of instantaneous frequency," in *Proc. IEEE Int'l. Conf. Acoust., Speech, Signal Proc.*, vol. 4, May 1989, pp. 2238–2241.

[18] ——, "Instantaneous bandwidth for signals and spectrogram," in *Proc. IEEE Int'l. Conf. Acoust., Speech, Signal Proc.*, vol. 5, Albuquerque, NM, Apr. 3-6, 1990, pp. 2451–2454.

[19] J. G. Daugman, "Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters," *J. Opt. Soc. Am. A*, vol. 2, no. 7, pp. 1160–1169, Jul. 1985.

[20] V. DeBrunner, M. Ozaydin, and T. Przebinda, "Resolution in time-frequency," *IEEE Trans. Signal Proc.*, vol. 47, no. 3, pp. 783–788, Mar. 1999.

[21] V. DeBrunner, M. Ozaydin, T. Przebinda, and J. Havlicek, "The optimal solutions to the continuous- and discrete-time versions of the Hirschman uncertainty principle," in *Proc. IEEE Int'l. Conf. Acoust., Speech, Signal Proc.*, vol. 1, Istanbul, Turkey, Jun. 5-9, 2000, pp. 81–84 vol.1.

[22] M. Felsberg and U. Kothe, "Get: The connection between monogenic scale-space and Gaussian derivatives," in *Proc. Scale-Space Conf.*, 2005, pp. 192–203.

[23] D. Feng, N. Wu, and B. Liu, "A region-growing phase unwrapping approach based on local frequency estimation for interferometric SAR," in *Proc. IEEE Int'l. Conf. Sig. Proc.*, vol. 4, 2006.

[24] T. J. Flynn, "Consistent 2-D phase unwrapping guided by a quality map," in *Proc. Int'l. Geoscience and Remote Sensing Symp.*, vol. 4, May 1996, pp. 2057–2059.

[25] G. Fornaro and E. Sansosti, "A two-dimensional region growing least squares phase unwrapping algorithm for interferometric SAR processing," *IEEE Trans. Geoscience and Remote Sensing*, vol. 37, no. 5, pp. 2215–2226, Sep. 1999.

[26] W. T. Freeman and E. H. Adelson, "The design and use of steerable filters," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 13, no. 9, pp. 891–906, Sep. 1991.

[27] D. Gabor, "Theory of communication," *J. Inst. Elect. Eng. London*, vol. 93, no. III, pp. 429–457, 1946.

[28] F. Gianfelici, G. Biagetti, P. Crippa, and C. Turchetti, "Multicomponent AM-FM representations: An asymptotically exact approach," *IEEE Trans. Audio, Speech, and Language Processing*, vol. 15, no. 3, pp. 823–837, Feb. 2007.

[29] J. P. Havlicek, "AM-FM image models," Ph.D. dissertation, The University of Texas at Austin, Dec. 1996.

[30] ——, "The evolution of modern texture processing," *Elektrik, Turkish Journal of Electrical Engineering and Computer Sciences*, vol. 5, no. 1, pp. 1–28, 1997.

[31] J. P. Havlicek and A. C. Bovik, "Multi-component AM-FM image models and wavelet-based demodulation with component tracking," in *Proc. IEEE Int'l. Conf. Image Proc.*, vol. 1, Austin, TX, Nov. 13-16, 1994, pp. 41–45.

[32] ——, "AM-FM models, the analytic image, and nonlinear demodulation techniques," in *Tech. Rept. TR-95-001, Center for Vision and Image Sciences*, The University of Texas at Austin, Mar. 1995.

[33] J. P. Havlicek, A. C. Bovik, M. D. Desai, and D. S. Harding, "The discrete quasi-eigenfunction approximation," in *Proc. Int'l. Conf. on Digital Signal Process.*, Limassol, Cyprus, June 26-28, 1995, pp. 747–752.

[34] J. P. Havlicek, A. C. Bovik, and P. Maragos, "Modulation models for image processing and wavelet-based image demodulation," in *Proc. 26th IEEE Asilomar Conf. Signals, Syst., Comput.*, Pacific Grove, CA, Oct. 26-28, 1992, pp. 805–810.

[35] J. P. Havlicek, D. S. Harding, and A. C. Bovik, "Multi-component signal demodulation and reconstruction using AM-FM modulation models," in *Proc. 1995 IEEE Workshop Nonlin. Signal and Image Process.*, Neos Marmaras, Halkidiki, Greece, June 20-22, 1995, pp. 41–45.

[36] ——, "Reconstruction from the multi-component AM-FM image representation," in *Proc. IEEE Int'l. Conf. Image Proc.*, vol. 2, Washington, DC, Oct. 23-26, 1995, vol. II, pp. 280–283.

[37] ——, "Computation of the multi-component AM-FM image representation," in *Tech. Rept. TR-96-001, Center for Vision and Image Sciences*, The University of Texas at Austin, May 1996.

[38] ——, "Discrete quasi-eigenfunction approximation for AM-FM image analysis," in *Proc. IEEE Int'l. Conf. Image Proc.*, Lausanne, Switzerland, Sep. 16-19, 1996, pp. 633–636 vol.1.

[39] ——, "Extracting essential modulated image structure," in *Proc. 30th IEEE Asilomar Conf. Signals, Syst., Comput.*, Pacific Grove, CA, Nov. 3-6, 1996, pp. 1014–1018 vol.2.

[40] ——, "The multicomponent AM-FM image representation," *IEEE Trans. Image Proc.*, vol. 5, no. 6, pp. 1094–1100, Jun. 1996.

[41] ——, "Multicomponent multidimensional signals," *Multidimensional Syst. and Signal Process.*, vol. 9, no. 4, pp. 391–398, Oct. 1998.

[42] ——, "Multidimensional quasi-eigenfunction approximations and multicomponent AM-FM models," *IEEE Trans. Image Proc.*, vol. 9, no. 2, pp. 227–242, Feb. 2000.

[43] J. P. Havlicek, D. S. Harding, N. D. Mamuya, and A. C. Bovik, "Wideband frequency excursions in computed AM-FM image models," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, Tucson, AZ, Apr. 6-7, 1998, pp. 211–216.

[44] J. P. Havlicek, J. W. Havlicek, and A. C. Bovik, "The analytic image," in *Proc. IEEE Int'l. Conf. Image Proc.*, Santa Barbara, CA, Oct. 26-29, 1997, pp. 446–449.

[45] J. P. Havlicek, J. W. Havlicek, N. D. Mamuya, and A. C. Bovik, "Skewed 2D Hilbert transforms and computed AM-FM models," in *Proc. IEEE Int'l. Conf. Image Proc.*, Chicago, IL, Oct. 4-7, 1998, pp. 602–606 vol.1.

[46] J. P. Havlicek, M. S. Pattichis, D. S. Harding, A. C. Christofides, and A. C. Bovik, "AM-FM image analysis techniques," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, Apr. 8-9, 1996, pp. 195–200.

[47] J. P. Havlicek, J. Tang, S. T. Acton, R. Antonucci, and F. N. Ouandji, "Modulation domain texture retrieval for CBIR in digital libraries," in *Proc. 37th IEEE Asilomar Conf. Signals, Syst., Comput.*, vol. 2, Pacific Grove, CA, Nov. 9-12, 2003, pp. 1580–1584.

[48] R. Ishii and K. Furukawa, "The uncertainty principle in discrete signals," *IEEE Transactions on Circuits and Systems*, vol. 33, no. 10, pp. 1032–1034, Oct. 1986.

[49] C. M. Johnston, N. A. Mould, and J. P. Havlicek, "Multichannel dual domain infrared target tracking for highly evolutionary target signatures," in *Proc. IEEE Int'l. Conf. Image Proc.*, Cairo, Egypt, Nov. 7-11, 2009, pp. 4117–4120.

[50] J. P. Jones and L. A. Palmer, "An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex," *J. Neurophysio.*, vol. 58, no. 6, pp. 1233–1258, 1987.

[51] ——, "The two-dimensional spatial structure of simple receptive fields in cat striate cortex," *J. Neurophysio.*, vol. 58, no. 6, pp. 1187–1211, 1987.

[52] J. P. Jones, A. Stepnoski, and L. A. Palmer, "The two-dimensional spatial structure of simple receptive fields in cat striate cortex," *J. Neurophysio.*, vol. 58, no. 6, pp. 1212–1232, 1987.

[53] B. Julesz, "Textons, the elements of texture perception, and their interactions," *Nature*, vol. 290, pp. 91–97, 1981.

[54] B. Julesz and J. R. Bergen, "Textons, the fundamental elements in preattentive vision and perception of textures," *Bell Syst. Tech. J.*, vol. 62, pp. 1619–1645, 1983.

[55] J. F. Kaiser, "On a simple algorithm to calculate the 'energy' of a signal," in *Proc. IEEE Int'l. Conf. Acoust., Speech, Signal Proc.*, vol. 1, Albuquerque, NM, Apr. 3-6, 1990, pp. 381–384.

[56] ——, "On Teager's energy algorithm and its generalization to continuous signals," in *Proc. 4th IEEE DSP Workshop*, New Paltz, NY, Sep. 1990.

[57] S. B. Kim and Y. S. Kim, "Least squares phase unwrapping in wavelet domain," *IEE Proc. Vis. Image Signal Proc.*, vol. 152, no. 3, pp. 261–267, Jun. 2005.

[58] N. Kitiyanan and J. P. Havlicek, "Modulation domain reference point detection for fingerprint recognition," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, Lake Tahoe, NV, Mar. 28-30, 2004, pp. 147–151.

[59] J. J. Kulikowski, S. Marcelja, and P. . Bishop, "Theory of spatial position and spatial frequency relations in the receptive field of simple cells in the visual cortex," *Biol. Cybern.*, vol. 43, pp. 187–198, 1982.

[60] K. Larkin, "Uniform estimation of orientation using local and nonlocal 2D energy operators," *Optics Express*, vol. 13, no. 20, pp. 8097–8121, Oct. 2005.

[61] S. Lu and P. C. Doerschuk, "Nonlinear modeling and processing of speech based on sums of AM-FM formant models," *IEEE Trans. Signal Proc.*, vol. 44, no. 4, pp. 773–782, Apr. 1996.

[62] S. G. Mallat, "A theory for multiresolution signal decomposition: The wavelet representation," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 11, no. 7, pp. 674–693, Jul. 1989.

[63] P. Maragos and A. C. Bovik, "Image demodulation using multidimensional energy separation," *J. Opt. Soc. Amer. A*, vol. 12, no. 9, pp. 1867–1876, Sep. 1995.

[64] P. Maragos, A. C. Bovik, and T. F. Quatieri, "A multidimensional energy operator for image processing," in *Proc. SPIE Symp. Visual Commun. Image Proc.*, Boston, MA, Nov. 16-18, 1992, pp. 177–186.

[65] P. Maragos, J. F. Kaiser, and T. F. Quatieri, "On separating amplitude from frequency modulations using energy operators," in *Proc. IEEE Int'l. Conf. Acoust., Speech, Signal Proc.*, vol. 2, Mar. 1992, pp. 1–4.

[66] ——, "Energy separation in signal modulations with applications to speech analysis," *IEEE Trans. Signal Proc.*, vol. 41, no. 10, pp. 3024–3051, Oct. 1993.

[67] ——, "On amplitude and frequency demodulation using energy operators," *IEEE Trans. Signal Proc.*, vol. 41, no. 4, pp. 1532–1550, Apr. 1993.

[68] P. Maragos, T. F. Quatieri, and J. F. Kaiser, "Speech nonlinearities, modulations, and energy operators," in *Proc. IEEE Int'l. Conf. Acoust., Speech, Signal Proc.*, vol. 1, May 1991, pp. 421–424.

[69] S. Marcelja, "Mathematical description of the responses of simple cortical cells," *J. Opt. Soc. Amer. A*, vol. 70, pp. 1297–1300, 1980.

[70] R. McAulay and T. Quatieri, "Magnitude-only reconstruction using a sinusoidal speech model," in *Proc. IEEE Int'l. Conf. Acoust., Speech, Signal Proc.*, vol. 9, Mar. 1984, pp. 441–444.

[71] J. A. Moorer, "Signal processing aspects of computer music: A survey," *Proceedings of the IEEE*, vol. 65, no. 8, pp. 1108–1137, Aug. 1977.

[72] N. A. Mould, C. T. Nguyen, and J. P. Havlicek, "Infrared target tracking with AM-FM consistency checks," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, Santa Fe, NM, Mar. 24-26, 2008, pp. 5–8.

[73] C. T. Nguyen, "Modulation domain image processing," Ph.D. dissertation, The University of Oklahoma, May 2012.

[74] C. T. Nguyen, P. A. Campbell, and J. P. Havlicek, "FM filters for modulation domain image processing," in *Proc. IEEE Int'l. Conf. Image Proc.*, Cairo, Egypt, Nov. 7-11, 2009, pp. 3973–3976.

[75] C. T. Nguyen and J. P. Havlicek, "Modulation domain features for discriminating infrared targets and backgrounds," in *Proc. IEEE Int'l. Conf. Image Proc.*, Atlanta, GA, Oct. 8-11, 2006, pp. 3245–3248.

[76] ——, "AM-FM image filters," in *Proc. IEEE Int'l. Conf. Image Proc.*, San Diego, CA, Oct. 12-15, 2008, pp. 789–792.

[77] ——, "Coherent texture decomposition using AM-FM model," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, Austin, TX, May 23-25, 2010, pp. 81–84.

[78] ——, "Modulation domain texture decomposition," in *Proc. IEEE Int'l. Conf. Image Proc.*, Hong Kong, Sep. 26-29, 2010, pp. 2741–2744.

[79] ——, "AM-FM models, partial Hilbert transform, and the monogenic signal," in *Proc. IEEE Int'l. Conf. Image Proc.*, Orlando, FL, Sep. 30 - Oct. 3, 2012, pp. 2337–2340.

[80] ——, "Gradient-based texture cartoon decomposition," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, Santa Fe, NM, Apr. 22-24, 2012, pp. 85–88.

[81] ——, "Linear adaptive infrared image fusion," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, San Diego, CA, Apr. 6-8, 2014, pp. 117–120.

[82] ——, "On the amplitude and phase computation of the AM-FM image model," in *Proc. IEEE Int'l. Conf. Image Proc.*, Paris, France, Oct. 27-30, 2014, pp. 4318–4322.

[83] ——, "Color to grayscale image conversion using modulation domain quadratic programming," in *Proc. IEEE Int'l. Conf. Image Proc.*, Quebec City, Canada, Sep. 27-30, 2015, pp. 4580–4584.

[84] C. T. Nguyen, R. A. Sivley, and J. P. Havlicek, "First results in perceptually-based AM-FM image filtering," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, Santa Fe, NM, Mar. 24-26, 2008, pp. 77–80.

[85] C. T. Nguyen, J. D. Williams, J. P. Havlicek, and M. Ozaydin, "FM processing with generalized amplitude & phase: Application to modulation domain geometric image transformations," in *Proc. IEEE Int'l. Conf. Image Proc.*, Brussels, Belgium, Sep. 11-14, 2011, pp. 81–84.

[86] A. V. Oppenheim and J. S. Lim, "The importance of phase in signals," *Proceedings of the IEEE*, vol. 69, no. 5, pp. 529–541, May 1981.

[87] S. Peddireddy, N. Mould, and J. P. Havlicek, "AM-FM picture carrier beat type noise filters," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, Santa Fe, NM, Apr. 22-24, 2012, pp. 209–212.

[88] M. Porat and Y. Y. Zeevi, "The generalized Gabor scheme of image representation in biological and machine vision," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 10, no. 4, pp. 452–468, Jul. 1988.

[89] ——, "Localized texture processing in vision: Analysis and synthesis in the Gaborian space," *IEEE Transactions on Biomedical Engineering*, vol. 36, no. 1, pp. 115–129, Jan. 1989.

[90] A. Potamianos and P. Maragos, "A comparison of energy operators and the Hilbert transform approach to signal and speech demodulation," *Signal Processing*, vol. 37, no. 1, May 1994.

[91] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C++: The Art of Scientific Computing*, 2nd ed. Cambridge, United Kingdom: Cambridge University Press, 2002.

[92] N. Ray, J. Havlicek, S. T. Acton, and M. Pattichis, "Active contour segmentation guided by AM-FM dominant component analysis," in *Proc. IEEE Int'l. Conf. Image Proc.*, vol. 1, Thessaloniki, Greece, Oct. 7-10, 2001, pp. 78–81, vol.1.

[93] T. C. Rearick, "A texture analysis algorithm inspired by a theory of preattentive vision," in *Proc. Int. Conf. Comput. Vision Patt. Recog.*, San Francisco, CA, 1985.

[94] M. Shook, J. Junger, N. Mould, and J. P. Havlicek, "Quantifying infrared target signature evolution using AM-FM features," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, Austin, TX, May 23-25, 2010, pp. 189–192.

[95] E. P. Simoncelli and W. T. Freeman, "The steerable pyramid: a flexible architecture for multi-scale derivative computation," in *Proc. IEEE Int'l. Conf. Image Proc.*, vol. 3, Washington, DC., Oct. 23-26, 1995, pp. 444–447.

[96] E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, "Shiftable multi-scale transform," *IEEE Trans. Info. Theory*, vol. 38, no. 2, pp. 587–607, Mar. 1992.

[97] R. A. Sivley, "Perfect reconstruction AM-FM image models," Master's thesis, The University of Oklahoma, 2006.

[98] R. A. Sivley and J. P. Havlicek, "Multidimensional phase unwrapping for consistent APF estimation," in *Proc. IEEE Int'l. Conf. Image Proc.*, vol. 2, Genoa, Italy, Sep. 11-14, 2005, pp. 458–461.

[99] ——, "Perfect reconstruction AM-FM image models," in *Proc. IEEE Int'l. Conf. Image Proc.*, Atlanta, GA, Oct. 8-11, 2006, pp. 2125–2128.

[100] ——, "A spline-based framework for perfect reconstruction AM-FM models," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, Denver, CO, Mar. 26-28, 2006, pp. 198–202.

[101] S. M. H. Song, S. Napel, N. J. Pelc, and G. H. Glover, "Phase unwrapping of MR phase images using Poisson equation," *IEEE Trans. Image Proc.*, vol. 4, no. 5, pp. 667–676, May 1995.

[102] U. Spagnolini, "2-D phase unwrapping and phase aliasing," *Geophys.*, vol. 58, no. 9, pp. 1324–1334, Sep. 1993.

[103] ——, "2-D phase unwrapping and instantaneous frequency estimation," *IEEE Trans. Geosci. Remote Sensing*, vol. 33, no. 3, pp. 579–589, May 1995.

[104] J. Strand, T. Taxt, and A. K. Jain, "Two-dimensional phase unwrapping using a block least-squares method," *IEEE Trans. Image Proc.*, vol. 8, no. 3, pp. 375–386, Mar. 1999.

[105] B. J. Super and A. C. Bovik, "Shape from texture using local spectral moments," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 17, no. 4, pp. 333–343, Apr. 1995.

[106] K. Suri and J. P. Havlicek, "Phase algorithm for blocking artifact reduction in reconstructions from analysis-only AM-FM models," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, Denver, CO, Mar. 26-28, 2006, pp. 6–10.

[107] H. Takajo and T. Takahashi, "Least-squares phase estimation from the phase difference," *J. Opt. Soc. Am. A*, vol. 5, no. 3, pp. 416–425, Mar. 1988.

[108] T. Tangsukson and J. P. Havlicek, "Modulation domain image segmentation," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, Austin, TX, Apr. 2-4, 2000, pp. 45–50.

[109] ——, "AM-FM image segmentationn," in *Proc. IEEE Int'l. Conf. Image Proc.*, Vancouver, BC, Canada, Sep. 10-13, 2000, vol. II, pp. 104–107.

[110] P. Tay, J. P. Havlicek, and V. DeBrunner, "Discrete wavelet transform with optimal joint localization for determining the number of image texture segments," in *Proc. IEEE Int'l. Conf. Image Proc.*, vol. 3, Rochester, NY, Sep. 22-25, 2002, pp. 281–284.

[111] ——, "A wavelet filter bank which minimizes a novel translation invariant discrete uncertainty measure," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, Santa Fe, NM, Apr. 7-9, 2002, pp. 173–177.

[112] P. C. Tay and J. P. Havlicek, "Joint uncertainty measures for maximally decimated M-channel prime factor cascaded wavelet filter banks," in *Proc. IEEE Int'l. Conf. Image Proc.*, vol. 1, Barcelona, Spain, Sep. 14-17, 2003, vol. I, pp. 1033–1036.

[113] ——, "Frequency implementation of discrete wavelet transforms," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, Lake Tahoe, NV, Mar. 28-30, 2004, pp. 167–171.

[114] M. Unser, "Splines: A perfect fit for signal and image processing," *IEEE Signal Proc. Mag.*, vol. 16, no. 6, pp. 22–38, Nov. 1999.

[115] M. Unser, A. Aldroubi, and M. Eden, "Fast B-spline transforms for continuous image representation and interpolation," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 13, no. 3, pp. 277–285, Mar. 1991.

[116] ——, "B-spline signal processing: Part I–theory," *IEEE Trans. Signal Proc.*, vol. 41, no. 2, pp. 821–833, Feb. 1993.

[117] ——, "B-spline signal processing: Part II–efficient design and applications," *IEEE Trans. Signal Proc.*, vol. 41, no. 2, pp. 834–848, Feb. 1993.

[118] D. Vakman, "On the analytic signal, the Teager-Kaiser energy algorithm, and other methods for defining amplitude and frequency," *IEEE Trans. Signal Proc.*, vol. 44, no. 4, pp. 791–797, Apr. 1996.

[119] M. Vetterli and C. Herley, "Wavelets and filter banks: Theory and design," *IEEE Trans. Signal Proc.*, vol. 40, no. 9, pp. 2207–2232, Sep. 1992.

[120] H. Voorhees and T. Poggio, "Detecting textons and texture boundaries in natural images," in *Proc. First Intl. Conf. Comput. Vision*, London, England, June 8-11, 1987.

[121] W. Xu and I. Cumming, "A region-growing algorithm for InSAR phase unwrapping," *IEEE Trans. Geoscience and Remote Sensing*, vol. 37, no. 1, pp. 124–134, May 1999.

[122] T. B. Yap, J. P. Havlicek, and V. DeBrunner, "Bayesian segmentation of AM-FM texture images," in *Proc. 35th IEEE Asilomar Conf. Signals, Syst., Comput.*, vol. 2, Pacific Grove, CA, Nov. 4-7, 2001, pp. 1156–1160, vol.2.

[123] T. B. Yap, T. Tangsukson, P. C. Tay, N. D. Mamuya, and J. P. Havlicek, "Unsupervised texture segmentation using dominant image modulations," in *Proc. 34th IEEE Asilomar Conf. Signals, Syst., Comput.*, Pacific Grove, CA, Oct. 29-Nov. 1, 2000, vol. II, pp. 911–915.

[124] H. Zhong, J. Tang, S. Zhang, and M. Chen, "An improved quality-guided phase-unwrapping algorithm based on priority queue," *Geoscience and Remote Sensing Letters, IEEE*, vol. 8, no. 2, pp. 364–368, Mar. 2011.