

BOSTON UNIVERSITY
COLLEGE OF ENGINEERING

Dissertation

**FEATURE-ENHANCED SYNTHETIC APERTURE RADAR
IMAGING**

by

MÜJDAT ÇETİN

B.S., Boğaziçi University, İstanbul, Turkey, 1993

M.S., University of Salford, Manchester, UK, 1995

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

2001

Approved by

First Reader: _____

W. Clement Karl, Ph.D.
Associate Professor of Electrical and Computer Engineering and
Associate Professor of Biomedical Engineering
Boston University

Second Reader: _____

David A. Castañon, Ph.D.
Professor of Electrical and Computer Engineering
Boston University

Third Reader: _____

S. Hamid Nawab, Ph.D.
Associate Professor of Electrical and Computer Engineering and
Associate Professor of Biomedical Engineering
Boston University

Fourth Reader: _____

Janusz Konrad, Ph.D.
Associate Professor of Electrical and Computer Engineering
Boston University

ACKNOWLEDGMENTS

The accomplishments in this work were made possible by the help and support of many individuals. First, and foremost, I would like to express my sincere gratitude to my advisor, Clem Karl, for his continuous guidance, suggestions, patience, and friendship during the course of my work. He has contributed tremendously to the technical content of this dissertation. Perhaps more importantly, he has provided me a perfect example by establishing a productive and enjoyable advisor-student relationship. I was very fortunate to have David Castañon as a continuous resource for my work. David's ability to provide insight on virtually any problem is absolutely astounding. This ability of his helped me tackle numerous hurdles on the way. I would also like to thank Clem and David for providing me continuous financial support as a graduate student. I want to thank the other members of my dissertation committee, Professors Hamid Nawab and Janusz Konrad for a careful evaluation of my work, and their useful feedback, which has significantly improved the quality of this document. I was very fortunate to have Hamid as my teacher, as well. His dedication to, and perfection in, teaching continue to serve as a valuable guide for myself.

I am obviously indebted to my friends at the MDSP lab, with whom I shared a considerable portion of my daily life. I thank John Kaufhold for having more trust in my work than I did, and for being the nice dude he is. I am thankful to Bob Weisenseel for patiently answering every question, that we, the foreigners, had about this country, and also for "volunteering" to be our system administrator (and a good one too). I thank Zhengrong Ying for acting as our online shopping consultant, for his tireless efforts in organizing the MDSP seminars, and for the Chinese goodies. I would like to thank Julia Pavlovich, for lending me her computer for my final defense, and for the delicious Russian chocolates. I thank Jeffrey O'Neill, for his help with my proposal, and for increasing the amount of social interaction in the MDSP lab, during the time he spent here. I am thankful to Ying-Jui Chen, for helping me produce CDs containing experimental results of my work. I would like to thank the

other MDSP members, Lingmin Meng, Haihua Feng, Yonggang Shi, Andrey Litvin, Vincenzo Galdi for many technical and non-technical discussions.

I appreciate the valuable collaboration of Gil Ettinger from Alphatech, in our work on automatic recognition. I would like to acknowledge The Scientific and Technical Research Council of Turkey (TÜBİTAK), for providing the initial financial support for my Ph.D.

My friends outside the MDSP lab were kind enough to keep me as a friend, despite my inability to spend much time with them. I am grateful to Alper Uzun for being an excellent friend for over two decades (yes, that long). I have shared a lot of work and play with Erol Şahin when he was in Boston, and he became one of my best friends. Hüseyin and Vildan Tanrıverdi have always provided food for thought, and enjoyable company. I thank Mutlu Özdoğan (my neighbor) for his support (even in the form of soup and meals). Murat Okatan has always been great in discussing and analyzing many shared questions in our minds. I thank Hikmet Günay for (having to suffer) being my roommate for two years. I would also like to thank Demet Haksever, Murat Canpolat, Izhak Shafran, Kağan Gökbayrak, Murat Kaynar, Ahmet Özcan, Filiz Aslan, Osman Murat Anlı, Rana Nomak and many others for their friendship in the last few years.

I am grateful to my parents, my wife's parents, and our families in Turkey, for their everlasting trust, encouragement, emotional support, and pure love. I thank my mom, dad, and brother for supporting my decision to spend this many years away from them.

I wish to recognize numerous anonymous individuals, whom I met personally, or through books, who have shaped my thoughts and ideals.

Finally, I would like to express my appreciation of my beloved wife, Remle. This work could not have been completed without her unconditional love and support. I am very happy to have her by my side to face many remaining intellectual challenges of life together.

model of the SAR observation process with prior information regarding the nature of the features of interest. We demonstrate the inclusion of prior information through a variety of non-quadratic potential functions. Efficient and robust numerical solution of the optimization problems posed in our framework is achieved through novel extensions of half-quadratic regularization methods to the complex-valued SAR problem.

We have established a methodology for quantitative evaluation of a SAR image formation technique based on recognition-oriented features. Through qualitative and quantitative analyses on large sets of real and synthetic SAR images, we have demonstrated the benefits provided by feature-enhanced imaging. These benefits include increased resolution, ease of region segmentation, sidelobe reduction, and speckle suppression, which are important attributes for automated decision-making.

Furthermore, we have demonstrated the potential of feature-enhanced SAR imaging to improve automated decision-making performance, through classification experiments on automatic target recognition (ATR) systems.

Contents

1	Introduction	1
1.1	The Synthetic Aperture Radar Imaging Problem	1
1.2	Current State of SAR Technology	2
1.3	The Need for Enhanced SAR Image Formation Techniques	5
1.4	Contributions of this Dissertation	7
1.5	Organization	8
2	Principles of Synthetic Aperture Radar	10
2.1	Overview of SAR Technology	10
2.2	Basic Spotlight-Mode SAR Signal Pre-processing	14
2.3	Range and Azimuth Resolution	19
2.4	Relation to Computed Tomography	22
2.5	Examples	25
3	Current SAR Image Reconstruction Methods	31
3.1	Polar Format Algorithm	31
3.2	Filtered Backprojection	32
3.3	Spectral Estimation-based Methods	33

3.4	Data Extrapolation-based Methods	35
3.5	Estimation-Theoretic and Entropy-based Methods	35
3.6	Conclusion	36
4	Regularization Methods in Image Processing	38
4.1	The Need for Regularization	38
4.2	Tikhonov Regularization	41
4.3	Non-Quadratic Regularization	42
5	A Regularization-based Framework for SAR Imaging	46
5.1	Discrete Observation Model	47
5.2	Properties of the SAR Projection Matrix	50
5.3	Objective Function for Feature-Enhanced Imaging	56
5.4	Enhancement of Point-based Features by $\ \mathbf{f}\ _k^k$	59
5.5	Enhancement of Region-based Features by $\ \mathbf{D} \mathbf{f}\ _k^k$	63
5.6	Statistical Interpretation	66
6	Efficient Solution of the Optimization Problem	72
6.1	Quasi-Newton-based Algorithm	73
6.2	Auxiliary Processes and Feature-Enhancement	76
6.3	Ties to Half-Quadratic Regularization	77
6.4	Experimental Results	79
6.4.1	Algorithm Initialization and Parameter Selection	80
6.4.2	Synthetic Scene Reconstructions	81
6.4.3	ADTS Data Reconstructions	85

6.4.4	MSTAR Data Reconstructions	92
6.4.5	URISD Reconstructions	97
6.4.6	Auxiliary Processes	97
6.4.7	Results of Smoothing Real and Imaginary Components	100
6.4.8	Comparison with Post-Processing	101
6.4.9	Bandwidth Extrapolation Property	102
6.4.10	Sensitivity to Initial Conditions	107
6.4.11	Behavior of the Iterative Scheme	107
6.4.12	Computational Complexity	114
7	Image Reconstruction with More General Potential Functions	117
7.1	Feature-Preserving Potential Functions	117
7.2	General Numerical Solution	119
7.3	Discussion on the Choice and Behavior of Potential Functions	120
7.4	Examples	126
7.4.1	Region-based Feature Enhancement	126
7.4.2	Point-based Feature Enhancement	126
8	Generalization to Other Variational Formulations	133
8.1	Variational Formulations based on the Mumford-Shah Energy	133
8.2	A Mumford-Shah-type Formulation for SAR Imaging	135
8.3	The Link to Previous Image Formation Approach	136
8.4	Example	138
9	Evaluation based on Recognition-Oriented Features	141

9.1	Evaluation Criteria for Point-Enhanced Images	142
9.1.1	Target-to-Clutter Ratio	142
9.1.2	Mainlobe Width	142
9.1.3	Peak Matching Accuracy	143
9.1.4	Average Associated Peak Distance	144
9.2	Evaluation Criteria for Region-Enhanced Images	144
9.2.1	Speckle Suppression	144
9.2.2	Segmentation Accuracy	145
9.2.3	Statistical Separability of Regions	146
9.3	Experimental Results	147
9.3.1	Experimental Setup	147
9.3.2	Point-Enhanced Imaging from Full-Resolution Data	149
9.3.3	Point-Enhanced Superresolution Imaging from Reduced-Resolution Data	150
9.3.4	Region-Enhanced Imaging	160
9.4	Summary	163
10	Recognition Tests with Feature-Enhanced Imagery	166
10.1	Template-based Classifier	167
10.2	Likelihood-based Classifier	172
10.3	Point-feature-based Classifier	173
10.4	Experimental Results	174
10.4.1	Experimental Setup	174
10.4.2	Template-based Classification Results	176

10.4.3	Likelihood-based Classification Results	181
10.4.4	Point-feature-based Classification Results	181
10.5	Summary	182
11	Application of Feature-Enhanced Reconstruction in HRR Radar	185
11.1	Introduction to HRR Radar	185
11.2	Superresolution HRR Signal Reconstruction	186
11.3	Examples	188
12	Conclusions and Future Directions	195
12.1	Summary and Conclusions	195
12.2	Topics for Future Research	199
12.2.1	Further Analysis and Evaluation of the Techniques Developed	199
12.2.2	Variations and Enhancements of the Techniques Developed . .	201
12.2.3	Relations and Extensions to Other Domains, Problems, Methods	206
A	Appendix	207
A.1	Discrete 2-D Derivative Operators	207
A.2	Gradient of the Objective Function for Image Reconstruction	210
A.3	Augmented Cost Function for Half-Quadratic Regularization	214
	Bibliography	217
	Vita	231

List of Tables

2.1	Sample SAR system parameters.	27
5.1	Example to demonstrate minimizing ℓ_1 -norm as opposed to ℓ_2 -norm results in the preference of a relatively sparse vector structure.	62
6.1	Computation times and values of the objective function at the solution point for the region-enhanced image reconstruction example of Figure 6.7(c) with various choices of the convergence tolerances.	114
6.2	Computation times and values of the objective function at the solution point for the point-enhanced image reconstruction example of the top row of Figure 6.5(b) with various choices of the convergence tolerances.	115
6.3	Computation times and values of the objective function at the solution point for the point-enhanced, superresolution image reconstruction example of the middle row of Figure 6.5(b) with various choices of the convergence tolerances.	116
6.4	Computation times and values of the objective function at the solution point for the point-enhanced, superresolution image reconstruction example of the bottom row of Figure 6.5(b) with various choices of the convergence tolerances.	116

7.1	Families of potential functions used. k is a parameter determining the shape of the functions. ϵ is a small smoothing constant.	118
7.2	The weighting functions associated with the potential functions. . . .	120
9.1	Average target-to-clutter ratios of images reconstructed from full-resolution data.	151
9.2	Average mainlobe widths of images reconstructed from reduced-resolution (50×50) data.	154
9.3	Average associated peak distances in images reconstructed from reduced-resolution (50×50) data.	156
9.4	Average associated peak distances in the synthetic T72 reconstructions.	156
9.5	Average speckle amplitude in the dB-valued reconstructed images. . .	160
9.6	Average segmentation accuracy for region-enhanced images, measured as the percentage of correctly classified pixels.	161
9.7	Bhattacharyya distances between Gaussian densities characterizing different regions in SAR images.	162
10.1	Composition of the MSTAR data set used in recognition experiments.	175
10.2	Confusion matrices summarizing the template-based classification results. The entry in row i , column j shows the number of images from vehicle type i classified as vehicle j	177
10.3	Confusion matrices summarizing the likelihood-based classification results. The entry in row i , column j shows the number of images from vehicle type i classified as vehicle j	182

10.4 Confusion matrices summarizing the point-feature-based classification results, from reduced-resolution data. The entry in row i , column j shows the number of images from vehicle type i classified as vehicle j . 183

10.5 Overall summary of the classification experiments, in terms of the probability of correct classification, P_{cc} 184

List of Figures

1.1	Simple illustration of data collection by synthetic aperture radar. (Image obtained from the web site of Sandia National Laboratories.) . . .	3
1.2	SAR image of a military vehicle with 0.3 m resolution.	4
1.3	Stages of a typical future SAR system.	5
2.1	Ground-plane geometry for data collection in spotlight-mode SAR. . .	13
2.2	Illustration of the approximation of an arc (which contains points equidistant to the platform at a particular observation angle) by a line.	16
2.3	Graphical representation of an annulus segment containing known samples of the Fourier transform of the reflectivity density.	20
2.4	Data collection geometry in CT.	23
2.5	An abstract illustration of the data support for CT and SAR in a limited angle scenario. The offset SAR data has higher bandwidth in cross-range, which provides a good cross-range resolution despite the limitation in angular diversity.	25

2.6	Simulated phase history data for single point scatterers. (a) Scene containing a point scatterer displaced from the scene center in the range direction. (b) Real part of the phase history data (in a rectangular grid) corresponding to the range-displaced scatterer. (c) Scene containing a point scatterer displaced from the scene center in the cross-range direction. (d) Real part of the phase history data (in a rectangular grid) corresponding to the cross-range-displaced scatterer.	29
2.7	Spatial frequency domain data support for SAR systems with various parameter choices. Top: $\omega_0 = 2\pi \times 10^{10}$ rad/s, $\Delta\theta = 2.3^\circ$. Middle: $\omega_0 = \pi \times 10^{10}$ rad/s, $\Delta\theta = 2.3^\circ$. Bottom: $\omega_0 = \pi \times 10^{10}$ rad/s, $\Delta\theta = 4.6^\circ$	30
2.8	Sample reconstructions of two synthetic point scatterers using the three parameter choices in Figure 2.7. (a) $\omega_0 = 2\pi \times 10^{10}$ rad/s, $\Delta\theta = 2.3^\circ$. (b) $\omega_0 = \pi \times 10^{10}$ rad/s, $\Delta\theta = 2.3^\circ$. (c) $\omega_0 = \pi \times 10^{10}$ rad/s, $\Delta\theta = 4.6^\circ$	30
5.1	The grayscale plot (black corresponds to the maximum value and white to the minimum) of the magnitude of the elements in the projection matrices for CT and SAR for a 32×32 field. (a) CT, $\Delta\theta = 2.3^\circ$. (b) SAR, $f_0 = 10$ GHz, $\Delta\theta = 2.3^\circ$. (c) CT, $\Delta\theta = 23^\circ$. (d) SAR, $f_0 = 1$ GHz, $\Delta\theta = 23^\circ$	51
5.2	The grayscale plot of the magnitude of the elements in the SAR high-resolution-field to low-resolution-data projection matrix for a 32×32 field. $f_0 = 5$ GHz, $\Delta\theta = 2.3^\circ$, pixel-spacing = 0.375 m, resolution = 0.75 m.	52

5.3	Magnitude of the rows of the SAR projection matrices, reshaped as images. These images indicate the points in the scene contributing to a particular data sample at observation angle θ . (a) \mathbf{T}_1 , $\theta = 0.7^\circ$. (b) \mathbf{T}_2 , $\theta = 0.7^\circ$. (c) \mathbf{T}_2 , $\theta = 10^\circ$	54
5.4	Magnitude of the columns of the SAR projection matrices, reshaped as images. These images indicate the contribution of a single pixel in the scene to the entire data. The vertical direction corresponds to the range bins, and the horizontal direction corresponds to the observation angles. Here, a cross-range displaced point scatterer is considered. (a) \mathbf{T}_1 . (b) \mathbf{T}_2	55
5.5	Real parts of the rows of the SAR projection matrices, reshaped as images, for a particular data sample at observation angle θ . (a) \mathbf{T}_1 , $\theta = 0.7^\circ$. (b) \mathbf{T}_2 , $\theta = 0.7^\circ$	56
5.6	Demonstration of the dependence of the cross-range resolving power of SAR on the central frequency, as reflected in the complex-valued structure of the projection matrix. (a) Scene containing a single point scatterer reconstructed from data generated by \mathbf{T}_1 ($\Delta\theta = 2.3^\circ$). (b) Same scene reconstructed from CT-type data generated by the magnitude of \mathbf{T}_1 for the same angular diversity.	57
5.7	Behavior of the function $ f ^k$ for various choices of k	60

5.8	Geometrical figures representing constant ℓ_k -norm points for $k = \{0.5, 1, 2\}$ for a 2-element vector \mathbf{f} . Given all points with a particular ℓ_k -norm, minimization of $\ell_{k'}$ -norms where $k' < k$ yields the coordinate axes as the minimizer (hence a sparse vector structure). On the other hand, given all points with a particular $\ell_{k'}$ -norm, minimization of ℓ_k -norms where $k > k'$ yields a solution where the two coordinates have identical magnitudes (hence a less sparse structure).	61
5.9	Roughness penalty functions: (a) on reflectivities: $\ \mathbf{D}\mathbf{f}\ _k^k$, (b) on reflectivity magnitudes: $\ \mathbf{D} \mathbf{f} \ _k^k$. Top: $k = 2$. Bottom: $k = 1$. (Here \mathbf{f} is real-valued, and has length two. $\mathbf{D} = [1 \ -1]$.)	65
5.10	Gaussian (top) and Laplacian (bottom) probability density functions for a complex-valued scalar f	70
5.11	1-D Gaussian and Laplacian probability density functions for a real-valued random variable.	71
6.1	Quasi-Newton-based iterative algorithm.	76
6.2	Mesh plots of synthetic point scatterer reconstructions. (a) Truth. (b) Conventional method. (c) Proposed method with $k = 0.8$, $\lambda_1 = 1$, $\lambda_2 = 0$. (d) Proposed method with $k = 0.1$, $\lambda_1 = 1$, $\lambda_2 = 0$	82
6.3	Mesh plots of synthetic point scatterer reconstructions from noisy data (SNR = 10 dB). (a) Conventional method. (b) Proposed method with $k = 0.8$, $\lambda_1 = 2$, $\lambda_2 = 0$	83

6.4	Reconstruction of the synthetic target image at the top, from data with SNR = 30 dB (middle row), and SNR = 10 dB (bottom row). (a) Original scene. (b) Conventional method. (c) Proposed method with $k = 1$, $\lambda_1 = 0$, $\lambda_2 = 1.4$. (d) Proposed method with $k = 0.7$, $\lambda_1 = 0$, $\lambda_2 = 1.4$. (e) Conventional method. (f) Proposed method with $k = 1$, $\lambda_1 = 1.4$, $\lambda_2 = 2.5$. (g) Proposed method with $k = 0.7$, $\lambda_1 = 1.4$, $\lambda_2 = 2$.	84
6.5	Enhancement of point-based features. Resolution: top: 0.3 m, middle: 0.6 m, bottom: 1.2 m. (a) Conventional method. (b) Proposed method with $k = 1$, $\lambda_2 = 0$, and top: $\lambda_1 = 14$, middle: $\lambda_1 = 7$, bottom: $\lambda_1 = 10$. (c) Proposed method with $k < 1$ and $\lambda_2 = 0$, with top: $\lambda_1 = 14$, $k = 0.8$, middle: $\lambda_1 = 7$, $k = 0.8$, bottom: $\lambda_1 = 10$, $k = 0.95$.	87
6.6	Enhancement of point-based features from low-quality data (SNR = 10 dB). Resolution: 0.3 m. (a) Conventional method. (b) Proposed method with $k = 1$, $\lambda_1 = 14$, $\lambda_2 = 0$.	88
6.7	Enhancement of region-based features. (a) Conventional method. (b) Tikhonov-type reconstruction (i.e. $k = 2$) with $\lambda_1 = 4$, $\lambda_2 = 22$. (c) Proposed method with $k = 1$, and $\lambda_1 = 5$, $\lambda_2 = 9$. (d) Proposed method with $k = 0.7$, and $\lambda_1 = 4$, $\lambda_2 = 6$.	89
6.8	Enhancement of region-based features from low-quality data (SNR = 8 dB). (a) Conventional method. (b) Proposed method with $k = 0.7$, and $\lambda_1 = 5$, $\lambda_2 = 6$.	90
6.9	Enhancement of region-based features. (a) Conventional method. (b) Tikhonov-type reconstruction (i.e. $k = 2$) with $\lambda_1 = 4$, $\lambda_2 = 22$. (c) Proposed method with $k = 1$, and $\lambda_1 = 4$, $\lambda_2 = 5$. (d) Proposed method with $k = 0.7$, and bottom: $\lambda_1 = 4$, $\lambda_2 = 4$.	91
6.10	Magnitude of the 2-D FFT of a sample MSTAR image.	93

6.11	A 100×100 2-D Taylor window with 35 dB sidelobe suppression, and a quality of approximation of $\bar{n} = 4$	94
6.12	Enhancement of point-based features. Resolution: top: 0.3 m, bottom: 0.6 m. (a) Conventional method. (b) Proposed method with $k = 0.8$, $\lambda_1 = 7$, $\lambda_2 = 0$	95
6.13	Results from data with Taylor windowing removed. Resolution: 0.3 m. (a) Conventional method. (b) Proposed method with $k = 0.8$, $\lambda_1 = 14$, $\lambda_2 = 0$	96
6.14	Enhancement of region-based features. (a) Conventional method. (b) Proposed method with $k = 0.8$, $\lambda_1 = 1.7$, $\lambda_2 = 2.5$	96
6.15	Results with the URISD. (a) CAD model of the fire truck. (b) Conventional reconstruction. (c) Enhancement of point-based features with $k = 0.8$, $\lambda_1 = 22$, $\lambda_2 = 0$. (d) Enhancement of region-based features with $k = 0.8$, $\lambda_1 = 5$, $\lambda_2 = 9$	98
6.16	Auxiliary processes representing the edges in the region-enhanced reconstructions of Figure 6.4. (a) SNR = 30 dB, $k = 1$, $\lambda_1 = 0$, $\lambda_2 = 1.4$. (b) SNR = 30 dB, $k = 0.7$, $\lambda_1 = 0$, $\lambda_2 = 1.4$. (c) SNR = 10 dB, $k = 1$, $\lambda_1 = 1.4$, $\lambda_2 = 2.5$. (d) SNR = 10 dB, $k = 0.7$, $\lambda_1 = 1.4$, $\lambda_2 = 2$	99
6.17	Comparison of smoothness constraints on the real and imaginary parts of a signal versus its magnitude, in the problem of denoising a random-phase signal with piecewise smooth magnitude. Top: noisy observations. Middle row: results of smoothing real and imaginary components. Bottom: results of smoothing the magnitudes directly. Left: magnitude of the signal. Middle column: real part of the signal. Right: imaginary part of the signal.	103

6.18	Comparison of smoothness constraints on the real and imaginary parts of the reflectivities versus their magnitude, for the example of Figure 6.7. (a) Reconstructed field by smoothing real and imaginary components. (b) Reconstructed field by smoothing the magnitudes directly.	104
6.19	Demonstration that simple post-processing techniques do not lead to robust feature enhancement. Top: Superresolution SAR image from Figure 6.5. Bottom: Results of post-processing the corresponding conventional image, based on (6.24), with $k = 0.8$. From left to right: $\lambda_1 = 1.4$, $\lambda_1 = 1.7$, $\lambda_1 = 2.2$	105
6.20	Bandwidth extrapolation property of point-enhanced, superresolution imaging, for the 1.2 m resolution example of Figure 6.5. Top: reconstructed images. Bottom: magnitude of the DFT coefficients (zero frequency at the center of the image). (a) Conventional method. (b) Point-enhanced, $k = 1$. (c) Point-enhanced, $k = 0.8$	106
6.21	Sensitivity of the SAR image reconstruction algorithm to initial conditions. Top: $k = 1$, and $\lambda_1 = 5$, $\lambda_2 = 9$. Bottom $k = 0.7$, and $\lambda_1 = 4$, $\lambda_2 = 6$. Left: adjoint-based initial condition. Right: zero initial condition. Values of the objective function at the solution point: (a) $J^* = 1.7248 \times 10^5$ (b) $J^* = 1.7972 \times 10^5$ (c) $J^* = 1.7671 \times 10^5$ (d) $J^* = 1.8998 \times 10^5$	109
6.22	Evolution of the objective function as a function of the iteration number n	110
6.23	Evolution of the ℓ_2 -norm of the gradient of the objective function as a function of the iteration number n	111
6.24	Evolution of the field $\hat{\mathbf{f}}^{(n)}$ as a function of the iteration number n . . .	112

6.25	Evolution of the field $\hat{\mathbf{f}}^{(n)}$ as a function of the iteration number n , with a zero initial condition.	113
6.26	Demonstration of the tradeoff between computation and accuracy for the superresolution reconstruction example of Figure 6.5(b), bottom row. (a) Reconstruction with $\delta = 10^{-6}$, $\delta_{CG} = 10^{-3}$, with a computation time of 14.40 minutes.(a) Reconstruction with $\delta = 10^{-4}$, $\delta_{CG} = 10^{-2}$, with a computation time of 1.53 minutes.	116
7.1	Families of potential functions used in the prior terms for feature-enhanced image reconstruction.	123
7.2	Scaled weighting functions corresponding to the potential functions in Figure 7.1.	124
7.3	The three potential functions for $k = 0.8$	125
7.4	Edge-preserving, region-enhanced reconstructions with various potential functions. (a) Conventional image. (b) ψ_1 , $p = 0.8$, $\lambda_1^2 = 30$, $\lambda_2^2 = 60$. (c) ψ_2 , $p = 0.8$, $\lambda_1^2 = 30$, $\lambda_2^2 = 70$. (d) ψ_3 , $p = 0.8$, $\lambda_1^2 = 30$, $\lambda_2^2 = 70$	128
7.5	Point-enhanced, superresolution imaging of a synthetic scene. (a) Original scene. (b) Conventional image. (c) ψ_1 , $p = 0.8$, $\lambda_1^2 = 50$. (d) ψ_2 , $p = 0.8$, $\lambda_1^2 = 25$. (e) ψ_3 , $p = 0.8$, $\lambda_1^2 = 25$	129
7.6	Contour plots of the images in Figure 7.5. (a) Original scene. (b) Conventional reconstruction. (c) ψ_1 , $p = 0.8$, $\lambda_1^2 = 50$. (d) ψ_2 , $p = 0.8$, $\lambda_1^2 = 25$. (e) ψ_3 , $p = 0.8$, $\lambda_1^2 = 25$	130
7.7	Point-enhanced, superresolution imaging from 0.6 m resolution data. (a) Conventional image. (b) ψ_1 , $p = 0.8$, $\lambda_1^2 = 50$. (c) ψ_2 , $p = 0.8$, $\lambda_1^2 = 100$. (d) ψ_3 , $p = 0.8$, $\lambda_1^2 = 100$	131

7.8	Point-enhanced, superresolution imaging from 1.2 m resolution data. (a) Conventional image. (b) ψ_1 , $p = 0.8$, $\lambda_1^2 = 150$. (c) ψ_2 , $p = 1$, $\lambda_1^2 = 200$. (d) ψ_3 , $p = 0.8$, $\lambda_1^2 = 150$	132
8.1	Synthetic example. (a) Original scene. (b) Conventional reconstruction from noisy data (SNR=15 dB).	139
8.2	Feature-enhanced reconstructions. (a) Reconstructed field using the objective function in (8.15) with $\lambda_2^2 = 20$. (b) Edge field \mathbf{e} corresponding to (a). (c) Reconstructed field using the Mumford-Shah-based objective function in (8.4) with $\lambda_2^2 = 20$, $\beta = 0.1$, $\lambda_e^2 = 0.4$, $k_e = 1$. (d) Edge field \mathbf{e} corresponding to (c).	140
9.1	Sample pictures of the three types of vehicles whose SAR images are used in the experiments.	148
9.2	128×128 sample images from the MSTAR public target data set. . .	148
9.3	100×100 sample images reconstructed from full-resolution (100×100) data. Top: conventional. Bottom: proposed point-enhanced, with $k = 0.8$, $\lambda_1 = 14$, $\lambda_2 = 0$	150
9.4	100×100 sample images reconstructed from reduced-resolution (50×50) data. Top: conventional. Bottom: proposed point-enhanced, with $k = 0.8$, $\lambda_1 = 14$, $\lambda_2 = 0$	152
9.5	Synthetic T72 image reconstruction example from reduced-resolution data. Left: ground truth. Middle: results from 50×50 data. Right: results from 25×25 data. Top row: conventional. Bottom row: proposed point-enhanced, with $k = 0.8$, $\lambda_2 = 0$, and middle: $\lambda_1 = 14$, right: $\lambda_1 = 7$	153

9.6	Sample peak extraction results for images reconstructed from 50×50 data. Circles indicate “reference” dominant scatterer locations extracted from full-resolution conventional images. Plus signs indicate peaks extracted from the reconstructed images. Top: conventional. Bottom: proposed point-enhanced, with $k = 0.8$, $\lambda_1 = 14$, $\lambda_2 = 0$. . .	155
9.7	Average number of peak matches in images reconstructed from reduced-resolution (50×50) data as a function of the radius of match declaration r	157
9.8	Sample peak extraction results for the synthetic T72 scenes. Circles indicate the scatterer locations in the synthetic scene. Plus signs indicate peaks extracted from the reconstructed images. Left: 50×50 data. Right: 25×25 data. Top: conventional. Bottom: proposed point-enhanced, with $k = 0.8$, $\lambda_1 = 7$, $\lambda_2 = 0$	158
9.9	Average number of peak matches for the synthetic T72 scenes as a function of the radius of match declaration r . Top: 50×50 data. Bottom: 25×25 data.	159
9.10	Region-enhanced image formation and segmentation. Top row: MSTAR images. Second row: region-enhanced reconstructions with $k = 1$, $\lambda_1 = 4$, $\lambda_2 = 6$. Third row: human segmentations. Fourth row: threshold-based segmentations of conventional imagery. Bottom row: threshold-based segmentations of region-enhanced imagery.	164
9.11	Gaussian probability density functions for regions in SAR images. (a) Conventional images. (b) Region-enhanced images.	165

10.1	Conventional templates for the T72 target at 17° depression angle. Each image shows the template for a different aspect angle, starting from 0° on the top left, and covering all 360° with 5° increments, and a 10° window for averaging.	169
10.2	Point-enhanced templates for the T72 target at 17° depression angle. Each image shows the template for a different aspect angle, starting from 0° on the top left, and covering all 360° with 5° increments, and a 10° window for averaging.	170
10.3	Region-enhanced templates for the T72 target at 17° depression angle. Each image shows the template for a different aspect angle, starting from 0° on the top left, and covering all 360° with 5° increments, and a 10° window for averaging.	171
10.4	Conventional templates for the T72 target at 17° depression angle, reconstructed from low-SNR data. Each image shows the template for a different aspect angle, starting from 0° on the top left, and covering all 360° with 5° increments, and a 10° window for averaging.	178
10.5	Point-enhanced templates for the T72 target at 17° depression angle, reconstructed from low-SNR data. Each image shows the template for a different aspect angle, starting from 0° on the top left, and covering all 360° with 5° increments, and a 10° window for averaging.	179
10.6	Region-enhanced templates for the T72 target at 17° depression angle, reconstructed from low-SNR data. Each image shows the template for a different aspect angle, starting from 0° on the top left, and covering all 360° with 5° increments, and a 10° window for averaging.	180

11.1	HRR profiles of fire truck at 0° orientation, reconstructed from X-band data. (a) Reference profile at 0.1 m resolution. (b) Conventional reconstruction from 0.4 m resolution data. (c) Point-enhanced, super-resolution reconstruction from 0.4 m resolution data ($k = 1, \lambda_1^2 = 10$).	191
11.2	Detail of the profiles in Figure 11.1 around the 110th range bin. Conventional reconstruction cannot resolve the two peaks, whereas point-enhanced reconstruction can.	192
11.3	HRR profiles of fire truck at 0° orientation, reconstructed from low SNR X-band data. (a) Reference profile at 0.1 m resolution. (b) Conventional reconstruction from 0.4 m resolution data. (c) Point-enhanced, superresolution reconstruction from 0.4 m resolution data ($k = 1, \lambda_1^2 = 10$).	193
11.4	HRR profiles of M1 tank at 0° orientation, reconstructed from X-band data. (a) Reference profile at 0.1 m resolution. (b) Conventional reconstruction from 0.4 m resolution data. (c) Point-enhanced, super-resolution reconstruction from 0.4 m resolution data ($k = 0.9, \lambda_1^2 = 3$).	194
12.1	Illustration of the idea of incorporating feedback from the final decision-making task in image formation.	203
12.2	Illustration of the multiresolution reconstruction concept for a SAR image. Each scale also corresponds to intermediate results at various stages of image formation.	204

Chapter 1

Introduction

This dissertation presents a new approach to the problem of synthetic aperture radar (SAR) image reconstruction. SAR is a sensor that has become increasingly popular in recent years in a variety of remote sensing applications. The purpose of this chapter is to: 1) introduce the problem addressed in this dissertation; 2) summarize the current state of SAR technology, and within this context discuss the needs for new image formation techniques; 3) provide a concise description of the approach taken in this work by pointing out the main contributions and an outline of the dissertation.

1.1 The Synthetic Aperture Radar Imaging Problem

In the past few decades, remote sensing of man-made objects or natural phenomena on the earth's surface has enjoyed a growing amount of attention due to a variety of emerging military and civilian objectives. Depending on the objective, remote sensing tasks differ in many ways. For example, the ground region of interest may be as small as the area around a military vehicle or as large as a whole continent.

The sensor may be mounted on an airplane flying a few kilometers above ground, or on a satellite in orbit. However, there are some common desirable properties of such a remote sensing device, such as day and night capability and all-weather operation. One of the most promising remote sensing modalities which possesses these properties is the imaging radar, known as synthetic aperture radar (SAR). SAR is an active sensor using its own “illumination,” hence it can operate any time of the day. Since it works in the microwave regime, it avoids weather-related limitations like cloud-cover or rainfall. The resolution achieved by a SAR sensor is largely independent of the distance to the target region of interest.

SAR uses a sensor carried on a platform (aircraft or satellite) which travels along a path transmitting microwave pulses towards the ground, as illustrated in Figure 1.1. Some of the transmitted microwave energy is reflected back towards the sensor where it is received as a signal. This signal first undergoes some pre-processing, involving demodulation. The SAR image formation problem is the problem of reconstruction of a spatial reflectivity distribution of the scene from the pre-processed SAR returns.

1.2 Current State of SAR Technology

Due to the expected key advantages offered by SAR in the realm of remote sensing, many experimental SAR systems have been built [1–3], and the hardware associated with these systems has improved considerably in the last two decades. One example is the Advanced Detection Technology Sensor (ADTS) [3], which is an airborne SAR sensor operated by the MIT Lincoln Laboratory. The SAR image of a military vehicle produced by the ADTS system is shown in Figure 1.2. The resolution of this image is 0.3 m. In order to achieve this level of precision and quality in SAR systems, much effort has been spent on the development and reliable operation of the radar device itself. One major concern has been resolution. Achieving




Figure 1.1: Simple illustration of data collection by synthetic aperture radar. (Image obtained from the web site of Sandia National Laboratories.)

high resolution depends on transmitting a high-frequency, high-bandwidth signal. Because of this fact, considerable effort has been spent on the choice and generation of appropriate waveforms [4]. Another concern has been reliably relating the collected data to physical locations on the ground. A SAR sensor is mounted on a moving platform, and reliable imaging depends on knowing the exact position of the sensor with respect to the imaged ground region at every instant of data collection. As a result, much effort has been spent on motion compensation for SAR, both through the use of inertial navigation systems and through data-driven autofocus techniques [5,6].

Recently, SAR technology has reached a mature level in terms of these concerns, and consequently the data collection has achieved high enough reliability and quality that the demand for using SAR in a variety of applications is increasing. For military purposes SAR is evolving to become an indispensable reconnaissance tool. There are a




Figure 1.2: SAR image of a military vehicle with 0.3 m resolution.

variety of potential earth-science related applications such as mapping and monitoring vegetation and sea-ice, finding minerals, and assessment of environmental damages [7]. NASA has used SAR on a number of missions including the Apollo 17 lunar mission, JPL's SEASAT oceanographic satellite, the Spaceborne Imaging Radar (SIR), and the Magellan mission to Venus [8]. New and interesting application areas for SAR are emerging every day. For example, in [9] the use of SAR for automatic aircraft landing has been proposed. The emergence of lower cost electronics is beginning to make SAR technology economical for small scale uses as well. Based on the current trend, it would not be unrealistic to say that in the very near future SAR systems will be ubiquitous, working on many different kinds of platforms and tasks. Such use of SAR will create an explosion in the amount of collected SAR data. Currently, the coverage rates of an airborne SAR system are capable of exceeding $1 \text{ km}^2/\text{s}$ at a resolution of 1 m^2 , thus producing over one million pixels each second. Furthermore, in most of the applications, the environment is not very structured and cooperating. As a result, the collected data will almost inevitably be noisy, and sometimes be incomplete due to limitations such as observation time. Hence the collected data will carry a *reduced* representation of the underlying scene, rather than a complete one.




Figure 1.3: Stages of a typical future SAR system.

Reduction in quality or quantity of the data will force systems to provide same or improved performance with worse data.

The anticipated high data rates and the time critical nature of emerging SAR tasks motivate the use of automated processing techniques in extracting information from a SAR image for an accurate and efficient interpretation of the scene. There is growing interest in such techniques, wherein features extracted from the formed imagery are used for tasks such as automatic target detection and recognition. With the inclusion of these techniques, most SAR systems will then contain the stages shown in the block diagram of Figure 1.3.

1.3 The Need for Enhanced SAR Image Formation Techniques

The preceding discussion about current directions in the use of SAR in reduced data environments, in conjunction with automated decision-making techniques triggers the question of whether the current signal processing methods used for SAR image formation can face the associated challenges.

As we mentioned before, up to the present most of the efforts in SAR systems have been targeted to the device, for solving “hardware” related issues, and considerable improvements have been achieved. However, in terms of the processing of SAR returns, improvements have only been minor. This is mostly due to the logical way of technology development: the primary concerns were to get the SAR systems

“working”, before using them in advanced tasks. Another observation regarding the SAR problem is that the field has been dominated by researchers in physics, electromagnetics and radar engineering for more than two decades. A signal processing view has been formulated only in the mid 80’s [10–12].

In current SAR systems, image formation is achieved through a Fourier transform-based algorithm. This technique has been developed with the clean and complete, rather than the reduced data situation in mind. Furthermore, this kind of processing does not take into account either any contextual information we have, or the final objectives of the SAR mission regarding the automated decisions or interpretations to be made. The conventional image formation method is basically only data driven and this limits the output quality.

The ADTS SAR image in Figure 1.2 has been formed by such processing. This image has been collected for the task of automatic target recognition (ATR). ATR is aimed at reducing the workload for human operators (e.g. image analysts, pilots) who are tasked with a large scope of other activities. Its need is dictated by large volumes of data requiring analysis and by the short timelines required by target acquisition scenarios [13]. ATR algorithms need to use certain features in the SAR image for recognition. Accurate and efficient extraction of such features from conventional SAR images is complicated by a number of factors. First, the resolution of the formed images is limited by the SAR system bandwidth. This complicates point scatterer localization for ATR tasks. In addition, such images suffer from speckle and sidelobe artifacts. These artifacts complicate region segmentation for shape-based recognition. Furthermore, there is no explicit mechanism in the image formation process to deal with limited quantity or quality of the observations.

Therefore, in order for SAR systems to meet the challenges of future tasks, new processing techniques which are geared towards the final objectives of the mission, and which are robust to reduced data domains are required.

1.4 Contributions of this Dissertation

The first major contribution of this dissertation is the formulation of the SAR imaging problem within a model-based, regularized, feature-enhanced image reconstruction framework. This framework provides mechanisms for achieving robustness to limited and noisy data, for taking into account the observation geometry and parameters, and for enhancing features in the scene which are important for the final objectives of the SAR mission. In this framework, the image formation problem is solved through the minimization of an objective function. This objective function requires a discrete observation model and potential functions for regularization and feature enhancement. For the observation model, we have exploited the tomographic structure of SAR and introduced the concept of a SAR projection matrix. We have demonstrated the use of a variety of potential functions in this framework.

The second major contribution of this dissertation is the extension of a robust regularization technique, known as half-quadratic regularization to complex-valued, random-phase fields. The resulting technique is an effective, computationally efficient iterative algorithm for complex-valued image formation. Although this algorithm has been developed primarily for SAR imaging, the technique is general enough to be used in other signal restoration and reconstruction problems involving complex-valued quantities, such as those arising in inverse scattering.

The third contribution of this dissertation is the extension of our SAR image formation framework to include more general variational formulations, such as the Mumford-Shah function [14] developed for segmentation applications.

The fourth contribution of this dissertation is the demonstration of the application of our SAR image formation approach in signal reconstruction problems arising in another type of radar, known as high range-resolution (HRR) radar.

The fifth contribution of this dissertation is the establishment of a framework for

quantitative evaluation of a SAR image formation technique based on recognition-oriented features. For this purpose, we have proposed new evaluation criteria and utilized previously existing criteria. We have carried out such evaluation experiments for our image reconstruction approach. Based on this analysis, practical outcomes of the feature-enhanced SAR imaging framework developed in this dissertation include superresolution, sidelobe and clutter reduction, speckle suppression, and easy-to-segment regions.

The final contribution of this dissertation is a practical one. We have constructed ATR systems to characterize the end-to-end impact of feature-enhanced SAR imaging, and have demonstrated its potential for improved recognition performance.

Overall, this dissertation proposes a new, effective and efficient framework for inverse problems with complex-valued data. The use of this approach in SAR results in images which lend themselves to robust, accurate, efficient feature-extraction. Consequently, the approach developed here is a promising SAR technology especially for tasks involving automated decision-making.

1.5 Organization

This dissertation is organized as follows. In Chapter 2, we review the principles of SAR. This includes an introduction of the main concepts involved, a mathematical description of the pre-processing of the collected data, the important issue of resolution, and a brief discussion of the similarity of SAR and computed tomography (CT), as well as synthetic examples highlighting some central issues in SAR. After this overview, Chapter 3 describes conventional, as well as recently proposed relevant SAR image reconstruction methods. In Chapter 4, we review regularization methods with an emphasis on their use in inverse problems arising in image processing. This is the final chapter covering background material in this dissertation.

Chapter 5 describes our framework for regularized, feature-enhanced SAR image formation. The chapter contains detailed descriptions of how we build the discrete observation model, and the potential functions for robust regularization, as well as a statistical interpretation of the resulting optimization problems. In Chapter 6, we propose an efficient numerical solution technique for the problem formulated in the previous chapter. This chapter also establishes the tie between this technique and half-quadratic regularization, which is a method that has been used in various domains for robust image restoration and reconstruction. Lastly, the chapter contains experimental results, where various aspects of the numerical scheme and the resulting reconstructions are investigated. Chapter 7 generalizes the approach developed in the previous two chapters to other feature-preserving potential functions, which have successfully been used in image restoration before. Chapter 8 provides yet another extension of feature-enhanced SAR imaging to include a class of objective functions which contain explicit terms for more effective feature preservation. In Chapter 9, we propose criteria for evaluating SAR images in terms of the degree of preservation of recognition-oriented features, and present the results of evaluation experiments performed on a large set of conventional and feature-enhanced images. Chapter 10 contains the description and results of recognition experiments based on three different classifiers to investigate the impact of feature-enhanced imaging on the success of automated decision-making. In Chapter 11, we demonstrate the application of the methods developed for SAR imaging in another sensor, the HRR radar, and illustrate the potential benefits. Finally, Chapter 12 summarizes the results we have obtained, and suggests a number of topics that emerge from this work as potential future research directions.

Chapter 2

Principles of Synthetic Aperture Radar

In this chapter, we provide background information on SAR, that will be useful in the rest of this dissertation. We describe the basic principles for SAR imaging, and present a mathematical description of data collection and pre-processing in SAR. We then discuss the resolution properties of SAR, and the relationship between SAR and computed tomography. Finally, we present some simulated examples to illustrate some of the concepts and issues discussed.

2.1 Overview of SAR Technology

Radars have the ability to distinguish between targets (reflectors) that are separated in range from the radar. Anyone who has shouted across a canyon and listened for the returning echo of his/her own voice has used the basic principle employed by radars for ranging. The concept of echo-ranging simply states that to know an echo signal's round trip flight time and its speed of propagation is to

know the range from the signal source to the target. In radar systems this basic principle is implemented by transmitting high-bandwidth pulses and then using pulse compression techniques [15–17]. In this way, points in a scene can be discriminated based on their distance. In order to visualize or distinguish targets in a 2-D scene however, having resolving power in the range direction only is not sufficient; we also need to be able to separate points in the scene that are at the same distance from the radar, but are in slightly different directions. To achieve such cross-range resolving capability, we - in principle - need to be able to transmit a narrow beam to illuminate a narrow strip of the ground. This problem is essentially the standard aperture problem also encountered in optical imaging. The discriminatory power of an optical system, in all directions in the scene, is proportional to the lens size. In radars, the resolution in the range direction in the scene is not related to the aperture problem, since it is based on echo-ranging as discussed before. On the other hand, cross-range resolution in radars and the resolution in optical systems (which is all cross-range-based) are analogous. In particular, both the cross-range resolution for conventional radars and the resolution for optical systems are on the order of

$$\rho = \frac{\lambda R}{d} \tag{2.1}$$

where λ is the wavelength of the illuminating source, R is the target range, and d is the width of the antenna aperture or the diameter of the lens. Given that the wavelength of microwaves is approximately 10^4 times that of visible light, a conventional radar would have to use an impractically large antenna in order to achieve the same level of resolution as that of an optical system. To demonstrate this, let us consider a radar operating at a nominal range to the ground patch of 50 kilometers. Suppose that we desire a radar image for which the resolution is 1 meter. This level of resolution would be required, for example, in order to count the number of vehicles (trucks, tanks, aircraft, etc.) present in a scene. Let us assume that we use a wavelength λ of

0.03 m, which is a typical wavelength for X-band radar. Then the required width of the antenna on board the aircraft would be

$$d = \frac{\lambda R}{\rho} = \frac{0.03 \cdot (50 \cdot 10^3)}{1} = 1500 \text{ m.} \quad (2.2)$$

The width of the physical antenna required is 1.5 kilometers! Because carrying such a structure on board is clearly impractical, an ordinary radar in this scenario is incapable of attaining the desired cross-range resolution of 1 meter.

Synthetic aperture radar solves this problem by sending multiple pulses from a number of observation points, and then focusing the received information coherently to obtain a high-resolution 2-D description of the scene. Hence it *synthesizes* the effect of a large antenna, using multiple observations from a small antenna. In order to resolve points separated in range, it uses the same techniques as an ordinary radar (which do not depend on aperture size). But its cross-range resolving capability makes it a promising means of constructing high resolution microwave images, using antennas of reasonable size. Actually current SAR systems can produce imagery with a spatial resolution that begins to approach that of remote optical imagers, while avoiding their shortcomings such as nighttime and cloud-cover limitations. Because of that, synthetic aperture radar imaging is rapidly becoming a key technology in the world of modern remote sensing.

There are two distinct modes in which a SAR imaging system can operate: stripmap-mode SAR, and spotlight-mode SAR. In stripmap-mode SAR, the antenna remains fixed with respect to the radar platform so that the antenna beam sweeps out a strip on the ground. In spotlight-mode, the antenna is steered to continuously illuminate a single spot of terrain. We will be focusing on this latter type of SAR. There is extensive literature on stripmap-mode SAR [12, 15, 18], which we will not be concentrating on here. Spotlight-mode SAR is able to provide higher resolution of




Figure 2.1: Ground-plane geometry for data collection in spotlight-mode SAR.

a more limited area than stripmap-mode SAR, because by steering the antenna, the same terrain portion can be observed through a wider range of angles as compared to that in stripmap-mode SAR.

The geometry for data collection in a spotlight-mode SAR is shown in Figure 2.1. The $x - y$ coordinate system (denoting *range* and *azimuth* coordinates respectively) is centered on a relatively small patch of ground illuminated by a narrow RF beam from the moving radar. As the radar traverses the flight path, the radar beam is continuously pointed in the direction of the ground patch. At points corresponding to equal increments of θ (the angle between the x -axis and u -axis in Figure 2.1), high-bandwidth pulses (such as linear FM) are transmitted to the ground patch and echoes are then received and processed.

As we will illustrate in the following sections, demodulated SAR returns at each observation point (after some pre-processing and certain approximations) are related to a particular projectional view of the underlying scene, and the full set of returns provide a band-limited spatial frequency domain description of the scene.

2.2 Basic Spotlight-Mode SAR Signal Pre-processing

In this section, we explain the pre-processing of the received echoes in spotlight-mode SAR, based on the derivation in [11]. Let us model the reflectivity density of the ground patch by the complex function $f(x, y)$, and assume that $f(x, y)$ is constant over the range of frequencies and the range of viewing angles θ employed by the radar. Note that this is an approximation, and there are cases where the dependence of the reflectivity on frequency or aspect angle is important and must be taken into account [19].

The most commonly transmitted pulses in SAR are linear FM chirp signals:

$$s(t) = \begin{cases} e^{j(\omega_0 t + \alpha t^2)}, & |t| \leq \frac{T_p}{2} \\ 0, & \text{otherwise} \end{cases} \quad (2.3)$$

where ω_0 is the carrier frequency and 2α is the chirp rate. Suppose the radar transmits the real part of such a signal, $\Re(s(t))$, at the instant when the angle between the u axis and the x axis in Figure 2.1 is θ . The return signal $\nu_{(\theta, x_0, y_0)}(t)$ from a differential area centered on the point (x_0, y_0) at a distance R_0 from the radar will be

$$\nu_{(\theta, x_0, y_0)}(t) = |f(x_0, y_0)| \cos \left(\omega_0 \left(t - \frac{2R_0}{c} \right) + \alpha \left(t - \frac{2R_0}{c} \right)^2 + \angle f(x_0, y_0) \right) dx dy \quad (2.4)$$

where c is the speed of light, and $2R_0/c$ accounts for the two-way travel time from radar to target. Here, we have completely neglected the effect of propagation attenuation, since it can in practice be compensated for, and therefore is not critical for our analysis. Note that the complex-valued nature of $f(x, y)$ captures both amplitude scaling and phase shifting of the transmitted waveform by the scatterers. The amplitude scaling occurs because only a fraction of the incident radiation is

reflected back to the radar. The phase shift of the reflected wave may be caused by several factors; foremost is the shift at the air/target interface due to the difference between the dielectric constants of air and the target material. The phase shift is also due to the tendency of the RF radiation to creep around target surfaces and its ability to penetrate soft objects and be reflected from within [11]. For most SAR scenes the phase of the reflectivity at a certain location can be modeled to be random, with a uniform probability density, and uncorrelated with the phase at other locations [20].

The return $\nu_{(\theta, x_0, y_0)}(t)$ can be more simply written as

$$\nu_{(\theta, x_0, y_0)}(t) = \Re \left\{ f(x_0, y_0) s \left(t - \frac{2R_0}{c} \right) \right\} dx dy. \quad (2.5)$$

Now let us consider the return from a continuum of scatterers which are at the same distance to the radar. The return from such scatterers will be received by the radar at precisely the same time. Let R be the distance from the radar to the center of the scene¹, and L be the radius of the ground region of interest, as shown in Figure 2.1. Points in the ground patch equidistant from the radar lie on an arc, but for a typical system $R \gg L$, so that this arc is nearly a straight line. This is illustrated in Figure 2.2. This inequality is related to two conditions that must be satisfied, so that we can assume points at the same range lie on a line (i.e. so that curvature of the wavefront can be neglected). First, the range error due to this assumption for any point in the ground patch must be less than a resolution cell:

$$\frac{L^2}{2R} < \frac{c}{2B} = \rho_x, \quad (2.6)$$

where ρ_x is the range resolution (to be introduced in Section 2.3), and B is the bandwidth of the transmitted waveform. Second, the range error due to this assumption

¹Note that the value of R actually depends on the particular observation angle θ , however we suppress this dependence for the sake of notational simplicity.




Figure 2.2: Illustration of the approximation of an arc (which contains points equidistant to the platform at a particular observation angle) by a line.

at a particular point must not vary much through the aperture:

$$\frac{L^2 \sin(2\theta_{max})}{R} \ll \frac{c}{2\omega_0}, \quad (2.7)$$

where θ_{max} is the maximum look angle. The derivations of these conditions can be found in [11]. We will assume that the combined return from such an “equidistant” set of scatterers is the sum of the returns that would be received from each individual scatterer. This assumption of superposition is a common and reasonable one, about which the discussion in [21] states: “When an arbitrary target complex is illuminated by a radar, the backscattered signal will consist of the superposition of the returns from a multitude of scattering centers.” Let us take $q_\theta(u)$ to be such a sum of reflectivities (i.e. a line integral) at distance $R + u$ to the radar, from observation angle θ . Then, we can write the relationship between the projection $q_\theta(u)$ and the field $f(x, y)$ as [22]:

$$q_\theta(u) = \iint_{x^2+y^2 \leq L^2} \delta(u - x \cos \theta - y \sin \theta) f(x, y) dx dy \quad (2.8)$$

This is the standard Radon transform. With this definition, the return signal from a differential line of scatterers normal to the u axis at $u = u_0$ is given by

$$\nu_{(\theta, u_0)}(t) = \Re \left\{ q_\theta(u_0) s \left(t - \frac{2(R + u_0)}{c} \right) \right\} du. \quad (2.9)$$

This is the contribution to the received signal of all scatterers at range $R + u_0$. Then the return from the entire ground patch (which is what the sensor actually receives) at observation angle θ is given by the integral of $\nu_{(\theta, u)}$ over u

$$\nu_\theta(t) = \Re \left\{ \int_{-L}^L q_\theta(u) s \left(t - \frac{2(R + u)}{c} \right) du \right\}. \quad (2.10)$$

Taking into account that $s(\cdot)$ is a chirp pulse, we have

$$\nu_\theta(t) = \Re \left\{ \int_{-L}^L q_\theta(u) \exp \left\{ j \left[\omega_0 \left(t - \frac{2(R + u)}{c} \right) + \alpha \left(t - \frac{2(R + u)}{c} \right)^2 \right] \right\} du \right\} \quad (2.11)$$

on the interval

$$-\frac{T_p}{2} + \frac{2(R + L)}{c} \leq t \leq \frac{T_p}{2} + \frac{2(R - L)}{c}. \quad (2.12)$$

Letting $\tau_0 = 2R/c$ be the round-trip delay to the center of the ground patch and mixing (multiplying) $\nu_\theta(t)$ with the reference chirp

$$\exp \left[-j \left(\omega_0(t - \tau_0) + \alpha(t - \tau_0)^2 \right) \right] \quad (2.13)$$

and then low-pass filtering yields the complex signal²

$$\bar{r}_\theta(t) = \int_{-L}^L q_\theta(u) \exp \left\{ j \frac{4\alpha u^2}{c^2} \right\} \cdot \exp \left\{ -j \frac{2}{c} (\omega_0 + 2\alpha(t - \tau_0)) u \right\} du. \quad (2.14)$$

So $\bar{r}_\theta(t)$ is the demodulated observation signal at platform position θ , as a function of

²We neglect a constant factor of 1/2 here.

time. In practice, the mixing operation described above is carried out by multiplying $\nu_\theta(t)$ with the *in phase* and *quadrature* (i.e. real and imaginary) components of the reference chirp separately. Also note, we assume here that τ_0 is known. In practice it is only known imperfectly and this makes it necessary to have a post-processing technique in SARs known as autofocus or automatic phase-error correction [17].

We will assume that the effect of the quadratic phase term $\exp\{j4\alpha u^2/c^2\}$ in (2.14) can be neglected. This is a reasonable approximation for most situations, since usually it is true that $4\alpha u^2 \ll c^2$. For a more detailed analysis of this approximation, see [17]. The observed signal, after this approximation is given by:

$$\begin{aligned} r_\theta(t) &= \int_{-L}^L q_\theta(u) \exp \left\{ -j \frac{2}{c} (\omega_0 + 2\alpha(t - \tau_0)) u \right\} du \\ &= \int_{-L}^L q_\theta(u) \exp \{ -j \Omega(t) u \} du. \end{aligned} \quad (2.15)$$

This signal can be identified as the Fourier transform of the projection $q_\theta(u)$ where the spatial frequency variable is $\Omega(t) = \frac{2}{c}(\omega_0 + 2\alpha(t - \tau_0))$. Note $\Omega(t)$ is limited to a finite spatial frequency interval, because the observation duration t is limited, and the chirp rate α is finite (equivalently $s(t)$ is narrow-band). Also $\Omega(t)$ is offset from the origin of the spatial frequency plane due to ω_0 . In summary, at least within the time interval considered, the processed return signal $r_\theta(t)$ carries band-pass information related to a particular line integral of the reflectivity field.

Now, let us derive the relationship between the field $f(x, y)$ and the demodulated observed signal $r_\theta(t)$. To this end, we substitute (2.8), in the observation relationship (2.15), to obtain

$$\begin{aligned}
r_\theta(t) &= \int_{|u| \leq L} \iint_{x^2+y^2 \leq L^2} \delta(u - x \cos \theta - y \sin \theta) f(x, y) \exp \{-j\Omega(t)u\} dx dy du \\
&= \iint_{x^2+y^2 \leq L^2} f(x, y) \exp \{-j\Omega(t)(x \cos \theta + y \sin \theta)\} dx dy. \tag{2.16}
\end{aligned}$$

Hence, $r_\theta(t)$ is a finite (i.e. band-limited) slice at angle θ from the 2-D Fourier transform of the field $f(x, y)$. Here $\Omega(t)$ serves as the radial spatial frequency. So, we have two interpretations for $r_\theta(t)$: a 1-D Fourier transform of the projections (based on (2.15)), and a slice through the 2-D Fourier transform of the field (based on (2.16)). This equivalence is essentially a band-limited version of the projection slice theorem [22] from computed tomography (CT). The data $r_\theta(t)$ collected from all observation angles, are usually called the phase histories.

2.3 Range and Azimuth Resolution

Since each return $r_\theta(t)$ is a radial slice at a particular angle in the frequency domain, the complete data from a pre-specified diversity of angles (i.e. full-aperture) lie in an annular region like that shown in Figure 2.3. Note that at this stage of our development, the data are still continuous in the radial direction, nevertheless we present a discrete illustration in Figure 2.3.

Now, let us try to relate the achievable resolution of the SAR image to the dimensions of this annulus, which is determined by system parameters. A definition of resolution in the image domain can be motivated by assuming this annulus can be approximated by a rectangle of width $\Delta\Omega_x$ and height $\Delta\Omega_y$. If we consider a point reflector in the scene and compute its Fourier transform limited to this rectangular region, and then compute an inverse Fourier transform, the output will be a two-




Figure 2.3: Graphical representation of an annulus segment containing known samples of the Fourier transform of the reflectivity density.

dimensional sinc function. The wider the support of the rectangular region, the narrower the mainlobe of this sinc, and hence the better the resolution of the formed image will be. More explicitly, the first zero crossings of the sinc occur at $2\pi/\Delta\Omega_x$ and $2\pi/\Delta\Omega_y$. Therefore, as a rough guide, the resolution of two point reflectors having equal reflectivity requires that the reflectors be separated by more than $\rho_x = 2\pi/\Delta\Omega_x$ in the x (range) dimension and $\rho_y = 2\pi/\Delta\Omega_y$ in the y (cross-range) dimension.³

First, we will consider range resolution. Let us take the width of the rectangle, $\Delta\Omega_x$, to be equal to the radial width of the annulus, which is essentially the spatial frequency bandwidth of each return. The lower and upper limits of the radial extent can be determined by substituting the limits for the observation time t from (2.12) into the definition of $\Omega(t)$ to find the lowest and highest spatial frequencies:

³Note that we leave the possible coherence between the two point reflectors out of consideration here.

$$\begin{aligned}\Omega_{x_l} &= \frac{2}{c} \left(\omega_0 - \alpha T_p + \frac{4\alpha L}{c} \right) \\ \Omega_{x_h} &= \frac{2}{c} \left(\omega_0 + \alpha T_p - \frac{4\alpha L}{c} \right).\end{aligned}\tag{2.17}$$

For a typical spotlight-mode SAR, we have $T_p \gg 4L/c$. So we can deduce that

$$\Delta\Omega_x = \Omega_{x_h} - \Omega_{x_l} \approx \frac{4\alpha T_p}{c} = \frac{4\pi B}{c},\tag{2.18}$$

where we have used the fact that the bandwidth of the transmitted pulse (in Hz) is given by $B = \alpha T_p/\pi$.

Next, we will consider the cross-range resolution, which is determined by $\Delta\Omega_y$. Let us use the geometry to find $\Delta\Omega_y$. From Figure 2.3:

$$\sin\left(\frac{\Delta\theta}{2}\right) \approx \frac{\Delta\Omega_y/2}{\Omega_0}.\tag{2.19}$$

Here $\Omega_0 = 2\omega_0/c$. Hence for $\Delta\theta \ll 1$, we have

$$\Delta\Omega_y \approx \frac{2\omega_0\Delta\theta}{c}.\tag{2.20}$$

Finally, noting that the wavelength of the transmitted pulse is given by $\lambda = 2\pi c/\omega_0$, the following range and cross-range resolution relationships for the system can be obtained:

$$\rho_x \approx \frac{c}{2B}\tag{2.21}$$

$$\rho_y \approx \frac{\pi c}{\omega_0\Delta\theta} = \frac{\lambda}{2\Delta\theta}\tag{2.22}$$

Hence the resolution in the range direction, ρ_x , depends on the bandwidth of the pulse used for transmission, whereas the azimuth resolution, ρ_y , depends on the angular diversity of observations and the central frequency (equivalently wavelength) of the transmitted chirp.

Here we have used the classical Rayleigh resolution criterion in our discussion, which is the most commonly used criterion for SAR imaging [17]. For a survey of alternative resolution criteria used in various applications, please see [23].

Before leaving the discussion on resolution, there is one final point we want to mention. In the radar community, cross-range resolving properties of SAR has usually been formulated with reference to Doppler shifts. This might lead to some confusion, and needs clarification. As the aircraft moves with respect to the target patch, the received signal undergoes a slight Doppler shift which varies according to the observation angle. The SAR imaging equations can be derived as a function of either the Doppler shift or the underlying change in viewing angle. It is important to emphasize, however, that the imaging principle employed in spotlight mode SAR is tomographic, rather than Doppler-based. That is, although the radar antenna must be moved from point to point to obtain different viewing angles, successful imaging is not dependent on a difference in relative velocity between the antenna and ground patch during pulse transmission and reception. As far as the imaging mechanism is concerned, the aircraft could completely stop at each transmission point in space and the SAR would still work properly [11].

2.4 Relation to Computed Tomography

Another imaging modality which constructs high-resolution images by processing data obtained from many different perspective views of a target area is computed tomography (CT). The CT scan is an X-ray technique which enables the imaging of

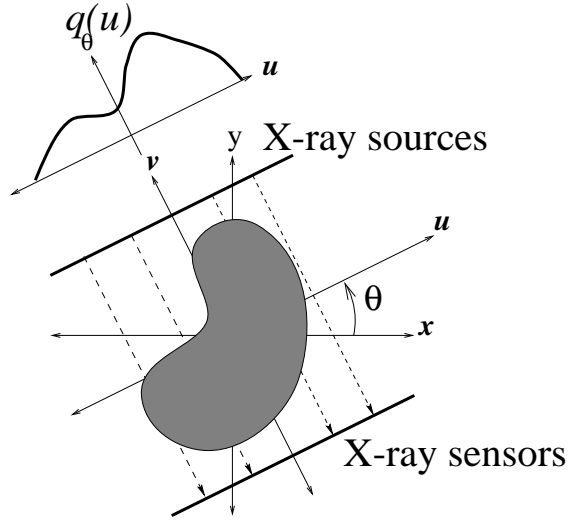


Figure 2.4: Data collection geometry in CT.

two-dimensional cross sections of solid objects. It has been shown that spotlight-mode SAR can be considered as a narrow-band version of CT [11]. Although this similarity has been known for some time now, its exploitation for the SAR problem has only been limited to unrealistic cases, e.g. involving real-valued fields. Throughout our work, we aim to exploit this close relationship between the two imaging modalities, and carry our tomographic insights to the SAR problem. However we should note, due to the differences in the observation geometries, the portion of the electromagnetic spectrum used for sensing, and the properties of the underlying physical objects present in the two modalities and also due to the fact that SAR is a coherent sensor, SAR presents itself as a demanding cousin of CT, rather than its exact replica.

The observation geometry for CT is shown in Figure 2.4. The relationship between the CT observations $q_\theta(u)$ and the corresponding underlying field $f(x, y)$ is again governed by the integral equation given in (2.8). Note that in an actual system the observations would be band-limited, sampled versions of these projections.

A major difference between the two modalities is the following. The SAR data obtained through an angular diversity of a few degrees are usually sufficient to reconstruct a high cross-range resolution SAR image. On the other hand, a practical discrete-data CT system requires a much larger angular range to form a satisfactory image. To better understand this difference, note that, if we define resolution in relation to the width of the mainlobe of the impulse response function as we did in Section 2.3, then resolution is purely a function of the bandwidth available in Fourier space. In order to compare the two modalities in terms of bandwidth, note that the SAR frequency region is offset from the origin as in Figure 2.3, but the CT frequency region is not. If we treat the Fourier space in Figure 2.3 as a circular pie, the observation angles determine the sides of a slice from the pie. The pie slice widens as we move away from the center, that is as we increase frequency. A wider piece of the slice means wider frequency support, hence better spatial resolution. Since SAR operates with a high offset frequency, sufficient bandwidth can be achieved with a small angular diversity. However, since CT operates at baseband, it must subtend a much wider range of viewing angles, in order to have a comparable bandwidth. Hence, observing projectional information at a high frequency rather than at the baseband is really the key behind our ability to obtain high cross-range resolution in SAR with a range of look angles of only a few degrees [24]. This notion is illustrated in Figure 2.5.

Independent of the discussion on bandwidth (or resolution), however there's the issue of image quality. It might seem intriguing that a Fourier-offset imaging system like SAR can produce high quality images. Indeed high-quality reconstructions of real-valued images are generally impossible from offset Fourier data. For example, a high-pass filter (an example of Fourier-offset reconstruction) applied to a real image produces a scene of poor quality consisting mainly of edges. Note that the field (attenuation) in CT is real-valued, but the field (reflectivity) in SAR is complex-valued. It has been shown in [20] that high-quality reconstructions from frequency-

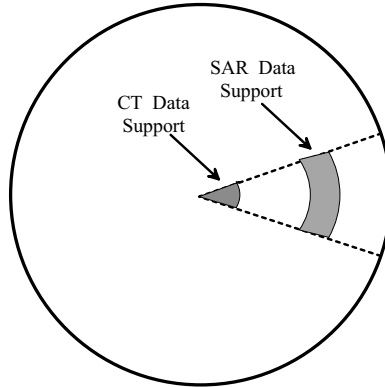


Figure 2.5: An abstract illustration of the data support for CT and SAR in a limited angle scenario. The offset SAR data has higher bandwidth in cross-range, which provides a good cross-range resolution despite the limitation in angular diversity.

offset data are possible for complex-valued fields so long as the phase of the field is random and spatially uncorrelated. It has been argued this is due to the fact that the uncorrelated phase will help distribute information about the magnitude of the field all over the frequency plane, and even a high frequency portion will carry magnitude information. Hence the quality of SAR reconstructions from Fourier-offset data is related to the random phase nature of the underlying fields.

2.5 Examples

We will present a number of simple synthetic examples to illustrate some basic properties of SAR. This includes the demonstration of the structure of the simulated sampled phase history data for point scatterers, and the effect of system parameters on the resolution of the images reconstructed from such data.

Figure 2.6 shows two scenes with point scatterers displaced from the scene center in the range and cross-range directions respectively, and the corresponding phase histories. For the phase histories, we plot only the real parts. The observation angles for this example are centered at 0° , and cover a width of 22° . This is a much larger

angular span than that will be used in most synthetic and real examples in this dissertation, however we would like to make the role of observation angles noticeable in our simulated phase histories here. First, consider the scene in Figure 2.6(a). The scatterer in this scene is located entirely in the range dimension of the imaged patch. If we took a rectangular 2-D discrete Fourier transform (DFT) of this image (using the scene center as the origin), we would obtain a complex sinusoid with frequency only in the range dimension. The SAR phase history data should be similar. However, as the observation angle in SAR changes, the range of the scatterer as measured by the sensor also changes slightly. This results in the type of oscillations shown in the phase history of Figure 2.6(b). Note that we created the phase history corresponding to this scatterer location and placed the sampled returns in a rectangular array on a pulse-by-pulse basis, without any polar-to-rectangular resampling here. Next, consider the cross-range-displaced scatterer in Figure 2.6(c). The corresponding phase history in Figure 2.6(d) has a predominant sinusoidal structure in the vertical (i.e. cross-range) dimension.

Now, let us examine the effect of system parameters on the data support in the spatial frequency domain, and consequently on the resolution of the reconstructed images. Table 2.1 contains a set of parameters we will use in this example, as well as in some of our other spotlight-mode SAR system simulations in this dissertation. With these parameters, the bandwidth of the transmitted signal is 0.4 GHz, implying a range resolution of 0.375 m. If we choose to set the azimuth resolution equal to the range resolution, this requires an angular observation interval of 2.3 degrees. The angular measurements can be taken symmetrically around $\theta = 0^\circ$, or at another angle. In the former case the SAR would be operating in a *side-looking mode*, and in the latter case it would be operating in a *squint mode*. We will take them around $\theta = 0^\circ$. The data sampling rate should be chosen according to $f_s \geq 4L\alpha/(\pi c)$ to prevent aliasing in the reconstructed image.

carrier frequency (ω_0)	$2\pi \times 10^{10}$ rad/s
chirp rate (2α)	$2\pi \times 10^{12}$ rad/s ²
pulse duration (T_p)	4×10^{-4} sec.
angular range ($\Delta\theta$)	2.3°

Table 2.1: Sample SAR system parameters.

With the above choice of parameters, the data lie in the spatial-frequency region shown in the top plot in Figure 2.7. Since the angular observation range is quite small in this example, the annulus is very close to a rectangle. In the same figure, we show the frequency support with two different parameter choices for comparison. When we decrease the chirp carrier frequency ω_0 to one half its original value, keeping everything else the same, the data samples lie in the region shown in the middle plot. Obviously this system would suffer from lower azimuth (cross-range) resolution. Now, if we want to gain back the original azimuth resolution with this choice of ω_0 , the angular observation range can be doubled, which results in the data support shown in the bottom plot.

Finally, let us observe the relationship between data support and resolution through some sample reconstructions. For image reconstruction, we will use the conventional techniques, as will be described in Chapter 3. Figure 2.8(a) shows the reconstruction of two point scatterers (separated in cross-range) by a simulated SAR system which has the data support shown in the uppermost plot in Figure 2.7. The bandwidth in this case is enough to resolve the scatterers. The reduced cross-range bandwidth (middle plot in Figure 2.7) on the other hand, is not sufficient to distinguish these scatterers, as shown in Figure 2.8(b). The reduction in resolution is due to the lower central frequency used by the SAR system in this case. Operating at that frequency, the required cross-range bandwidth can be recovered through a wider range of observation angles as previously shown in the bottom plot of Figure 2.7.

Consequently, the reconstruction with this increased angular observation interval, can resolve the scatterers, as shown in Figure 2.8(c).




Figure 2.6: Simulated phase history data for single point scatterers. (a) Scene containing a point scatterer displaced from the scene center in the range direction. (b) Real part of the phase history data (in a rectangular grid) corresponding to the range-displaced scatterer. (c) Scene containing a point scatterer displaced from the scene center in the cross-range direction. (d) Real part of the phase history data (in a rectangular grid) corresponding to the cross-range-displaced scatterer.




Figure 2.7: Spatial frequency domain data support for SAR systems with various parameter choices. Top: $\omega_0 = 2\pi \times 10^{10}$ rad/s, $\Delta\theta = 2.3^\circ$. Middle: $\omega_0 = \pi \times 10^{10}$ rad/s, $\Delta\theta = 2.3^\circ$. Bottom: $\omega_0 = \pi \times 10^{10}$ rad/s, $\Delta\theta = 4.6^\circ$.




Figure 2.8: Sample reconstructions of two synthetic point scatterers using the three parameter choices in Figure 2.7. (a) $\omega_0 = 2\pi \times 10^{10}$ rad/s, $\Delta\theta = 2.3^\circ$. (b) $\omega_0 = \pi \times 10^{10}$ rad/s, $\Delta\theta = 2.3^\circ$. (c) $\omega_0 = \pi \times 10^{10}$ rad/s, $\Delta\theta = 4.6^\circ$.

Chapter 3

Current SAR Image

Reconstruction Methods

The problem of SAR image reconstruction is to obtain an estimate of the reflectivity density $f(x, y)$ based on the observed, pre-processed SAR data. In this chapter, we briefly discuss conventional, as well as recently proposed relevant SAR image reconstruction methods.

3.1 Polar Format Algorithm

As discussed in Section 2.2, the spotlight-mode SAR observation kernel consists of a band-pass Fourier transform of the reflectivity field to be imaged. As a consequence of this, the standard image formation algorithm has been the polar format algorithm [10, 25] based on the two-dimensional fast Fourier transform (FFT). There is no known fast FFT-type algorithm for computing approximate samples of $f(x, y)$ from polar samples of its Fourier transform. As a result, in the case of SAR, the known data samples are first interpolated to a Cartesian grid, assuming unknown samples

(those outside the annulus segment in Figure 2.3 to be zero. After interpolation, an inverse 2-D FFT is employed and the magnitude of the reconstructed complex image is displayed for viewing. Before FFT processing, the data can be windowed to reduce sidelobe levels. This is the algorithm used in essentially all SAR systems today.

3.2 Filtered Backprojection

Another image reconstruction method, suggested by the tomographic formulation of SAR [11], is the filtered backprojection (FBP) algorithm [11, 22, 26]. FBP can be derived by writing the 2-D inverse Fourier transform in polar coordinates. The radial slices in the frequency domain are then recognized as 1-D Fourier transforms of the projections of the field at the corresponding angle, by virtue of the projection slice theorem [22]. This reduces the double Fourier integral to two sequential operations: first the data at each observation angle are filtered by a ramp (i.e. $|\Omega|$, where Ω denotes frequency) filter, and then the results are backprojected to obtain a reconstruction. FBP is the algorithm that is currently used in commercial CT scanners.

One might ask why FBP is the preferred technique for CT, whereas the 2-D FFT-based polar format algorithm has been the choice for SAR. Comparing CT and SAR reconstruction problems, there is a considerable difference between the shapes of the Fourier sampling grids. In CT, the sensor moves all around the object to be imaged. This produces Fourier samples in a disk shaped region rather than an annular region as in SAR. The nearly Cartesian nature of the grid in SAR makes interpolation more accurate than that in the case of CT, so that the polar format algorithm is attractive in SAR. In CT, however, interpolation is much more difficult, due to the circular distribution of the data, hence the FBP algorithm which does not require interpolation in the Fourier domain is preferred. In [27], it has been shown that for the CT problem, the computational effort required by an FFT-based method with

interpolation is larger than that required by the FBP method for a comparable level of quality. It has also been argued that for the SAR problem the polar format algorithm, even with simple interpolation methods, produces high-quality reconstructions.

Both the polar format and the FBP algorithms are based on the inverse operator for the case when perfect data are available throughout the spatial frequency domain. These methods have no explicit mechanism to counter any imperfection in the data. Although there are algorithmic differences between the two methods, the reconstructions they produce are very similar. We will call these methods the *conventional* methods for SAR image formation.

3.3 Spectral Estimation-based Methods

Recently there have been steps towards alternative algorithms for SAR image formation. The motivation for a class of these new approaches has been increasing the resolution of SAR images beyond the Fourier limit.

In conventional methods, image resolution is limited by the system bandwidth. Let us consider the observed signal from a synthetic aperture, which is a collection of received returns $r_\theta(t)$, as shown in Figure 2.3. Note that this 2-D signal is in the spatial frequency domain, hence its spectrum is in the spatial domain. Since peaks in the spectrum of this signal correspond to strong point scatterers in the scene, one idea to overcome the resolution limit of conventional techniques is to use modern 2-D spectral estimation methods [28] rather than a Fourier transform for forming the SAR image [29–34] (assuming that polar to rectangular resampling has already been done). Algorithms based on this idea are usually termed “superresolution” methods, and have become increasingly popular. For example, the Lincoln Laboratory baseline ATR system for SAR data currently uses a superresolution technique known as high-definition vector imaging (HDVI) [32, 35]. In addition to resolution

improvement, other motivations suggested for the application of these methods are to remove sidelobe artifacts, and to reduce speckle [29]. When applied to SAR imaging, spectral estimation-based methods are quite successful in preserving gain on ideal point scatterers, reducing sidelobes and reducing mainlobe widths, however most spectral estimation-based methods reduce gain on non-pointlike scatterers such as trees, and they usually fail to improve the quality of images containing objects with distributed features.

Spectral estimation-based methods implicitly incorporate some prior information into their processing, and can produce better results by virtue of that. The problem with these approaches is that they are not very flexible in the form of the prior information provided to the algorithm. Hence, these methods provide no explicit means to accentuate one type of feature over another. A comprehensive comparison of various spectral estimation methods in SAR can be found in [29].

Spectral estimation-based SAR imaging methods require the estimation of a correlation matrix of the collected radar returns. This estimation is the most computationally intensive part of the methods. Since the amount of data from a typical SAR scene is quite big, this is done in the following manner. First a SAR image is formed using the 2-D FFT processing, then small chips from that image are transformed back to the Fourier domain, and a correlation matrix is estimated for each of these downsampled signal histories corresponding to each image chip. Finally the imaging algorithm is applied on these signal histories, and the formed images are mosaiced together to form the full scene. The method generally used for the correlation matrix estimate is the modified covariance method [36].

3.4 Data Extrapolation-based Methods

Another way of increasing the resolution beyond the Fourier limit is to expand the support of the data in the frequency domain, by estimating samples outside the annular data region by linear prediction filters [37–39]. There has also been some recent limited attempt to compare the performance of spectral estimation-based methods with data extrapolation-based methods. For example in [40], it has been reported that spectral estimation based techniques have a degraded performance with real world targets (unlike with point targets), and may cause some loss of information about the target, whereas data extrapolation techniques offer increased resolution and better overall performance in these cases. On the other hand, according to [41], data extrapolation-based methods do not yield particularly good results, since they introduce significant amounts of noise.

3.5 Estimation-Theoretic and Entropy-based Methods

There also exists some limited previous work taking an estimation-theoretic approach to the problem. These methods are closest in spirit to our perspective. In [42], a regularized inversion method has been proposed for stripmap-mode SAR, which involves deconvolution of the projections of the field by Tikhonov-type regularization, followed by backprojection. In [43], an estimation-theoretic ℓ_1 -norm-based approach has been proposed for imaging closely-spaced multiple moving scatterers over a given spatial region. In [44] and [45], a stochastic inverse approach which seeks the best linear estimate of the reflectivity field in a stochastic least-squares sense has been proposed. This method starts from blurred image domain observations rather than projectional SAR data. A method similar in spirit to that in [44, 45], which considers

incorporating prior knowledge about the property of targets being limited in spatial extent has been introduced in [46]. Finally, a class of approaches for SAR imaging based on entropy methods has been developed [47–51]. These methods appear to offer good noise suppression properties, however the experimental results in published work are too limited to show other possible advantages over conventional methods. Rather than realistic SAR scenes, most results involve simple examples, such as two isolated point scatterers [47], a synthetic scene with reflectivity $1 - j$ in the target region and zero in the background (hence not random phase) [49,50], and a small scene consisting of straight lines and isolated point scatterers [51]. The method in [49,50] also requires post-processing (median filtering) to reduce some of the artifacts. The recent entropy-based work in [51] considers the problem of imaging extended (distributed) targets, however the method uses smoothness constraints which are built on the real and imaginary components of the field, rather than on the magnitudes directly, unlike our approach, as will be described in subsequent chapters.

3.6 Conclusion

In this chapter, we have provided brief descriptions of the conventional and some of the recently proposed techniques for SAR image reconstruction. A number of these recent techniques share some of our objectives for SAR imaging, such as increasing the resolution, or decreasing sidelobe and speckle artifacts. However, most of these techniques achieve (a subset of) such objectives by the implicit structure of the image formation algorithm. In contrast, our goal is to pose the image reconstruction problem in such a way to explicitly express the (potentially multiple and variable) feature-enhancement objectives up-front. This can be done in a regularized image reconstruction framework. The discussion in [52] regarding some current issues in SAR image reconstruction has also underlined the need for such regularized approaches

for the complex-valued inverse problems arising in SAR. In Chapter 4, we provide a general introduction to the use of regularization methods in image processing. We then propose a regularization-based framework for SAR imaging in Chapter 5.

Chapter 4

Regularization Methods in Image Processing

In this chapter, we provide a brief overview of regularization methods for inverse problems in image processing, such as restoration and reconstruction. The discussion here is not a comprehensive one, but is just aimed at presenting some preliminaries relevant to the approach that will be taken in the following portions of this dissertation for SAR imaging. Our coverage is mostly based on the reviews of the subject matter in [53] and [54].

4.1 The Need for Regularization

Restoration and reconstruction of images can be defined as the general problem of estimating a 2-D field $f(x, y)$ from some form of indirect observations related to this field. Although the problem arises in various branches of engineering and applied physics, most image restoration and reconstruction methods have a common estimation structure, and they present some common practical limitations. By virtue of

these commonalities, regularization has emerged as a central concept in the solution of image restoration and reconstruction problems in various fields and sensor modalities.

We will concentrate on problems where the mathematical relationship between the measurements and the field $f(x, y)$ is governed by a linear integral equation. The discretized version of such a relationship, taking into account the measurement noise as well, is given by:

$$\mathbf{g} = \mathbf{T}\mathbf{f} + \mathbf{w} \quad (4.1)$$

where \mathbf{g} and \mathbf{f} are vectors representing the sampled data, and the unknown field, respectively, \mathbf{w} is the noise, and \mathbf{T} is a known matrix which models the measurement or degradation mechanism. Although this is a simple observation model, it captures many situations of engineering interest. The image restoration or reconstruction problem is then one of solving for the unknown vector \mathbf{f} , given the observation vector \mathbf{g} . One may initially think that a simple matrix inverse can be used to find an estimate $\hat{\mathbf{f}}$ of \mathbf{f} , but that certainly is not possible in general. There are four basic issues that must be taken into account before any inversion can be done [53]:

1. Due to the presence of noise, an exact solution to the linear set of equations may not exist.
2. There may be more than one field which satisfy these equations, hence the solution may not be unique.
3. It is desirable to have a “stable” inversion mechanism, i.e. $\hat{\mathbf{f}}$ should be relatively insensitive to perturbations in the data.
4. It is desirable to include any *a priori* information about \mathbf{f} in the inversion process.

A solution to the first problem mentioned above can be obtained through a least-squares solution:

$$\hat{\mathbf{f}}_{\text{LS}} = \arg \min_{\mathbf{f}} \|\mathbf{g} - \mathbf{T}\mathbf{f}\|_2^2 \quad (4.2)$$

where $\|\cdot\|_2$ denotes the ℓ_2 -norm. However, when the null-space of \mathbf{T} is not empty, the least-squares solution is not unique. A common approach to address this problem is to choose the field with minimum norm, among the set of least-squares solutions, as the estimate of the true field \mathbf{f} :

$$\hat{\mathbf{f}}^+ = \arg \min_{\hat{\mathbf{f}}_{\text{LS}}} \|\hat{\mathbf{f}}_{\text{LS}}\|_2 \quad (4.3)$$

The resulting estimate $\hat{\mathbf{f}}^+$ is called the generalized solution. The generalized solution provides, a simple, reasonable way to deal with the first two issues mentioned in the list above. However, it does not directly address the third and fourth issues. The generalized solution is unstable in the face of perturbations to the data. Basically, this behavior is due to the fact that the main focus of the generalized solution is to reduce the data fit error. Naturally, when the data is noisy, the generalized solution will unfortunately fit to the noise components. Since \mathbf{T} is most often ill-conditioned, these noise components will usually be unacceptably amplified in the generalized solution. A detailed explanation of the reasons of this behavior can be found in [53] or [54]. The generalized solution does not allow the inclusion of any prior information about \mathbf{f} , either.

These difficulties, which the generalized solution cannot address, can be resolved by what is known as regularization. The purpose of regularization is to allow the inclusion of prior information to stabilize the solution in the presence of noisy data, and allow reasonable estimates. The basic idea is to constrain the solution such that the amplified noise effects are avoided.

4.2 Tikhonov Regularization

Tikhonov regularization [55, 56] is probably the most commonly used regularization method. Tikhonov method incorporates prior information about the field \mathbf{f} by augmenting the least-squares cost function with an additional term as follows:

$$\hat{\mathbf{f}}_{\text{Tik}} = \arg \min_{\mathbf{f}} \|\mathbf{g} - \mathbf{T}\mathbf{f}\|_2^2 + \lambda^2 \|\mathbf{P}\mathbf{f}\|_2^2 \quad (4.4)$$

where \mathbf{P} is a matrix, and λ is a scalar parameter.¹ The first term in (4.4) is a data fidelity term. The second term is called the “regularizer,” and prior information about \mathbf{f} is incorporated through this term. The regularization parameter λ controls the tradeoff between the two terms.

The simplest choice for \mathbf{P} is an identity matrix. In this case, the role of the regularizer term is to penalize large values in the reconstruction and hence reduce potential noise amplification (as happens in the generalized solution). Another common choice for \mathbf{P} is a 2-D derivative (gradient) operator. For a discussion of the structure of such discrete derivative operators, see Appendix A.1. In this case, the regularizer imposes a roughness penalty on the solution, and thus captures the prior belief that the field we want to reconstruct is smooth.

When we take the gradient of (4.4) with respect to \mathbf{f} , and set it equal to zero, we obtain the following set of linear equations for the Tikhonov solution:

$$(\mathbf{T}^T \mathbf{T} + \lambda^2 \mathbf{P}^T \mathbf{P}) \hat{\mathbf{f}}_{\text{Tik}} = \mathbf{T}^T \mathbf{g}. \quad (4.5)$$

When the null spaces of \mathbf{T} and \mathbf{P} are distinct, there exists a unique, closed-form solution to (4.5).

¹The variable λ we use for the regularization parameter here, and the variable λ we have used for the wavelength in Chapter 2 are not related, and should not be confused.

An important problem in Tikhonov and other regularization methods is the choice of the regularization parameter λ , as in (4.4), and any other hyperparameters that might appear in other regularization methods. Automated choice of these parameters based on data is still an open research problem. For a survey of current techniques for hyperparameter choice, please see [53, 57].

4.3 Non-Quadratic Regularization

The Tikhonov cost function (4.4) is a quadratic function of \mathbf{f} . Consequently, the optimality condition (4.5) is a linear function of \mathbf{f} . This leads to linear processing of the data \mathbf{g} for image restoration or reconstruction. While such linear processing is desirable, since it leads to straightforward and reasonably efficient computation methods, it is also limiting, in that far more powerful results are possible if non-linear methods are allowed [53]. To this end, let us consider more general problems of the following form:

$$\hat{\mathbf{f}}_{\text{NQ}} = \arg \min_{\mathbf{f}} \|\mathbf{g} - \mathbf{Tf}\|_2^2 + \lambda^2 \sum_{i=1}^M \psi((\mathbf{Pf})_i) \quad (4.6)$$

where M is the length of the vector \mathbf{Pf} , and $(\mathbf{Pf})_i$ denotes its i -th element. Note that when $\psi(x) = x^2$, (4.6) reduces to the Tikhonov cost function in (4.4), however in general $\psi(x)$ is non-quadratic. The formulation of (4.6) includes well-known regularization approaches such as the maximum entropy [58], and the total variation [59, 60] methods.

Unlike the Tikhonov case, (4.6) does not lead to a closed-form solution in general, so numerical techniques must be used to find $\hat{\mathbf{f}}_{\text{NQ}}$. Here, we will consider a particular approach, known as half-quadratic regularization [61, 62], for this solution. The reason we consider this particular method is that our approach for SAR image reconstruction

is related to it, as we will show in Chapter 6. Let us represent the cost function in (4.6) as:

$$J(\mathbf{f}) = \|\mathbf{g} - \mathbf{T}\mathbf{f}\|_2^2 + \lambda^2 \sum_{i=1}^M \psi((\mathbf{P}\mathbf{f})_i). \quad (4.7)$$

The basic idea in half-quadratic regularization is to introduce a new cost function, which has the same minimum as $J(\mathbf{f})$, but one which can be manipulated with linear algebraic methods. To this end, let us consider a new function $K(\mathbf{f}, \mathbf{b})$, which is quadratic in \mathbf{f} (hence the name half-quadratic), and where \mathbf{b} is an auxiliary vector, such that

$$\inf_{\mathbf{b}} K(\mathbf{f}, \mathbf{b}) = J(\mathbf{f}). \quad (4.8)$$

One particular way to construct the augmented half-quadratic cost function $K(\mathbf{f}, \mathbf{b})$ is the following [61]:

$$K(\mathbf{f}, \mathbf{b}) = \|\mathbf{g} - \mathbf{T}\mathbf{f}\|_2^2 + \lambda^2 \sum_{i=1}^M [\mathbf{b}_i ((\mathbf{P}\mathbf{f})_i^2) + \eta(\mathbf{b}_i)] \quad (4.9)$$

where $\psi(\cdot)$ of (4.6), and $\eta(\cdot)$ of (4.9) are related through convex duality relations as follows:

$$\begin{aligned} \psi(x) &= \inf_{\omega} (\omega x^2 + \eta(\omega)) \\ \eta(\omega) &= \sup_x (\psi(x) - \omega x^2) \end{aligned} \quad (4.10)$$

Minimizing $K(\mathbf{f}, \mathbf{b})$ rather than $J(\mathbf{f})$ may have some structural advantages. In particular, one can benefit from the half-quadratic structure by using alternating updates of \mathbf{f} given \mathbf{b} , and of \mathbf{b} given \mathbf{f} . This can be done through an iterative block

coordinate descent method as follows [63]:

$$\hat{\mathbf{b}}^{(n+1)} = \arg \min_{\mathbf{b}} K(\hat{\mathbf{f}}^{(n)}, \mathbf{b}) \quad (4.11)$$

$$\hat{\mathbf{f}}^{(n+1)} = \arg \min_{\mathbf{f}} K(\mathbf{f}, \hat{\mathbf{b}}^{(n+1)}) \quad (4.12)$$

where n denotes the iteration number. The aim is to take advantage of such coordinate descent steps, where each of these steps is a simple enough operation. We know that when \mathbf{b} is fixed, $K(\mathbf{f}, \mathbf{b})$ is quadratic in \mathbf{f} , hence (4.12) is simple. To investigate a typical structure for (4.11), let us consider a specific non-quadratic cost function. A non-quadratic regularizer that has achieved popularity in recent years is the total variation measure [59, 60]. The idea is to use $\psi(x) = |x|$, i.e. penalize the absolute value of the argument, unlike the Tikhonov approach which penalizes its square. The matrix \mathbf{P} used in total variation is a derivative operator. The total variation of a signal is just the total amount of change the signal goes through and can be thought as a measure of signal variability. One important advantage of total variation regularization over Tikhonov regularization is that the total variation solutions can contain localized steep gradients, so that edges are better preserved than that in Tikhonov solutions. To overcome the problem of non-differentiability of the absolute value function around the origin, a smooth approximation to the total variation regularizer can be used: $\psi(x) = \sqrt{x^2 + \epsilon}$, where ϵ is a small smoothing constant. This yields the following pair of functions for total variation, based on the convex duality relationships of (4.10):

$$\begin{aligned} \psi(x) &= (x^2 + \epsilon)^{1/2} \\ \eta(\omega) &= \omega\epsilon + \frac{1}{4\omega} \end{aligned} \quad (4.13)$$

Then the cost function $J(\mathbf{f})$ for the total variation problem, and its augmented, half-quadratic version $K(\mathbf{f}, \mathbf{b})$ are as follows:

$$J(\mathbf{f}) = \|\mathbf{g} - \mathbf{T}\mathbf{f}\|_2^2 + \lambda^2 \sum_{i=1}^M [(\mathbf{P}\mathbf{f})_i^2 + \epsilon]^{1/2} \quad (4.14)$$

$$K(\mathbf{f}, \mathbf{b}) = \|\mathbf{g} - \mathbf{T}\mathbf{f}\|_2^2 + \lambda^2 \sum_{i=1}^M \left[\mathbf{b}_i ((\mathbf{P}\mathbf{f})_i^2 + \epsilon) + \frac{1}{4\mathbf{b}_i} \right] \quad (4.15)$$

Using (4.11),(4.12) we then obtain the following coordinate descent steps for the minimization of $K(\mathbf{f}, \mathbf{b})$ in (4.15):

$$\hat{\mathbf{b}}_i^{(n+1)} = \frac{1}{2 \left[(\mathbf{P}\hat{\mathbf{f}}^{(n)})_i^2 + \epsilon \right]^{1/2}} \quad (4.16)$$

$$\hat{\mathbf{f}}^{(n+1)} = \left[2\mathbf{T}^T\mathbf{T} + \lambda^2\mathbf{P}^T \text{diag} \left\{ \frac{1}{\left[(\mathbf{P}\hat{\mathbf{f}}^{(n)})_i^2 + \epsilon \right]^{1/2}} \right\} \mathbf{P} \right]^{-1} (2\mathbf{T}^T\mathbf{g}) \quad (4.17)$$

where $\text{diag}\{\cdot\}$ is a diagonal matrix whose i -th diagonal element is given by the expression inside the brackets. Hence both of the coordinate descent steps required are simple operations. Note that we have substituted the result of (4.16) into (4.17) above, so the algorithm required for total variation regularization consists of just the fixed-point iterative scheme of (4.17).

Our discussion of half-quadratic regularization was aimed only at providing a flavor of the method, so it is not comprehensive. Consequently, some issues on the reader's mind may have remained unresolved. Please see [61–64] for a more detailed treatment of this topic.

Chapter 5

A Regularization-based Framework for SAR Imaging

In this chapter, we describe our framework for regularized, feature-enhanced SAR image formation. The first major contribution of this chapter is an explicit projection-type complex-valued forward operator for SAR. This operator (projection matrix) not only opens the door to exploiting the similarities with the projectional operator in CT, but also provides a new perspective to explore and interpret interesting properties of the SAR phenomenology. The second major contribution of this chapter is the proposal and setup of a specific optimization problem for SAR imaging. This optimization-based approach addresses the issue of robustness to noisy or limited data, through a regularized objective function in a model-based framework. This objective function extends the use of feature-preserving potential functions to the SAR imaging problem. In doing this extension, we focus on particular features important for automated interpretation of SAR imagery. Unlike standard image processing problems, SAR involves *complex-valued* and *random-phase* reflectivities, as we have seen in Chapter 2. This complex-valued, random-phase nature of SAR scenes

makes extension and application of real-valued regularization methods challenging. We formulate the objective function in such a way to address these challenges. We also provide a statistical interpretation of our approach.

5.1 Discrete Observation Model

Most standard SAR image formation methods have only an implicit dependence on the SAR observation relationship developed in Section 2.2, i.e. an inverse relationship based on the continuous, complete-data problem formulation is developed, and then applied to the particular discrete problem. An explicit discrete model of the particular SAR sensor and observation geometry is generally not used in practice. Using an explicit discrete forward model however, can facilitate incorporation of information about properties of the SAR sensor and measurement parameters into the processing. A model-based approach may offer a number of advantages for the SAR problem:

- It lets us handle limitations in data quantity more effectively. Examples of such limitations are angular diversity limitations (e.g. due to sensor re-tasking), resolution limitations, and missing observations.
- The model-based approach ties readily into the statistical processing methods. This lets us handle limitations in data quality more effectively, through a noisy observation model.
- When the particular data relationship deviates significantly from the Fourier model, a model-based approach can readily take this into account. An example for this situation is when the quadratic phase term in (2.14) cannot be neglected.
- Since the model can be built for the polar-sampled data, polar to rectangular interpolation in the data domain, as required for example in the polar format

algorithm, is not required.

In this section we will briefly explain the discrete observation model we use, and also demonstrate how we can obtain a projection-like data relationship. Let us consider the second relationship in (2.16), which we repeat here for convenience:

$$r_\theta(t) = \iint_{x^2+y^2 \leq L^2} f(x, y) \exp(-j\Omega(t)(x \cos \theta + y \sin \theta)) dx dy \quad (5.1)$$

where $r_\theta(t)$ is the demodulated received signal at observation angle θ , $f(x, y)$ is the underlying unknown reflectivity field to be reconstructed, $\Omega(t)$ is the spatial frequency in the radial direction, and L is the radius of the ground patch to be imaged. We can compactly write the above relationship as $r_\theta(t) = (\mathcal{C}_\theta f(x, y))(t)$, where \mathcal{C}_θ is the continuous observation kernel.

In practice, the observations at the i -th observation angle θ_i are samples $r_{\theta_i}(t_j)$ of the continuous received signal $r_{\theta_i}(t)$ at sampling times t_j . This sampling in the time domain results in sampling of the spectrum of the underlying reflectivity field. Sampling places a limit on the maximum allowable scene size that can be imaged without aliasing (in the spatial domain).

Let \mathbf{r}_{θ_i} be the vector of these observed samples, \mathbf{C}_{θ_i} be a discretized approximation to the continuous observation kernel \mathcal{C}_{θ_i} and \mathbf{f} be a vector representing the unknown sampled reflectivity image. Then, overall, we can write:

$$\underbrace{\begin{bmatrix} \mathbf{r}_{\theta_1} \\ \mathbf{r}_{\theta_2} \\ \vdots \\ \mathbf{r}_{\theta_P} \end{bmatrix}}_{\mathbf{r}} = \underbrace{\begin{bmatrix} \mathbf{C}_{\theta_1} \\ \mathbf{C}_{\theta_2} \\ \vdots \\ \mathbf{C}_{\theta_P} \end{bmatrix}}_{\mathbf{C}} \mathbf{f} \quad (5.2)$$

where P is the total number of angular observation points. The data in \mathbf{r} are the sampled phase histories, and are confined to an annular region in the spatial frequency plane as shown in Figure 2.3. In practice, we form \mathbf{C} by calculating the sampled phase history corresponding to unit reflectivity point targets in the scene. In doing this, we take into account the particular system parameters, such as those shown in Table 2.1 before.

In Section 2.2, we have seen that the observed signal $r_\theta(t)$ can also be identified as a band-pass filtered Fourier transform of the projections $q_\theta(u)$ of the field:

$$r_\theta(t) = \int_{|u| \leq L} q_\theta(u) \exp \{-j\Omega(t)u\} du \quad (5.3)$$

We will now exploit this relationship to obtain a discrete projectional data relationship, between the sampled field \mathbf{f} and band-pass filtered projections, similar to the one in CT data collection. In particular, we can obtain samples of the band-pass projectional information by an inverse discrete Fourier transform (IDFT) of the sampled data \mathbf{r}_{θ_i} , $i \in \{1, \dots, P\}$, at each observation angle:

$$\begin{bmatrix} \bar{\mathbf{q}}_{\theta_1} \\ \bar{\mathbf{q}}_{\theta_2} \\ \vdots \\ \bar{\mathbf{q}}_{\theta_P} \end{bmatrix} = \begin{bmatrix} \mathbf{F}^{-1} & & & \\ & \mathbf{F}^{-1} & & \\ & & \ddots & \\ & & & \mathbf{F}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{r}_{\theta_1} \\ \mathbf{r}_{\theta_2} \\ \vdots \\ \mathbf{r}_{\theta_P} \end{bmatrix}$$

$$\begin{aligned}
\begin{bmatrix} \bar{\mathbf{q}}_{\theta_1} \\ \bar{\mathbf{q}}_{\theta_2} \\ \vdots \\ \bar{\mathbf{q}}_{\theta_P} \end{bmatrix} &= \begin{bmatrix} \mathbf{F}^{-1} & & & \\ & \mathbf{F}^{-1} & & \\ & & \ddots & \\ & & & \mathbf{F}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{C}_{\theta_1} \\ \mathbf{C}_{\theta_2} \\ \vdots \\ \mathbf{C}_{\theta_P} \end{bmatrix} \mathbf{f} \\
\underbrace{\hspace{1.5cm}}_{\bar{\mathbf{q}}} &= \underbrace{\hspace{1.5cm}}_{\mathbf{T}} \hspace{1.5cm} \mathbf{f}
\end{aligned} \tag{5.4}$$

Here \mathbf{F} is the DFT matrix, and \mathbf{T} represents a complex-valued discrete ‘‘SAR projection operator’’. The data $\bar{\mathbf{q}}$ obtained in this way are the range profiles. Note that in the formation of $\bar{\mathbf{q}}_\theta$, we have taken into account the band-pass structure of the observations as well, so $\bar{\mathbf{q}}_\theta$ is not just a baseband sampled version of $q_\theta(u)$ of (2.8).

In the presence of noise, our model of the SAR range profile observations becomes:

$$\mathbf{g} = \mathbf{T}\mathbf{f} + \mathbf{w} \tag{5.5}$$

where \mathbf{w} accounts for additive measurement noise. This is the observation relationship we will use throughout our development in this chapter. A similar observation relationship could also be written in terms of the phase histories \mathbf{r} , and the matrix \mathbf{C} , however we will mostly use (5.5) in our experiments, since the sparser nature of \mathbf{T} offers computational advantages. Note that, since this system model relates reflectivities to measurements directly, we will not require polar to rectangular resampling.

5.2 Properties of the SAR Projection Matrix

The matrix \mathbf{T} obtained as described in the previous section resembles the projection matrix commonly used in CT. The CT projection matrix would be a discrete approximation to the kernel in (2.8). The similarity of the structure of the discrete CT and SAR projection operators for two SAR data collection geometries is shown




Figure 5.1: The grayscale plot (black corresponds to the maximum value and white to the minimum) of the magnitude of the elements in the projection matrices for CT and SAR for a 32×32 field. (a) CT, $\Delta\theta = 2.3^\circ$. (b) SAR, $f_0 = 10$ GHz, $\Delta\theta = 2.3^\circ$. (c) CT, $\Delta\theta = 23^\circ$. (d) SAR, $f_0 = 1$ GHz, $\Delta\theta = 23^\circ$.

in Figure 5.1. The images in this figure are grayscale plots of the magnitudes of the elements of the corresponding projection matrices. Here an underlying field of 32×32 pixels, and projectional data of length 32 at each of 32 angular observation points have been assumed. Figure 5.1(a) and (b) contain the CT and SAR projection matrices respectively, for an angular observation range of $\Delta\theta = 2.3^\circ$. For SAR, the central frequency is $f_0 = \omega_0/2\pi = 10$ GHz. Based on (2.22), this results in a cross-range (azimuth) resolution of 0.375 m. We choose the bandwidth of the chirp pulse, based on (2.21) such that range resolution is also 0.375 m. We will call this specific




Figure 5.2: The grayscale plot of the magnitude of the elements in the SAR high-resolution-field to low-resolution-data projection matrix for a 32×32 field. $f_0 = 5$ GHz, $\Delta\theta = 2.3^\circ$, pixel-spacing = 0.375 m, resolution = 0.75 m.

SAR projection matrix \mathbf{T}_1 . Next, we form projection matrices with different sensor parameters. We reduce the SAR central frequency to $f_0 = 1$ GHz, and keep the previous cross-range resolution. This then requires an angular observation range of $\Delta\theta = 23^\circ$. The resulting SAR projection matrix, \mathbf{T}_2 , is shown in Figure 5.1(d). The CT projection matrix for $\Delta\theta = 23^\circ$ is shown in Figure 5.1(c) for comparison. The structural similarity of the magnitude of the CT and SAR projection matrices is evident. For the SAR projection matrices in this figure, we have assumed that the pixel spacing in the field is equal to the resolution supported by the data, and that the number of data samples is equal to the number of pixels (hence a square projection matrix). Neither of these has to be the case. In Figure 5.2, we show a SAR projection matrix, \mathbf{T}_3 , where the pixel spacing is 0.375 m, the (range and cross-range) resolution is 0.75 m, the angular range is $\Delta\theta = 2.3^\circ$, the central frequency is $f_0 = 5$ GHz, and the number of data samples is half the number of pixels.

We will provide some insight into the tomographic behavior of these matrices by examining their rows and columns. One row of \mathbf{T} represents the contribution of all points in the scene to one particular data sample. When a row is reshaped and its magnitude is viewed as an image, the bright points in this image correspond to locations in the scene which have a large contribution to that data sample. In the context of SAR, this indicates a line integral over points which are equidistant to the

sensor. Figure 5.3 shows the magnitudes of such images for one of the data samples at an observation angle of $\theta = 0.7^\circ$, for the projection matrix \mathbf{T}_1 in (a), and for \mathbf{T}_2 in (b). These two images are identical, since the line integral path is determined by the observation angle. In Figure 5.3(c), we show a similar image for \mathbf{T}_2 , this time at an observation angle of $\theta = 10^\circ$. As expected, the line integral is now over a tilted path in the scene, based on the observation angle of 10° .

A different kind of intuition can be obtained by inspecting the columns of \mathbf{T} . The columns indicate the contribution of a single pixel in the scene to the entire data collected through the synthetic aperture. We can reshape the columns and display them as images, where the vertical direction corresponds to the range bins (hence data samples collected at a particular observation angle), and the horizontal direction corresponds to the observation angles. The resulting image is a trace of the varying range of a location in the scene as a function of the observation angle. Such plots are analogous to sinograms used in CT, where rotating the sensor around the scene to collect data through a 360° range of observation angles produces one period of a sinusoid in such a plot. In SAR, the range of observation angles is limited, hence these plots produce what looks like portions of a sinusoid. Such plots are shown in Figure 5.4 for a cross-range displaced point in the scene. Note that the plot for \mathbf{T}_2 shows a greater variability than that of \mathbf{T}_1 in the angular direction. This is simply because the angular observation range for \mathbf{T}_2 is ten times that of \mathbf{T}_1 (23° versus 2.3°).

So far, we have investigated the magnitude of the SAR projection matrices. However, the SAR projection matrices are in general complex-valued, whereas the CT projection matrices are real-valued. To understand this, note that the components of the SAR projection matrix \mathbf{T} at each angle are composed of two cascade operations: finding an offset Fourier transform slice of the rectangularly sampled field, followed by a 1-D inverse DFT. Due to the radial frequency offset inherent in the first operation, the generated Fourier domain signal is not necessarily conjugate symmetric at all




Figure 5.3: Magnitude of the rows of the SAR projection matrices, reshaped as images. These images indicate the points in the scene contributing to a particular data sample at observation angle θ . (a) \mathbf{T}_1 , $\theta = 0.7^\circ$. (b) \mathbf{T}_2 , $\theta = 0.7^\circ$. (c) \mathbf{T}_2 , $\theta = 10^\circ$.

observation points, consequently the inverse DFT does not yield a real-valued signal. Let us now revisit Figure 5.3(a) and (b), where we have observed that the magnitude of a row of \mathbf{T}_1 and \mathbf{T}_2 at a particular observation angle produced identical plots of a line integral path. Note that the matrices \mathbf{T}_1 and \mathbf{T}_2 were based on sensors with different central frequencies (10 GHz versus 1 GHz). This information can be brought out by examining the real and imaginary parts of the matrices. In Figure 5.5, we show the mesh plots of the real parts of the reshaped rows corresponding to the magnitude images shown in Figure 5.3(a) and (b). The difference in the structure of \mathbf{T}_1 and \mathbf{T}_2 can now be observed.

As we mentioned in Section 2.4, CT and SAR differ in their angular diversity requirements: SAR requires a much smaller angular observation range than CT for a satisfactory (cross-range) resolution. This is related to the high central frequencies used in SAR, as explained in Section 2.4. In our projection matrices, the information regarding the central frequency is carried in the full complex structure of the matrices, as we have observed in Figure 5.5. The magnitudes of the SAR projection matrices




Figure 5.4: Magnitude of the columns of the SAR projection matrices, reshaped as images. These images indicate the contribution of a single pixel in the scene to the entire data. The vertical direction corresponds to the range bins, and the horizontal direction corresponds to the observation angles. Here, a cross-range displaced point scatterer is considered. (a) \mathbf{T}_1 . (b) \mathbf{T}_2 .

roughly correspond to a sensor operating at baseband, and are similar to the CT projection matrices, as we have seen in Figure 5.1. Now, we will show on a simple example that although a regular SAR projection matrix has a resolving power with a small range of observation angles, the magnitude of such a matrix (i.e. a CT projection scenario) does not. Figure 5.6(a) shows a scene with a single point scatterer, reconstructed from data generated by \mathbf{T}_1 ($\Delta\theta = 2.3^\circ$). Let us now generate simulated data and reconstruct this scene using a CT-type projection operator, using the magnitude of \mathbf{T}_1 . The resulting reconstruction from such data is shown in Figure 5.6(b). Clearly, the cross-range resolving power has been lost.




Figure 5.5: Real parts of the rows of the SAR projection matrices, reshaped as images, for a particular data sample at observation angle θ . (a) \mathbf{T}_1 , $\theta = 0.7^\circ$. (b) \mathbf{T}_2 , $\theta = 0.7^\circ$.

5.3 Objective Function for Feature-Enhanced Imaging

We formulate the SAR image reconstruction problem as the following optimization problem:

$$\hat{\mathbf{f}} = \arg \min_{\mathbf{f}} J(\mathbf{f}) \tag{5.6}$$

where we choose $J(\mathbf{f})$ to be an objective function of the following form:

$$J(\mathbf{f}) = \|\mathbf{g} - \mathbf{T}\mathbf{f}\|_2^2 + \Psi(\mathbf{f}) \tag{5.7}$$

where $\Psi(\mathbf{f})$ is a function from \mathbb{C}^N to $\mathbb{R}^+ \cup \{0\}$. The structure of the above objective function is that of the regularization functions we have seen in Chapter 4. The first term is a data fidelity term minimizing the squared error between the actual observations and the observations that would be produced by the reconstructed field based on our system model. The second term is the regularization term reflecting




Figure 5.6: Demonstration of the dependence of the cross-range resolving power of SAR on the central frequency, as reflected in the complex-valued structure of the projection matrix. (a) Scene containing a single point scatterer reconstructed from data generated by \mathbf{T}_1 ($\Delta\theta = 2.3^\circ$). (b) Same scene reconstructed from CT-type data generated by the magnitude of \mathbf{T}_1 for the same angular diversity.

the prior information we would like to impose. The reconstruction naturally depends on the kind of constraint used, hence the choice of $\Psi(\mathbf{f})$ is critical. We would like to choose $\Psi(\mathbf{f})$ such that it helps both in reducing undesired artifacts, and also in enhancing the features in the image that are critical for further processing based on its intended use.

The simplest and most common choice in many regularization problems is to set $\Psi(\mathbf{f})$ to be a quadratic function of \mathbf{f} , which leads to Tikhonov regularization [55,56], as we have seen in Section 4.2. This choice results in computationally easy optimization problems, however it has been observed that in many imaging problems this constraint may suppress useful features in the image. Recently, considerable effort has been spent in designing alternative, non-quadratic constraints which are general enough to be used in a class of problems, and yet preserve the critical features for the particular task. In our work, we extend the use of such methods to complex-valued SAR imaging.

In particular we choose $\Psi(\mathbf{f})$ as follows:

$$\Psi(\mathbf{f}) = \lambda_1^2 \|\mathbf{f}\|_k^k + \lambda_2^2 \|\mathbf{D}|\mathbf{f}|\|_k^k, \quad (5.8)$$

which results in the following objective function:

$$J(\mathbf{f}) = \|\mathbf{g} - \mathbf{T}\mathbf{f}\|_2^2 + \lambda_1^2 \|\mathbf{f}\|_k^k + \lambda_2^2 \|\mathbf{D}|\mathbf{f}|\|_k^k. \quad (5.9)$$

Here $\|\cdot\|_k$ denotes the ℓ_k -norm, \mathbf{D} is a discrete approximation to the 2-D derivative operator (gradient), $|\mathbf{f}|$ denotes the vector of magnitudes of the complex-valued vector \mathbf{f} , and λ_1, λ_2 are scalar parameters. For a description of the discrete 2-D derivative operators used in this dissertation, please see Appendix A.1. The formulation of (5.6),(5.9) starts from the observed range profiles and is not simply a post-processing of a formed image.

The first term in the objective function (5.9) incorporates the tomographic SAR observation model (5.5), and thus information about the observation geometry. The second and third terms in (5.9) incorporate prior information regarding both the behavior of the field \mathbf{f} , and the nature of the features of interest in the resulting reconstructions. In general, the values of k used for the norms in the two prior information terms do not have to be identical. Here we use the same norm for both terms for the sake of simplicity. These terms are aimed at enhancing point-based and region-based features respectively. The relative magnitudes of the parameters λ_1 and λ_2 determine the relative emphasis on these two types of features.

In Sections 5.4 and 5.5, we will discuss in more detail the reasoning behind our particular choices for these prior information terms.

5.4 Enhancement of Point-based Features by $\|\mathbf{f}\|_k^k$

Many object recognition methods rely on locations of dominant point scatterers extracted from SAR images [65, 66]. Extraction of these locations is however complicated due to the limited resolution achievable by the SAR system (hence widening or merging individual peaks) and due to sidelobe artifacts. One of our objectives is to produce images in which such point-based features are enhanced. In applications such as nuclear magnetic resonance (NMR) spectroscopy [67] and astronomical imaging [68], similar objectives have previously been achieved by using maximum entropy methods. These approaches provide reconstructions with good energy concentration (i.e. most elements are small and a few are very large). It has been shown that similar behavior can be obtained using minimum ℓ_1 -norm reconstruction [69]. In spectral analysis, ℓ_k -norm constraints, where $k < 2$, have been shown to result in higher resolution spectral estimates compared to the ℓ_2 -norm case (which is proportional to the periodogram) [70]. Based on these observations, we use a prior term of the form $\|\mathbf{f}\|_k^k$ with $k \leq 1$. This function imposes an energy-type constraint on the solution, and aims to suppress artifacts and increase the resolvability of scatterers.

Now, let us get some insight about the role and effect of k . First, let us consider a scalar, real-valued version of $\|\mathbf{f}\|_k^k$, i.e. $|f|^k$, and plot it versus f . Figure 5.7 shows such plots for a number of choices of k . When we view these plots as penalty functions, we deduce that as the value of k gets smaller, the relative penalty on large values of f reduces. We will show that in the SAR imaging problem, this effect helps the preservation and enhancement of strong scatterers in the scene, while still suppressing artifacts. Note also that for $k < 1$, the penalty functions in Figure 5.7 are concave, hence minimization of objective functions containing such terms can be challenging.

Next, let us consider the effect of k on the structure of the vector which minimizes $\|\mathbf{f}\|_k^k$. For simplicity, let us consider vectors with two elements. Figure 5.8(a) shows




Figure 5.7: Behavior of the function $|f|^k$ for various choices of k .

curves representing unit ℓ_k -norm points for $k = \{0.5, 1, 2\}$, where the coordinates are the elements of the vector \mathbf{f} . Consider the points that lie on the circle in Figure 5.8(a), all of which have the same ℓ_2 -norm, hence an ℓ_2 -norm minimization would not distinguish between these points. On the other hand, if we performed an ℓ_1 -norm minimization among these points, the optimal points would be the ones on the coordinate axes, where the equi-norm curves for $k = 2$ and $k = 1$ intersect. Similar observations can be made as we move on to $k = 0.5$, where the coordinate axes preference is even stronger. Hence we can say, as k gets smaller, a sparser structure in the solution is favored. What we mean by sparsity in a vector is that, there are a small number of dominant elements, and many other elements have a magnitude close to zero. We can do this exercise in the reverse direction as well. Consider the points lying on the constant $\ell_{0.5}$ -norm curve in Figure 5.8(b). By drawing equi-norm curves

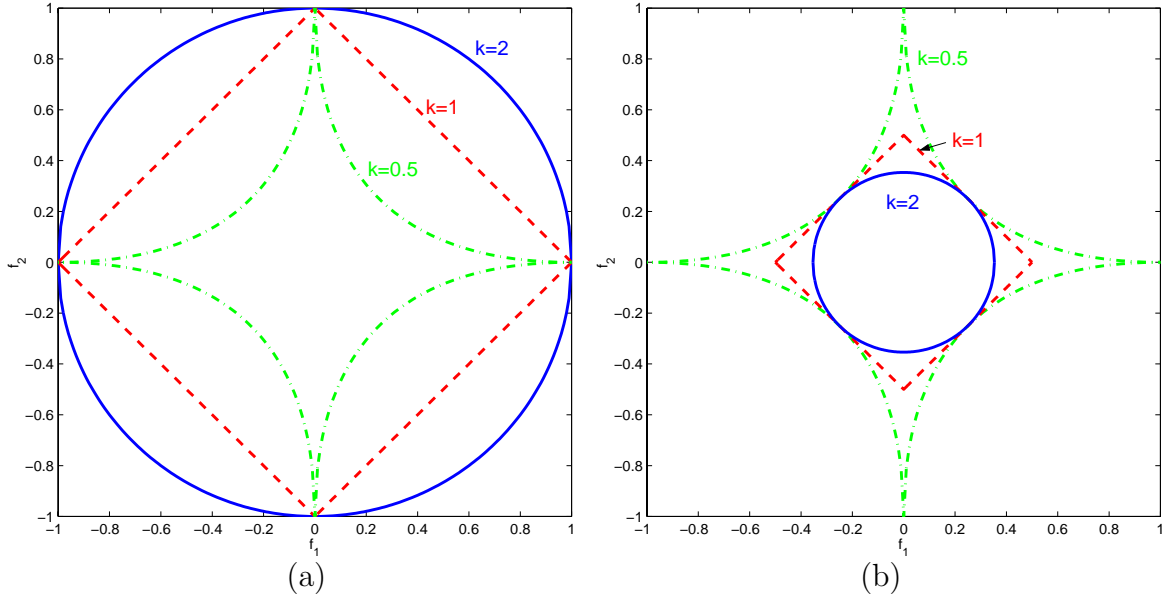


Figure 5.8: Geometrical figures representing constant ℓ_k -norm points for $k = \{0.5, 1, 2\}$ for a 2-element vector \mathbf{f} . Given all points with a particular ℓ_k -norm, minimization of $\ell_{k'}$ -norms where $k' < k$ yields the coordinate axes as the minimizer (hence a sparse vector structure). On the other hand, given all points with a particular $\ell_{k'}$ -norm, minimization of ℓ_k -norms where $k > k'$ yields a solution where the two coordinates have identical magnitudes (hence a less sparse structure).

for larger values of k , which are circumscribed by the $\ell_{0.5}$ -norm curve, we can see that the intersection points are now at orientations of 45° and its multiples. Hence, among the points with identical $\ell_{0.5}$ -norms, the ones with the two coordinates having equal magnitudes are favored by minimization of norms with a larger k . So, as k gets larger, a less sparse structure is favored. The same idea can be observed from the example in Table 5.1. Given two vectors with the same ℓ_2 -norm (v_1 and v_2 in Table 5.1), the vector with the smaller ℓ_1 -norm is v_2 , which has a sparser structure. On the other hand, given two vectors with the same ℓ_1 -norm (v_1 and v_3), ℓ_2 -norm minimization prefers v_1 , which has a less sparse structure.

Overall, the effect of a small k relative to a large k then is the favoring of a field with a smaller number of dominant scatterers, and better preservation of the scatterers and their magnitudes.

	$v_1 = (3, 4)$	$v_2 = (5, 0)$	$v_3 = (7, 0)$
ℓ_2 -norm	5	5	7
ℓ_1 -norm	7	5	7

Table 5.1: Example to demonstrate minimizing ℓ_1 -norm as opposed to ℓ_2 -norm results in the preference of a relatively sparse vector structure.

Apart from the implicit stabilizing effect due to regularization, the $\|\mathbf{f}\|_k^k$ term then serves two closely related, but still distinct goals:

1. Signal-to-noise enhancement (scatterer enhancement)
2. Superresolution

These effects are obtained by virtue of the non-quadratic nature of the objective functions (and the non-linearity of the resulting inversion algorithms) we use. The mechanism by which these effects are achieved by non-linear inversion schemes like maximum entropy and minimum ℓ_1 -norm minimization, and the conditions required for the occurrence of such effects have been discussed in detail in [69]. One of the major conclusions of that study is that these effects occur if and only if the image to be recovered is *nearly black* - nearly zero in all but a small fraction of samples. In SAR imaging, this condition can correspond to the existence of a small number of dominant scatterers in the scene.

The term superresolution refers to the ability to resolve better than the Rayleigh resolution limit, which has been discussed in Section 2.3. A superresolving effect in the reconstructions can be sought when the data carries limited-resolution information about the underlying truth, in which case the forward operator \mathbf{T} maps a high-resolution field to lower-resolution data.

5.5 Enhancement of Region-based Features by $\|\mathbf{D}|\mathbf{f}|\|_k^k$

SAR recognition algorithms also use region-based shape features [71–73]. As an example, in military target recognition applications, the image is sometimes first segmented into three regions: target, shadow and background, and features related to the shapes of these regions are then used for classification of the targets. With conventional SAR images, segmentation of such regions is particularly difficult due to speckle. We are thus interested in reducing variability in homogeneous regions, while preserving discontinuities at region boundaries. Such behavior has been obtained in real-valued image restoration and reconstruction problems by using constraints of the form $\|\mathbf{D}\mathbf{f}\|_k^k$ with $k \approx 1$ [60, 74]. However, straightforward application of such a term to the complex-valued, random-phase SAR case is problematic, since it would impose smoothness separately on the real and imaginary parts of the complex field \mathbf{f} . In Section 6.4.7, we will experimentally demonstrate how the use of such a term produces poor SAR reconstructions. The correlation in a homogeneous region of \mathbf{f} in SAR is due to the similarity of backscatter power, which is better represented in the magnitude of \mathbf{f} than its real and imaginary parts. As a result, for region-based SAR imaging, we propose using the prior term $\|\mathbf{D}|\mathbf{f}|\|_k^k$. The resulting optimization problem is made much more difficult by the substitution of the term $\mathbf{D}|\mathbf{f}|$ for $\mathbf{D}\mathbf{f}$, since $|\mathbf{f}|$ is a non-linear function of the real and imaginary parts of \mathbf{f} . In order to visualize some of the complexities resulting from the use of $\mathbf{D}|\mathbf{f}|$, let us consider a simple example, and observe the convexity properties. Let \mathbf{f} be a two-element, real-valued vector, and let $\mathbf{D} = [1 \ -1]$. Figure 5.9(a) shows a 3-D mesh plot of $\|\mathbf{D}\mathbf{f}\|_k^k$ as a function of the first and second elements of \mathbf{f} for $k = 2$ (top) and for $k = 1$ (bottom). These functions are convex. The corresponding plots for $\|\mathbf{D}|\mathbf{f}|\|_k^k$ are shown in Figure 5.9(b). Contrary to the curves in Figure 5.9(a), these are non-convex functions. Efficient and robust solution of (5.9) thus becomes a major challenge. In

Chapter 6, we overcome this limitation by providing an effective, efficient algorithm.

Note that the insight we have obtained in Section 5.4 regarding the behavior of the ℓ_k -norms through Figure 5.7 and Table 5.1 is also valid for region-based feature-enhancement. This time the argument holds not for the reflectivities themselves, but for the derivatives of the reflectivity magnitudes. So, as the role of the point-enhancement term was to preserve strong scatterers, the role of the region-enhancement term is to preserve strong reflectivity gradients, such as object boundaries, while serving as a regularizer.

Note that the advantage of region-based feature enhancement during image reconstruction rather than by post-processing a conventionally formed image is that the former is more effective in suppressing potential artifacts resulting from limitations (e.g. partial aperture) or imperfections in the data.




Figure 5.9: Roughness penalty functions: (a) on reflectivities: $\|\mathbf{D}\mathbf{f}\|_k^k$, (b) on reflectivity magnitudes: $\|\mathbf{D}|\mathbf{f}|\|_k^k$. Top: $k = 2$. Bottom: $k = 1$. (Here \mathbf{f} is real-valued, and has length two. $\mathbf{D} = [1 \ -1]$.)

5.6 Statistical Interpretation

The reconstruction problem we have described in Section 5.3 can also be obtained through Bayesian means. With a statistical perspective, let us seek a reconstructed reflectivity field as the solution to a maximum *a posteriori* (MAP) estimation problem, which defines the estimate of the random field \mathbf{f} , based on the observed data \mathbf{g} , as follows:

$$\begin{aligned}\hat{\mathbf{f}}_{\text{MAP}} &= \arg \max_{\mathbf{f}} [\log (p_{\mathbf{f}|\mathbf{g}}(\mathbf{f}|\mathbf{g}))] \\ &= \arg \max_{\mathbf{f}} [\log (p_{\mathbf{g}|\mathbf{f}}(\mathbf{g}|\mathbf{f})) + \log (p_{\mathbf{f}}(\mathbf{f}))]\end{aligned}\quad (5.10)$$

where $\log(\cdot)$ denotes the natural logarithm. Note that maximizing the posterior density $p_{\mathbf{f}|\mathbf{g}}(\mathbf{f}|\mathbf{g})$, or its logarithm are equivalent, due to the monotonicity property of the logarithm. Now, we need the pieces of the posterior density $p_{\mathbf{f}|\mathbf{g}}(\mathbf{f}|\mathbf{g})$, namely the likelihood $p_{\mathbf{g}|\mathbf{f}}(\mathbf{g}|\mathbf{f})$, and the prior $p_{\mathbf{f}}(\mathbf{f})$.¹

To find an expression for the likelihood, let us consider the noisy observation model (5.5), and assume that the observation noise \mathbf{w} is independent identically distributed (i.i.d.) complex Gaussian noise:

$$p_{\mathbf{w}}(\mathbf{w}) \propto \exp \left(-\frac{\|\mathbf{w}\|_2^2}{2\sigma^2} \right), \quad (5.11)$$

where σ^2 is proportional to the noise power. This is the most commonly used statistical model for radar measurement noise [47, 75]. Then, for the likelihood, we have:

¹The posterior density also depends on $p_{\mathbf{g}}(\mathbf{g})$, but we do not need this piece since the MAP estimation problem does not depend on it.

$$p_{\mathbf{g}|\mathbf{f}}(\mathbf{g}|\mathbf{f}) \propto \exp\left(-\frac{1}{2\sigma^2}\|(\mathbf{g} - \mathbf{T}\mathbf{f})\|_2^2\right) \quad (5.12)$$

Let the prior probability density function for the field \mathbf{f} be given by:

$$p_{\mathbf{f}}(\mathbf{f}) \propto \exp(-\mu\Psi(\mathbf{f})), \quad (5.13)$$

where $\Psi(\mathbf{f})$ is as in (5.8), and μ is a constant.

Now, we first substitute (5.12) and (5.13) into (5.10), and then convert the maximization to a minimization through a sign change. The MAP estimation problem then reduces to the following:

$$\hat{\mathbf{f}}_{\text{MAP}} = \arg \min_{\mathbf{f}} \left[\frac{1}{2\sigma^2}\|(\mathbf{g} - \mathbf{T}\mathbf{f})\|_2^2 + \mu\Psi(\mathbf{f}) \right]$$

Now, letting $\mu = 1/2\sigma^2$, and substituting for $\Psi(\mathbf{f})$, we obtain:

$$\hat{\mathbf{f}}_{\text{MAP}} = \arg \min_{\mathbf{f}} \left[\|\mathbf{g} - \mathbf{T}\mathbf{f}\|_2^2 + \lambda_1^2\|\mathbf{f}\|_k^k + \lambda_2^2\|\mathbf{D}|\mathbf{f}|\|_k^k \right]. \quad (5.14)$$

Note that the function to be minimized for the MAP solution is nothing but the objective function in (5.9).

We see that the second and third terms in (5.9) and (5.14) pertain to the inclusion of prior information from a regularized reconstruction point of view, and from an estimation-theoretic point of view respectively. In the former view, they act as feature-enhancement constraints as components of the regularizer, and in the latter they are tied into the prior probability density function $p_{\mathbf{f}}(\mathbf{f})$. Now, let us provide some insight for the expected impact of the use of these terms from a statistical standpoint. Let us

assume we use only the point-based feature enhancement term $\|\mathbf{f}\|_k^k$ (hence $\lambda_2 = 0$) in $\Psi(\mathbf{f})$. Then, we have:

$$p_{\mathbf{f}}(\mathbf{f}) \propto \exp(-\mu\lambda_1^2\|\mathbf{f}\|_k^k) \quad (5.15)$$

As seen in this expression, the use of a single prior term based on $\|\mathbf{f}\|_k^k$, would be equivalent to a prior model on \mathbf{f} which assumes independent identically distributed pixels with a circular generalized Gaussian density [74]. Note that with this prior, phase is assumed to be uniformly distributed. Two particular cases of the density would be, $k = 2$ leading to a Gaussian prior, and $k = 1$ leading to a Laplacian prior for each complex-valued pixel. Figure 5.10 illustrates such densities for a complex-valued scalar f . Figure 5.11 shows 1-D versions of these densities for ease of interpretation. Note that the Gaussian distribution assigns very low probabilities to large pixel amplitudes, hence it prevents large noise artifacts in the image, but while doing that, it also penalizes large values which might actually be due to objects present in the scene. A heavier-tailed distribution can assign a bigger probability for large reflectivity magnitudes, as shown in Figures 5.10 and 5.11 for the Laplacian case. Note that by choosing $k < 1$ in our framework, we can obtain even heavier-tailed densities than Laplacian, and strengthen this effect. As we will experimentally demonstrate, such a choice can be helpful for the objective of noise suppression and scatterer enhancement. Finally, we should also mention that, if ground truth is available for a particular type of scene, the shape of the probability density function estimated from such data can also be used in the determination of the prior model to be used for reconstruction purposes.

Similar to the point-based feature-enhancement case, we can also consider statistical models for the region-based feature enhancement term $\|\mathbf{D}\mathbf{f}\|_k^k$. In this case we would have an i.i.d. Gaussian distribution for the derivatives of the field magnitude

when $k = 2$ and a heavier-tailed density when $k < 2$. The $k = 2$ case may impose a field where large values of the gradient are severely penalized. This can result in suppression or blurring of the edges where there actually are large deviations in the underlying scene. As we use a smaller value of k however, the penalty on very large deviations is reduced while still keeping a considerable amount of penalty on small deviations. Hence, in effect, formation of edges is allowed, and a smoothness penalty is imposed in regions where there are no edges.





Figure 5.10: Gaussian (top) and Laplacian (bottom) probability density functions for a complex-valued scalar f .




Figure 5.11: 1-D Gaussian and Laplacian probability density functions for a real-valued random variable.

Chapter 6

Efficient Solution of the Optimization Problem

In this chapter, we present a numerical solution method for the image reconstruction problem we have formulated in Chapter 5. In particular, we solve the optimization problem in (5.9), where the objective is to find the field \mathbf{f} , which minimizes the function:

$$J(\mathbf{f}) = \|\mathbf{g} - \mathbf{T}\mathbf{f}\|_2^2 + \lambda_1^2 \|\mathbf{f}\|_k^k + \lambda_2^2 \|\mathbf{D}|\mathbf{f}|\|_k^k. \quad (6.1)$$

We propose an efficient iterative algorithm, whose structure gives some insight regarding the feature-enhancement properties of our approach. We then show that this algorithm is an extension of half-quadratic regularization techniques described in Section 4.3. The final portion of this chapter contains experimental results, demonstrating various aspects of feature-enhanced imaging.

6.1 Quasi-Newton-based Algorithm

In order to avoid problems due to non-differentiability of the ℓ_k -norm around the origin when $k \leq 1$, we will use the following smooth approximation to the ℓ_k -norm in (6.1) [60]:

$$\|\mathbf{z}\|_k^k \approx \sum_{i=1}^K (|(\mathbf{z})_i|^2 + \epsilon)^{k/2} \quad (6.2)$$

where $\epsilon \geq 0$ is a small constant, K is the length of the complex vector \mathbf{z} , and $(\mathbf{z})_i$ denotes its i -th element. For numerical purposes, we thus will use the following slightly modified cost function:

$$J_\epsilon(\mathbf{f}) = \|\mathbf{g} - \mathbf{T}\mathbf{f}\|_2^2 + \lambda_1^2 \sum_{i=1}^N (|(\mathbf{f})_i|^2 + \epsilon)^{k/2} + \lambda_2^2 \sum_{i=1}^M (|(\mathbf{D}|\mathbf{f}|)_i|^2 + \epsilon)^{k/2}. \quad (6.3)$$

Note that $J_\epsilon(\mathbf{f}) \rightarrow J(\mathbf{f})$ as $\epsilon \rightarrow 0$. A simple, closed-form solution for the minimizer of $J(\mathbf{f})$ or $J_\epsilon(\mathbf{f})$ does not exist in general, so numerical optimization techniques must be used.

Standard numerical optimization techniques like Newton's method or quasi-Newton methods with a conventional Hessian update scheme, such as Davidon-Fletcher-Powell [76] or Broyden-Fletcher-Goldfarb-Shanno [76] methods have been shown to perform poorly in optimization problems involving non-quadratic constraints, which are special cases of (6.3) [77], and we have observed this behavior as well. This precludes the use of such standard methods here. The additional presence of a constraint on the magnitude of \mathbf{f} in our case makes the problem even more difficult. To overcome these obstacles, we develop a new and efficient quasi-Newton method to solve problems of the form (6.3) for complex-valued fields. This algorithm is based on a new Hessian update scheme, which can be obtained by extending ideas from half-quadratic regularization [61] to account for the complex-valued nature of the SAR problem and the associated prior information terms, as will be discussed in

Section 6.3. This new Hessian approximation and update strategy is matched to the structure of the SAR problem in (6.3). The resulting new optimization algorithm is a non-trivial extension of existing numerical schemes, and provides an efficient, robust solution.

In order to develop our scheme, we use a structure which effectively deals with both the complex-valued nature of \mathbf{f} and the non-linearity associated with $|\mathbf{f}|$. We first take the gradient of (6.3) with respect to the real and imaginary parts of \mathbf{f} . This yields a gradient vector of length $2N$. We then put this vector into a compact form, by defining a complex-valued gradient vector of length N , whose real and imaginary components contain the derivatives with respect to the real and imaginary parts of \mathbf{f} respectively. This requires substantial manipulation, which is described in Appendix A.2. The resulting compact gradient can be placed in the following convenient form:

$$\nabla J_\epsilon(\mathbf{f}) = \mathbf{H}(\mathbf{f})\mathbf{f} - 2\mathbf{T}^H \mathbf{g} \quad (6.4)$$

where:

$$\mathbf{H}(\mathbf{f}) \triangleq 2\mathbf{T}^H \mathbf{T} + k\lambda_1^2 \mathbf{\Lambda}_1(\mathbf{f}) + k\lambda_2^2 \mathbf{\Phi}^H(\mathbf{f}) \mathbf{D}^T \mathbf{\Lambda}_2(\mathbf{f}) \mathbf{D} \mathbf{\Phi}(\mathbf{f}) \quad (6.5)$$

$$\begin{aligned} \mathbf{\Lambda}_1(\mathbf{f}) &\triangleq \text{diag} \left\{ \frac{1}{(|(\mathbf{f})_i|^2 + \epsilon)^{1-k/2}} \right\} \\ \mathbf{\Lambda}_2(\mathbf{f}) &\triangleq \text{diag} \left\{ \frac{1}{(|(\mathbf{D}\mathbf{f})_i|^2 + \epsilon)^{1-k/2}} \right\} \\ \mathbf{\Phi}(\mathbf{f}) &\triangleq \text{diag} \{ \exp(-j\phi[(\mathbf{f})_i]) \} \end{aligned} \quad (6.6)$$

Here $\phi[(\mathbf{f})_i]$ denotes the phase of the complex number $(\mathbf{f})_i$, $(\cdot)^H$ denotes the Hermitian of a matrix, and $\text{diag}\{\cdot\}$ is a diagonal matrix whose i -th diagonal element is given by the expression inside the brackets.

Examining the gradient expression (6.4), the term $\mathbf{H}(\mathbf{f})$ resembles a “coefficient” matrix multiplying \mathbf{f} . As a result, we use $\mathbf{H}(\mathbf{f})$ as an approximation to the Hessian. Note that this Hessian approximation depends on \mathbf{f} itself. We use this approximate Hessian $\mathbf{H}(\mathbf{f})$ in the following quasi-Newton iteration:

$$\hat{\mathbf{f}}^{(n+1)} = \hat{\mathbf{f}}^{(n)} - \gamma \left[\mathbf{H} \left(\hat{\mathbf{f}}^{(n)} \right) \right]^{-1} \nabla J_\epsilon(\hat{\mathbf{f}}^{(n)}) \quad (6.7)$$

where γ is the step size. After substituting (6.4) into (6.7) and rearranging, we obtain our new iterative algorithm:

$$\mathbf{H} \left(\hat{\mathbf{f}}^{(n)} \right) \hat{\mathbf{f}}^{(n+1)} = (1 - \gamma) \mathbf{H} \left(\hat{\mathbf{f}}^{(n)} \right) \hat{\mathbf{f}}^{(n)} + \gamma 2\mathbf{T}^H \mathbf{g} \quad (6.8)$$

We run the iteration (6.8) until

$$\frac{\|\hat{\mathbf{f}}^{(n+1)} - \hat{\mathbf{f}}^{(n)}\|_2^2}{\|\hat{\mathbf{f}}^{(n)}\|_2^2} < \delta, \quad (6.9)$$

where $\delta > 0$ is a small constant.

At the $(n+1)$ -st iteration, the only unknown in (6.8) is $\hat{\mathbf{f}}^{(n+1)}$. Hence (6.8) defines the iterate $\hat{\mathbf{f}}^{(n+1)}$ implicitly as the solution of a linear set of equations, in the form $\mathbf{H}(\hat{\mathbf{f}}^{(n)}) \hat{\mathbf{f}}^{(n+1)} = \mathbf{v}$, where $\mathbf{v} \triangleq (1 - \gamma) \mathbf{H}(\hat{\mathbf{f}}^{(n)}) \hat{\mathbf{f}}^{(n)} + \gamma 2\mathbf{T}^H \mathbf{g}$. The coefficient matrix $\mathbf{H}(\hat{\mathbf{f}}^{(n)})$ of this set is sparse, Hermitian and positive semi-definite, and hence these equations may themselves be efficiently solved using iterative approaches. We use the conjugate gradient (CG) algorithm for the solution of this linear problem. We terminate the CG algorithm when the ratio of the ℓ_2 -norm of the residual $\mathbf{H}(\hat{\mathbf{f}}^{(n)}) \hat{\mathbf{f}}^{(n+1)} - \mathbf{v}$ to the ℓ_2 -norm of the right-hand-side vector \mathbf{v} becomes smaller than a threshold $\delta_{CG} > 0$ [78]. A conceptual diagram of this algorithm is shown in Figure 6.1.

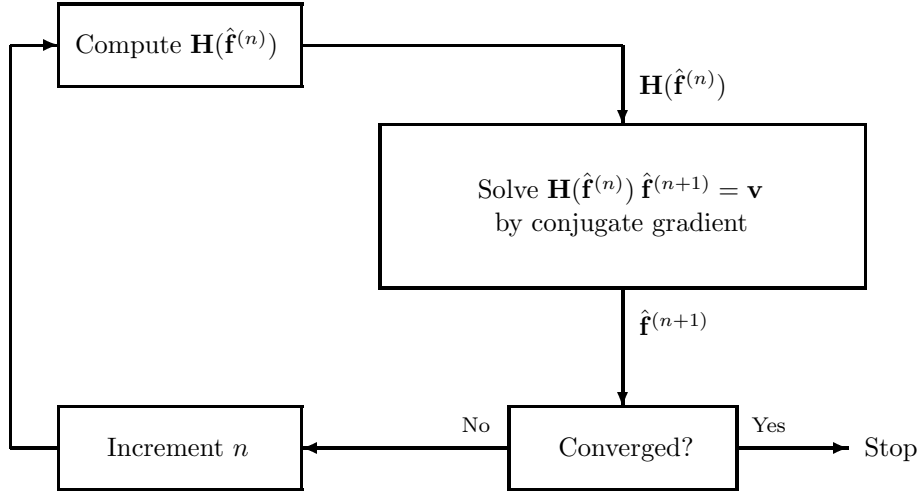


Figure 6.1: Quasi-Newton-based iterative algorithm.

6.2 Auxiliary Processes and Feature-Enhancement

The structure of the above algorithm provides some insight into the expected feature-preserving behavior of the approach. For the sake of simplicity, let us assume that $\gamma = 1$. Then, the solution of (6.8) is also the minimizer of the following quadratic function with respect to \mathbf{f} :

$$\|\mathbf{g} - \mathbf{T}\mathbf{f}\|_2^2 + \frac{k\lambda_1^2}{2}\mathbf{f}^H\mathbf{\Lambda}_1(\hat{\mathbf{f}}^{(n)})\mathbf{f} + \frac{k\lambda_2^2}{2}\mathbf{f}^H\mathbf{\Phi}^H(\hat{\mathbf{f}}^{(n)})\mathbf{D}^T\mathbf{\Lambda}_2(\hat{\mathbf{f}}^{(n)})\mathbf{D}\mathbf{\Phi}(\hat{\mathbf{f}}^{(n)})\mathbf{f} \quad (6.10)$$

In this quadratic problem, the quantities $\mathbf{\Lambda}_j(\hat{\mathbf{f}}^{(n)})$ ($j \in \{1, 2\}$) act as spatially varying weighting matrices. The diagonal elements of these matrices correspond to auxiliary processes in the context of half-quadratic regularization. First consider the role of $\mathbf{\Lambda}_2(\hat{\mathbf{f}}^{(n)})$. At a location where there is an edge in the field, $|(\mathbf{D}|\hat{\mathbf{f}}^{(n)}|)_i|^2$ will be large, hence the corresponding element of $\mathbf{\Lambda}_2(\hat{\mathbf{f}}^{(n)})$ will be small, essentially suppressing the derivative penalty at that location through (6.10), and preventing smoothing across the edge. On the other hand, smaller fluctuations in the field (assumed to be the result of noise and artifacts and not physical edges) will be suppressed. Next, consider the

role of $\Lambda_1(\hat{\mathbf{f}}^{(n)})$. At a location where there is a strong scatterer in the field, $|(\hat{\mathbf{f}}^{(n)})_i|^2$ will be large, hence the corresponding element of $\Lambda_1(\hat{\mathbf{f}}^{(n)})$ will be small, essentially removing the energy-type penalty at that location through (6.10), and preventing the suppression of the scatterer. On the other hand, at locations where $|(\hat{\mathbf{f}}^{(n)})_i|^2$ is not large enough to be “declared” as a physical scatterer, the energy-type penalty is in effect.

Overall, the auxiliary processes associated with Λ_1 and Λ_2 can be viewed as foreground/background and edge maps respectively, and may be useful for interpretation of the formed image. Our algorithm generates these processes during the iteration process without any additional cost.

6.3 Ties to Half-Quadratic Regularization

We now establish the connection between the quasi-Newton-based algorithm of Section 6.1 and half-quadratic regularization methods discussed in Section 4.3. Note that there are two major differences between the type of cost functions considered in Chapter 4, and our cost function in (6.3). First, (6.3) contains *complex-valued* vectors and matrices, as opposed to the real-valued formulation of standard image processing problems. Second, the cost function in (6.3) contains roughness penalties on the *magnitude* of the field through the terms $\mathbf{D}|\mathbf{f}|$, which are more complicated than standard linear constraints of the form $\mathbf{D}\mathbf{f}$.

In establishing the link to half-quadratic regularization in this section, we consider a special case of the cost function (6.3) by setting $k = 1$ and $\lambda_1 = 0$:

$$J_\epsilon(\mathbf{f}) = \|\mathbf{g} - \mathbf{T}\mathbf{f}\|_2^2 + \lambda_2^2 \sum_{i=1}^M (|(\mathbf{D}|\mathbf{f}|)_i|^2 + \epsilon)^{1/2}. \quad (6.11)$$

This simplification is aimed at keeping the algebra in this section simple, and it does

not cause any loss of generality. Now, let us develop a half-quadratic approach to minimize $J_\epsilon(\mathbf{f})$ of (6.11). To this end, consider the following augmented cost function:

$$K(\mathbf{f}, \mathbf{b}, \mathbf{s}) = \|\mathbf{g} - \mathbf{T}\mathbf{f}\|_2^2 + \lambda_2^2 \sum_{i=1}^M \left[\mathbf{b}_i (|(\mathbf{D}\mathbf{S}\mathbf{f})_i|^2 + \epsilon) + \frac{1}{4\mathbf{b}_i} \right] \quad (6.12)$$

where

$$\mathbf{S} = \text{diag}\{\exp(-js_l)\}, \quad (6.13)$$

with s_l being the l -th element of the vector \mathbf{s} . In this formulation, \mathbf{b} and \mathbf{s} act as auxiliary vectors. Note that $K(\mathbf{f}, \mathbf{b}, \mathbf{s})$ is a quadratic function of \mathbf{f} , hence it is half-quadratic. In Appendix A.3, we show the following relationship between $K(\mathbf{f}, \mathbf{b}, \mathbf{s})$ of (6.12), and $J_\epsilon(\mathbf{f})$ of (6.11):

$$\inf_{\mathbf{b}, \mathbf{s}} K(\mathbf{f}, \mathbf{b}, \mathbf{s}) = J_\epsilon(\mathbf{f}). \quad (6.14)$$

Based on (6.14), $J_\epsilon(\mathbf{f})$ and $K(\mathbf{f}, \mathbf{b}, \mathbf{s})$ share the same minima in \mathbf{f} . Then, we can use a block coordinate descent scheme on $K(\mathbf{f}, \mathbf{b}, \mathbf{s})$, to find the field \mathbf{f} that minimizes $J_\epsilon(\mathbf{f})$:

$$\hat{\mathbf{s}}^{(n+1)} = \arg \min_{\mathbf{s}} K(\hat{\mathbf{f}}^{(n)}, \hat{\mathbf{b}}^{(n)}, \mathbf{s}) \quad (6.15)$$

$$\hat{\mathbf{b}}^{(n+1)} = \arg \min_{\mathbf{b}} K(\hat{\mathbf{f}}^{(n)}, \mathbf{b}, \hat{\mathbf{s}}^{(n+1)}) \quad (6.16)$$

$$\hat{\mathbf{f}}^{(n+1)} = \arg \min_{\mathbf{f}} K(\mathbf{f}, \hat{\mathbf{b}}^{(n+1)}, \hat{\mathbf{s}}^{(n+1)}) \quad (6.17)$$

Using the results of Appendix A.3, we obtain:

$$\hat{\mathbf{s}}_i^{(n+1)} = \phi[(\hat{\mathbf{f}}^{(n)})_i] \quad (6.18)$$

$$\hat{\mathbf{b}}_i^{(n+1)} = \frac{1}{2 \left[(\mathbf{D}\hat{\mathbf{S}}^{(n+1)}\hat{\mathbf{f}}^{(n)})_i^2 + \epsilon \right]^{1/2}} \quad (6.19)$$

$$\left[2\mathbf{T}^H\mathbf{T} + 2\lambda_2^2(\hat{\mathbf{S}}^{(n+1)})^H\mathbf{D}^T \text{diag} \left\{ \hat{\mathbf{b}}_i^{(n+1)} \right\} \mathbf{D}\hat{\mathbf{S}}^{(n+1)} \right] \hat{\mathbf{f}}^{(n+1)} = 2\mathbf{T}^H\mathbf{g} \quad (6.20)$$

Note that $\hat{\mathbf{S}}^{(n+1)} = \Phi(\hat{\mathbf{f}}^{(n)})$, where $\Phi(\cdot)$ is as defined in (6.6). Now, substituting (6.18), and (6.19) into (6.20) results in the following fixed point iteration for $\hat{\mathbf{f}}^{(n+1)}$:

$$\left[2\mathbf{T}^H\mathbf{T} + \lambda_2^2\Phi^H(\hat{\mathbf{f}}^{(n)})\mathbf{D}^T \text{diag} \left\{ \frac{1}{\left[(\mathbf{D}|\hat{\mathbf{f}}^{(n)})_i \right]^2 + \epsilon} \right\} \mathbf{D}\Phi(\hat{\mathbf{f}}^{(n)}) \right] \hat{\mathbf{f}}^{(n+1)} = 2\mathbf{T}^H\mathbf{g}. \quad (6.21)$$

Note that this iterative algorithm obtained through half-quadratic regularization is nothing but the quasi-Newton scheme of (6.8) with $\gamma = 1$, for the special case of $k = 1$, $\lambda_1 = 0$, as we have assumed in this section. Hence, our iterative scheme for SAR image formation is an extension of half-quadratic regularization methods.

This association is important for two reasons. First, it allows the use of previous results concerning e.g. the convergence properties of half-quadratic regularization methods, as in [63, 79], for our algorithm as well. Second, it extends the use of half-quadratic regularization methods to problems involving complex-valued, random-phase fields.

6.4 Experimental Results

We demonstrate the effectiveness of our method on synthetic and real SAR scenes. We present examples for both point-based feature enhancement and region-based feature enhancement, and compare these results to conventional reconstructions.

6.4.1 Algorithm Initialization and Parameter Selection

For all the examples, we initialize our algorithm with $\hat{\mathbf{f}}^{(0)} = a\mathbf{T}^H\mathbf{g}$. Here a is a normalization factor, and we choose a to be the reciprocal of the value that the diagonal elements of the matrix $\mathbf{T}^H\mathbf{T}$ take. Note that this initialization is obtained by application of the adjoint of the observation matrix to the data. From a tomographic standpoint, this can also be viewed as the backprojection operation. Another simple initialization is the conventional polar format reconstruction. The difference between these two initialization choices is minimal. Sensitivity of our algorithm to initial conditions will be discussed in Section 6.4.10.

Although the iterative scheme in (6.8) allows a variable step size, we use a fixed step size of $\gamma = 1$ in our examples. In our experience, the algorithm has always converged with this choice, hence varying (reducing) the step size has not been necessary. We choose the approximation parameter ϵ in (6.2) between 10^{-7} and 10^{-5} , depending on the particular data set, so that it is small enough not to affect the behavior of the solution. For the termination conditions of (6.8), we use $\delta = 10^{-6}$ and a CG tolerance of $\delta_{CG} = 10^{-3}$. We choose the feature accentuation parameters λ_1 and λ_2 in (6.3) based on subjective qualitative assessment of the formed imagery, coupled with our imaging goals, as described in subsequent sections. Our experience on a large database of SAR images composed of similar scenes is that one set of parameters chosen on a single image can be used for the entire data set. Automatic selection of these regularization parameters is beyond the scope of our work and information on this topic can be found in [53, 57]. In all the examples, we show the magnitude (in dB) of the reconstructed complex-valued field. We increase the sparsity of $\mathbf{H}(\hat{\mathbf{f}}^{(n)})$ by neglecting elements in $\mathbf{T}^H\mathbf{T}$ whose magnitudes are smaller than 1% of the largest element.

6.4.2 Synthetic Scene Reconstructions

6.4.2.1 Point-based Feature Enhancement

First, we demonstrate the superresolution capability of our method on a simple synthetic scene composed of eight single-pixel scatterers with unit reflectivity magnitude and random, uniform phase. The 3-D mesh plot of the magnitude of this 16×16 pixel scene is shown in Figure 6.2(a). We simulate SAR returns from this ideal scene such that the bandwidth of the data supports a resolution cell of 2×2 pixels. The conventional SAR reconstruction in Figure 6.2(b) cannot resolve four of the scatterers falling into one resolution cell, and suffers from sidelobes. In this example, we want to accentuate points, hence we set $\lambda_2 = 0$ in (6.3). Figure 6.2(c) and (d) show the results of our method with two different choices of k in (6.3). In these reconstructions, all the scatterers are resolved, background is suppressed, and peak reflectivity magnitudes are preserved (0.9552 in (c) and 0.9947 in (d)). Next, we show the performance of our method in the face of degradation in data quality. To this end, we add complex Gaussian noise to the simulated SAR returns, so that the signal-to-noise ratio (SNR) is 10 decibels (dB). We take the SNR to be the variance ratio of the noise-free data to noise in dB. The resulting conventional and point-enhanced reconstructions are shown in Figure 6.3. Our method is still able to resolve the scatterers and suppress the effect of noise.

6.4.2.2 Region-based Feature Enhancement

We now demonstrate region-based feature enhancement in our framework through the reconstruction of the synthetic scene of Figure 6.4(a). This scene represents an object (e.g. a vehicle) in a uniform background (the gray area in the image). The white portions of the scene are where there is scattering from the object, and the black area is where the shadow of the object falls. Here, we display the magnitude (in

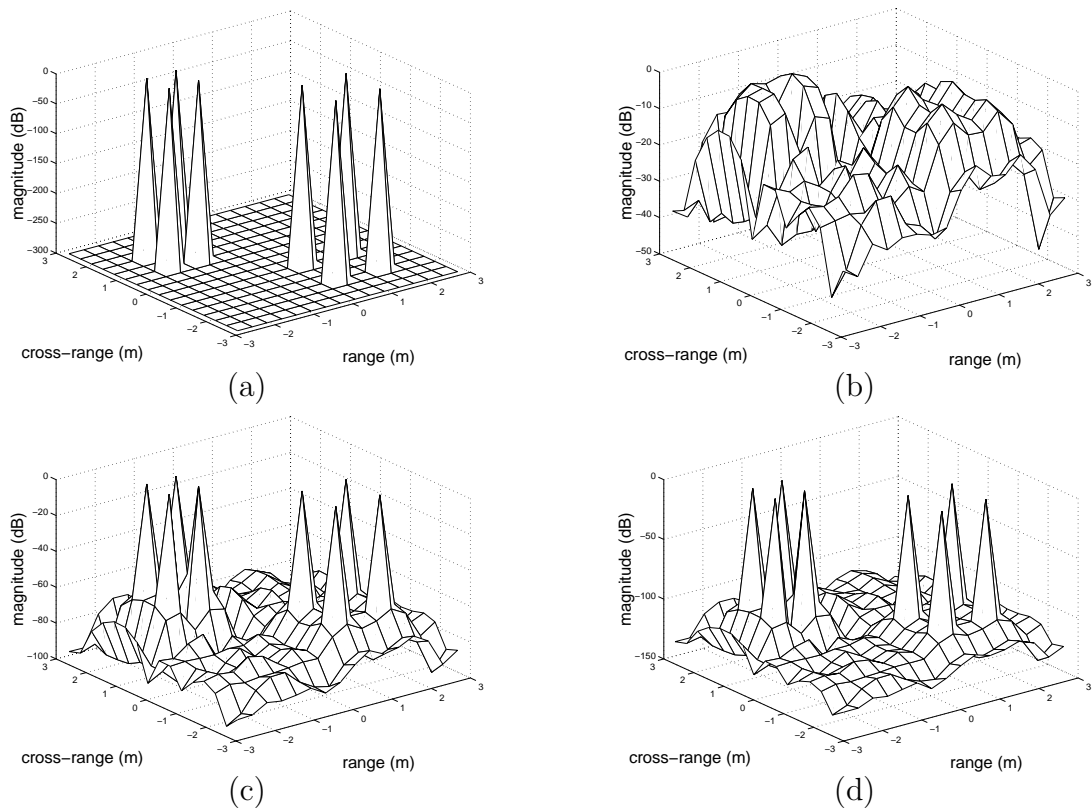


Figure 6.2: Mesh plots of synthetic point scatterer reconstructions. (a) Truth. (b) Conventional method. (c) Proposed method with $k = 0.8$, $\lambda_1 = 1$, $\lambda_2 = 0$. (d) Proposed method with $k = 0.1$, $\lambda_1 = 1$, $\lambda_2 = 0$.

dB) of the synthetic reflectivities only, however the reflectivities are complex-valued, with random, uniform phase.

We simulate the SAR returns from this scene and corrupt the observations by additive complex Gaussian noise. We consider two noise levels of $\text{SNR} = 30$ dB and $\text{SNR} = 10$ dB. The conventional and region-enhanced reconstructions from such observations are shown in the middle and bottom rows of Figure 6.4. The conventional reconstructions suffer from artifacts. The region-enhanced reconstructions provide a better representation of the scene through homogeneous regions and preserved region

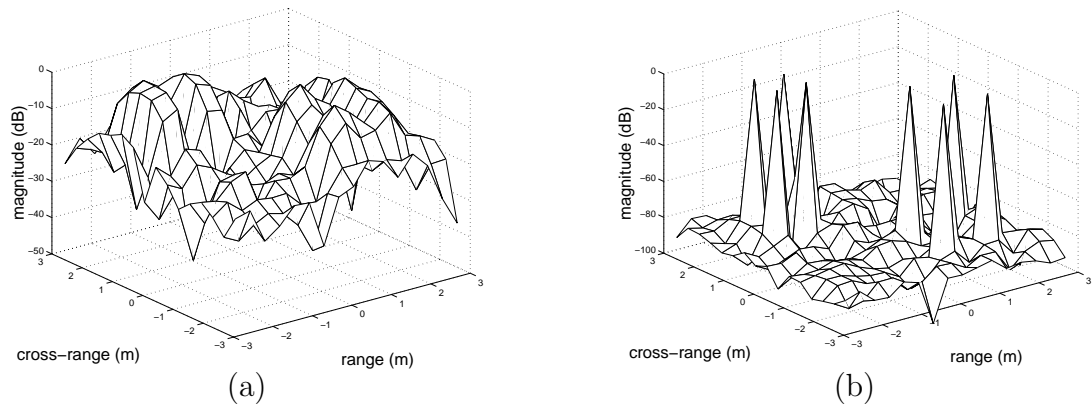


Figure 6.3: Mesh plots of synthetic point scatterer reconstructions from noisy data (SNR = 10 dB). (a) Conventional method. (b) Proposed method with $k = 0.8$, $\lambda_1 = 2$, $\lambda_2 = 0$.

boundaries.




Figure 6.4: Reconstruction of the synthetic target image at the top, from data with $\text{SNR} = 30$ dB (middle row), and $\text{SNR} = 10$ dB (bottom row). (a) Original scene. (b) Conventional method. (c) Proposed method with $k = 1$, $\lambda_1 = 0$, $\lambda_2 = 1.4$. (d) Proposed method with $k = 0.7$, $\lambda_1 = 0$, $\lambda_2 = 1.4$. (e) Conventional method. (f) Proposed method with $k = 1$, $\lambda_1 = 1.4$, $\lambda_2 = 2.5$. (g) Proposed method with $k = 0.7$, $\lambda_1 = 1.4$, $\lambda_2 = 2$.

6.4.3 ADTS Data Reconstructions

We now show results on data from the MIT Lincoln Laboratory Advanced Detection Technology Sensor (ADTS) data set [3, 80]. Since the ADTS data set provides formed imagery only, we generate synthetic radar returns by computing Fourier transform samples on a polar grid and then using the resulting range profiles as the input to our reconstruction algorithm.

6.4.3.1 Point-based Feature Enhancement

First, we show point-based feature enhancement results. In this example, all of the reconstructed images consist of 64×64 complex-valued pixels. Since we want to accentuate point features, we set $\lambda_2 = 0$ in (6.3). The top row in Figure 6.5 contains 0.3 m resolution reconstructions of a scene containing an M48 tank. Our reconstructions appear to produce images with accentuated dominant peaks. Next, we reduce the bandwidth of the data equally in range and cross-range, and attempt to generate superresolution reconstructions. The middle and bottom rows in Figure 6.5 contain reconstructions where the resolution has been reduced to 0.6 m and 1.2 m respectively. Although precise superresolution arguments are not as easy for this complicated real SAR scene as for the synthetic scene of Figure 6.2, the peaks still appear to be better localized by our approach. Note that a particular parameter choice in our method (such as that associated with the image in the bottom row of Figure 6.5(b)) produces reconstructions which are visually very similar to the imagery obtained by the spectral estimation-based superresolution method of [32]. In Figure 6.6, we show reconstructions from lower quality data for the 0.3 m resolution case. Lower quality data are obtained by adding complex Gaussian noise to the ADTS SAR returns used for the example in Figure 6.5. If we treat those ADTS returns as the clean signal, the SNR in this case is 10 dB. The quality of the point-enhanced

reconstruction is not significantly affected by noise. For a quantitative analysis of the improvements in scattering center locating accuracy provided by our technique, please see Chapter 9.

6.4.3.2 Region-based Feature Enhancement

Now, we demonstrate formation of images with enhanced region-based features. Since we want to accentuate homogeneous regions, we set $\lambda_2 \geq \lambda_1$ in (6.3). We do not necessarily set $\lambda_1 = 0$, since we have observed that its presence can help in preserving the shadow regions in some cases. The images of this example consist of 128×128 complex-valued pixels. Figure 6.7 contains ADTS images of an M48 tank reconstructed by using the conventional method and by our proposed scheme for different choices of k . By choosing $k = 2$, our algorithm can produce reconstructions analogous to standard Tikhonov regularization, which we show for comparison. When $k \leq 1$, our method produces images where background fluctuations are suppressed, in contrast to the conventional image. Furthermore, this is achieved without compromising the sharp boundaries, unlike Tikhonov-type reconstructions. In Figure 6.8, we show reconstructions from lower quality data. The region-enhanced image is again able to provide smooth homogeneous regions with clear boundaries.

Similar observations apply to the reconstructions of a natural scene consisting of trees, corner reflectors, fields and a road, shown in Figure 6.9. Our method forms images in which the tree shapes and shadows are very distinguishable, and the background is quite smooth, whereas the conventional SAR image suffers from considerable amount of speckle.




Figure 6.5: Enhancement of point-based features. Resolution: top: 0.3 m, middle: 0.6 m, bottom: 1.2 m. (a) Conventional method. (b) Proposed method with $k = 1$, $\lambda_2 = 0$, and top: $\lambda_1 = 14$, middle: $\lambda_1 = 7$, bottom: $\lambda_1 = 10$. (c) Proposed method with $k < 1$ and $\lambda_2 = 0$, with top: $\lambda_1 = 14$, $k = 0.8$, middle: $\lambda_1 = 7$, $k = 0.8$, bottom: $\lambda_1 = 10$, $k = 0.95$.




Figure 6.6: Enhancement of point-based features from low-quality data (SNR = 10 dB). Resolution: 0.3 m. (a) Conventional method. (b) Proposed method with $k = 1$, $\lambda_1 = 14$, $\lambda_2 = 0$.

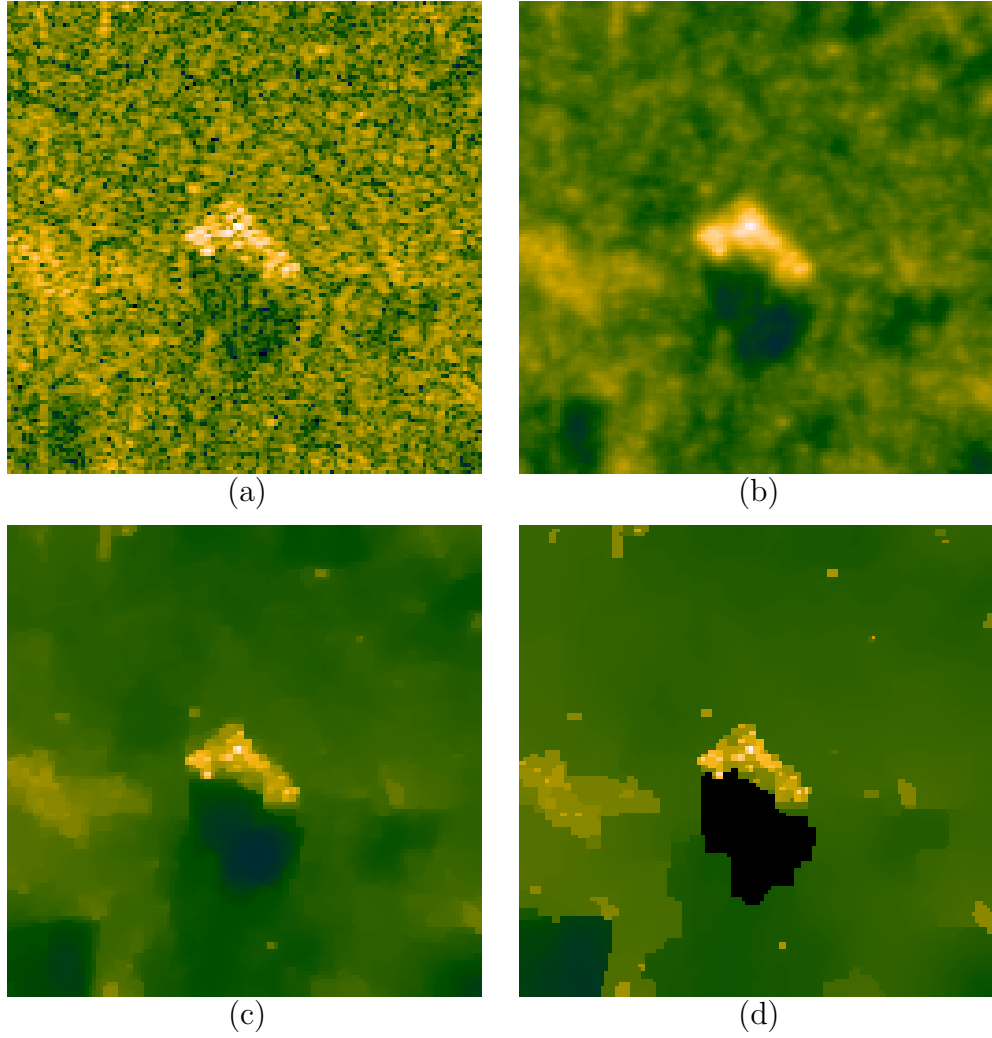


Figure 6.7: Enhancement of region-based features. (a) Conventional method. (b) Tikhonov-type reconstruction (i.e. $k = 2$) with $\lambda_1 = 4$, $\lambda_2 = 22$. (c) Proposed method with $k = 1$, and $\lambda_1 = 5$, $\lambda_2 = 9$. (d) Proposed method with $k = 0.7$, and $\lambda_1 = 4$, $\lambda_2 = 6$.




Figure 6.8: Enhancement of region-based features from low-quality data (SNR = 8 dB). (a) Conventional method. (b) Proposed method with $k = 0.7$, and $\lambda_1 = 5$, $\lambda_2 = 6$.




Figure 6.9: Enhancement of region-based features. (a) Conventional method. (b) Tikhonov-type reconstruction (i.e. $k = 2$) with $\lambda_1 = 4$, $\lambda_2 = 22$. (c) Proposed method with $k = 1$, and $\lambda_1 = 4$, $\lambda_2 = 5$. (d) Proposed method with $k = 0.7$, and bottom: $\lambda_1 = 4$, $\lambda_2 = 4$.

6.4.4 MSTAR Data Reconstructions

We now show results on data from the Moving and Stationary Target Acquisition and Recognition (MSTAR) public target data set [81].

6.4.4.1 Obtaining the Projectional Data

Similar to the ADTS data, the MSTAR data set also provides formed imagery only. In order to apply our image formation technique, we need the projectional data (phase histories or range profiles). The MSTAR data contain some information about how the images are formed. Hence we can attempt to undo the final steps of MSTAR image formation using that information to obtain a closer approximation to the raw collected data, as done in [82].

We will now highlight the main blocks of this processing. MSTAR images are formed by taking a 2-D inverse FFT of the Taylor-windowed, zero-padded phase history data on a rectangular grid. To undo this image formation process, we can first take the 2-D FFT of the 128×128 images, and shift the transformed signal such that small frequencies are at the center. A mesh plot of the magnitude of the resulting 2-D signal for a sample MSTAR scene is shown in Figure 6.10. One can notice a band of points with significantly small values near the borders of the 2-D signal. We can then assume, as in [82] that these are due to zero-padding, and remove a band of width 28 around this signal to obtain a 100×100 signal. The next step is to undo the Taylor windowing. From the MSTAR file headers, we know that a 35 dB Taylor window [15] has been used, and we assume that the quality of approximation of the Taylor window used is $\bar{n} = 4$. The mesh plot of such a Taylor window is shown in Figure 6.11. The 100×100 data can then be divided by this 2-D window to obtain the unwindowed phase histories. Finally by a rectangular-to-polar resampling, the polar formatted phase histories can be obtained. Range profiles can be obtained through




Figure 6.10: Magnitude of the 2-D FFT of a sample MSTAR image.

1-D Fourier transforms.

6.4.4.2 Point-based Feature Enhancement

Figure 6.12 shows 100×100 conventional and point-enhanced MSTAR images. Images in the top row have been reconstructed from 100×100 data (resolution = 0.3 m), and images in the bottom row have been reconstructed from 50×50 data (resolution = 0.6 m). Point-enhanced imaging is able to localize the dominant peaks despite the resolution loss. In forming these images, we have used the data without removing the Taylor windowing. For comparison, Figure 6.13 shows the reconstructions for the 0.3 m resolution case, using the data after the removal of the Taylor window.




Figure 6.11: A 100×100 2-D Taylor window with 35 dB sidelobe suppression, and a quality of approximation of $\bar{n} = 4$.

6.4.4.3 Region-based Feature Enhancement

Figure 6.14 shows a conventional and a region-enhanced MSTAR image. Qualitatively, the behavior is very similar to what we have observed with the ADTS data. Region-enhanced imaging suppresses fluctuations in homogeneous regions, without destroying the region boundaries.




Figure 6.12: Enhancement of point-based features. Resolution: top: 0.3 m, bottom: 0.6 m. (a) Conventional method. (b) Proposed method with $k = 0.8$, $\lambda_1 = 7$, $\lambda_2 = 0$.




Figure 6.13: Results from data with Taylor windowing removed. Resolution: 0.3 m. (a) Conventional method. (b) Proposed method with $k = 0.8$, $\lambda_1 = 14$, $\lambda_2 = 0$.




Figure 6.14: Enhancement of region-based features. (a) Conventional method. (b) Proposed method with $k = 0.8$, $\lambda_1 = 1.7$, $\lambda_2 = 2.5$.

6.4.5 URISD Reconstructions

Our final examples are from the XPATCH-generated [83] University Research Initiative Synthetic Dataset (URISD) [84]. The URISD provides phase histories and range profiles, which we directly use as the input to our algorithm. Figure 6.15 contains the CAD model of a fire truck used for data generation, and the corresponding reconstructed images. The conventional image in Figure 6.15(b) suffers from large sidelobes. Sidelobes can be suppressed by windowing the data prior to image formation, as we have seen for the MSTAR data in Section 6.4.4, however this may reduce the effective resolution in the formed image. Our reconstruction with a point-based prior is shown in Figure 6.15(c), and achieves sidelobe suppression, as well as increased resolvability of point scatterers. Our method with a region-based prior, on the other hand, produces an image with an enhanced object shape, as shown in Figure 6.15(d).

6.4.6 Auxiliary Processes

So far, we have shown only the images reconstructed by our method. However, as discussed in Section 6.2, our algorithm also produces auxiliary processes in addition to the reconstructed field. Here, we will consider one of these processes, consisting of the diagonal elements of the matrix Λ_2 , as defined in (6.6). Note that this process acts as an indicator of the edges in the scene. In Figure 6.16 we show such edge fields corresponding to the region-enhanced reconstructions of Figure 6.4. These images are displayed in a dB scale.




Figure 6.15: Results with the URISD. (a) CAD model of the fire truck. (b) Conventional reconstruction. (c) Enhancement of point-based features with $k = 0.8$, $\lambda_1 = 22$, $\lambda_2 = 0$. (d) Enhancement of region-based features with $k = 0.8$, $\lambda_1 = 5$, $\lambda_2 = 9$.




Figure 6.16: Auxiliary processes representing the edges in the region-enhanced reconstructions of Figure 6.4. (a) SNR = 30 dB, $k = 1$, $\lambda_1 = 0$, $\lambda_2 = 1.4$. (b) SNR = 30 dB, $k = 0.7$, $\lambda_1 = 0$, $\lambda_2 = 1.4$. (c) SNR = 10 dB, $k = 1$, $\lambda_1 = 1.4$, $\lambda_2 = 2.5$. (d) SNR = 10 dB, $k = 0.7$, $\lambda_1 = 1.4$, $\lambda_2 = 2$.

6.4.7 Results of Smoothing Real and Imaginary Components

When choosing a region-based feature enhancement term in Section 5.5, we have argued that due to the random-phase nature of the SAR reflectivities, the smoothing constraint should be explicitly imposed on the magnitude of the field through $\mathbf{D}|\mathbf{f}|$, rather than on the real and imaginary components of the field through $\mathbf{D}\mathbf{f}$.

We experimentally demonstrate the results of using these two types of smoothness constraints now. First, we illustrate this phenomenon on a 1-D signal denoising problem. We consider a complex signal with a piecewise smooth magnitude and random, uniform phase. We add complex Gaussian noise to this signal to obtain the noisy signal, \mathbf{g} , shown in the top row of Figure 6.17. Figure 6.17(a) shows the magnitude of the noisy signal, while Figure 6.17(b) and (c) show the real and imaginary components respectively. Treating this signal as the observation, our objective then is to denoise it to reconstruct the original piecewise constant true signal magnitude. We present two approaches. The first approach finds the estimate $\hat{\mathbf{f}}$ by minimizing the following objective function:

$$J_\epsilon(\mathbf{f}) = \|\mathbf{g} - \mathbf{f}\|_2^2 + \lambda_2^2 \sum_{i=1}^M (|(\mathbf{D}\mathbf{f})_i|^2 + \epsilon)^{k/2} \quad (6.22)$$

Note that the smoothness constraint is based on $\mathbf{D}\mathbf{f}$. The results of this approach, with $k = 1$, are shown in the middle row of Figure 6.17. Clearly, the objective of recovering the piecewise smooth signal has not been achieved. There is some smoothing on the real and imaginary components, which can be observed by comparing the middle portions of the reconstructed signals in (e) and (f) with their counterparts in the noisy observation of (b) and (c).

The alternative approach involves constraints explicitly on the magnitude of the

field, and finds the estimate $\hat{\mathbf{f}}$ by minimizing the following objective function:

$$J_\epsilon(\mathbf{f}) = \|\mathbf{g} - \mathbf{f}\|_2^2 + \lambda_2^2 \sum_{i=1}^M (|\mathbf{D}\mathbf{f}|_i|^2 + \epsilon)^{k/2} \quad (6.23)$$

The resulting denoised signal, with $k = 1$, is shown in the bottom row of Figure 6.17. This reconstructed signal magnitude provides the piecewise smooth signal, as desired.

Next, we show a similar comparison for a SAR image reconstruction example. Figure 6.18(b) contains the magnitude of an image reconstructed by our technique. This is the example shown in Figure 6.7(c) before. Now, if we replace the term $\mathbf{D}|\mathbf{f}|$ in (6.3) with $\mathbf{D}\mathbf{f}$, and use the resulting objective function in image formation, we obtain the reconstructed field in Figure 6.18(a). The parameter values used for this reconstruction are the same as those used for the one in Figure 6.18(b). This example demonstrates that if the smoothing constraint is not explicitly imposed on the magnitude of the field, the resulting reconstruction suffers from artifacts caused by the random-phase nature of the reflectivities.

6.4.8 Comparison with Post-Processing

Our technique starts from the projectional SAR observations, and forms feature-enhanced images. Here we illustrate that such feature enhancement cannot in general be achieved alternatively by post-processing operations on a conventionally reconstructed image. By “post-processing,” we refer to techniques that do not take into account a proper observation model for the data at hand.

Let us consider the point enhanced superresolution image in the middle row of Figure 6.5(c), repeated in the top row of Figure 6.19. The question we pose here is whether we can obtain such localization of point scatterers starting from the

conventional image $\hat{\mathbf{f}}_{\text{CONV}}$ (shown in the middle row of Figure 6.5(a)) through the minimization of an objective function of the following form:

$$J_{\epsilon}(\mathbf{f}) = \|\hat{\mathbf{f}}_{\text{CONV}} - \mathbf{f}\|_2^2 + \lambda_1^2 \sum_{i=1}^N (|\mathbf{f}_i|^2 + \epsilon)^{k/2} \quad (6.24)$$

In the bottom row of Figure 6.19, we show the results of such post-processing of the conventional image. The three images correspond to small, medium, and large regularization parameter choices. Neither of these images can provide the scatterer localization power of the image in the top row.

However, we should note that, if an observation model taking into account the SAR data collection and image formation processes is incorporated in the processing, objectives similar to ours can be achieved by acting on a conventional image as well.

6.4.9 Bandwidth Extrapolation Property

In Section 2.3, we have seen that the resolution of a SAR image is determined by the bandwidth of the spatial frequency domain data available. As discussed in Section 3.4, a class of recently proposed SAR image formation methods aim to produce superresolution images directly by extrapolating these available data samples in the spatial Fourier domain. Our approach to superresolution was different. However, naturally, our method also implicitly performs extrapolation. This is demonstrated in Figure 6.20. The top row in Figure 6.20 shows the conventional and superresolution images for the 1.2 m resolution example in Figure 6.5. The bottom row in Figure 6.20 displays the magnitude of the centered 2-D DFT coefficients of these images. The support of the significant DFT coefficients for the superresolution images is larger than that for the conventional image.




Figure 6.17: Comparison of smoothness constraints on the real and imaginary parts of a signal versus its magnitude, in the problem of denoising a random-phase signal with piecewise smooth magnitude. Top: noisy observations. Middle row: results of smoothing real and imaginary components. Bottom: results of smoothing the magnitudes directly. Left: magnitude of the signal. Middle column: real part of the signal. Right: imaginary part of the signal.




Figure 6.18: Comparison of smoothness constraints on the real and imaginary parts of the reflectivities versus their magnitude, for the example of Figure 6.7. (a) Reconstructed field by smoothing real and imaginary components. (b) Reconstructed field by smoothing the magnitudes directly.




Figure 6.19: Demonstration that simple post-processing techniques do not lead to robust feature enhancement. Top: Superresolution SAR image from Figure 6.5. Bottom: Results of post-processing the corresponding conventional image, based on (6.24), with $k = 0.8$. From left to right: $\lambda_1 = 1.4$, $\lambda_1 = 1.7$, $\lambda_1 = 2.2$.




Figure 6.20: Bandwidth extrapolation property of point-enhanced, superresolution imaging, for the 1.2 m resolution example of Figure 6.5. Top: reconstructed images. Bottom: magnitude of the DFT coefficients (zero frequency at the center of the image). (a) Conventional method. (b) Point-enhanced, $k = 1$. (c) Point-enhanced, $k = 0.8$.

6.4.10 Sensitivity to Initial Conditions

In previous sections, we have initialized our iterative image reconstruction scheme as described in Section 6.4.1, with an adjoint-based initial condition, which was also very similar to a conventional image. Here we address the issue of sensitivity of the reconstructions to the choice of the initial conditions, by showing the results of initializing our algorithm with zero initial conditions. We consider the examples of Figure 6.7(c) and (d). Note that for the reconstruction in Figure 6.7(c), we used $k = 1$, and for the reconstruction in Figure 6.7(d), we used $k = 0.7$.

In Figure 6.21, we show these images, as well as the corresponding reconstructions obtained by using zero initial conditions. For the $k = 1$ case, the difference between the reconstructions with different initial conditions is not visually noticeable. The values of the objective function at the solution point for the two initializations are also very close: $J^* = 1.7248 \times 10^5$ for the adjoint-based initialization, versus $J^* = 1.7972 \times 10^5$ for initialization with zeros. For the $k = 0.7$ case, the difference caused by the different initialization is more noticeable, but the two reconstructions are still quite similar. The values of the objective function in this case are $J^* = 1.7671 \times 10^5$, for the adjoint-based initial condition, and $J^* = 1.8998 \times 10^5$ for the zero initial condition.

6.4.11 Behavior of the Iterative Scheme

So far, we have shown only the final results of our iterative scheme. However, the path to this final solution is also of interest. Here, we consider the image reconstruction example of Figure 6.7(c), and examine various aspects of the iterative structure.

The first issue is the behavior in which the objective function progresses towards its minimum. Figure 6.22 shows the value of the objective function $J_\epsilon(\mathbf{f})$ of (6.3) as

a function of the iteration number. For this example, the algorithm has converged in 26 steps. The convergence tolerances used are $\delta = 10^{-6}$, and $\delta_{CG} = 10^{-3}$, as stated in Section 6.4.1. Based on Figure 6.22, the objective function quickly approaches its minimum value. In fact, 98.6% of the path from the initial value of the objective function to its final value is taken in the first three iterations.

Next, we consider the behavior of the ℓ_2 -norm of the gradient of the objective function $\nabla J_\epsilon(\mathbf{f})$ of (6.4), as a function of the iteration number. Figure 6.23 shows the evolution of the gradient norm with the iterations. Again, this plot confirms the fast approach of the algorithm to the convergence point.

Finally, we are interested in the behavior of the field $\hat{\mathbf{f}}^{(n)}$ itself, as a function of the iterations. Again, we consider the example of Figure 6.7(c). In Figure 6.24, we show the intermediate stages of the reconstructed field after the iterations $n = 1, 2, 3, 4, 13, 26$. The last image corresponds to the final reconstruction. We can observe that it is possible to obtain reconstructions which closely resemble the final solution, with only three or four iterations of the algorithm. Figure 6.25 shows the results after the same number of iterations as in Figure 6.24, for the case of starting the algorithm from a zero initial condition.




Figure 6.21: Sensitivity of the SAR image reconstruction algorithm to initial conditions. Top: $k = 1$, and $\lambda_1 = 5$, $\lambda_2 = 9$. Bottom $k = 0.7$, and $\lambda_1 = 4$, $\lambda_2 = 6$. Left: adjoint-based initial condition. Right: zero initial condition. Values of the objective function at the solution point: (a) $J^* = 1.7248 \times 10^5$ (b) $J^* = 1.7972 \times 10^5$ (c) $J^* = 1.7671 \times 10^5$ (d) $J^* = 1.8998 \times 10^5$.




Figure 6.22: Evolution of the objective function as a function of the iteration number n .




Figure 6.23: Evolution of the ℓ_2 -norm of the gradient of the objective function as a function of the iteration number n .




Figure 6.24: Evolution of the field $\hat{\mathbf{f}}^{(n)}$ as a function of the iteration number n .




Figure 6.25: Evolution of the field $\hat{\mathbf{f}}^{(n)}$ as a function of the iteration number n , with a zero initial condition.

	$\delta = 10^{-6}$		$\delta = 10^{-4}$	
	time (seconds)	J^*	time (seconds)	J^*
$\delta_{CG} = 10^{-3}$	93	1.7247×10^5	37	1.7319×10^5
$\delta_{CG} = 10^{-2}$	51	1.7250×10^5	23	1.7322×10^5

Table 6.1: Computation times and values of the objective function at the solution point for the region-enhanced image reconstruction example of Figure 6.7(c) with various choices of the convergence tolerances.

6.4.12 Computational Complexity

In this section, we provide some information about the computational cost of our algorithm. We have done all the implementation presented here on a Sun Ultra 60 system with a 350 MHz processor using non-optimized MATLAB code. We have used simple diagonal pre-conditioning [78] for the conjugate gradient algorithm, more advanced pre-conditioning techniques can be used to reduce the computational cost further.

First, let us consider the region-enhanced reconstruction of Figure 6.7(c). As shown in Table 6.1, it takes our algorithm about 1.5 minutes to produce this image, with the convergence tolerances of $\delta = 10^{-6}$, and $\delta_{CG} = 10^{-3}$. However, we have noticed that considerable computational savings can be achieved, without compromising the final result significantly, by relaxing these convergence tolerances. Table 6.1 contains the running times of the algorithm with the relaxed tolerances, as well. The images produced with these choices of the tolerances are visually hardly distinguishable from the reconstruction of Figure 6.7(c). The values of the objective function at the solution point for each of these cases are also very close to each other, as shown in Table 6.1. Here we have used a fixed value for δ_{CG} , for a complete run of the algorithm. Another approach might be to start with a relaxed value of δ_{CG} , and reduce it at each iteration n .

	$\delta = 10^{-6}$		$\delta = 10^{-4}$	
	time (seconds)	J^*	time (seconds)	J^*
$\delta_{CG} = 10^{-3}$	14	5.6308×10^4	5.6	5.6410×10^4
$\delta_{CG} = 10^{-2}$	12	5.6308×10^4	4.7	5.6412×10^4

Table 6.2: Computation times and values of the objective function at the solution point for the point-enhanced image reconstruction example of the top row of Figure 6.5(b) with various choices of the convergence tolerances.

Next, we consider the point-enhanced reconstructions in Figure 6.5(b). Table 6.2 contains the computation times, and the values of the objective function at the solution point with various choices of the convergence tolerances. Note that this was a non-superresolution example.

For the superresolution reconstruction problems, the computation time significantly increases, as shown in Table 6.3 for the 2-to-1 resolution loss example of Figure 6.5(b), middle row. This is mainly due to the less sparse structure of the matrix $\mathbf{T}^H \mathbf{T}$ in the case of a high-resolution-field to low-resolution-data forward operator \mathbf{T} .

Finally, in Table 6.4, we present similar results for the 4-to-1 resolution loss example of Figure 6.5(b), bottom row. In this case, the reconstructions with the relaxed convergence tolerances are visually distinguishable from the more accurate reconstruction obtained by the use of tight tolerances. For that reason, we display the solutions with tight and relaxed tolerances in Figure 6.26, to demonstrate the computation versus accuracy tradeoff in this case. Note that the image in Figure 6.26(b) is obtained by only one tenth of the computation required for the image in Figure 6.26(a), but still provides a reasonable approximation.

	$\delta = 10^{-6}$		$\delta = 10^{-4}$	
	time (seconds)	J^*	time (seconds)	J^*
$\delta_{CG} = 10^{-3}$	296	8.5160×10^3	56	8.5313×10^3
$\delta_{CG} = 10^{-2}$	71	8.5257×10^3	51	8.5325×10^3

Table 6.3: Computation times and values of the objective function at the solution point for the point-enhanced, superresolution image reconstruction example of the middle row of Figure 6.5(b) with various choices of the convergence tolerances.

	$\delta = 10^{-6}$		$\delta = 10^{-4}$	
	time (minutes)	J^*	time (minutes)	J^*
$\delta_{CG} = 10^{-3}$	14.40	1.4958×10^4	2.10	1.5006×10^4
$\delta_{CG} = 10^{-2}$	2.60	1.5007×10^4	1.53	1.5034×10^4

Table 6.4: Computation times and values of the objective function at the solution point for the point-enhanced, superresolution image reconstruction example of the bottom row of Figure 6.5(b) with various choices of the convergence tolerances.




Figure 6.26: Demonstration of the tradeoff between computation and accuracy for the superresolution reconstruction example of Figure 6.5(b), bottom row. (a) Reconstruction with $\delta = 10^{-6}$, $\delta_{CG} = 10^{-3}$, with a computation time of 14.40 minutes. (b) Reconstruction with $\delta = 10^{-4}$, $\delta_{CG} = 10^{-2}$, with a computation time of 1.53 minutes.

Chapter 7

Image Reconstruction with More General Potential Functions

In Chapter 5, we have formulated the SAR image reconstruction problem using a particular family of functions (ℓ_k -norms) for the prior information terms. In this chapter, we generalize our framework and iterative algorithm to incorporate a wider range of potentially useful choices. We demonstrate feature-enhanced imaging through the use of a number of classes of potential functions. We also compare these functions in terms of their expected behavior in image reconstruction.

7.1 Feature-Preserving Potential Functions

For the prior information term in (5.7), we now assume the following structure:

$$\Psi(\mathbf{f}) = \lambda_1^2 \sum_i \psi(|(\mathbf{f})_i|) + \lambda_2^2 \sum_i \psi((\mathbf{D}\mathbf{f})_i). \quad (7.1)$$

$\psi_1(x)$	$(x^2 + \epsilon)^{k/2}$
$\psi_2(x)$	$\frac{(x^2 + \epsilon)^{k/2}}{1 + (x^2 + \epsilon)^{k/2}}$
$\psi_3(x)$	$\log(1 + (x^2 + \epsilon)^{k/2})$

Table 7.1: Families of potential functions used. k is a parameter determining the shape of the functions. ϵ is a small smoothing constant.

The role of the first term in (7.1) is to put an energy-type constraint on the solution, and this term should be chosen in such a way to suppress artifacts and increase the resolvability of scatterers. The second term is a piecewise smoothness penalty which should be chosen in such a way that it provides the required smoothing while preserving the edges and hence the shapes of the objects. Our method does not require the potential functions ψ used for the two terms in (7.1) to be identical, however we use identical potential functions for simplicity here. As before, the relative magnitudes of λ_1 and λ_2 reflect the emphasis on point-based feature enhancement (superresolution imaging) and region-based feature-enhancement (edge-preserving smoothing) respectively.

Now, let us discuss the choice of the potential function ψ . There exist some recent studies in image restoration which have focused on the problem of finding appropriate potential functions both from a superresolution [70,85] and from an edge-preservation [63,86] point of view. Conditions to be satisfied for such potential functions have also been proposed [63]. We will not go into the details of such discussions, but rather pick a number of such potential functions from image restoration literature, and demonstrate their use for SAR image reconstruction in our framework. The three particular classes of feature-preserving functions ψ we consider in this chapter are shown in Table 7.1. For all of these functions, ϵ is a small smoothing parameter, hence $(x^2 + \epsilon)^{k/2} \approx x^k$. One might subtract an appropriate constant from each potential function to set $\psi_l(0) = 0$ ($l \in \{1, 2, 3\}$), however we have chosen not to do

so in Table 7.1, to keep the expressions simpler. Note that these potential functions can more generally be expressed in terms of x/Δ , where Δ is a scaling parameter. This parameter Δ can also be used to control the shape, hence the behavior of the potential function. Here, we will use a fixed Δ , and omit it in our analysis for notational simplicity.

First note that the use of ψ_1 in (7.1) leads to constraints in terms of approximate ℓ_k -norms. This is nothing new for us: it results exactly in the prior terms used in (6.3). The other functions, ψ_2 and ψ_3 , are new within our framework. The potential function ψ_2 is based on previous work in [86]. Special cases of ψ_2 for $k = 1$ and $k = 2$ yield the potential functions used in [61] and [63] respectively. Finally, ψ_3 is a generalized version of the potential function proposed in [87].

7.2 General Numerical Solution

The only modification required for the generalization of the iterative algorithm developed in Section 6.1, to incorporate more general potential functions is the structure of the matrix $\mathbf{H}(\mathbf{f})$. Let us choose $\Psi(\mathbf{f})$ in (5.7) so that it contains one of the potential functions ψ_l ($l \in \{1, 2, 3\}$) from Table 7.1. Then, based on the gradient of the cost function (5.7), the Hessian approximation to be used for each l is as follows:

$$\mathbf{H}(\mathbf{f}) \triangleq 2\mathbf{T}^H\mathbf{T} + \lambda_1^2\mathbf{V}_l(|\mathbf{f}|) + \lambda_2^2\Phi^H(f)\mathbf{D}^T\mathbf{V}_l(\mathbf{D}|\mathbf{f}|)\mathbf{D}\Phi(\mathbf{f}) \quad (7.2)$$

$$\begin{aligned} \mathbf{V}_l(\mathbf{z}) &\triangleq \text{diag}\{v_l(z_i)\} \\ v_l(z_i) &= \frac{\psi'_l(z_i)}{z_i} \end{aligned}$$

$v_1(x)$	$\frac{k}{(x^2+\epsilon)^{1-k/2}}$
$v_2(x)$	$\frac{k}{(x^2+\epsilon)^{1-k/2} [(x^2+\epsilon)^{k/2}+1]^2}$
$v_3(x)$	$\frac{k}{(x^2+\epsilon)^{1-k/2} [(x^2+\epsilon)^{k/2}+1]}$

Table 7.2: The weighting functions associated with the potential functions.

The weighting functions $v_l(\cdot)$, for the three potential functions $\psi_l(\cdot)$ considered here, are shown in Table 7.2. With this $\mathbf{H}(\mathbf{f})$, the iterative algorithm (6.8) remains valid for the general case.

7.3 Discussion on the Choice and Behavior of Potential Functions

Although we will not attempt to answer the question of which potential function and choice of parameter k is “best” for a reconstruction problem, we will briefly discuss the behavior of the potential functions of Table 7.1. Figure 7.1 shows plots of these functions for a number of choices of k .

One important issue in the choice of a function to be used in an optimization problem is its convexity properties. Convex functions are known to result in easier minimization problems. From an image restoration and reconstruction standpoint however, non-convex (or relatively “less” convex) potential functions for the prior terms have been observed to be more feature-preserving (more resolvable or more edge-preserving) [79, 88]. All of the potential functions in Figure 7.1, except ψ_1 for $k > 1$, are non-convex.

We will now highlight the feature preservation properties of each potential function, first as a function of the parameter k , and then relative to the other potential

functions. Note that what we mean by feature preservation here is just the degree to which the potential function suppresses the penalization of large-valued quantities.

Let us consider the weighting functions v_l corresponding to the potential functions ψ_l . Figure 7.2 plots the scaled weighting functions $v_l(x)/k$ on a logarithmic scale.¹ Scaling is done so as to compare the structure of the functions more easily. Note also that scaling can always be adjusted through the regularization parameters.

Consider the first weighting function $v_1(x)$, and note that $v_1(x)/k$ is constant for $k = 2$. Hence this function imposes uniform penalties (energy-type or smoothness) everywhere in the scene, and does not lead to feature enhancement. When $k < 2$ however, the weighting $v_1(x)$ decreases, as x gets larger. This results in the preservation of large-valued reflectivities or gradients. This effect becomes more pronounced as k decreases, hence we will say v_1 (hence (ψ_1)) becomes “more feature-preserving” as k decreases.

The structure of v_2 as shown in Figure 7.2 is more interesting. First, note that² $v_2(x)/k \approx 1/(x^{2-k}(1+x^k)^2)$. If we just focus on small values of the magnitude of x (e.g. smaller than 1), the behavior of v_2 , as observed from the plots, is similar to that of v_1 , that is, as k gets smaller, the rate of decay of the weighting function increases. This is because v_2 behaves like x^{k-2} in that region, which is also the behavior of v_1 . On the other hand, for large x , v_2 decays more rapidly as k increases, unlike the case for v_1 . This is because for large x , the weighting function v_2 behaves like $x^{-(2+k)}$, hence for a large k the non-uniform weighting effect becomes more dramatic. Hence the effect of k on the feature preservation property of v_2 depends on the range of values of x of interest.

The dependence of the third weighting function $v_3(x)/k \approx 1/(x^{2-k}(1+x^k))$ on k for small x is similar to the first two weighting functions. For large x , v_3 behaves

¹Note that $v_l(x)$ also depends on k .

²Considering positive x in the expressions here, without loss of generality.

like x^{-2} , hence different choices of k do not lead to very different behavior.

Now, let us compare the three potential functions in terms of the “feature-preservation” effect for a given k . For the three weighting functions we have:

$$v_2(x) \approx \frac{v_3(x)}{(1+x^k)} \approx \frac{v_1(x)}{(1+x^k)^2} \quad (7.3)$$

Note that the term $(1+x^k)$ is always larger than 1, and monotonically increasing, so v_2 has the fastest rate of decay. Hence the second function has the most, and the first function has the least feature-preserving property for fixed k . This can also be observed from the plots of potential functions themselves for a fixed k , as shown for $k = 0.8$ in Figure 7.3. Here we can observe that $\psi_2(x)$ has the least penalty on large values of x as compared to the other two potential functions.




Figure 7.1: Families of potential functions used in the prior terms for feature-enhanced image reconstruction.




Figure 7.2: Scaled weighting functions corresponding to the potential functions in Figure 7.1.




Figure 7.3: The three potential functions for $k = 0.8$.

7.4 Examples

To demonstrate the performance of our image reconstruction method, generalized to incorporate all three potential functions described in this chapter, we use images from the MIT Lincoln Laboratory Advanced Detection Technology Sensor (ADTS) data set [3]. For all the results presented here, we have chosen the values of λ_1 and λ_2 based on subjective qualitative assessment of the formed imagery. We use $\Delta = 1$ for all the examples.

7.4.1 Region-based Feature Enhancement

First, we consider region-based feature enhancement through edge-preserving image formation in our framework. For this task, the dominant prior information term in (7.1) should be the smoothness constraint (hence we need to set $\lambda_2 > \lambda_1$). Images of a military vehicle obtained by the use of various potential functions, as well as the conventional method are shown in Figure 7.4. The images produced by our scheme exhibit reduced speckle, and clear object and shadow boundaries. The three potential functions produce very similar reconstructions. If we compare the image in Figure 7.4(c) with the one in Figure 7.4(b), we can say that the more “feature-preserving” potential function ψ_2 produces slightly sharper edges around the target area as compared to ψ_1 . However it also accentuates some pieces of clutter in the background.

7.4.2 Point-based Feature Enhancement

We now demonstrate superresolution imaging. For this task, we set $\lambda_2 = 0$ in (7.1). First, we show a synthetic example. We generate a synthetic scene by extracting the brightest 30 peaks from a real SAR image, preserving the complex

reflectivities of these points, and setting everything else in the scene to zero. This provides a simple scene where we have a small number of scatterers, and no clutter. Visualizing the resolution properties of the resulting reconstructions should then be simpler than that in a real, complicated scene. We generate simulated SAR data from this scene, by a 2-to-1 resolution reduction. We will show the reconstructions both as intensity images, and as contour plots for ease of interpretation. Figure 7.5 contains the images, and Figure 7.6 contains the contour plots for the synthetic scene, and the conventional, as well as the point-enhanced reconstructions. Point-enhanced reconstructions with each of the three potential functions provide a better visual representation of the underlying scene than the conventional image.

Next, we demonstrate results on real data. Figure 7.7 shows images of the vehicle reconstructed from 0.6 m resolution data. Figure 7.8 contains similar results for an even more reduced resolution level of 1.2 m. Images formed by our approach exhibit improved resolution, reduced sidelobes, narrower mainlobes, hence higher resolution and better dominant scatterer localization than the conventional images for both the synthetic and the real data cases. For a quantitative analysis of the improvements in scattering center locating accuracy provided by our technique (for a subset of the potential functions used here) please see Chapter 9.

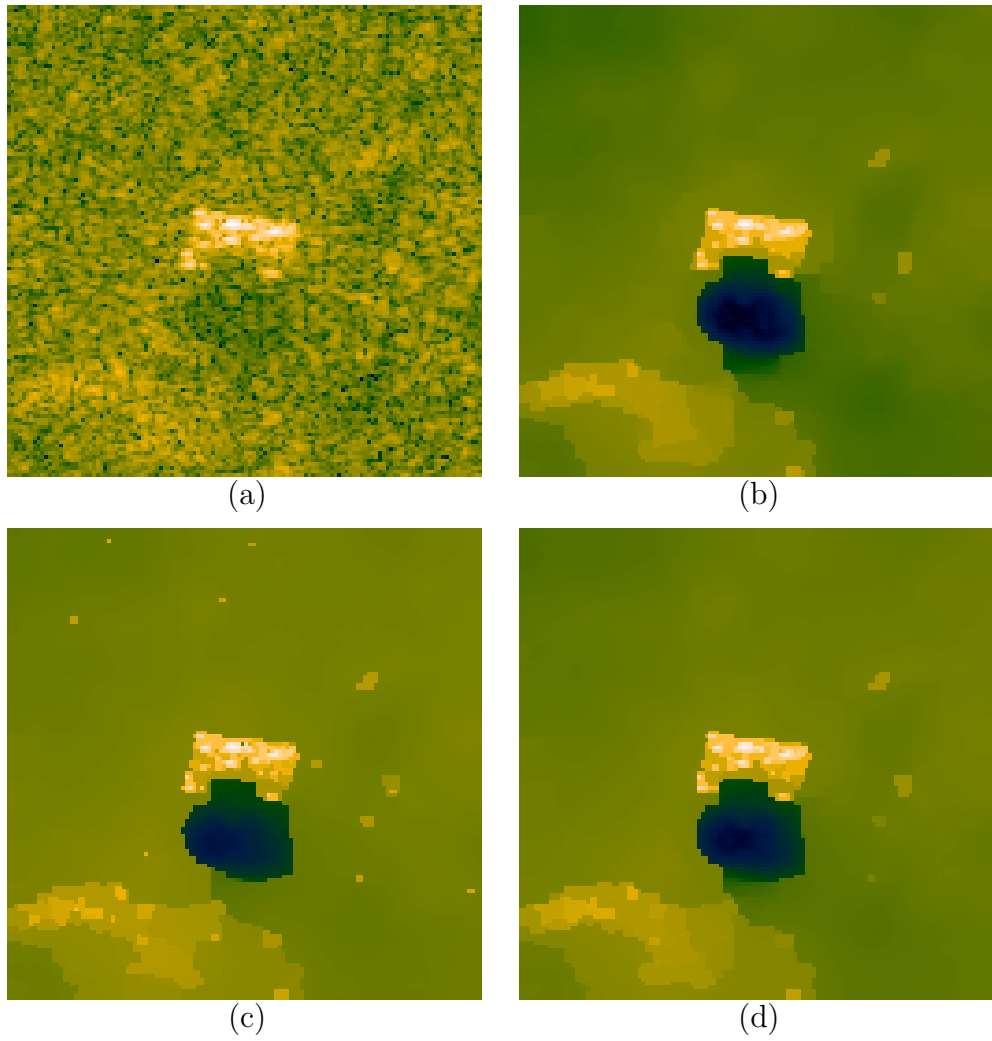


Figure 7.4: Edge-preserving, region-enhanced reconstructions with various potential functions. (a) Conventional image. (b) ψ_1 , $p = 0.8$, $\lambda_1^2 = 30$, $\lambda_2^2 = 60$. (c) ψ_2 , $p = 0.8$, $\lambda_1^2 = 30$, $\lambda_2^2 = 70$. (d) ψ_3 , $p = 0.8$, $\lambda_1^2 = 30$, $\lambda_2^2 = 70$.




Figure 7.5: Point-enhanced, superresolution imaging of a synthetic scene. (a) Original scene. (b) Conventional image. (c) $\psi_1, p = 0.8, \lambda_1^2 = 50$. (d) $\psi_2, p = 0.8, \lambda_1^2 = 25$. (e) $\psi_3, p = 0.8, \lambda_1^2 = 25$.

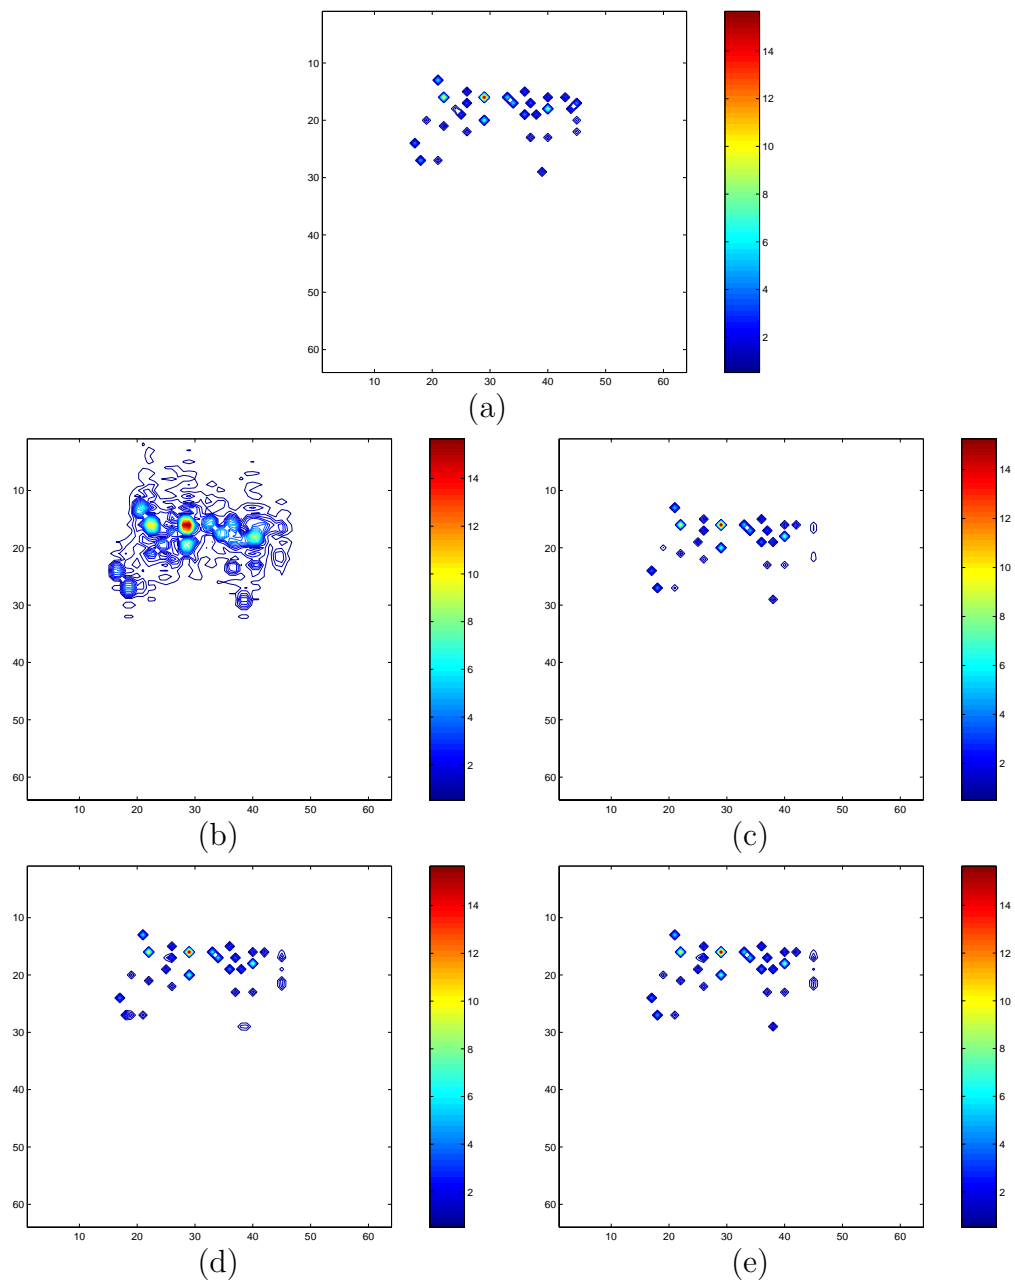


Figure 7.6: Contour plots of the images in Figure 7.5. (a) Original scene. (b) Conventional reconstruction. (c) ψ_1 , $p = 0.8$, $\lambda_1^2 = 50$. (d) ψ_2 , $p = 0.8$, $\lambda_1^2 = 25$. (e) ψ_3 , $p = 0.8$, $\lambda_1^2 = 25$.




Figure 7.7: Point-enhanced, superresolution imaging from 0.6 m resolution data. (a) Conventional image. (b) ψ_1 , $p = 0.8$, $\lambda_1^2 = 50$. (c) ψ_2 , $p = 0.8$, $\lambda_1^2 = 100$. (d) ψ_3 , $p = 0.8$, $\lambda_1^2 = 100$.




Figure 7.8: Point-enhanced, superresolution imaging from 1.2 m resolution data. (a) Conventional image. (b) ψ_1 , $p = 0.8$, $\lambda_1^2 = 150$. (c) ψ_2 , $p = 1$, $\lambda_1^2 = 200$. (d) ψ_3 , $p = 0.8$, $\lambda_1^2 = 150$.

Chapter 8

Generalization to Other Variational Formulations

In this chapter, we demonstrate the use of a different variational energy formulation for SAR imaging. This formulation, based on the Mumford-Shah energy [14], has been commonly used in image restoration and segmentation problems before. Our work here provides an extension of our previous framework for SAR imaging to incorporate the richer structure of this variational formulation. On the other hand, it also extends the use of Mumford-Shah-type formulations to complex-valued, random-phase fields, with non-trivial observation models.

8.1 Variational Formulations based on the Mumford-Shah Energy

Many variational approaches for image segmentation problems are based on the well-known Mumford-Shah energy [14]. Such approaches share our goal in region-enhanced imaging of smoothing the field where there are no edges, while preserving

the discontinuities. Our philosophy in Chapters 5 and 7 to solve this problem was to use non-quadratic potential functions, which implicitly produced the desired effect. In contrast, the Mumford-Shah-type variational formulations start from an objective function which is explicitly aimed at smoothing the homogeneous regions in an image, but simultaneously preserving the discontinuities. One such function, which is a relaxed version of the Mumford-Shah energy, is the Ambrosio-Tortorelli energy [89]. A discretized version of this objective function is:

$$E(\mathbf{f}, \mathbf{e}) = \|\mathbf{g} - \mathbf{f}\|_2^2 + \lambda^2 \left\{ \left[\sum_{i=1}^M (1 - \mathbf{e}_i)^2 (\mathbf{Df})_i^2 \right] + \frac{1}{2\beta} \|\mathbf{e}\|_2^2 + \frac{\beta}{2} \|\mathbf{De}\|_2^2 \right\} \quad (8.1)$$

where \mathbf{g} denotes the observed image data, \mathbf{f} denotes the piecewise smooth approximating field, \mathbf{e} denotes the continuous-valued edge field, \mathbf{D} is a discrete derivative operator, and λ, β are scalar parameters. The subscript i denotes the i -th element of a vector. Presence and absence of an edge at a particular location make the edge field \mathbf{e} approach 1 and 0 respectively at that location. The first term in (8.1) is a data fidelity term. Note that since segmentation problems use an observed *image* as the data, the observation model is just the identity. The second of the four terms is a smoothness constraint. Note that this term is essentially a weighted ℓ_2 -norm of \mathbf{Df} where the weights are given by $(1 - \mathbf{e}_i)^2$. Due to the edge-dependent weighting, smoothing is done only in locations where the edge field \mathbf{e} is small. Such a constraint makes sense, however this term by itself favors placing edges everywhere in the scene. This potential problem is prevented by the last two terms in (8.1), which penalize the norm and roughness of the edge field \mathbf{e} .

8.2 A Mumford-Shah-type Formulation for SAR Imaging

By adapting the Ambrosio-Tortorelli energy in (8.1), we propose the following cost function for SAR imaging:

$$J(\mathbf{f}, \mathbf{e}) = \|\mathbf{g} - \mathbf{Tf}\|_2^2 + \lambda^2 \left\{ \left[\sum_{i=1}^M (1 - \mathbf{e}_i)^2 (|(\mathbf{D}|\mathbf{f})_i|^2) \right] + \frac{1}{2\beta} \|\mathbf{e}\|_2^2 + \frac{\lambda_e^2 \beta}{2} \|\mathbf{De}\|_{k_e}^{k_e} \right\} \quad (8.2)$$

Note that this cost function takes into account the SAR observation model \mathbf{T} , and imposes a smoothness constraint on the *magnitude* of the complex-valued SAR reflectivity field \mathbf{f} . To solve this optimization problem, we take a half-quadratic approach. The half-quadratic energy $\tilde{K}(\mathbf{f}, \mathbf{e}, \mathbf{s})$ must satisfy:

$$\inf_{\mathbf{s}} \tilde{K}(\mathbf{f}, \mathbf{e}, \mathbf{s}) = J(\mathbf{f}, \mathbf{e}). \quad (8.3)$$

Such a half-quadratic augmented cost function is given by:

$$\tilde{K}(\mathbf{f}, \mathbf{e}, \mathbf{s}) = \|\mathbf{g} - \mathbf{Tf}\|_2^2 + \lambda^2 \left\{ \left[\sum_{i=1}^M (1 - \mathbf{e}_i)^2 (|(\mathbf{DSf})_i|^2) \right] + \frac{1}{2\beta} \|\mathbf{e}\|_2^2 + \frac{\lambda_e^2 \beta}{2} \|\mathbf{De}\|_{k_e}^{k_e} \right\} \quad (8.4)$$

where

$$\mathbf{S} = \text{diag}\{\exp(-j\mathbf{s}_l)\}, \quad (8.5)$$

with \mathbf{s}_l being the l -th element of the vector \mathbf{s} .

We minimize $\tilde{K}(\mathbf{f}, \mathbf{e}, \mathbf{s})$ of (8.4) by the following coordinate descent approach:

$$\hat{\mathbf{s}}^{(n+1)} = \arg \min_{\mathbf{s}} \tilde{K}(\hat{\mathbf{f}}^{(n)}, \hat{\mathbf{e}}^{(n)}, \mathbf{s}) \quad (8.6)$$

$$\hat{\mathbf{e}}^{(n+1)} = \arg \min_{\mathbf{e}} \tilde{K}(\hat{\mathbf{f}}^{(n)}, \mathbf{e}, \hat{\mathbf{s}}^{(n+1)}) \quad (8.7)$$

$$\hat{\mathbf{f}}^{(n+1)} = \arg \min_{\mathbf{f}} \tilde{K}(\mathbf{f}, \hat{\mathbf{e}}^{(n+1)}, \hat{\mathbf{s}}^{(n+1)}) \quad (8.8)$$

Based on the structure of $\tilde{K}(\mathbf{f}, \mathbf{e}, \mathbf{s})$, we obtain the following coordinate descent steps:

$$\hat{\mathbf{s}}_i^{(n+1)} = \phi[(\hat{\mathbf{f}}^{(n)})_i] \quad (8.9)$$

$$\hat{\mathbf{e}}^{(n+1)} = \arg \min_{\mathbf{e}} \left\{ \left[\sum_{i=1}^M (1 - \mathbf{e}_i)^2 \left(|(\mathbf{D}\hat{\mathbf{S}}^{(n+1)}\hat{\mathbf{f}}^{(n)})_i|^2 \right) \right] + \frac{1}{2\beta} \|\mathbf{e}\|_2^2 + \frac{\lambda_e^2 \beta}{2} \|\mathbf{D}\mathbf{e}\|_{k_e}^{k_e} \right\} \quad (8.10)$$

$$\left[2\mathbf{T}^H \mathbf{T} + 2\lambda^2 (\hat{\mathbf{S}}^{(n+1)})^H \mathbf{D}^T \text{diag} \left\{ \left(\hat{\mathbf{e}}_i^{(n+1)} \right)^2 \right\} \mathbf{D} \hat{\mathbf{S}}^{(n+1)} \right] \hat{\mathbf{f}}^{(n+1)} = 2\mathbf{T}^H \mathbf{g} \quad (8.11)$$

Note that (8.9), and (8.11) provide what are conceptually closed-form solutions for $\hat{\mathbf{s}}^{(n+1)}$ and $\hat{\mathbf{f}}^{(n+1)}$ respectively (although an iterative scheme can be used in practice for $\hat{\mathbf{f}}^{(n+1)}$). On the other hand (8.10) gives a closed-form solution for $\hat{\mathbf{e}}^{(n+1)}$ only if $k_e = 2$, in which case this is a quadratic minimization problem. When $k_e \neq 2$, a numerical method, such as the half-quadratic methods we have been discussing, must be used for (8.10).

8.3 The Link to Previous Image Formation Approach

In this section, we illuminate the similarity between the function $\tilde{K}(\mathbf{f}, \mathbf{e}, \mathbf{s})$ of Section 8.2, and the half-quadratic objective functions we have considered in Chapters 5 and 7. In particular, we show that we may view the Mumford-Shah-type variational

approach of this chapter as a generalization of our previous framework, with a specific potential function. In our development, we follow the line of thought in [88].

Consider the special case of the objective functions of Chapters 5 and 7, where we take $\lambda_1 = 0$:

$$J(\mathbf{f}) = \|\mathbf{g} - \mathbf{Tf}\|_2^2 + \lambda_2^2 \sum_{i=1}^M \psi((\mathbf{Df})_i). \quad (8.12)$$

In Section 6.3, we have seen that minimization of such functions $J(\mathbf{f})$ can be done by the minimization of a half-quadratic energy, $K(\mathbf{f}, \mathbf{b}, \mathbf{s})$.

Now, suppose we choose $\psi(x) = x^2/(1 + x^2)$ in (8.12). Note that this is a particular case of the potential function ψ_2 from Table 7.1, with $k = 2$ and $\epsilon = 0$. With this potential function in (8.12), we can show that the augmented half-quadratic cost function $K(\mathbf{f}, \mathbf{b}, \mathbf{s})$ is given by:

$$K(\mathbf{f}, \mathbf{b}, \mathbf{s}) = \|\mathbf{g} - \mathbf{Tf}\|_2^2 + \lambda_2^2 \sum_{i=1}^M \left[\mathbf{b}_i (|(\mathbf{DSf})_i|^2) + (\sqrt{\mathbf{b}_i} - 1)^2 \right] \quad (8.13)$$

Let us make a change of variable and replace \mathbf{b}_i with $(1 - \mathbf{e}_i)^2$. This yields:

$$K(\mathbf{f}, \mathbf{e}, \mathbf{s}) = \|\mathbf{g} - \mathbf{Tf}\|_2^2 + \lambda_2^2 \sum_{i=1}^M \left[(1 - \mathbf{e}_i)^2 (|(\mathbf{DSf})_i|^2) + (\mathbf{e}_i)^2 \right]. \quad (8.14)$$

Finally, recognizing the last term in (8.14) can be expressed as an ℓ_2 -norm, we obtain:

$$K(\mathbf{f}, \mathbf{e}, \mathbf{s}) = \|\mathbf{g} - \mathbf{Tf}\|_2^2 + \lambda_2^2 \left\{ \left[\sum_{i=1}^M (1 - \mathbf{e}_i)^2 (|(\mathbf{DSf})_i|^2) \right] + \|\mathbf{e}\|_2^2 \right\} \quad (8.15)$$

The cost function $K(\mathbf{f}, \mathbf{e}, \mathbf{s})$ of (8.15) is a particular case of the type of functions we have used in Chapters 5 and 7. Now, let us compare this function to the Mumford-Shah-based cost function $\tilde{K}(\mathbf{f}, \mathbf{e}, \mathbf{s})$ of (8.4). One can easily establish a correspondence between the terms in these two functions. We can observe that $K(\mathbf{f}, \mathbf{e}, \mathbf{s})$ can be viewed as a special case of $\tilde{K}(\mathbf{f}, \mathbf{e}, \mathbf{s})$ with $\lambda = \lambda_2$, $\lambda_e = 0$, $\beta = 0.5$. Hence the Mumford-Shah-based cost function $\tilde{K}(\mathbf{f}, \mathbf{e}, \mathbf{s})$ can be viewed as a generalization of our previous formulation in a particular way.

8.4 Example

We now illustrate some benefits that may be obtained by the richness of the variational formulations considered in this chapter. Consider the synthetic scene in Figure 8.1(a). We simulate SAR returns from this scene, and add Gaussian noise to the simulated noise-free returns so that SNR is 15 dB. We then consider the reconstruction problem from such noisy data. Figure 8.1(b) shows the conventional reconstruction.

We now reconstruct this scene first using the approach of Chapter 7, and then using the variational formulation presented in this chapter. For the discrete derivative operators, we use the second approach described in Appendix A.1 here. Figure 8.2(a) shows the reconstructed field obtained by the method of Chapter 7 with ψ_2 , $k = 2$. Note that this corresponds to minimizing the function in (8.15). Figure 8.2(b) shows the associated reconstructed edge field \mathbf{e} . Using the same regularization parameter, we now form an image by minimizing the function in (8.4) using the iterative scheme of Section 8.2. The reconstructed field and the edge field are shown in Figure 8.2(c) and (d). Note that the edge field in (d) suppresses the artifacts much better than the edge field of (b).




Figure 8.1: Synthetic example. (a) Original scene. (b) Conventional reconstruction from noisy data (SNR=15 dB).




Figure 8.2: Feature-enhanced reconstructions. (a) Reconstructed field using the objective function in (8.15) with $\lambda_2^2 = 20$. (b) Edge field \mathbf{e} corresponding to (a). (c) Reconstructed field using the Mumford-Shah-based objective function in (8.4) with $\lambda_2^2 = 20$, $\beta = 0.1$, $\lambda_e^2 = 0.4$, $k_e = 1$. (d) Edge field \mathbf{e} corresponding to (c).

Chapter 9

Evaluation based on Recognition-Oriented Features

While the images produced by our methods appear visually good, the performance of such images in automated processing algorithms is our main interest. Therefore, in this chapter, we propose quantitative criteria for evaluating the images produced by the SAR image formation technique developed in Chapters 5 and 6, based on recognition-oriented features. We run experiments on the MSTAR public target data set [81] to compare the SAR images formed by feature-enhanced imaging to conventional images in terms of these quantitative measures. The criteria we use to evaluate point-based features are target-to-clutter ratio, mainlobe width, peak matching accuracy, and average associated peak distance. The metric of peak matching accuracy is particularly useful for testing the superresolution properties of an image formation method. The criteria we use to evaluate region-based features are speckle suppression, segmentation accuracy, and statistical separability of different regions from one another. The results of our experiments indicate that feature-enhanced images exhibit superresolution and improved localization accuracy for dominant scat-

terers, and improved separability for different regions in the scene.

9.1 Evaluation Criteria for Point-Enhanced Images

In this section, we propose measures focusing on point-based features, for evaluating the quality of reconstructed SAR images. Many of these criteria have appeared in the literature before, and they are mostly directed towards quantities to be used in target recognition tasks.

9.1.1 Target-to-Clutter Ratio

As a measure of accentuation of the target pixels with respect to the background, we will use the target-to-clutter ratio in dB, defined as [32]:

$$\text{Target-to-clutter ratio} = 20 \log_{10} \left(\frac{\max_{(i,j) \in \mathcal{T}} (|\hat{\mathbf{f}}_{ij}|)}{\frac{1}{N_C} \sum_{(i,j) \in \mathcal{C}} |\hat{\mathbf{f}}_{ij}|} \right) \quad (9.1)$$

where the pair (i, j) denotes the pixel indices, $\hat{\mathbf{f}}$ is the reconstructed image, $\hat{\mathbf{f}}_{ij}$ is the reconstructed reflectivity at location (i, j) , \mathcal{T} denotes the target region, \mathcal{C} denotes a clutter patch in the image, and N_C denotes the number of pixels in the clutter patch.

9.1.2 Mainlobe Width

As one of the measures of the effective resolution of an image, we will use the 3-dB mainlobe width of the strong scatterers. To obtain an estimate of the mainlobe width, we concentrate on the target region. In each row and column in the target region of the reconstructed image, we find the first point near the maximum where the reflectivity magnitude is more than 3 dB below the maximum value. We then obtain a better estimate of the 3-dB distance by means of a linear interpolation between

pixels. Finally, we average the distances obtained from each row and column in the target region to find an overall estimate of the 3-dB lobe width for a particular image.

9.1.3 Peak Matching Accuracy

Locations of dominant point scatterers extracted from a target image are important characteristics for recognition [65, 66]. Loss of resolution manifests itself by merging and moving such characteristic points, and this makes the accurate localization of these points in the scene more difficult. Thus, we evaluate the superresolution properties of our method by measuring how well the dominant scatterers are preserved when we use reduced-resolution data to form the image.

The procedure we use for measuring peak matching accuracy is as follows. The first step is to extract the locations of N_p brightest scatterers from a reconstructed image. For this purpose, we first find all the peaks in the scene. The peaks are taken to be the points where the discrete spatial derivatives of the reflectivity magnitude in both the x and the y directions change sign from positive to negative. Once the peaks are found, we order them based on their magnitudes, and pick the largest N_p of them. In our experiments, we will use $N_p = 20$.

We then evaluate how well the coordinates of these N_p peaks match those of the N_p “reference” locations of the scatterers. These reference positions may be obtained either from the ground truth, in case that is available, or from the locations of the scatterers extracted from a higher resolution image, otherwise. Once we have two sets of N_p peak locations, we need to determine how many of the extracted locations “match” a reference location. The simplest matching criterion would be to count the exact position matches between the two sets of locations. However, this criterion may be too restrictive. For this task, we use the feature matching method used in [90]. This method allows a match declaration between two peaks, if the estimated

peak location is within a radius r of the “reference” peak location. Hence it is more powerful than counting only the exact matches, with r used as a variable parameter ($r = 0$ corresponds to counting the exact matches). A one-to-one association of the peaks is made such that the sum of the squared distances between the locations of the “reference” peaks and the corresponding matched peaks from the image is minimized. We can then count the number of matched peaks out of N_p , to see how well the peaks are preserved.

We apply the above procedure to the conventional images and the feature-enhanced images reconstructed from the same reduced-resolution data. We compute the number of matched peaks for a large set of such conventional and feature-enhanced images, and compare the average number of matched peaks, parameterized by r .

9.1.4 Average Associated Peak Distance

Another criterion based on peak locations that we use is the average distance between the two sets of N_p matched peak coordinates. To compute this measure, we relax the matching radius r of Section 9.1.3, so that each of the N_p peaks from the reconstructed image is matched to one of the “reference” peaks. We then find the average of the distances between these associated peaks.

9.2 Evaluation Criteria for Region-Enhanced Images

9.2.1 Speckle Suppression

Speckle complicates intensity-based region description in conventional SAR images. One measure that has been used for speckle amplitude is the standard deviation

of a clutter patch in the dB-valued SAR images [32]. We use this measure to quantitatively compare our reconstructions with conventional images in terms of speckle suppression.

9.2.2 Segmentation Accuracy

It is of interest to obtain accurate segmentations of SAR images for effective use of region-based shape features in target recognition. Recently there has been much interest in the development of segmentation algorithms for conventional SAR images. Our region-enhanced images would appear to provide easier-to-segment regions as compared to conventional SAR images. We demonstrate this property by segmenting our reconstructions to target, shadow and background regions by simple adaptive thresholding, where the thresholds for a particular image under consideration depend on the statistics of that image.

To determine the thresholds, we find the mean μ and the standard deviation σ of the dB-valued pixel magnitudes in the image. Then, we apply the following decision rule at each pixel:

$$\begin{aligned} 20 \log_{10}(\hat{\mathbf{f}}_{ij}) < \mu - c_1\sigma &\implies \hat{\mathbf{f}}_{ij} \in \mathcal{S} \\ \mu - c_1\sigma \leq 20 \log_{10}(\hat{\mathbf{f}}_{ij}) < \mu + c_2\sigma &\implies \hat{\mathbf{f}}_{ij} \in \mathcal{B} \\ \mu + c_2\sigma \leq 20 \log_{10}(\hat{\mathbf{f}}_{ij}) &\implies \hat{\mathbf{f}}_{ij} \in \mathcal{T} \end{aligned}$$

where $\mathcal{T}, \mathcal{S}, \mathcal{B}$ denote the target, shadow and background regions respectively and c_1, c_2 are two constants that are fixed beforehand. Hence, this is really a “histogram-based” threshold.

From a statistical standpoint, it would make more sense to develop a decision metric based on the statistics of particular regions. However, our objective here is not to develop the best decision metric, but rather to show that we can obtain reasonable segmentations of the region-enhanced images even by simple suboptimal processing.

The above procedure produces a segmented image. Next, we need to evaluate the accuracy of this segmentation. For such evaluation, we need the ground truth. Ground truth for this problem consists of the assignment of each location in the scene to one of the three sets $\mathcal{T}, \mathcal{S}, \mathcal{B}$. Such ground truth information is not readily available in general, as is the case with the MSTAR data set. We use segmentations done by humans as the ground truth. The measure of accuracy we use is the percent of pixels in the automatically segmented image, which are classified correctly (i.e. assigned to the same set as the ground truth assignment of the pixel).

9.2.3 Statistical Separability of Regions

Segmentation accuracy provides a good flavor of the degree of separability of different regions from one another in the SAR image. However it does not provide a statistical measure for the similarity of different regions, $\mathcal{T}, \mathcal{S}, \mathcal{B}$. To obtain such a measure, we first assume that all the (dB-valued) reflectivity magnitudes in a particular region of the reconstructed SAR target image are drawn from the same Gaussian distribution. We also assume that these region-based probability density functions are independent of the target type. We then estimate the mean and variance of such a density for each region, $\mathcal{T}, \mathcal{S}, \mathcal{B}$, using a large set of SAR images. Note that the actual distribution of those elements may not be close to a Gaussian, but such a simple model is sufficient for our objective of measuring region separability.

As a measure of the similarity of two Gaussian densities corresponding to classes (regions) i and j , we use the Bhattacharyya distance d_{ij} [91]:

$$d_{ij} = \frac{1}{8}(\mu_i - \mu_j)^T \left(\frac{\Sigma_i + \Sigma_j}{2} \right)^{-1} (\mu_i - \mu_j) + \frac{1}{2} \ln \left(\frac{|\frac{\Sigma_i + \Sigma_j}{2}|}{|\Sigma_i|^{1/2} |\Sigma_j|^{1/2}} \right) \quad (9.2)$$

where μ_i, μ_j are the mean vectors and Σ_i, Σ_j are the covariance matrices for class

i and class j . Note that, in our case the means and covariances are just scalars. The distance d_{ij} is a common measure of the separability of classes characterized by multi-dimensional Gaussian distributions and gives an upper bound on the Bayesian error for classification of two classes. We compute Bhattacharyya distances for each pair of different regions, $\mathcal{T}, \mathcal{S}, \mathcal{B}$, and compare such distances for conventional and region-enhanced images. A larger Bhattacharyya distance means that the regions are easier to separate from each other.

9.3 Experimental Results

9.3.1 Experimental Setup

The MSTAR public target data set provides SAR images of various military vehicles. We use images of T72 tanks, BMP2 tanks, and BTR70 armored personnel carriers from this data set to evaluate the performance of our reconstructed images in terms of the criteria described in Sections 9.1 and 9.2. Sample pictures of these three types of vehicles are shown in Figure 9.1. We use 72 images for each vehicle type, all at 17° depression angle, and evenly spaced in azimuth (approximately 5°) to cover 360° . Depression angle is the angle between the horizontal plane and the line that connects the radar platform to the imaged scene. Azimuth or aspect angle is the angle between the major axis of the vehicle and the line between the radar platform and the vehicle in the ground plane, hence this angle basically shows the orientation of the vehicle with respect to the sensor. Figure 9.2 shows the magnitude of an example MSTAR image in dB for each target type. We will use such sample images to display the nature of our reconstructions and the associated feature extraction results for each target type. As we will describe, we have also carried out some synthetic scene reconstruction experiments to make some evaluations where ground truth is exactly



Figure 9.1: Sample pictures of the three types of vehicles whose SAR images are used in the experiments.




Figure 9.2: 128×128 sample images from the MSTAR public target data set.

known. We obtain the projectional SAR data from the MSTAR images as described in Section 6.4.4.1.

When applying our image reconstruction method, we choose the feature accentuation parameters λ_1 and λ_2 in (6.3) based on subjective qualitative assessment of one image from the entire data set, and use those values for all the images for that particular experiment.

9.3.2 Point-Enhanced Imaging from Full-Resolution Data

In this section, we report the results of the experiments using the 100×100 phase history samples to form point-enhanced images. The resolution supported by this data is 0.3 m. We form “critically-sampled” (i.e. 100×100) images with our technique. Therefore, in order to have conventional SAR images of this size for comparison, we first form 100×100 Taylor-windowed Fourier images. Samples of such reconstructions are shown in the top row of Figure 9.3. Naturally these are very similar to their oversampled versions in Figure 9.2.

We form point-enhanced images with $k = 0.8$, and $\lambda_2 = 0$ in (6.3). We do not apply any windowing to the data before processing, since our method is able to suppress sidelobes considerably even with rectangular weighting. However, if desired, the method can be used with windowing, as it was demonstrated in Section 6.4.4. The bottom row in Figure 9.3 shows the reconstructions obtained. The dominant scatterers appear to be accentuated as compared to the conventional images at the top row.

9.3.2.1 Target-to-Clutter Ratio

We will quantify the enhancement of the target pixels in the full-resolution data reconstructions by means of their target-to-clutter ratio. We compute the target-to-clutter ratio as defined in (9.1), by using the bottom 20 rows (2000 pixels) of the reconstructed images as the clutter region. This region is big enough to give a reliable estimate of the mean reflectivity magnitude, and is safe to use, since target and shadow appear to be located outside this region for the entire data set. Table 9.1 shows the average target-to-clutter ratio achieved by the conventional and the point-enhanced reconstructions over the 72 images for each target type. These results indicate a clear improvement of the target-to-clutter ratio by our proposed image




Figure 9.3: 100×100 sample images reconstructed from full-resolution (100×100) data. Top: conventional. Bottom: proposed point-enhanced, with $k = 0.8$, $\lambda_1 = 14$, $\lambda_2 = 0$.

formation method.

9.3.3 Point-Enhanced Superresolution Imaging from Reduced-Resolution Data

In this section, we report the results of experiments on two sets of data: the actual MSTAR data, and the synthetic point scatterer scenes constructed using the MSTAR images. The reason for using synthetic examples is to demonstrate the superresolution properties of our method in a situation where the ground truth is

Average Target-to-Clutter Ratio	T72	BMP2	BTR70
Conventional	31.88 dB	28.92 dB	26.92 dB
Proposed	88.28 dB	85.38 dB	82.62 dB

Table 9.1: Average target-to-clutter ratios of images reconstructed from full-resolution data.

exactly known.

We present the mainlobe width results for only the actual MSTAR reconstructions. We present the peak matching accuracy and the average associated peak distance results for both actual and synthetic images. We do not present the target-to-clutter ratio results in this section, since they are very similar to the full-resolution target-to-clutter ratio results of Table 9.1.

For experiments on actual MSTAR data, we form images from a 50×50 subset of the 100×100 phase history samples previously used. This results in a 2-to-1 resolution loss in the range and cross-range directions. Hence, the resolution supported by such reduced data is 0.6 m. All the images we present in this section are composed of 100×100 pixels. The top row in Figure 9.4 shows Taylor weighted Fourier images from the reduced-resolution, 50×50 data. The resolution loss in these images is evident when they are compared to their high-resolution counterparts in Figure 9.3. We now form point-enhanced images with $k = 0.8$, and $\lambda_2 = 0$ in (6.3), samples of which are shown in the bottom row of Figure 9.4.

We now consider the synthetic examples. To generate synthetic scenes, we find the 20 peaks with the largest magnitude in each of the 72 100×100 Taylor-windowed T72 Fourier images, and form a synthetic scene by placing simulated point-scatterers at the locations of these peaks, with the original complex reflectivities, and zeros in the rest of the scene. An example contour plot of the magnitude of such a synthetic scene is shown in the left third of Figure 9.5. We then generate simulated phase histories




Figure 9.4: 100×100 sample images reconstructed from reduced-resolution (50×50) data. Top: conventional. Bottom: proposed point-enhanced, with $k = 0.8$, $\lambda_1 = 14$, $\lambda_2 = 0$.

from this scene. The reconstructed conventional Taylor-windowed image from 50×50 phase history samples is shown at the top row, middle column of Figure 9.5. The loss of resolution is easy to observe. The corresponding point-enhanced image produced by our method is shown at the bottom row, middle column of the same figure, and we can visually observe that most of the scatterers that were merged by the conventional reconstruction are now resolved. The images in the rightmost column demonstrate similar results for the 25×25 data case. Although the amount of data we use here is only one sixteenth of the data required for full resolution, our method is able to localize most of the scatterers.

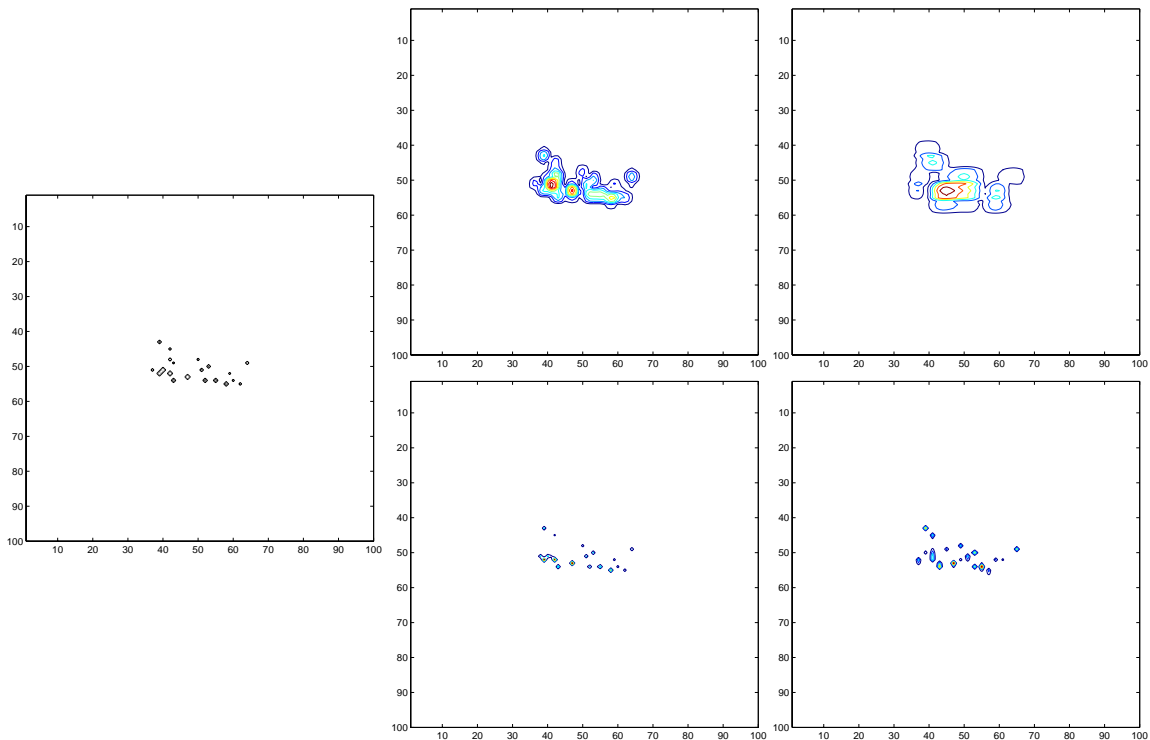


Figure 9.5: Synthetic T72 image reconstruction example from reduced-resolution data. Left: ground truth. Middle: results from 50×50 data. Right: results from 25×25 data. Top row: conventional. Bottom row: proposed point-enhanced, with $k = 0.8$, $\lambda_2 = 0$, and middle: $\lambda_1 = 14$, right: $\lambda_1 = 7$.

In the remainder of this section, we quantitatively demonstrate the resolution improvement achieved by the images presented here.

9.3.3.1 Mainlobe Width

We compute the average 3-dB mainlobe width as described in Section 9.1.2 for all the 216 reconstructed MSTAR scenes. The results in Table 9.2 for 50×50 data reconstructions show that our proposed scheme is able to reduce the mainlobe width considerably. To put these numbers in perspective, note that the resolution supported by the data is 0.6 m in this experiment.

Average Mainlobe Width	T72	BMP2	BTR70
Conventional	0.447 m	0.466 m	0.459 m
Proposed	0.098 m	0.097 m	0.094 m

Table 9.2: Average mainlobe widths of images reconstructed from reduced-resolution (50×50) data.

9.3.3.2 Peak Matching Accuracy

We now evaluate how the locations of the dominant peaks are preserved in reduced-resolution data situations by the conventional reconstructions and by our point-enhanced images. For the MSTAR examples, we use the locations of the 20 peaks extracted from the Taylor-windowed image reconstructed from full-resolution data, as the “reference” locations. Figure 9.6 provides a visual comparison of the peak locating accuracy of the reconstructions from 50×50 data. The circles indicate the “reference” locations of the 20 dominant scatterers, and the plus signs indicate the peaks extracted from the reconstructed reduced-resolution images. The top row contains the results for the conventional images, while the bottom row contains those for the point-enhanced images. The clear observation we can make out of these results is that, since conventional image formation causes peaks to merge, some of the peaks in the target area are lost, and peaks outside this area may become dominant. We now evaluate the peak matching accuracy of our method by using the criterion described in Section 9.1.3. In Figure 9.7, we plot the average number of peak matches for the images formed by the conventional and the proposed methods as a function of the radius r within which a match declaration is allowed. The standard deviation of this estimate of the mean is very small, hence we do not show error bars on these plots. The peak matching accuracy of our images appear to be higher than that of the conventional images. Note that our analysis is based on finding peaks all around the scene. Alternatively, the search for peaks can be done in a pre-determined target




Figure 9.6: Sample peak extraction results for images reconstructed from 50×50 data. Circles indicate “reference” dominant scatterer locations extracted from full-resolution conventional images. Plus signs indicate peaks extracted from the reconstructed images. Top: conventional. Bottom: proposed point-enhanced, with $k = 0.8$, $\lambda_1 = 14$, $\lambda_2 = 0$.

region only.

We now report the results of similar experiments for the synthetic T72 scenes. Figure 9.8 shows the peaks extracted from the point-enhanced images and those extracted from the conventional images using 50×50 and 25×25 data. The improved accuracy provided by our method is easy to observe in these plots. Similarly, the peak matching accuracy results from the entire data, shown in Figure 9.9, verify the superresolution properties of our imaging scheme.

9.3.3.3 Average Associated Peak Distance

We now compute the average distance between the true 20 peaks and the 20 peaks extracted from the reconstructed images as described in Section 9.1.4. Tables 9.3 and 9.4 illustrate the average associated peak distances for the real and synthetic scene experiments respectively. These results indicate a clear reduction in peak distances by point-enhanced imaging.

Average Associated Peak Distance	T72	BMP2	BTR70
Conventional	3.13 m	4.32 m	4.30 m
Proposed	0.82 m	1.06 m	1.25 m

Table 9.3: Average associated peak distances in images reconstructed from reduced-resolution (50×50) data.

Average Associated Peak Distance (synthetic T72)	50×50 data	25×25 data
Conventional	1.22 m	1.35 m
Proposed	0.07 m	0.61 m

Table 9.4: Average associated peak distances in the synthetic T72 reconstructions.




Figure 9.7: Average number of peak matches in images reconstructed from reduced-resolution (50×50) data as a function of the radius of match declaration r .




Figure 9.8: Sample peak extraction results for the synthetic T72 scenes. Circles indicate the scatterer locations in the synthetic scene. Plus signs indicate peaks extracted from the reconstructed images. Left: 50×50 data. Right: 25×25 data. Top: conventional. Bottom: proposed point-enhanced, with $k = 0.8$, $\lambda_1 = 7$, $\lambda_2 = 0$.




Figure 9.9: Average number of peak matches for the synthetic T72 scenes as a function of the radius of match declaration r . Top: 50×50 data. Bottom: 25×25 data.

Average Speckle Amplitude	T72	BMP2	BTR70
Conventional	5.919 dB	5.921 dB	5.898 dB
Proposed	2.261 dB	2.283 dB	2.269 dB

Table 9.5: Average speckle amplitude in the dB-valued reconstructed images.

9.3.4 Region-Enhanced Imaging

We now compare our region-enhanced images with conventional ones in terms of the criteria described in Section 9.2. Here, we form 128×128 images from 128×128 Taylor-windowed phase history samples. In our image reconstruction method, we use $k = 1$, and set $\lambda_2 > \lambda_1$ in (6.3) to enhance regions. The second row of Figure 9.10 shows sample reconstructions using the proposed method. In contrast to the conventional images in the top row, these reconstructions reduce variability in homogeneous regions, while preserving discontinuities at region boundaries.

9.3.4.1 Speckle Suppression

We quantify the speckle amplitude in images as described in Section 9.2.1, by using the bottom 20 rows (2560 pixels) of the reconstructed images as the clutter region. The results in Table 9.5 illustrate the speckle suppression achieved by the region-enhanced reconstructions.

9.3.4.2 Segmentation Accuracy

We now demonstrate that our region-enhanced images simplify segmentation of the images into target, shadow and background regions. In our evaluation, we use human segmentations of the MSTAR images as the ground truth. These segmentations have been obtained by having a number of volunteers draw the region

	T72	BMP2	BTR70
Average Segmentation Accuracy	96.85 %	97.28 %	97.64 %

Table 9.6: Average segmentation accuracy for region-enhanced images, measured as the percentage of correctly classified pixels.

boundaries on conventional MSTAR images based on their own judgment. Samples of such human segmentations are shown in the third row of Figure 9.10. Naturally, these segmentations are themselves not perfect.

We segment our region-enhanced images by simple adaptive thresholding, as described in Section 9.2.2, using $c_1 = 1.2$ and $c_2 = 2.5$. Sample results of such processing are shown in the bottom row of Figure 9.10. These results show that segmentation is considerably simplified by our reconstruction method. If such thresholding-based segmentation were applied to conventional images, the result would be dominated by fluctuations in homogeneous regions, as shown in the fourth row of Figure 9.10.

In Table 9.6, we present the average percentage of accurately classified pixels in segmentations of region-enhanced images using the entire data set. We should note that the major error contributing to our results is due to the gap between the target and shadow regions in the segmentations. This is a systematic error and may be improved upon by incorporation of additional information during segmentation. Our error analysis in Table 9.6 has the limitation that the human segmentations, which we use as the truth, are really not perfect. We should also note that the feature-enhancement parameters λ_1 , λ_2 , and the thresholding parameters c_1 , c_2 have not been optimized for best segmentation performance, but rather picked based on visual assessment of one image, and applied to the processing of the entire data set.

Bhattacharyya Distances	Targ.-Backg.	Targ.-Shad.	Backg.-Shad.
Conventional	0.28	0.76	0.25
Region-Enhanced	1.48	1.81	0.45

Table 9.7: Bhattacharyya distances between Gaussian densities characterizing different regions in SAR images.

9.3.4.3 Statistical Separability of Regions

We now provide a measure of the separability of different regions from each other in conventional and region-enhanced images by using Bhattacharyya distances as described in Section 9.2.3. To this end, we treat the human segmented images as the truth, and extract region labels from these segmentations for each location in the scene. We assume that all the pixels in a particular region are drawn from an identical distribution. We also assume that the probability distribution for a particular type of region is the same for all target types. We can then model each region by a Gaussian density after estimating the mean and variance from the reconstructed dB-valued images. Figure 9.11 illustrates the Gaussian densities we thus obtain for the target, shadow and background regions of the conventional and region-enhanced images. By visual inspection, we can claim that the densities for the region-enhanced images are farther from each other than those for the conventional ones. We can quantify this claim by computing the Bhattacharyya distances between region pairs. Table 9.7 contains the results of this computation, which demonstrate that the Bhattacharyya distances between all region pairs are larger for the region-enhanced images than the conventional images.

9.4 Summary

The results of the study in this chapter show that emphasizing point-based features through the proposed method yields images with higher resolution and better dominant scatterer localization than conventional images. Our experiments indicate that the method is able to produce accurate superresolution reconstructions from considerably reduced amounts of data. Emphasizing region-based features on the other hand, results in enhanced anomaly and speckle suppression in homogeneous regions, and hence, easier-to-segment images. Our results should be compared to those in similar analyses carried out for other enhanced SAR image formation techniques [32, 82, 92].

In this chapter, the focus was mostly on how well the reconstructed images represent the underlying truth, in terms of recognition-oriented features. However, the actual issue of how this behavior is reflected to recognition performance has not been addressed. The impact of feature-enhanced SAR imaging on automatic target recognition is evaluated in Chapter 10.




Figure 9.10: Region-enhanced image formation and segmentation. Top row: MSTAR images. Second row: region-enhanced reconstructions with $k = 1$, $\lambda_1 = 4$, $\lambda_2 = 6$. Third row: human segmentations. Fourth row: threshold-based segmentations of conventional imagery. Bottom row: threshold-based segmentations of region-enhanced imagery.

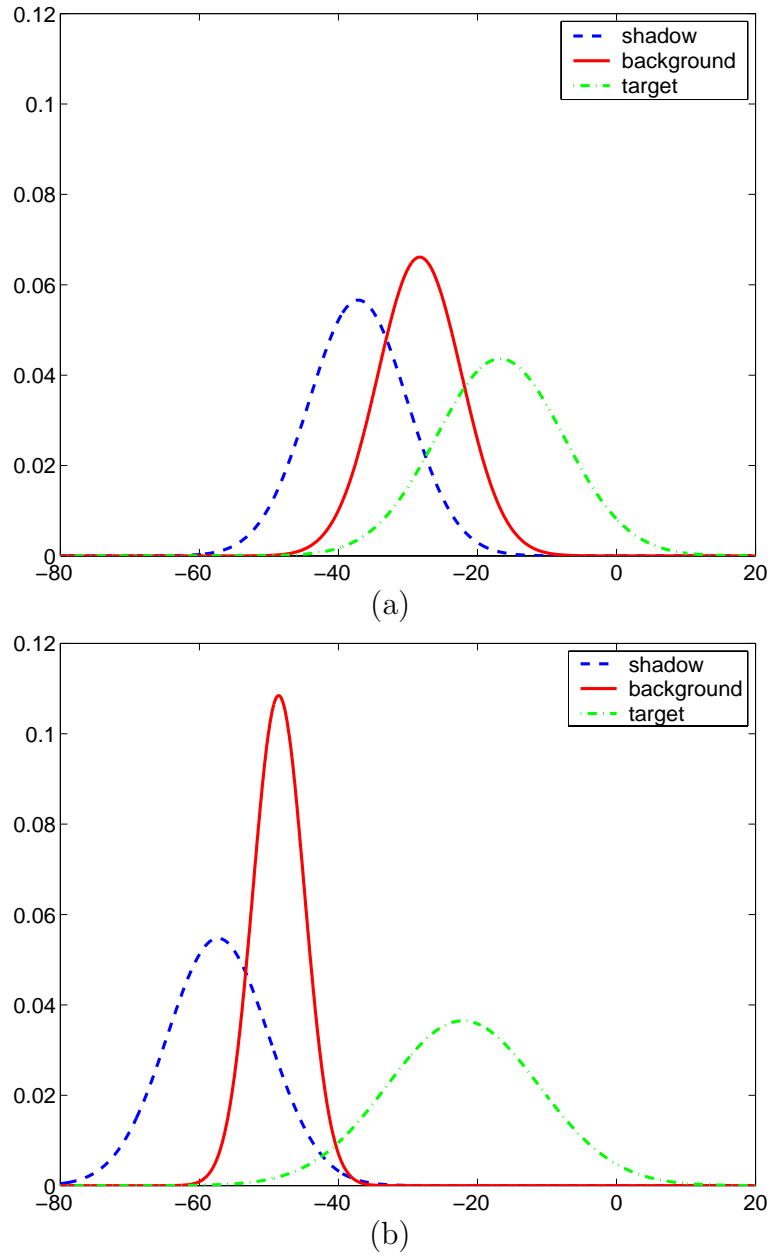


Figure 9.11: Gaussian probability density functions for regions in SAR images. (a) Conventional images. (b) Region-enhanced images.

Chapter 10

Recognition Tests with Feature-Enhanced Imagery

In this chapter, we present an evaluation of the impact of feature-enhanced imaging on automatic recognition performance. We use the MSTAR targets, and three different classifiers, to make decisions about the target type given conventional or feature-enhanced images, and compare the results for these two types of images. The first classifier is based on template-matching, which is a conventional approach for SAR ATR. The second classifier is based on conditionally Gaussian models for reflectivities, and performs classification through a likelihood test. The third classifier is aimed at the use of point-based features in superresolution imaging, and performs classification by measuring how well the dominant scatterer locations are preserved in the face of resolution loss in data. We present the recognition results in the form of confusion matrices. These experiments demonstrate that feature-enhanced imaging can offer higher probability of correct classification than conventional imaging.

10.1 Template-based Classifier

The idea in template-based classification is to first measure how well a given test image matches reference images, called templates, which represent the hypotheses in the problem, and then declare the test image to be from the class giving the best matching score. A common metric used for the degree of match is the mean-squared error (MSE) between the test image, and the template. This classification approach has been used in the classifier component of a series of ATR systems developed by Lincoln Laboratory [35, 93].

The classification problem in ATR is to determine the type of the vehicle in a given scene. In SAR, the image of a particular type of vehicle at a particular orientation (azimuth, aspect angle) with respect to the radar platform is usually not equivalent to a rotated version of the image of the same vehicle at a significantly different orientation. Hence for template-based classification, we need templates of the targets at a reasonable number of different orientations. Let $\mathbf{t}(\vartheta, a)$ be such a template (stacked as a vector) for vehicle type a , at orientation $\vartheta \in [0^\circ, 360^\circ)$, and let $\hat{\mathbf{f}}_{\text{dB}}$ be a normalized (to have unit ℓ_2 -norm) test image (in dB) stacked as a vector. Then the template-based minimum MSE classification can be done as follows:

$$\hat{a}_{\text{MSE}} = \arg \min_a \left(\min_{\vartheta} \|\hat{\mathbf{f}}_{\text{dB}} - \mathbf{t}(\vartheta, a)\|_2^2 \right) \quad (10.1)$$

The classifier in (10.1) needs a stored set of templates, which in practice can be obtained from training data. Usually we do not have many training images at exactly the same orientation. Then, the template for vehicle type a , at orientation ϑ_k can be constructed by first aligning (in angle) the training images that have an orientation close to ϑ_k , and then finding an average of these images, as follows:

$$\mathbf{t}(\vartheta_k, a) = \frac{1}{N_k} \sum_{\vartheta \in W_k} \hat{\mathbf{f}}_{\text{dB}}^r(\vartheta, a) \quad (10.2)$$

where W_k denotes the range of angular orientations that are used in training for the orientation ϑ_k , N_k is the number of available training images that have an orientation within this range, and $\hat{\mathbf{f}}_{\text{dB}}^r(\vartheta, a)$ is a training image which has an original orientation of ϑ , and which is rotated $\vartheta_k - \vartheta$ degrees, so that it is aligned with other training images used for the computation of $\mathbf{t}(\vartheta_k, a)$. Such templates for T72 tanks at 17° depression angle and various orientations for the conventional, point-enhanced (non-superresolution), and region-enhanced images are shown in Figures 10.1, 10.2, and 10.3. The orientation ϑ for each image used in the computation of these templates is obtained from the corresponding MSTAR file header. In the construction of these templates, we have used $\vartheta_k = 5(k - 1)$ degrees, and $W_k = [\vartheta_k - 5, \vartheta_k + 5)$, where $k \in \{1, 2, \dots, 72\}$. Hence the sets of images used in the computation of templates for subsequent orientations intersect each other, and each training image is used in the construction of two templates. In our experimental work, described in Section 10.4, we use images at 17° depression angle for training the classifiers, and an independent set of images at 15° depression angle to test the classification performance.




Figure 10.1: Conventional templates for the T72 target at 17° depression angle. Each image shows the template for a different aspect angle, starting from 0° on the top left, and covering all 360° with 5° increments, and a 10° window for averaging.




Figure 10.2: Point-enhanced templates for the T72 target at 17° depression angle. Each image shows the template for a different aspect angle, starting from 0° on the top left, and covering all 360° with 5° increments, and a 10° window for averaging.




Figure 10.3: Region-enhanced templates for the T72 target at 17° depression angle. Each image shows the template for a different aspect angle, starting from 0° on the top left, and covering all 360° with 5° increments, and a 10° window for averaging.

10.2 Likelihood-based Classifier

We now describe a likelihood-based classifier, proposed in [94] for SAR ATR. This approach starts from a statistical signal model, in which the underlying SAR image $\mathbf{f}(\vartheta, a)$, where ϑ and a are as defined in Section 10.1, is assumed to be a complex Gaussian vector, and the observed image $\hat{\mathbf{f}}$ is given by the following observation model:

$$\hat{\mathbf{f}} = \mathbf{f}(\vartheta, a) + \mathbf{w} \quad (10.3)$$

where \mathbf{w} is i.i.d. complex Gaussian noise with mean $\mathbf{0}$ and covariance $\sigma^2\mathbf{I}$. The classifier is based on a generalized likelihood test as follows:

$$\hat{a}_{\text{GLRT}} = \arg \max_a \left(\max_{\vartheta} l(\hat{\mathbf{f}}|\vartheta, a) \right) \quad (10.4)$$

where $l(\hat{\mathbf{f}}|\vartheta, a)$ denotes the logarithm of the likelihood that the test image is of target type a , at orientation ϑ . Let the signal $\mathbf{f}(\vartheta, a)$ have mean $\mathbf{m}(\vartheta, a)$ and covariance $\mathbf{Q}(\vartheta, a)$. Then the probability density function $p(\hat{\mathbf{f}}|\vartheta, a)$ is Gaussian with the following mean and covariance:

$$E\{\hat{\mathbf{f}}|\vartheta, a\} = \mathbf{m}(\vartheta, a) \quad (10.5)$$

$$E\{[\hat{\mathbf{f}} - \mathbf{m}][\hat{\mathbf{f}} - \mathbf{m}]^H|\vartheta, a\} = \mathbf{Q}(\vartheta, a) + \sigma^2\mathbf{I} \quad (10.6)$$

Let us define $\mathbf{K}(\vartheta, a) \triangleq \mathbf{Q}(\vartheta, a) + \sigma^2\mathbf{I}$. Then the log-likelihood is given by:

$$l(\hat{\mathbf{f}}|\vartheta, a) \propto -\log |\mathbf{K}(\vartheta, a)| - [\hat{\mathbf{f}} - \mathbf{m}(\vartheta, a)]^H (\mathbf{K}(\vartheta, a))^{-1} [\hat{\mathbf{f}} - \mathbf{m}(\vartheta, a)]. \quad (10.7)$$

We will assume that $\mathbf{m}(\vartheta, a) = \mathbf{0}$, as in [94]. This is a reasonable assumption due to the random phase nature of SAR reflectivities. The second simplifying assumption made in [94] is that of statistical independence of different pixels in the SAR image, which results in a diagonal $\mathbf{Q}(\vartheta, a)$, hence a diagonal $\mathbf{K}(\vartheta, a)$. This in turn simplifies the computation of the likelihood to a simple summation:

$$l(\hat{\mathbf{f}}|\vartheta, a) \propto \sum_i \left[-\log(\mathbf{K}_{i,i}(\vartheta, a)) - \frac{|\hat{\mathbf{f}}_i|^2}{\mathbf{K}_{i,i}(\vartheta, a)} \right] \quad (10.8)$$

The classifier in (10.4) with the likelihood function in (10.8) requires the model variances $\mathbf{K}_{i,i}(\vartheta, a)$. These variances can be estimated from training data as follows:

$$\mathbf{K}_{ii}(\vartheta_k, a) = \frac{1}{N_k} \sum_{\vartheta \in W_k} \left| \left(\hat{\mathbf{f}}^r(\vartheta, a) \right)_i \right|^2 \quad (10.9)$$

where $(\hat{\mathbf{f}}^r(\vartheta, a))_i$ denotes the i -th pixel of an aligned training image, and N_k, W_k are as defined in Section 10.1. The operation in (10.9) produces variance “images”, which in spirit are very similar to the template images of Figures 10.1, 10.2, and 10.3.

10.3 Point-feature-based Classifier

We now concentrate on the reconstruction problem from reduced-resolution data, and propose a classifier specifically aimed at evaluating the ability of the point-enhanced, superresolution images to preserve the locations of dominant scatterers in the scene. Similar peak-based classification schemes have previously been proposed in [66] and [73]. Such feature-based classification techniques, rather than the pixel-based classification schemes of the previous two sections, have been an important component of recent research efforts such as those in DARPA’s Moving and Stationary Target Acquisition and Recognition (MSTAR) program [73].

The classifier works as follows. Given a test image, the locations of the largest N_p peaks are extracted. The peak extraction is done as described in Section 9.1.3. Then a one-to-one association is established between these peaks and the true peak locations for the hypothesis under test. Next, an average associated peak distance between these two sets of peaks is computed as described in Section 9.1.4. The class assignment is achieved by finding the vehicle type which yields the minimum associated peak distance over all orientations:

$$\hat{a}_{\text{PEAK}} = \arg \min_a \left(\min_{\vartheta} \frac{1}{N_p} \|\mathbf{d}_{\text{PEAK}}(\hat{\mathbf{f}}, \mathbf{f}(\vartheta, a))\|_1 \right) \quad (10.10)$$

where $\mathbf{d}_{\text{PEAK}}(\cdot, \cdot)$ is a vector of Euclidean distances between associated peak pairs extracted from the two argument images. Note that the ℓ_1 -norm in (10.10) is just equivalent to a summation of the distances between individual peak pairs. The classifier in (10.10) requires the “true” peak locations for each target type, at a reasonable number of orientations ϑ , extracted from the reference scenes $\mathbf{f}(\vartheta, a)$. One way to obtain the reference peak locations is through the use of 3-D CAD models of the vehicles together with an electromagnetic signature prediction tool such as XPATCH [83] to simulate the expected reference scene. A second option is to use conventional SAR images obtained from higher resolution data as the reference scenes, and extract the “true” peak locations from such images. In our experiments, we use the latter approach.

10.4 Experimental Results

10.4.1 Experimental Setup

We now evaluate the performance of the classifiers described in the previous sections, given MSTAR images produced by feature-enhanced versus conventional

Target	Training set		Test set	
	Depression	No. of images	Depression	No. of images
T72	17°	232	15°	196
BMP2	17°	233	15°	195
BTR70	17°	233	15°	196

Table 10.1: Composition of the MSTAR data set used in recognition experiments.

imaging. The MSTAR data set consists of a large assortment of military targets. In our recognition experiments, we use the T72, BMP2 and BTR70 targets. Sample pictures of these vehicles were shown in Figure 9.1. All the images in this chapter (both conventional and feature-enhanced) have been formed from Taylor-windowed data. For all three classifiers, our training set is composed of images at 17° depression angle, and our test set is composed of an independent set of images at 15° depression angle. The numbers of each type of target images in each of these sets are shown in Table 10.1. In training and testing with all three classifiers, we extract and use a near-central portion of the SAR image which contains the target. In all cases, we normalize the images so that they have the same ℓ_2 -norm, before using them in the classifiers.

We present the results of our evaluations in the form of classifier confusion matrices, which show the number of correct and incorrect classifications achieved on test inputs of each type. A single number characterizing the classifier’s ability to recognize test inputs can be obtained through the probability of correct classification, P_{cc} , which is defined as the fraction of all target test inputs that were correctly classified.

10.4.2 Template-based Classification Results

We now present the recognition performance of the conventional, point-enhanced (non-superresolution) and region-enhanced images, when they are used as inputs to the template-based classifier, described in Section 10.1. We use 72 templates for each target type, with each template representing an orientation 5° apart from the subsequent template. In the construction of each template, we use a training window W_k of 10 degrees, as described in Section 10.1.

We have initially conducted recognition experiments with images reconstructed from high SNR data. By high SNR data, we mean the SAR data obtained by undoing the image formation steps for the images in the MSTAR data set, without any additional measurement noise. The templates for such images were shown in Figures 10.1, 10.2, 10.3. We have observed that for this high-SNR, 3-target classification problem both conventional and feature-enhanced images result in a high recognition rate, and the performance difference is not significant. Next, we investigate the recognition performance of these images in the face of degraded data, which may provide a better representation of a practical situation. To this end, we have included additive complex Gaussian noise in the projectional SAR data. Treating the original data before the addition of this noise as the clean signal, the SNR of the corrupted data we use is -4 dB. This represents a very noisy scenario. The templates of the T72 target for the conventional and feature-enhanced images in this case are shown in Figures 10.4, 10.5, and 10.6.

We now present the recognition results for this low-SNR case. Table 10.2 shows the confusion matrices for the classification of conventional and feature-enhanced images. The conventional images result in an overall probability of correct classification of 69.85%. Point-enhanced and region-enhanced imaging increase this rate to 88.93% and 96.76% respectively. So degradation in recognition performance caused by the

noise is much less with our feature-enhanced images, as compared to the conventional images. Hence feature-enhanced imaging is more robust to limitations in data quality.

	T72	BMP2	BTR70
T72	159	30	7
BMP2	53	120	22
BTR70	33	32	131

(a) Conventional images. $P_{cc} = 69.85\%$

	T72	BMP2	BTR70
T72	176	13	7
BMP2	8	173	14
BTR70	3	20	173

(b) Point-enhanced images. $P_{cc} = 88.93\%$

	T72	BMP2	BTR70
T72	184	11	1
BMP2	3	191	1
BTR70	1	2	193

(c) Region-enhanced images. $P_{cc} = 96.76\%$

Table 10.2: Confusion matrices summarizing the template-based classification results. The entry in row i , column j shows the number of images from vehicle type i classified as vehicle j .




Figure 10.4: Conventional templates for the T72 target at 17° depression angle, reconstructed from low-SNR data. Each image shows the template for a different aspect angle, starting from 0° on the top left, and covering all 360° with 5° increments, and a 10° window for averaging.




Figure 10.5: Point-enhanced templates for the T72 target at 17° depression angle, reconstructed from low-SNR data. Each image shows the template for a different aspect angle, starting from 0° on the top left, and covering all 360° with 5° increments, and a 10° window for averaging.




Figure 10.6: Region-enhanced templates for the T72 target at 17° depression angle, reconstructed from low-SNR data. Each image shows the template for a different aspect angle, starting from 0° on the top left, and covering all 360° with 5° increments, and a 10° window for averaging.

10.4.3 Likelihood-based Classification Results

We now present the results of running the low-SNR data used in Section 10.4.2 through the likelihood-based classifier described in Section 10.2. For each target type, we train 72 probability density functions, each representing an orientation 5° apart from the subsequent orientation. As in the construction of the templates in Section 10.4.2, we use a training window W_k of 10 degrees, in estimating the covariance matrices.

The classification results are shown in Table 10.3. The classifier has a correct classification rate of 87.05 % with conventional images, 94.38 % with point-enhanced images, and 99.15 % with region-enhanced images.

10.4.4 Point-feature-based Classification Results

We now report the results of recognition experiments with images reconstructed from data supporting a lower resolution than those used in the previous recognition tests in this chapter. The data used in Sections 10.4.2 and 10.4.3 had a resolution of 0.3 m. In this section, we use 50×50 test data with a resolution of 0.6 m. The data used for the low resolution experiments do not contain any noise in addition to what is already present in the MSTAR data. Hence, we consider a high-SNR scenario here.

We could certainly use a template-based, or likelihood-based classifier for the images reconstructed from such data, too. However, our objective in this section is to present the results of a feature-based classification scheme, rather than a pixel-based one. The classifier we use has been described in Section 10.3. We obtain the true peak locations for all targets at all the orientations available, using conventionally formed SAR images at 0.3 m resolution, and 17° depression angle. We then run the classifier on conventional and point-enhanced, superresolution images at 15° depression angle, reconstructed from data supporting a resolution of 0.6 m.

	T72	BMP2	BTR70
T72	170	18	8
BMP2	7	161	27
BTR70	4	12	180

(a) Conventional images. $P_{cc} = 87.05\%$

	T72	BMP2	BTR70
T72	190	5	1
BMP2	10	178	7
BTR70	8	2	186

(b) Point-enhanced images. $P_{cc} = 94.38\%$

	T72	BMP2	BTR70
T72	193	3	0
BMP2	1	194	0
BTR70	1	0	195

(c) Region-enhanced images. $P_{cc} = 99.15\%$

Table 10.3: Confusion matrices summarizing the likelihood-based classification results. The entry in row i , column j shows the number of images from vehicle type i classified as vehicle j .

The classification results presented in Table 10.4 demonstrate a clear improvement in recognition performance through the use of point-enhanced, superresolution images as compared to the case of conventional images. The correct classification rates for the conventional and point-enhanced images are 44.80%, and 81.43%, respectively.

10.5 Summary

In this chapter we have defined a 3-class target recognition problem to compare the recognition performance of feature-enhanced images and conventional images.

	T72	BMP2	BTR70
T72	38	84	74
BMP2	14	114	67
BTR70	4	81	111

(a) Conventional images. $P_{cc} = 44.80\%$

	T72	BMP2	BTR70
T72	162	23	11
BMP2	28	147	20
BTR70	10	17	169

(b) Point-enhanced, superresolution images. $P_{cc} = 81.43\%$

Table 10.4: Confusion matrices summarizing the point-feature-based classification results, from reduced-resolution data. The entry in row i , column j shows the number of images from vehicle type i classified as vehicle j .

We have constructed two pixel-based classifiers, and one feature-based classifier. With high-SNR data we have not observed significant performance difference between the conventional and feature-enhanced images. With reduced quality data however, feature-enhanced images resulted in higher recognition rates. We have also conducted reduced-resolution data experiments, and used the feature-based classifier to test the performance of our superresolution images. We have observed recognition improvements as compared to the conventional images in this case, as well. Table 10.5 summarizes the results obtained in this chapter in terms of the probability of correct classification.

The recognition systems, and the posed recognition problem in this chapter were only a first step in understanding the recognition behavior of feature-enhanced images. There is definitely need for more extensive studies on this topic. In Section 12.2.1, we suggest a number of such research directions.

Low-SNR experiment			
	Conventional	Point-Enhanced	Region-Enhanced
Template-based classifier	69.85 %	88.93 %	96.76 %
Likelihood-based classifier	87.05 %	94.38 %	99.15 %

Reduced-resolution, high-SNR experiment		
	Conventional	Point-Enhanced, Superresolution
Point-feature-based classifier	44.80 %	81.43 %

Table 10.5: Overall summary of the classification experiments, in terms of the probability of correct classification, P_{cc} .

Chapter 11

Application of Feature-Enhanced Reconstruction in HRR Radar

In this chapter, we demonstrate the application of the feature-enhanced SAR image reconstruction technique of Chapters 5 and 6 in another radar operating mode, known as high range resolution (HRR) radar. We illustrate the potential of our approach to produce superresolution reconstruction of HRR range profiles.

11.1 Introduction to HRR Radar

HRR radars have the same data collection mechanism, as in SAR. The difference between the two modes of operation is the way the collected phase history data are processed and used. In SAR, data at multiple observation angles, and within the extent of a synthetic aperture, are combined to form an image of the scene. In HRR radar, data at each observation angle are used to form a complex HRR range profile.

HRR radar uses high bandwidth pulses so that it achieves a high resolution, which is sufficient to profile a target. As in SAR, the information contained in the

profile is the magnitude of the radar scattering as a function of range, along the line of sight of the radar. This in turn provides information about the length of the object in the scene.

As SAR has been successfully used for imaging stationary scenes, the main objective of HRR radar has been the detection and tracking of vehicles in moving target indication (MTI) systems. An HRR/MTI system provides vehicle target length measurements and can be used as a cueing mechanism for the detection of critical targets through such measurements. This cue can then be handed off to a SAR system for imaging and recognition. However, the high resolution nature of the HRR profiles suggests more advanced uses for this sensor. Consequently, there has recently been much interest in using HRR profiles directly for recognition of moving targets, and both template-based [95–97], and feature-based [98] approaches to the problem have been proposed.

11.2 Superresolution HRR Signal Reconstruction

The primary difficulty associated with the HRR sensor for ATR is that the HRR signatures exhibit a high degree of variability. Yet, characteristic, robust features are required for successful recognition. In [98] locations and amplitudes of the dominant peaks in the HRR profile have been proposed as such features. Superresolution processing techniques for HRR profiles can help in the accurate extraction of such features, and there has been some recent work in this direction. In [99,100], a relaxation-based algorithm has been presented for superresolution target feature extraction, whereas in [101], the SAR imaging technique of [32] has been used for HRR radar. Furthermore, the study in [101] has also demonstrated that superresolution techniques can improve HRR ATR performance. With these motivations, we have extended the use of our point-enhanced, superresolution signal reconstruction technique to the

formation of HRR profiles.

From a mathematical standpoint, the observed data in HRR is equivalent to the SAR data received at a single observation angle. Hence the inverse problem we formulate here is one of obtaining a complex HRR profile from the received, pre-processed HRR phase history signal. To this end, let \mathbf{q} be the sampled HRR profile, and \mathbf{h} be the noisy sampled phase history data at a particular observation angle. Then, we have the following observation model:

$$\mathbf{h} = \mathbf{F}\mathbf{q} + \mathbf{w} \quad (11.1)$$

where \mathbf{w} is measurement noise, and \mathbf{F} is a high resolution to low resolution DFT matrix. This definition of \mathbf{F} reflects the belief that the underlying object (hence its profile) possesses high-frequency features that are not captured by the resolution supported by the data. The conventional way to reconstruct the HRR profile is through an inverse DFT, which, in this framework can be represented by $\hat{\mathbf{q}}_{\text{CONV}} = \mathbf{F}^H \mathbf{h}$, with appropriate normalization. In contrast, we formulate the HRR profile reconstruction problem as the following optimization problem:

$$\hat{\mathbf{q}} = \arg \min_{\mathbf{q}} J(\mathbf{q}) \quad (11.2)$$

where we choose $J(\mathbf{q})$ to be an objective function of the following form:

$$J(\mathbf{q}) = \|\mathbf{h} - \mathbf{F}\mathbf{q}\|_2^2 + \lambda_1^2 \|\mathbf{q}\|_k^k \quad (11.3)$$

where k and λ_1 are scalar parameters. Apart from the nature of the observation matrix \mathbf{F} , and the 1-D as opposed to 2-D structure of the signals, the optimization problem in (11.3) is in the same form as the optimization problem defined for SAR imaging in Chapter 5. Hence the algorithm described in Chapter 6 can be readily

used for the solution of this problem.

We end this section by pointing out the relationship between our technique for superresolution HRR profile reconstruction and the field of adaptive signal representation [102–104]. Adaptive signal representation addresses the problem of finding optimal representations of signals as combination of elements from an overcomplete dictionary. Such techniques have been used for feature extraction from HRR profiles [104]. One adaptive signal representation technique, called basis pursuit denoising [103], finds an optimal representation by minimizing an objective function of the same mathematical form as (11.3) with $k = 1$. Hence, we can interpret our signal reconstruction method as one of finding the optimal basis pursuit denoised representation of the observed HRR data, in terms of the complex exponential dictionary elements.

11.3 Examples

In our examples, we use the University Research Initiative Synthetic Dataset (URISD) [84]. The URISD provides a collection of simulated range profiles and the associated phase histories, produced by the HRR simulator XPATCH [83]. The URISD includes data for four ground vehicles (two tank models, school bus, fire truck) over three frequency bands (UHF, L, X). In our examples, we use the X-band data which has a central frequency of 10 GHz, and a bandwidth of 1.4775 GHz. This implies a resolution of around 0.1 m.

We use conventionally reconstructed 0.1 m resolution profiles as the “reference” profiles. In our experiments, we use limited data, supporting a resolution of 0.4 m to obtain conventional and point-enhanced reconstructions. We compare the resulting reconstructions to the reference profile in terms of the preservation of point-based features.

Figure 11.1(a) shows the reference profile for a fire truck at 0° orientation. The conventional reconstruction from reduced resolution data is shown in Figure 11.1(b). Note that the two most dominant peaks in the reference profile are merged into one in the conventional reconstruction. The point-enhanced reconstruction in Figure 11.1(c) however, resolves these dominant peaks. Note that the point-enhanced reconstruction also preserves the amplitudes of the peaks better and appears to provide a good visual approximation to the reference profile. We can quantify this degree of similarity by calculating the MSE between the reference and reconstructed profiles. The conventional profile in this case yields an MSE of $9.4 \cdot 10^{-3}$, and the point-enhanced profile yields an MSE of $6.3 \cdot 10^{-3}$, hence the latter is a better approximation of the reference profile in the MSE sense, as well. Figure 11.2 provides the detail of the profiles around the region where the dominant peaks of the reference profile are located. This plot demonstrates the capability of the point-enhanced method in resolving and accurately locating the dominant peaks. Figure 11.3 shows the profiles in the case of lower SNR data. To obtain such data, we have added complex Gaussian noise with a standard deviation of 0.3 to the phase history observations. If we treat the original phase histories before the addition of this noise as the clean signal, then this level of noise is equivalent to an SNR of 36 dB. The conventional profile in Figure 11.3(b) suffers from random oscillations in the background caused by the noise, whereas the point-enhanced reconstruction in Figure 11.3(c) is able to suppress those artifacts.

Next, we show similar reconstructions for another vehicle, the M1 tank. Figure 11.4 contains the M1 profiles, where similar observations to the fire truck example can be made. In particular, the two leftmost dominant peaks in the reference profile of Figure 11.4(a) are merged into one by the conventional reconstruction of Figure 11.4(b), but they are resolved by the point-enhanced reconstruction of Figure 11.4(c).

Our work in this chapter was aimed at suggesting that our techniques for SAR

image reconstruction may prove to be useful for HRR radar, as well. However, more extensive studies are required to investigate the practical significance of such techniques for HRR radar.




Figure 11.1: HRR profiles of fire truck at 0° orientation, reconstructed from X-band data. (a) Reference profile at 0.1 m resolution. (b) Conventional reconstruction from 0.4 m resolution data. (c) Point-enhanced, superresolution reconstruction from 0.4 m resolution data ($k = 1$, $\lambda_1^2 = 10$).




Figure 11.2: Detail of the profiles in Figure 11.1 around the 110th range bin. Conventional reconstruction cannot resolve the two peaks, whereas point-enhanced reconstruction can.




Figure 11.3: HRR profiles of fire truck at 0° orientation, reconstructed from low SNR X-band data. (a) Reference profile at 0.1 m resolution. (b) Conventional reconstruction from 0.4 m resolution data. (c) Point-enhanced, superresolution reconstruction from 0.4 m resolution data ($k = 1$, $\lambda_1^2 = 10$).




Figure 11.4: HRR profiles of M1 tank at 0° orientation, reconstructed from X-band data. (a) Reference profile at 0.1 m resolution. (b) Conventional reconstruction from 0.4 m resolution data. (c) Point-enhanced, superresolution reconstruction from 0.4 m resolution data ($k = 0.9$, $\lambda_1^2 = 3$).

Chapter 12

Conclusions and Future Directions

12.1 Summary and Conclusions

In this dissertation, we have contributed to the area of SAR imaging and to the area of regularized reconstruction of complex-valued fields. We have developed a novel mathematical foundation and associated algorithms for feature-enhanced imaging of complex-valued SAR reflectivity fields. We have provided mathematical extensions of robust regularization techniques to be used in complex-valued, random-phase problems. We have presented methods for quantitative assessment of the proposed techniques, and have shown potential benefits of the new perspective of this dissertation through extensive evaluation experiments on real SAR data.

In Chapter 5, our framework for feature-enhanced, regularized SAR imaging was established. At the heart of our formulation was an objective function to be minimized for image formation. This objective function consisted of a data fidelity term and prior information terms. The data fidelity term required the use of a discrete observation model. To be used as the observation model, the concept of a SAR projection matrix was introduced, and various interesting properties of this matrix were illustrated.

The prior information terms were aimed at preserving and enhancing certain types of features in the scene, and took into account the complex-valued, random phase nature of the SAR reflectivities. Our specific choices for these terms, to enhance point-based and region-based features, and their anticipated role in image reconstruction were discussed. Finally, a statistical interpretation of the posed problem was presented.

The optimization problem posed in Chapter 5 was challenging, due to both the random-phase nature of the reflectivities, and the presence of non-quadratic functions of the field, which were needed for effective, robust feature enhancement. Chapter 6 provided a robust and efficient numerical solution technique for this type of optimization-based image formation problems. This technique was based on a specific quasi-Newton method matched to the structure of the SAR image formation problem. The technique was then shown to be a generalization of half-quadratic regularization methods. Various aspects of the iterative scheme were discussed, including the mechanism by which feature-enhancement is achieved and the role of auxiliary processes. The rest of the chapter presented experimental results. Through reconstructions of a variety of real and synthetic scenes, the effectiveness of the proposed approach in forming SAR images with enhanced features and suppressed artifacts was demonstrated. Particular benefits included images with higher resolution than that supported by the measured data, as well as images with reduced variability of reflectivity magnitudes within homogeneous regions and preserved region boundaries. Experimental analyses were carried out to illuminate and quantify various aspects of the numerical solution method.

Chapter 7 extended the formulation of Chapter 5 to allow the use of different classes of potential functions in the prior information terms, and considered the use of three particular classes of functions. The generalization of the optimization problem, and of the numerical solution method to include such new potential functions was presented. The anticipated behavior of the three classes of potential functions

considered were discussed and compared. Finally, experimental results demonstrating the effectiveness of SAR imaging with the new potential functions were presented.

Chapter 8 provided an extension of our framework to the type of variational formulations based on the Mumford-Shah energy, commonly used in image segmentation problems. This extension involved a generalized version of the Ambrosio-Tortorelli objective function to be used in problems involving complex-valued, random-phase fields. This extension provided a richer and more flexible structure in the formulation of the optimization problems for SAR image formation. Potential benefits of this richness were illuminated through examples.

In Chapter 9, quantitative criteria for evaluating SAR images based on recognition-oriented features were introduced. While the experimental results in previous chapters had provided promising results based on subjective visual assessment, the aim of this chapter was to present a quantitative, feature-based analysis. The quantitative criteria used for point-enhanced images included the target-to-clutter ratio, mainlobe width, peak matching accuracy, and average associated peak distance. The criteria used for evaluating region-enhanced images were speckle suppression, segmentation accuracy, and statistical separability of regions. Extensive experiments on large sets of real and synthetic SAR images indicated that feature-enhanced SAR images exhibited superresolution and improved localization accuracy for dominant scatterers, and improved separability and ease of segmentation for different region types, which are all important attributes for automated decision-making from SAR images.

Chapter 10 demonstrated the impact of feature-enhanced SAR imaging on an automatic target recognition task. Three different ATR algorithms, including a template-based classifier, a likelihood-based classifier, and a point-feature-based classifier were implemented. Conventional and feature-enhanced SAR images were used as inputs to these decision-making systems. The recognition results, presented in the form of confusion matrices and the probability of correct classification, demonstrated

potential improvements provided by feature-enhanced SAR imaging in the face of limited data quality and reduced resolution.

In Chapter 11, extension of the use of feature-enhanced reconstruction methods to HRR radar was considered. The main focus here was superresolution, and resolution improvements achieved through feature-enhanced signal reconstruction were demonstrated by experimental results.

Overall, this dissertation has presented a new perspective to complex-valued image reconstruction problems, in particular those arising in SAR. At the center of the new philosophy were goal-directed processing based on intended use of the images, effective use of prior knowledge, and robustness to data limitations. The methods developed in this dissertation, resulting from such a philosophy, produce imagery from which accurate features can be efficiently extracted. This improves the feature extraction process which makes our approach a promising technology, especially in applications where automated recognition or decision-making algorithms are to be applied to the formed imagery.

12.2 Topics for Future Research

12.2.1 Further Analysis and Evaluation of the Techniques Developed

Further Analysis on Potential Function Choices

In Chapter 7, we have demonstrated the use of a number of classes of potential functions in constructing objective functions for SAR imaging. However, we have not addressed the question of which potential function is the “best” in some sense for a particular type of scene or objective. Although all potential functions considered in this work are aimed at similar objectives, further analysis of the relationship between the potential function and the resulting reconstruction can help in automated or semi-automated choice of both the class of potential functions, and the particular function (through the choice of parameter k) within a class to be used in image formation.

Comparison to Other Superresolution Methods

In Chapter 3, we have discussed a number of recently proposed SAR image formation techniques, which share some of the objectives of our approach to the problem. In particular, the published results on the spectral-estimation-based methods resemble the reconstructions obtained by our point-enhanced, superresolution method. We believe a qualitative and quantitative comparative evaluation of the two approaches, when applied to the same data set would be valuable.

Extensive Experiments on a Feature-based Recognition System

In Chapter 10, we have presented some recognition experiments to compare the performance of feature-enhanced and conventional images. Although this analysis

has provided a flavor of the impact of feature-enhanced imaging, more extensive experimentation on a more general setting would be an important contribution for a number of reasons. First, the recognition problem posed in Chapter 10 contained three classes. This is too small to represent a practical setting where the decisions usually involve a larger number of vehicle types. Furthermore, our experiments did not involve *confuser* vehicles. Confuser vehicles are test inputs, for which no training is done, and the ATR system should ideally respond as “unknown” to data containing such vehicles. This is an important practical issue, and a more realistic recognition evaluation should involve confuser vehicles. Finally, although the pixel-based classifiers of Chapter 10 were based on real recognition systems in use, there are more advanced, feature-based systems under development, such as that associated with the MSTAR program, as described in [73]. Currently, we are collaborating with the authors of [73], to test our images in their recognition system, which will hopefully address all the issues mentioned above.

Recognition Tests on HRR Radar Data

In Chapter 11, we have demonstrated a superresolution signal reconstruction technique for HRR radar. However, we have not investigated the impact of such reconstructions on the problem of target recognition from HRR radar profiles. Such an analysis would be valuable in evaluating the significance of our results.

12.2.2 Variations and Enhancements of the Techniques Developed

Use of Priors for Other Features

The potential functions in this dissertation, introduced in Chapters 5 and 7, used either the field reflectivities themselves, or the gradients of the reflectivity magnitudes, as their arguments. This was based on the relationship of these quantities to the types of features we wanted to preserve. Our framework however is not limited to these choices. One example of a different prior, used in surface and shape estimation is the so-called thin-plate prior model, which involves the second derivatives of the field. With a statistical perspective, this would be equivalent to a different correlational structure in the prior probability density function for the field. Our framework can be easily extended to allow the use of various functions of the field reflectivities as arguments of the potential functions.

More Physical/High Level/Statistically-based Prior Models

One of the key elements of our philosophy for image reconstruction was the inclusion of prior information about the scene. The information incorporated in this way was at a particular level of abstraction. It was not at a physical level of electromagnetic interactions and scattering. On the other extreme, it did not carry high level characterizations such as the shape, size etc. of the physical objects expected to be present in the scene. In a variety of applications, there is ongoing work in both ends of the spectrum: in physics-based signal processing, and in pattern theory. Use of such different levels of information for the SAR image formation and automated scene interpretation problems is an open area with interesting philosophical and practical questions.

On a different note, the optimization problems posed for image formation in this dissertation had statistical interpretations, as discussed in Section 5.6. However, the resulting statistical prior models themselves were really not based on the “statistics of the truth” itself. Work in the direction of incorporating more statistically-based information can add value to the estimation-theoretic perspective of our approach to SAR imaging.

Automated Choice of Hyperparameters

Our approach to image reconstruction involves scalar parameters, which need to be picked appropriately for successful imaging. In this dissertation, we have not attempted to choose those parameters automatically, but have picked them manually. In regularization problems, automated choice of hyperparameters based on data is still an open and important research problem. Development and incorporation of an efficient automated parameter choice technique producing reasonable results would add value to feature-enhanced SAR imaging.

Closing the Loop between Recognition and Image Formation

One of the elements of our image formation philosophy was to tailor the image formation process to the final objectives of the mission, by enhancing features important for automated decisions, such as recognition. Within our framework however, there are certain choices, which can still have an impact on the recognition performance. So far, we have made choices, such as the class of potential functions, the particular potential function through the selection of k , and the regularization parameters in an open-loop fashion, without optimizing recognition performance over these choices. Constructing a mechanism to include the feedback from the decision-making stage in driving the image formation process, as sketched in Figure 12.1, would increase




Figure 12.1: Illustration of the idea of incorporating feedback from the final decision-making task in image formation.

the effectiveness of our SAR imaging approach. In fact, if performed accurately and efficiently, this would be the ideal way to tune the image formation process for a particular decision-making task. An approach based on such a closed-loop recognition philosophy has recently been proposed in [105, 106] for integrated segmentation and recognition of optical images.

Multiresolution Image Formation

In our reconstruction framework, we have constructed the prior models directly on the field. In a number of similar domains however, researchers have obtained various advantages by constructing the prior on a transformation of the field. An example of this philosophy is multiscale statistical models, where the priors are built on a multiscale representation of the field. One advantage obtained through these kinds of models is multiresolution image formation, which naturally leads to an incremental refinement property for the reconstructed image. Figure 12.2 demonstrates the idea on a small SAR image reconstruction example. Multiscale models have also led to computational advantages via reasonable approximations in the transform domain. Extension of our framework to include such ideas is a future research topic.




Figure 12.2: Illustration of the multiresolution reconstruction concept for a SAR image. Each scale also corresponds to intermediate results at various stages of image formation.

Choosing Priors in a Projectional Transform Domain

Another direction for choosing priors not directly on the field is to choose them in a domain closer to the projectional observations. Such an idea has previously been applied to the CT problem [107], and has been shown to offer computational advantages. Similar approaches may be investigated for the SAR problem, as well.

Incorporation of Non-Point-Scattering Effects

In our SAR observation model, we have assumed a point-scattering mechanism, that is, we have assumed the reflected signal from a scatterer to be the same pure delay and scaling of the transmitted signal, for all observation regimes and angles. This in turn assumes that the scatterers are isotropic and frequency-independent, i.e. the reflectivities do not change as a function of the observation angle of the sensor

and frequency of the transmitted signal. These are assumptions that are also made by all conventional SAR imaging techniques. Physics of electromagnetic scattering however show that these may not be very realistic assumptions, and different physical structures can lead to different scattering behaviors. There is some recent work which attempts to take into account the aspect or frequency dependence present in SAR [19, 108, 109]. In our formulation, such an attempt would require the use of an observation model $\mathbf{T}(\mathbf{f})$, which depends on the field \mathbf{f} itself. This would be a challenging and interesting research direction for our framework.

12.2.3 Relations and Extensions to Other Domains, Problems, Methods

Application of the Techniques in Other Observation Modes or Domains

The model-based nature of our approach simplifies its use in problems requiring an observation model that is different from the one considered in this dissertation.

The first example of such an extension is regarding the approximations we have made in constructing an observation model. These include the sensor-scene distance being much larger than the scene radius, or the quadratic phase term being very small, as discussed in Section 2.2. It is straightforward to build and incorporate in our method, a more general observation model not requiring these assumptions, in case they are violated.

Secondly, there are other modes of SAR operation, such as stripmap-mode SAR, where the geometry of data collection is different from that in spotlight-mode SAR. Extension of our techniques to such operating modes only requires the use of an appropriate discrete observation model for that particular sensing scenario.

Finally, there are other inverse problems, especially in inverse scattering, which are geometrically, and physically similar to SAR. Examination of the extension and usefulness of our methods to such domains for data inversion, especially involving complex-valued fields is another research path.

Extension to Moving-Target Scenes

SAR imaging has so far been successful for stationary scenes, and that is what we have assumed in our work, as well. Imaging of scenes containing moving targets has recently started to attract attention [110]. Using feature-enhanced SAR imaging for moving scenes is another challenging research direction.

Appendix A

Appendix

A.1 Discrete 2-D Derivative Operators

In our methods, we use smoothness constraints on a field, which require the spatial derivatives of the field. We use the horizontal and vertical first order difference operators in approximating such derivatives. Derivatives of the field in other directions, such as the diagonals may be used as well, however we have found the use of only horizontal and vertical derivatives sufficient. Consider a real-valued, sampled field \mathbf{z} , column stacked as a vector of length $N = N_x N_y$, where N_x and N_y denote the number of rows and columns respectively in the 2-D field. We can compute first order differences of this field, $\mathbf{D}_x \mathbf{z}$ and $\mathbf{D}_y \mathbf{z}$, in the horizontal and vertical directions respectively, where the discrete derivative operators are given by:

$$\mathbf{D}_x = \begin{bmatrix} -\mathbf{I} & \mathbf{I} & & & \\ & \ddots & \ddots & & \\ & & & -\mathbf{I} & \mathbf{I} \end{bmatrix} \quad (\text{A.1})$$

and

$$\mathbf{D}_y = \begin{bmatrix} \mathbf{D}_1 & & & \\ & \mathbf{D}_1 & & \\ & & \ddots & \\ & & & \mathbf{D}_1 \end{bmatrix}, \quad (\text{A.2})$$

with

$$\mathbf{D}_1 = \begin{bmatrix} -1 & 1 & & \\ & \ddots & \ddots & \\ & & -1 & 1 \end{bmatrix}. \quad (\text{A.3})$$

Note that, since we take first order differences between neighboring pixels, it is appropriate to define the discrete derivatives between the locations of the adjacent pixels. With the above definitions, \mathbf{D}_x has a size of $N_y(N_x - 1) \times N_x N_y$, and \mathbf{D}_y has a size of $N_x(N_y - 1) \times N_x N_y$. Hence, these are non-square operators. However, if the use of square derivative operators is desired, the above definitions can be augmented by derivatives defined at the boundary of the field to make these operators square. This may be preferred, for example, when one wants to associate the derivatives to pixel locations.

We now describe two ways to compute the smoothness constraint terms of the form $\|\mathbf{D}\mathbf{z}\|_k^k$, that appear in objective functions such as (5.9). The discussion can easily be generalized to smoothness constraints with other potential functions, such as those considered in Chapter 7.

The first approach is based on treating the horizontal and vertical derivatives separately in imposing a smoothness constraint. This is achieved by defining the 2-D discrete derivative operator \mathbf{D} as follows:

$$\mathbf{D} \triangleq \begin{bmatrix} \mathbf{D}_x \\ \mathbf{D}_y \end{bmatrix}. \quad (\text{A.4})$$

With this definition, we can write $\|\mathbf{Dz}\|_k^k$ as

$$\|\mathbf{Dz}\|_k^k = \sum_{i=1}^M |(\mathbf{Dz})_i|^k = \sum_{i=1}^{M_x} |(\mathbf{D}_x\mathbf{z})_i|^k + \sum_{i=1}^{M_y} |(\mathbf{D}_y\mathbf{z})_i|^k = \|\mathbf{D}_x\mathbf{z}\|_k^k + \|\mathbf{D}_y\mathbf{z}\|_k^k \quad (\text{A.5})$$

where $M_x \triangleq N_y(N_x - 1)$, $M_y \triangleq N_x(N_y - 1)$, and $M = M_x + M_y$.

The second approach is based on treating the gradient at each pixel location as a two-element vector $[(\mathbf{D}_x\mathbf{z})_i \ (\mathbf{D}_y\mathbf{z})_i]^T$, composing of the horizontal and vertical gradients, and using the ℓ_2 -norms of such gradients at all locations in the field for the computation of the overall ℓ_k -norm:

$$\|\mathbf{Dz}\|_k^k \triangleq \sum_{i=1}^N [|(\mathbf{D}_x\mathbf{z})_i|^2 + |(\mathbf{D}_y\mathbf{z})_i|^2]^{k/2} \quad (\text{A.6})$$

Two things must be noted here. First, the use of a *linear* operator \mathbf{D} is only conceptual in this case, because no such explicit matrix exists. Second, this approach requires a one-to-one correspondence between horizontal and vertical derivatives at each location in the scene, hence in this case we use square ($N \times N$) derivative operators \mathbf{D}_x , \mathbf{D}_y .

In our methods, we have used the first approach described above, unless stated otherwise. All the mathematical expressions in the body of this dissertation are also based on the first approach. Note that when $k = 2$, the two approaches are identical, with the use of square derivative operators. To make the association between the two approaches clear, let us consider square derivative operators, and examine the first approach in this case:

$$\|\mathbf{Dz}\|_k^k = \sum_{i=1}^N |(\mathbf{D}_x\mathbf{z})_i|^k + \sum_{i=1}^N |(\mathbf{D}_y\mathbf{z})_i|^k \quad (\text{A.7})$$

$$= \sum_{i=1}^N [|(\mathbf{D}_x\mathbf{z})_i|^k + |(\mathbf{D}_y\mathbf{z})_i|^k] \quad (\text{A.8})$$

Let us compare this expression to the second approach given in (A.6). In (A.6), ℓ_2 -norm of the gradient vector at each location is used in the computation of the overall ℓ_k -norm. In contrast, the first approach, as shown in (A.8), corresponds to using an ℓ_k -norm for the gradient vector $[(\mathbf{D}_x \mathbf{z})_i \ (\mathbf{D}_y \mathbf{z})_i]^T$ at each location. This association lets us compare the consequences of using the two approaches. For example, when $k < 2$, the first approach used in a smoothness constraint would favor horizontal/vertical edges more than diagonal edges, relative to the second approach.

A.2 Gradient of the Objective Function for Image Reconstruction

Let us first define:

$$\mathbf{g}_b \triangleq \begin{bmatrix} \Re(\mathbf{g}) \\ \Im(\mathbf{g}) \end{bmatrix}, \quad \mathbf{f}_b \triangleq \begin{bmatrix} \Re(\mathbf{f}) \\ \Im(\mathbf{f}) \end{bmatrix}, \quad \mathbf{T}_b \triangleq \begin{bmatrix} \Re(\mathbf{T}) & -\Im(\mathbf{T}) \\ \Im(\mathbf{T}) & \Re(\mathbf{T}) \end{bmatrix} \quad (\text{A.9})$$

where $\Re(\cdot)$ and $\Im(\cdot)$ denote the real and imaginary components respectively. These matrices are real-valued versions of the original complex-valued matrices formed by stacking the real and imaginary parts appropriately.

When we take the gradient of (6.3) with respect to the real and imaginary parts of \mathbf{f} , we obtain the following gradient expression:

$$\nabla J_\epsilon^1(\mathbf{f}) = 2\mathbf{T}_b^T \mathbf{T}_b \mathbf{f}_b + k\lambda_1^2 \mathbf{Q}_1(\mathbf{f}) \mathbf{f}_b + k\lambda_2^2 \mathbf{\Phi}_{b1}^T(\mathbf{f}) \mathbf{D}^T \mathbf{\Lambda}_2(\mathbf{f}) \mathbf{D} |\mathbf{f}| - 2\mathbf{T}_b^T \mathbf{g}_b \quad (\text{A.10})$$

where

$$\mathbf{Q}_1(\mathbf{f}) \triangleq \begin{bmatrix} \mathbf{\Lambda}_1(\mathbf{f}) & \mathbf{0} \\ \mathbf{0} & \mathbf{\Lambda}_1(\mathbf{f}) \end{bmatrix} \quad (\text{A.11})$$

$$\mathbf{\Phi}_{b1}(\mathbf{f}) \triangleq \begin{bmatrix} \mathbf{c} & \mathbf{s} \end{bmatrix} \quad (\text{A.12})$$

$$\mathbf{c} \triangleq \text{diag} \{ \cos(\phi[(\mathbf{f})_i]) \} \quad (\text{A.13})$$

$$\mathbf{s} \triangleq \text{diag} \{ \sin(\phi[(\mathbf{f})_i]) \} \quad (\text{A.14})$$

$(\mathbf{f})_i$ denotes the i -th element of \mathbf{f} , $\phi[z]$ denotes the phase of the complex number z , and $\mathbf{\Lambda}_1(\mathbf{f})$, $\mathbf{\Lambda}_2(\mathbf{f})$ are as defined in (6.6). Note that $\nabla J_\epsilon^1(\mathbf{f})$ is a $2N$ -vector, where N is the length of the complex vector \mathbf{f} .

We would like to rewrite (A.10) in a form that resembles a linear function of \mathbf{f}_b . The reason is that we would like to use ideas from half-quadratic regularization. The relationship between our method and half-quadratic regularization is explained in Section 6.3, however let us provide a brief, qualitative explanation for desiring a linear-like gradient here. As discussed in Section 4.3, half-quadratic regularization uses an augmented cost function which is quadratic in the unknown field. Consequently, the gradient of the augmented cost function is linear in the unknown field (when the auxiliary variable is fixed). It is this linear structure that we are after here. To this end, let us concentrate on $\mathbf{\Phi}_{b1}^T(\mathbf{f})\mathbf{D}^T\mathbf{\Lambda}_2(\mathbf{f})\mathbf{D}|\mathbf{f}|$ which is the only term that needs to be manipulated for this purpose. Let us first define:

$$\mathbf{\Phi}_b(\mathbf{f}) \triangleq \begin{bmatrix} \mathbf{c} & \mathbf{s} \\ -\mathbf{s} & \mathbf{c} \end{bmatrix} \quad (\text{A.15})$$

$$\mathbf{Q}_2(\mathbf{f}) \triangleq \begin{bmatrix} \mathbf{D}^T\mathbf{\Lambda}_2(\mathbf{f})\mathbf{D} & \mathbf{0} \\ \mathbf{0} & \mathbf{D}^T\mathbf{\Lambda}_2(\mathbf{f})\mathbf{D} \end{bmatrix} \quad (\text{A.16})$$

and then do the following manipulation:

$$\begin{aligned} \Phi_{b1}^T(\mathbf{f})\mathbf{D}^T\Lambda_2(\mathbf{f})\mathbf{D}|\mathbf{f}| &= \begin{bmatrix} \mathbf{c}\mathbf{D}^T\Lambda_2(\mathbf{f})\mathbf{D}|\mathbf{f}| \\ \mathbf{s}\mathbf{D}^T\Lambda_2(\mathbf{f})\mathbf{D}|\mathbf{f}| \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{c} & -\mathbf{s} \\ \mathbf{s} & \mathbf{c} \end{bmatrix} \begin{bmatrix} \mathbf{D}^T\Lambda_2(\mathbf{f})\mathbf{D}|\mathbf{f}| \\ \mathbf{0} \end{bmatrix} \end{aligned} \quad (\text{A.17})$$

$$= \begin{bmatrix} \mathbf{c} & -\mathbf{s} \\ \mathbf{s} & \mathbf{c} \end{bmatrix} \begin{bmatrix} \mathbf{D}^T\Lambda_2(\mathbf{f})\mathbf{D} & \mathbf{0} \\ \mathbf{0} & \mathbf{D}^T\Lambda_2(\mathbf{f})\mathbf{D} \end{bmatrix} \begin{bmatrix} |\mathbf{f}| \\ \mathbf{0} \end{bmatrix} \quad (\text{A.18})$$

Now, noting that

$$\begin{bmatrix} |\mathbf{f}| \\ \mathbf{0} \end{bmatrix} = \Phi_b(\mathbf{f})\mathbf{f}_b \quad (\text{A.19})$$

and using the definitions in (A.15) and (A.16), we obtain:

$$\Phi_{b1}^T(\mathbf{f})\mathbf{D}^T\Lambda_2(\mathbf{f})\mathbf{D}|\mathbf{f}| = \Phi_b^T(\mathbf{f})\mathbf{Q}_2(\mathbf{f})\Phi_b(\mathbf{f})\mathbf{f}_b. \quad (\text{A.20})$$

Now, defining

$$\mathbf{H}_b(\mathbf{f}) \triangleq 2\mathbf{T}_b^T\mathbf{T}_b + k\lambda_1^2\mathbf{Q}_1(\mathbf{f}) + k\lambda_2^2\Phi_b^T(\mathbf{f})\mathbf{Q}_2(\mathbf{f})\Phi_b(\mathbf{f}) \quad (\text{A.21})$$

we can write the gradient as:

$$\nabla J_\epsilon^1(\mathbf{f}) = \mathbf{H}_b(\mathbf{f})\mathbf{f}_b - 2\mathbf{T}_b^T\mathbf{g}_b. \quad (\text{A.22})$$

This is in the form that we have desired: it resembles a linear function of \mathbf{f}_b , when the \mathbf{f} -dependence of the matrix $\mathbf{H}_b(\mathbf{f})$ is ignored.

We could develop our iterative scheme using the gradient expression in (A.22), which involves real-valued vectors of length $2N$. However, we choose to formulate

the algorithm in terms of complex-valued N -vectors. The two representations are equivalent. The reasons why we choose the latter are the following. First, it provides a compact, easier to interpret structure. Second, we have noticed some computational advantages of dealing with small, complex-valued, rather than large, real-valued matrices in MATLAB. To this end, first note:

$$\mathbf{H}_b(\mathbf{f}) \triangleq \begin{bmatrix} \Re(\mathbf{H}(\mathbf{f})) & -\Im(\mathbf{H}(\mathbf{f})) \\ \Im(\mathbf{H}(\mathbf{f})) & \Re(\mathbf{H}(\mathbf{f})) \end{bmatrix} \quad (\text{A.23})$$

where

$$\mathbf{H}(\mathbf{f}) \triangleq 2\mathbf{T}^H\mathbf{T} + k\lambda_1^2\mathbf{\Lambda}_1(\mathbf{f}) + k\lambda_2^2\mathbf{\Phi}^H(\mathbf{f})\mathbf{D}^T\mathbf{\Lambda}_2(\mathbf{f})\mathbf{D}\mathbf{\Phi}(\mathbf{f}) \quad (\text{A.24})$$

$$\mathbf{\Phi}(\mathbf{f}) \triangleq \text{diag} \{ \exp(-j\phi[(\mathbf{f})_i]) \} \quad (\text{A.25})$$

We note the fact that for any complex-valued matrices $\mathbf{M}_1, \mathbf{M}_2, \mathbf{M}_3$, such that $\mathbf{M}_3 = \mathbf{M}_1\mathbf{M}_2$, we have:

$$\begin{bmatrix} \Re(\mathbf{M}_3) & -\Im(\mathbf{M}_3) \\ \Im(\mathbf{M}_3) & \Re(\mathbf{M}_3) \end{bmatrix} = \begin{bmatrix} \Re(\mathbf{M}_1) & -\Im(\mathbf{M}_1) \\ \Im(\mathbf{M}_1) & \Re(\mathbf{M}_1) \end{bmatrix} \begin{bmatrix} \Re(\mathbf{M}_2) & -\Im(\mathbf{M}_2) \\ \Im(\mathbf{M}_2) & \Re(\mathbf{M}_2) \end{bmatrix}. \quad (\text{A.26})$$

We use the property in (A.26) to deduce:

$$\nabla J_\epsilon^1(\mathbf{f}) = \begin{bmatrix} \Re(\nabla J_\epsilon(\mathbf{f})) \\ \Im(\nabla J_\epsilon(\mathbf{f})) \end{bmatrix} \quad (\text{A.27})$$

where

$$\nabla J_\epsilon(\mathbf{f}) = \mathbf{H}(\mathbf{f})\mathbf{f} - 2\mathbf{T}^H\mathbf{g}. \quad (\text{A.28})$$

We use the complex-valued ‘‘gradient’’ $\nabla J_\epsilon(\mathbf{f})$ in (6.4).

A.3 Augmented Cost Function for Half-Quadratic Regularization

The objective of this appendix is to prove the relationship (6.14), which we repeat here:

$$\inf_{\mathbf{b}, \mathbf{s}} K(\mathbf{f}, \mathbf{b}, \mathbf{s}) = J_\epsilon(\mathbf{f}). \quad (\text{A.29})$$

This relationship shows that $K(\mathbf{f}, \mathbf{b}, \mathbf{s})$ of (6.12), is a valid augmented cost function to be used in half-quadratic regularization for the function $J_\epsilon(\mathbf{f})$ of (6.11).

To keep the derivation simple, we will consider a 1-D signal \mathbf{f} , rather than a 2-D field in this appendix. The results however can easily be extended to the 2-D case. We will assume the following structure for the discrete 1-D derivative operator \mathbf{D} :

$$\mathbf{D} = \begin{bmatrix} -1 & 1 & & \\ & \ddots & \ddots & \\ & & -1 & 1 \end{bmatrix} \quad (\text{A.30})$$

which simply consists of two-element differences.

Let us now find \mathbf{s} and \mathbf{b} which minimize $K(\mathbf{f}, \mathbf{b}, \mathbf{s})$. First consider \mathbf{s} . The portion of $K(\mathbf{f}, \mathbf{b}, \mathbf{s})$ which depends on \mathbf{s} is the following:

$$\sum_{i=1}^M \mathbf{b}_i |(\mathbf{D}\mathbf{S}\mathbf{f})_i|^2. \quad (\text{A.31})$$

Based on the structures of \mathbf{D} in (A.30) and \mathbf{S} in (6.13), we have

$$(\mathbf{D}\mathbf{S}\mathbf{f})_i = -e^{-j\mathbf{s}_i} \mathbf{f}_i + e^{-j\mathbf{s}_{i+1}} \mathbf{f}_{i+1} \quad (\text{A.32})$$

and consequently

$$|(\mathbf{D}\mathbf{S}\mathbf{f})_i|^2 = |\mathbf{f}_i|^2 + |\mathbf{f}_{i+1}|^2 - 2\Re \{ |\mathbf{f}_i| |\mathbf{f}_{i+1}| e^{j(\phi[(\mathbf{f})_i] - \phi[(\mathbf{f})_{i+1}] + \mathbf{s}_{i+1} - \mathbf{s}_i)} \} \quad (\text{A.33})$$

Here $\phi[(\mathbf{f})_i]$ denotes the phase of the complex number \mathbf{f}_i . The sum in (A.31) takes its minimum value, when the term inside the brackets in (A.33) has a zero imaginary part for all i . Hence the minimizing \mathbf{s} satisfies:

$$\mathbf{s}_{i+1} - \mathbf{s}_i + \phi[(\mathbf{f})_i] - \phi[(\mathbf{f})_{i+1}] = 0 \quad (\text{A.34})$$

We could have obtained this result by the following qualitative argument as well. We want to minimize (A.31), which is a weighted sum of squared norms of the differences between complex number pairs of the form $z_i = e^{-j\mathbf{s}_i} \mathbf{f}_i$. The variables we have for optimization are \mathbf{s}_i for all i , hence we can essentially choose the phase of each complex number. Naturally, the minimum is obtained when the complex numbers z_i have identical phase, since this makes the norm of the difference between two complex numbers as small as possible. This is exactly what the condition in (A.34) implies: the optimum \mathbf{s}_i should “rotate” \mathbf{f}_i in such a way that the resulting z_i have the same phase for all i . Note that we still have a freedom in choosing what that identical phase is. If we simply choose it to be 0, then we have the following optimal \mathbf{s} :

$$\mathbf{s}_i = \phi[(\mathbf{f})_i], \quad \forall i \quad (\text{A.35})$$

Note that with this \mathbf{s} , we have $\mathbf{S}\mathbf{f} = |\mathbf{f}|$. Hence,

$$\inf_{\mathbf{s}} K(\mathbf{f}, \mathbf{b}, \mathbf{s}) = \|\mathbf{g} - \mathbf{T}\mathbf{f}\|_2^2 + \lambda_2^2 \sum_{i=1}^M \left[\mathbf{b}_i (|(\mathbf{D}|\mathbf{f}|)_i|^2 + \epsilon) + \frac{1}{4\mathbf{b}_i} \right] \quad (\text{A.36})$$

Next, let us consider \mathbf{b} . Differentiating the term inside the summation in (A.36), and setting it equal to zero, we obtain the following condition for the minimizing \mathbf{b} :

$$\mathbf{b}_i = \frac{1}{2(|(\mathbf{D}\mathbf{f})_i|^2 + \epsilon)^{1/2}} \quad (\text{A.37})$$

Substituting (A.37) in $K(\mathbf{f}, \mathbf{b}, \mathbf{s})$, we obtain the result we desired:

$$\inf_{\mathbf{b}, \mathbf{s}} K(\mathbf{f}, \mathbf{b}, \mathbf{s}) = \|\mathbf{g} - \mathbf{T}\mathbf{f}\|_2^2 + \lambda_2^2 \sum_{i=1}^M (|(\mathbf{D}\mathbf{f})_i|^2 + \epsilon)^{1/2} = J_\epsilon(\mathbf{f}) \quad (\text{A.38})$$

which shows that (6.14) is true.

Bibliography

- [1] J. H. G. Ender. The airborne experimental multi-channel SAR-system AER-II. In *European Conference on Synthetic Aperture Radar*, pages 49–52, March 1996.
- [2] D. Sicard. RAMSES airborne SAR development status & experiments. In *European Conference on Synthetic Aperture Radar*, pages 45–48, March 1996.
- [3] J. C. Henry. The Lincoln Laboratory 35 GHz airborne polarimetric SAR imaging system. In *IEEE National Telesystems Conference*, pages 353–358, Atlanta, GA, USA, March 1991.
- [4] D. R. Bromaghim and J. P. Perry. A wideband linear FM ramp generator for the long-range imaging radar. *IEEE Transactions on Microwave Theory and Techniques*, MTT-26(5):322–325, May 1978.
- [5] J. C. Kirk Jr. Motion compensation for synthetic aperture radar. *IEEE Transactions on Aerospace and Electronic Systems*, AES-11(3):338–348, May 1975.
- [6] C. J. Oliver. Synthetic-aperture radar imaging. *Journal of Physics D - Applied Physics*, 22(7):871–890, July 1989.

- [7] A. Freeman, D. Evans, and J. J. van Zyl. SAR applications in the 21st century. In *European Conference on Synthetic Aperture Radar*, pages 25–30, March 1996.
- [8] D. A. Ausherman, A. Kozma, J. L. Walker, H. M. Jones, and E. C. Poggio. Developments in radar imaging. *IEEE Transactions on Aerospace and Electronic Systems*, AES-20(4):363–400, July 1984.
- [9] M. Soumekh. Automatic aircraft landing using interferometric inverse synthetic aperture radar imaging. *IEEE Transactions on Image Processing*, 5(9):1335–1345, September 1996.
- [10] J. Walker. Range-Doppler imaging of rotating objects. *IEEE Transactions on Aerospace and Electronic Systems*, AES-16(1):23–52, January 1980.
- [11] D. C. Munson Jr., J. D. O’Brien, and W. K. Jenkins. A tomographic formulation of spotlight-mode synthetic aperture radar. *Proceedings of the IEEE*, 71:917–925, August 1983.
- [12] D. C. Munson Jr. and R. L. Visentin. A signal processing view of strip-mapping synthetic aperture radar. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 37:2131–2147, December 1989.
- [13] B. Bhanu, D. E. Dudgeon, E. G. Zelnio, A. Rosenfeld, D. Casasent, and I. S. Reed. Introduction to the special issue on automatic target detection and recognition. *IEEE Transactions on Image Processing*, 6(1):1–6, January 1997.
- [14] D. Mumford and J. Shah. Boundary detection by minimizing functionals, I. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 22–26, 1985.
- [15] J. C. Curlander and R. N. McDonough. *Synthetic Aperture Radar: Systems and Signal Processing*. John Wiley & Sons, Inc., New York, 1991.

- [16] J. P. Fitch. *Synthetic Aperture Radar*. Springer-Verlag, New York, 1988.
- [17] C. V. Jakowatz Jr., D. E. Wahl, P. H. Eichel, D. C. Ghiglia, and P. A. Thompson. *Spotlight-mode Synthetic Aperture Radar: a Signal Processing Approach*. Kluwer Academic Publishers, Norwell, MA, 1996.
- [18] C. Wu, K. Y. Liu, and M. Jin. Modeling and a correlation algorithm for spaceborne SAR signals. *IEEE Transactions on Aerospace and Electronic Systems*, AES-18:563–575, 1982.
- [19] R. D. Chaney, A. S. Willsky, and L. M. Novak. Coherent aspect-dependent SAR image formation. In *Algorithms for Synthetic Aperture Radar Imagery*, pages 256–274. SPIE, April 1994.
- [20] D. C. Munson Jr. and J. L. C. Sanz. Image reconstruction from frequency-offset Fourier data. *Proceedings of the IEEE*, 72:661–669, June 1984.
- [21] A. W. Rihaczek. *Principles of High-Resolution Radar*. IEEE Press, New York, 1988.
- [22] A. C. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging*. IEEE Press, New York, 1988.
- [23] A. J. den Dekker and A. van den Bos. Resolution: a survey. *Journal of the Optical Society of America*, 14(3):547–557, March 1997.
- [24] A. Di Cenzo. A comparison of resolution for spotlight synthetic-aperture radar and computer-aided tomography. *Proceedings of the IEEE*, 74:1165–1166, August 1986.
- [25] W. G. Carrara, R. S. Goodman, and R. M. Majewski. *Spotlight Synthetic Aperture Radar: Signal Processing Algorithms*. Artech House, Boston, 1995.

- [26] M. D. Desai and W. K. Jenkins. Convolution backprojection image reconstruction for spotlight mode synthetic aperture radar. *IEEE Transactions on Image Processing*, 1(4):505–517, October 1992.
- [27] H. Choi and D. C. Munson Jr. Direct-Fourier reconstruction in tomography and synthetic aperture radar. *International Journal of Imaging Systems and Technology*, 9:1–13, January 1998.
- [28] C. W. Therrien. *Discrete Random Signals and Statistical Signal Processing*. Prentice-Hall, Englewood Cliffs, NJ, 1992.
- [29] S. R. Degraaf. SAR imaging via modern 2-D spectral estimation methods. *IEEE Transactions on Image Processing*, 7(5):729–761, May 1998.
- [30] J. Li and P. Stoica. An adaptive filtering approach to spectral estimation and SAR imaging. *IEEE Transactions on Signal Processing*, 44(6):1469–1484, June 1996.
- [31] G. R. Benitz. Adaptive high-definition imaging. In D. A. Giglio, editor, *Algorithms for Synthetic Aperture Radar Imagery*, volume 2230 of *Proceedings of SPIE*, pages 106–119, Orlando, FL, USA, 1994.
- [32] G. R. Benitz. High-definition vector imaging. *Lincoln Laboratory Journal*, 10(2):147–170, 1997.
- [33] S. Barbarossa, L. Marsili, and G. Mungari. SAR super-resolution imaging by signal subspace projection techniques. In *European Conference on Synthetic Aperture Radar*, pages 267–270, March 1996.
- [34] J. A. C. Lee and D. C. Munson Jr. Spatially variant apodization for image reconstruction from partial Fourier data. *IEEE Transactions on Image Processing*, 9(11):1914–1924, November 2000.

- [35] G. J. Owirka, S. M. Verbout, and L. M. Novak. Template-based SAR ATR performance using different image enhancement techniques. In E. G. Zelnio, editor, *Algorithms for Synthetic Aperture Radar Imagery VI*, volume 3721 of *Proceedings of SPIE*, pages 302–319, Orlando, FL, USA, April 1999.
- [36] J. H. McClellan. Parametric signal modeling. In J. S. Lim and A. V. Oppenheim, editors, *Advanced Topics in Signal Processing*, pages 1–57. Prentice-Hall, Englewood Cliffs, NJ, 1988.
- [37] D. C. Munson Jr. and E. A. Ullman. Support-limited extrapolation of offset Fourier data (synthetic aperture radar). In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages 2483–2486, April 1986.
- [38] H. Li, N. H. Farhat, and Y. Shen. A new iterative algorithm for extrapolation of data available in multiple restricted regions with application to radar imaging. *IEEE Transactions on Antennas and Propagation*, AP-35(5):581–588, May 1987.
- [39] I. J. Gupta, M. J. Beals, and A. Moghaddar. Data extrapolation for high resolution radar imaging. *IEEE Transactions on Antennas and Propagation*, 42(11):1540–1545, November 1994.
- [40] I. J. Gupta and A. Gandhe. Comparison of various enhanced radar imaging techniques. In *Algorithms for Synthetic Aperture Radar Imagery V*, volume 3370 of *Proceedings of SPIE*, pages 250–260, April 1998.
- [41] S. R. Degraaf. SAR imaging via modern 2-D spectral estimation methods. In *Algorithms for Synthetic Aperture Radar Imagery*, volume 2230 of *Proceedings of SPIE*, pages 36–47, April 1994.
- [42] J. A. C. Lee, O. Arikan, and D. C. Munson Jr. Formulation of a general imaging algorithm for high-resolution synthetic aperture radar. In *IEEE International*

- Conference on Acoustics, Speech and Signal Processing*, volume 4, pages 2092–2095, September 1996.
- [43] C. L. Logan. *An Estimation-Theoretic Technique for Motion-Compensated Synthetic-Aperture Array imaging*. PhD thesis, Massachusetts Institute of Technology, February 2000.
- [44] S. P. Lutrell. Prior knowledge and object reconstruction using the best linear estimate technique. *Optica Acta*, 32:703–716, 1985.
- [45] L. M. Delves, G. C. Pryde, and S. P. Lutrell. A super-resolution algorithm for SAR images. *Inverse Problems*, 4:681–703, 1988.
- [46] S. P. Lutrell and C. J. Oliver. Prior knowledge in synthetic-aperture radar processing. *Journal of Physics D - Applied Physics*, 19:333–356, 1986.
- [47] B. Borden. Maximum entropy regularization in inverse synthetic aperture radar imagery. *IEEE Transactions on Signal Processing*, 40(4):969–973, April 1992.
- [48] B. Borden. Cross-entropy regularization and complex-valued image analysis applicable to ISAR. In *Radar Processing, Technology, and Applications*, volume 2845 of *Proceedings of SPIE*, pages 175–182, August 1996.
- [49] B. R. Frieden and A. T. Bajkova. Entropic reconstruction of complex images. *Optics Communications*, 102(5-6):515–522, October 1993.
- [50] B. R. Frieden and A. T. Bajkova. Bayesian cross-entropy reconstruction of complex images. *Applied Optics*, 33(2):219–226, January 1994.
- [51] B. Borden. Regularization of noisy ISAR images containing extended features. *IEEE Transactions on Image Processing*, 8(1):124–127, January 1999.

- [52] B. Borden. Some issues in inverse synthetic aperture radar image reconstruction. *Inverse Problems*, 13:571–584, 1997.
- [53] W. C. Karl. Regularization in image restoration and reconstruction. In A. Bovik, editor, *Handbook of Image and Video Processing*, pages 141–160. Academic Press, 2000.
- [54] G. Demoment. Image reconstruction and restoration: Overview of common estimation structures and problems. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37(12):2024–2036, December 1989.
- [55] D. L. Phillips. A technique for the numerical solution of certain integral equations of the first kind. *Journal of the Association for Computing Machinery*, 9(1):84–97, January 1962.
- [56] A. N. Tikhonov. Solution of incorrectly formulated problems and the regularization method. *Soviet Mathematics - Doklady*, 4:1035–1038, 1963.
- [57] M Hanke and P. C. Hansen. Regularization methods for large-scale problems. *Surveys on Mathematics for Industry*, 3:253–315, 1993.
- [58] B. R. Frieden. Restoring with maximum likelihood and maximum entropy. *Journal of the Optical Society of America*, 62:511–518, 1972.
- [59] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60(1-4):259–268, November 1992.
- [60] C. R. Vogel and M. E. Oman. Fast, robust total variation-based reconstruction of noisy, blurred images. *IEEE Transactions on Image Processing*, 7(6):813–824, June 1998.

- [61] D. Geman and G. Reynolds. Constrained restoration and the recovery of discontinuities. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(3):367–383, March 1992.
- [62] D. Geman and C. Yang. Nonlinear image recovery with half-quadratic regularization. *IEEE Transactions on Image Processing*, 4(7):932–946, July 1995.
- [63] P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud. Deterministic edge-preserving regularization in computed imaging. *IEEE Transactions on Image Processing*, 6(2):298–310, February 1997.
- [64] J. Idier. Convex half-quadratic criteria and interacting auxiliary variables for image restoration. *IEEE Transactions on Image Processing*. Submitted for publication.
- [65] L. C. Potter and R. L. Moses. Attributed scattering centers for SAR ATR. *IEEE Transactions on Image Processing*, 6(1):79–91, January 1997.
- [66] G. Jones, III and B. Bhanu. Recognition of articulated and occluded objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(7):603–613, July 1999.
- [67] S. Sibisi, J. Skilling, R. G. Brereton, E. D. Laue, and J. Staunton. Maximum entropy signal processing in practical NMR spectroscopy. *Nature*, 311:446–447, October 1984.
- [68] R. Narayan and R. Nityananda. Maximum entropy restoration in astronomy. *Annual Review of Astronomy and Astrophysics*, 24:127–170, 1986.
- [69] D. L. Donoho, I. M. Johnstone, J. C. Koch, and A. S. Stern. Maximum entropy and the nearly black object. *Journal of the Royal Statistical Society, Series B*, 54(1):41–81, 1992.

- [70] P. Ciuciu, J. Idier, and J. F. Giovannelli. Markovian high resolution spectral analysis. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, volume 3, pages 1601–1604, Phoenix, AZ, USA, March 1999.
- [71] B. Bhanu. Automatic target recognition: state of the art survey. *IEEE Transactions on Aerospace and Electronic Systems*, AES-22(4):364–379, July 1986.
- [72] J. A. Ratches, C. P. Walters, and R. G. Buser. Aided and automatic target recognition based upon sensory inputs from image forming systems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(9):1004–1019, September 1997.
- [73] W. W. Irving and G. J. Ettinger. Classification of targets in synthetic aperture radar imagery via quantized grayscale matching. In E. G. Zelnio, editor, *Algorithms for Synthetic Aperture Radar Imagery VI*, volume 3721 of *Proceedings of SPIE*, pages 320–331, Orlando, FL, USA, April 1999.
- [74] C. Bouman and K. Sauer. A generalized Gaussian image model for edge-preserving MAP estimation. *IEEE Transactions on Image Processing*, 2(3):296–310, July 1993.
- [75] D. L. Mensa. *High Resolution Radar Imaging*. Artech House, Dedham, MA, 1981.
- [76] D. G. Luenberger. *Linear and Nonlinear Programming*. Addison Wesley, Reading, MA, 1989.
- [77] C. R. Vogel and M. E. Oman. Iterative methods for total variation denoising. *SIAM Journal on Scientific Computing*, 17(1):227–238, January 1996.

- [78] G. H. Golub and C. F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, 1996.
- [79] A. H. Delaney and Y. Bresler. Globally convergent edge-preserving regularized reconstruction: an application to limited-angle tomography. *IEEE Transactions on Image Processing*, 7(2):204–221, February 1998.
- [80] Air Force Research Laboratory, Model Based Vision Laboratory, Sensor Data Management System ADTS Web Page: <http://www.mbvlab.wpafb.af.mil/public/sdms/datasets/adts/>.
- [81] Air Force Research Laboratory, Model Based Vision Laboratory, Sensor Data Management System MSTAR Web Page: <http://www.mbvlab.wpafb.af.mil/public/sdms/datasets/mstar/>.
- [82] A. E. Brito, S. H. Chan, and S. D. Cabrera. SAR image formation using 2-D re-weighted minimum norm extrapolation. In E. G. Zelnio, editor, *Algorithms for Synthetic Aperture Radar Imagery VI*, volume 3721 of *Proceedings of SPIE*, pages 78–90, 1999.
- [83] M. Hazlett, D. J. Andersh, S. W. Lee, H. Ling, and C. L. Yu. XPATCH: a high-frequency electromagnetic scattering prediction code using shooting and bouncing rays. In W. R. Watkins and D. Clement, editors, *Targets and Backgrounds: Characterization and Representation*, volume 2469 of *Proceedings of SPIE*, pages 266–275, Orlando, FL, USA, April 1995.
- [84] Center for Imaging Science URISD Web Page: http://cis.jhu.edu/wu_sensor_data/urisd/.
- [85] M. D. Sacchi, T. J. Ulrych, and C. J. Walker. Interpolation and extrapolation using a high-resolution discrete Fourier transform. *IEEE Transactions on Signal Processing*, 46(1):31–38, January 1998.

- [86] S. Geman and D. E. McClure. Statistical methods for tomographic image reconstruction. In *Proceedings of the 46th Session of the International Statistical Institute*, volume 52 of *Bulletin of the ISI*, pages 5–21, 1987.
- [87] T. Hebert and R. Leahy. A generalized EM algorithm for 3-D Bayesian reconstruction from Poisson data using Gibbs priors. *IEEE Transactions on Medical Imaging*, 8:194–202, June 1989.
- [88] S. Teboul, L. Blanc-Féraud, G. Aubert, and M. Barlaud. Variational approach for edge-preserving regularization using coupled PDE's. *IEEE Transactions on Image Processing*, 7(3):387–397, March 1998.
- [89] L. Ambrosio and V. M. Tortorelli. Approximation of functionals depending on jumps by elliptic functionals via Γ -convergence. *Communications on Pure and Applied Mathematics*, 43(8):999–1036, December 1990.
- [90] Z. Ying and D. A. Castañón. Statistical model for occluded object recognition. In *IEEE International Conference on Information, Intelligence and Systems*, pages 324–327, 1999.
- [91] K. Fukunaga. *Introduction to Statistical Pattern Recognition*. Academic Press, New York, 1990.
- [92] W. Phillips, S. DeGraaf, and R. Chellappa. Enhanced segmentation of SAR images using non-Fourier imaging. In *IEEE International Conference on Image Processing*, volume 1, pages 583–586, Chicago, IL, USA, 1998.
- [93] L. M. Novak, G. J. Owirka, W. S. Brower, and A. L. Weaver. The automatic target-recognition system in SAIP. *Lincoln Laboratory Journal*, 10(2):187–202, 1997.

- [94] J. A. O’Sullivan, M. D. Devore, V. Kedia, and M. I. Miller. SAR ATR performance using a conditionally Gaussian model. *IEEE Transactions on Aerospace and Electronic Systems*. Submitted for publication.
- [95] S. P. Jacobs and J. A. O’Sullivan. High resolution radar models for joint tracking and recognition. In F. A. Sadjadi, editor, *Automatic Target Recognition VII*, volume 3069 of *Proceedings of SPIE*, pages 94–105, 1997.
- [96] S. P. Jacobs and J. A. O’Sullivan. Automatic target recognition using sequences of high resolution radar range-profiles. *IEEE Transactions on Aerospace and Electronic Systems*, 36(2):364–382, April 2000.
- [97] R. Williams, J. Westerkamp, J. Gross, and A Palomino. Automatic target recognition of time critical moving targets using 1D high range resolution (HRR) radar. *IEEE Aerospace and Electronic Systems Magazine*, 15(4):37–43, April 2000.
- [98] R. A. Mitchell and J. J. Westerkamp. Robust statistical feature based aircraft identification. *IEEE Transactions on Aerospace and Electronic Systems*, 35(3):1077–1094, July 1999.
- [99] Z. Bi, R. Wu, J. Li, and R Williams. Joint super-resolution moving target feature extraction and stationary clutter suppression. *IEE Proceedings - Radar, Sonar and Navigation*, 147(1):23–29, February 2000.
- [100] G. Liu, H. Li, and J. Li. Moving target feature extraction with polarisation diversity in the presence of arbitrary range migration and phase errors. *IEE Proceedings - Radar, Sonar and Navigation*, 147(4):208–216, August 2000.
- [101] D. H. Nguyen, G. R. Benitz, J. H. Kay, and R. H. Whiting. Super-resolution HRR ATR performance with HDVI. In F. A. Sadjadi, editor, *Automatic Target*

- Recognition X*, volume 4050 of *Proceedings of SPIE*, Orlando, FL, USA, April 2000.
- [102] S. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Transactions on Signal Processing*, 41(12):3397–3415, 1993.
- [103] S. Chen. *Basis Pursuit*. PhD thesis, Stanford University, 1995.
- [104] S. Jaggi, W. C. Karl, H. Krim, S. Mallat, and A. S. Willsky. Feature extraction through high-resolution pursuit. *Applied and Computational Harmonic Analysis*, 5, 1998.
- [105] J. Peng and B. Bhanu. Closed-loop object recognition using reinforcement learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(2):139–154, February 1998.
- [106] B. Bhanu and J. Peng. Adaptive integrated image segmentation and object recognition. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 30(4):427–441, November 2000.
- [107] M. Bhatia, W. C. Karl, and A. S. Willsky. Tomographic reconstruction and estimation based on multiscale natural-pixel bases. *IEEE Transactions on Image Processing*, 6(3):463–478, March 1997.
- [108] A. J. Kim, J. W. Fisher III, A. S. Willsky, and P. Viola. Nonparametric estimation of aspect dependence for ATR. In E. G. Zelnio, editor, *Algorithms for Synthetic Aperture Radar Imagery VI*, volume 3721 of *Proceedings of SPIE*, pages 332–342, Orlando, FL, USA, April 1999.
- [109] A. J. Kim, J. W. Fisher III, R. L. Moses, and A. S. Willsky. Attributing scatterer anisotropy for model based ATR. In E. G. Zelnio, editor, *Algorithms*

for Synthetic Aperture Radar Imagery VII, volume 4053 of *Proceedings of SPIE*, Orlando, FL, USA, April 2000. SPIE.

- [110] S. Barbarossa and A. Farina. Space-time-frequency processing of synthetic aperture radar signals. *IEEE Transactions on Aerospace and Electronic Systems*, 30(2):341–358, April 1994.

Vita

Müjdat Çetin was born in İzmit, Kocaeli, Turkey, in 1971. He received the B.S. degree from Boğaziçi University, İstanbul, Turkey, in 1993, and the M.S. degree from the University of Salford, Manchester, UK, in 1995, both in electrical and electronic engineering. He was a member of the Multi-Dimensional Signal Processing Laboratory at Boston University, while he was pursuing the Ph.D. degree. He then joined the Stochastic Systems Group at MIT, as a post-doctoral researcher. His research interests include statistical signal and image processing with an emphasis on inverse problems arising in remote sensing and medical imaging, statistical pattern recognition, and detection and estimation theory. He can be reached by e-mail at `mcetin@mit.edu`.