

Inference Methods for Synthetic Aperture Radar
and Array Processing With Lattices

Dissertation

Presented in Partial Fulfillment of the Requirements for the Degree
Doctor of Philosophy in the Graduate School of The Ohio State
University

By

David Tucker, B.S.

Graduate Program in Electrical and Computer Engineering

The Ohio State University

2024

Dissertation Committee:

Lee C. Potter, Advisor

Emre Ertin

Kiryung Lee

© Copyright by

David Tucker

2024

In reference to IEEE copyrighted material which is used with permission in this thesis, the IEEE does not endorse any of Ohio State University's products or services. Internal or personal use of this material is permitted. If interested in reprinting/republishing IEEE copyrighted material for advertising or promotional purposes or for creating new collective works for resale or redistribution, please go to [Right-sLink](#) to learn how to obtain a License from RightsLink. If applicable, University Microfilms and/or ProQuest Library, or the Archives of Canada may supply single copies of the dissertation.

Abstract

This dissertation addresses two distinct topics: inference from synthetic aperture radar (SAR) data and the application of lattice theory to array signal processing. First, we propose a series of methods that mitigate the effect of speckle in inference tasks involving multichannel SAR data. Coherent imaging methods like SAR are subject to speckle, and the suppression of this noise-like quality is often a prerequisite to image interpretation. We propose a technique that recovers the per-pixel multichannel SAR covariance matrix and incorporates a statistical model of speckle and *a priori* knowledge of the varieties of clutter present in the scene. In this approach, an expectation-maximization algorithm is made computationally tractable by a graph-coloring probing technique. We next address the problem of coherent change detection in repeat-pass SAR data. A Bayesian change detection approach is given that assigns prior distributions to the unobserved model variables to exploit spatial structure both in the geophysical scattering qualities of the scene and among the scene disturbances that take place between the passes. We also give a polarimetric SAR (PolSAR) despeckling method based on convolutional neural networks (CNNs). An invertible transformation involving a matrix logarithm is used to facilitate CNN processing of the PolSAR data. A residual learning strategy is adopted, in which the CNN is trained to identify the speckle component which is then removed from the corrupted image.

The latter portion of this dissertation is concerned with the application of lattice theory to array signal processing. We consider the problem of maximum likelihood parameter estimation in mixed integer linear models and provide two polynomial-time solution methods for special cases of this problem. These approaches extend the prior art by allowing for multivariate real-valued unknowns and more general linear models. We then provide a generally applicable alternative solution method that uses sphere decoding.

We next consider the design and analysis of nonuniform arrays in one, two, and three dimensions. We give simply tested necessary and sufficient conditions for an array of sensors to unambiguously determine the direction of arrival for a specified set of possible directions of arrival. The new results facilitate the design of nonuniform arrays, allowing for configurations with widely separated sensors or increased apertures without an increase in the number of sensors. The unambiguous region is shown to be a parallelotope, a property which admits simple geometric interpretation and facilitates array design.

Acknowledgments

First, I thank the faculty of Ohio State’s electrical and computer engineering department. In particular, I thank my advisor Lee Potter for his mentorship and patience. I also thank Emre Ertin, Kiryung Lee, Bradley Clymer, and Phil Schniter for serving on my exam committees. I am grateful to the Center for Surveillance Research, an NSF Industry and University Cooperative Research Center (IUCRC), and AFRL for funding in part the work in this dissertation. I also thank my coauthors for their insight and collaboration. Finally, I thank my family and friends for their support.

Vita

December 2013	B.S., Electrical and Computer Engineering, The Ohio State University
August 2014–June 2016	Radar systems engineer, MDA Information Systems
September 2016–Present	Graduate research associate, The Ohio State University.

Publications

D. Tucker and L. C. Potter, “Approximate maximum likelihood estimation in split-dwell pulse-Doppler radar,” in *IEEE Radar Conf.*, to be published.

D. Tucker, S. Zhao, and L. C. Potter, “Maximum likelihood estimation in mixed integer linear models,” *IEEE Signal Process. Lett.*, vol. 30, pp. 1557–1561, 2023.

D. Tucker, J. N. Ash, and L. C. Potter, “SAR coherent change detection with variational expectation maximization,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 59, no. 3, pp. 2163–2175, 2023.

D. Tucker, S. Zhao, R. Ahmad, and L. C. Potter, “Alias-free arrays,” *IEEE Signal Process. Lett.*, vol. 29, pp. 2457–2461, 2022.

S. Zhao, R. Ahmad, L. Potter, and D. Tucker “Venc design and velocity estimation for phase-contrast MRI,” US Patent 63 324 133, filed Mar. 27, 2022.

D. Tucker and L. C. Potter, “Polarimetric SAR despeckling with convolutional neural networks,” *IEEE Trans. Geosci. Remote Sens.*, vol. 60, pp. 1–12, 2022.

D. Tucker and L. C. Potter, “Speckle suppression in multi-channel coherent imaging: a tractable Bayesian approach,” *IEEE Trans. Comput. Imag.*, vol. 6, pp. 1429–1439, 2020.

Fields of Study

Major Field: Electrical and Computer Engineering

Specialization: Signal Processing

Table of Contents

	Page
Abstract	ii
Acknowledgments	iv
Vita	v
List of Tables	xii
List of Figures	xiv
1. Introduction	1
1.1 Inference Methods for Synthetic Aperture Radar	1
1.2 Array Processing With Lattices	3
2. Speckle Suppression in Multichannel SAR Imaging	6
2.1 Introduction	6
2.1.1 Motivating Applications	7
2.1.2 Ameliorating Speckle by Post-Processing	8
2.1.3 Regularized Imaging	9
2.2 Imaging Model	11
2.2.1 Spotlight-Mode SAR Imaging	11
2.2.2 Multichannel SAR	12
2.2.3 Modeling Speckle in Multichannel SAR	12
2.2.4 Markov Random Field Prior	13
2.3 MAP Estimation	14
2.3.1 MAP Objective Surrogate Function	16
2.3.2 Evaluating the Conditional Posterior Covariance	17
2.3.3 Summary of MAP Estimation Approach	21

2.4	Experimental Results	24
2.4.1	Evaluation of Algorithm Performance	24
2.4.2	PolSAR data	25
2.4.3	Simulated Data Results	26
2.4.4	Measured Data Results	30
2.4.5	Discussion	33
2.5	Conclusion	34
2.5.1	Summary	34
2.5.2	Future Work	35
3.	Coherent Change Detection in SAR Imaging	36
3.1	Introduction	36
3.2	Background	39
3.2.1	Variational Inference	39
3.2.2	Gamma Markov Random Fields	41
3.3	Coherent Change Detection	43
3.3.1	Bayesian Change Detection Model	43
3.3.2	Inference	45
3.3.3	Parameter Estimation	47
3.3.4	Summary	51
3.4	Experiments	52
3.4.1	Simulated Data Experiments	52
3.4.2	Measured Data Experiments	60
3.5	Conclusion	66
3.5.1	Summary	66
3.5.2	Future Work	67
4.	Polarimetric SAR Despeckling	68
4.1	Introduction	68
4.2	Background	71
4.2.1	Principles of Polarimetric SAR	71
4.2.2	Polarimetric SAR Data Notation	72
4.2.3	Related Work	73
4.3	Proposed Despeckling Approach	75
4.3.1	Approach Details	76
4.3.2	Network Architecture	77
4.3.3	Approach Summary	78
4.4	Experimental Procedure	78
4.4.1	Training	78
4.4.2	Conventional Despeckling Methods	79

4.4.3	MuLoG Despeckling Framework	80
4.4.4	PolSAR Despeckling Performance Assessment	81
4.5	Results	86
4.5.1	Simulated Data Experiments	86
4.5.2	Measured Data Experiments	91
4.6	Conclusion	97
4.6.1	Summary	97
4.6.2	Future Work	98
5.	Maximum Likelihood Estimation in Mixed Integer Linear Models	100
5.1	Introduction	100
5.2	Prior Art	102
5.3	Background	104
5.3.1	Lattices	104
5.3.2	Hermite Normal Form	105
5.3.3	Chinese Remainder Theorems	105
5.3.4	Negative Log-Likelihood Function	108
5.4	ML Estimator for Restricted Encoding Matrices	109
5.4.1	Preliminaries	109
5.4.2	Estimating the Integer Remainders	112
5.4.3	Estimating the Real-Valued Common Remainder	113
5.4.4	Constructing the ML Solution	116
5.4.5	Computational Complexity	117
5.4.6	Comment on the $n = 1$ Case	118
5.5	ML Estimator for General Encoding Matrices	118
5.6	ML Estimator for General Mixed Integer Linear Models	122
5.6.1	Estimation Task	122
5.6.2	Lattice Basis Construction	123
5.6.3	Lattice Basis Reduction	124
5.6.4	Babai Point	125
5.6.5	Estimation Procedure	126
5.6.6	Estimator Distribution	126
5.6.7	Direction of Arrival Estimation	129
5.7	Examples	131
5.7.1	Computation Time Comparison	132
5.7.2	Encoding Matrices in Application	133
5.7.3	Direction of Arrival Estimation With Planar Array	133
5.7.4	Magnetic Resonance Flow Imaging	135
5.8	Summary	137
5.9	Future Work	138

6.	Aliasing in Nonuniform Arrays	140
6.1	Introduction	140
6.2	Signal Model	142
6.3	Alias-free	142
6.4	Necessary and Sufficient Condition	146
6.4.1	Linear Arrays	147
6.4.2	Nonlinear Arrays	148
6.4.3	Invariances and Scaling	148
6.5	Examples	149
6.5.1	Linear Array	149
6.5.2	Planar Arrays	151
6.5.3	3D Array	152
6.6	Summary	153
6.7	Future Work	154
	Appendices	156
A.	Variational Inference Details for Chapter 3	156
B.	Proofs for Chapter 5	158
B.1	Proofs for Section 5.4	158
B.2	Correspondence Between Hyperplane Intersections and Chambers .	164
B.3	Constructing χ	167
B.3.1	Constructing χ with $n \geq 1$	167
B.3.2	Comment on the $n = 1$ Case	169
B.4	Bounds on Problem Size for Restricted Encoding Matrices	170
B.4.1	Upper Bound on N	170
B.4.2	Bound on Number of Wrapped Chambers	171
B.5	Bounds on Problem Size for General Encoding Matrices	172
B.5.1	Bound on N_{NC}	173
B.5.2	Comparison Between N and N_{NC}	174
B.5.3	Bound on the Number of Wrapped Chambers	174
B.5.4	Comparison Between K and K_{NC}	175
B.6	Proofs for Section 5.6	176
C.	Phase Noise Model	177

Bibliography	181
------------------------	-----

List of Tables

Table	Page
2.1 Recovery metrics from simulated fully polarimetric SAR data experiment.	29
2.2 Stockbridge ADTS data recovery metrics.	31
2.3 Recovery metrics from simulated dual-pol SAR data experiment. . . .	34
3.1 Error rates from the simulated data experiments using Scenarios 1–4. The error rates for the coherence estimates are computed using a threshold which minimizes the error rate.	58
3.2 Error rates from the simulated data experiments using Scenario 4 with various values for γ . The error rates for the coherence estimates are computed using a threshold which minimizes the error rate.	60
4.1 Quantitative despeckling assessment results from simulated single-look ($L = 1$) PolSAR data. The mean results over ten independent realizations of the speckle process from each of 18 ground truth covariance images of 256×256 pixels are shown with standard deviations in gray.	87
4.2 Quantitative despeckling assessment results from simulated four-look ($L = 4$) PolSAR data. The mean results over ten independent realizations of the speckle process from each of 18 ground truth covariance images of 256×256 pixels are shown with standard deviations in gray.	89
4.3 Quantitative polarimetric information preservation assessment results from simulated PolSAR data. The mean results over ten independent realizations of the speckle process from each of 18 ground truth covariance images of 256×256 pixels are shown.	90
4.4 Quantitative despeckling assessment results from the measured PolSAR data ($L = 1$).	95

4.5	Quantitative despeckling assessment results from the measured Pol-SAR data ($L = 4$).	95
5.1	Algorithms and problem sizes for examples of the linear model (5.1) appearing in application.	134

List of Figures

Figure		Page
2.1	Factor graph depicting the conditional probabilities comprising the assumed signal model. The text at the right describes the physical characteristics captured in the model, depicted as factor nodes in the graph.	15
2.2	Simulated SAR image estimation results: (a,f) the span of the ground truth covariance, (b,g) the single-look empirical covariance, (c,h) the boxcar-filtered estimate, (d,i) the MuLoG-filtered estimate, and (e,j) the MAP estimate with the top row involving dual-pol data and the bottom fully-polarimetric data.	28
2.3	Measured fully polarimetric Ka-band SAR data estimation results from two image chips: (a,f) the span of the full-data empirical covariance, (b,g) the single-look empirical covariance, (c,h) the boxcar-filtered estimate, (d,i) the MuLoG-filtered estimate, and (e,j) the MAP estimate. High-quality color graphics of reconstruction results have been made available as supplementary material.	32
3.1	The neighborhood structure of a 4-connected gamma MRF showing the vertical and horizontal neighbors for r_i . The auxiliary variables are shown in red.	42
3.2	Factor graph representation of the proposed change detection model. The squares denote factor nodes and the circles variable nodes.	45
3.3	Values used to simulate Scenarios 1–4 (columns 1–4, respectively). (a–d) Normalized mean backscatter power (in dB). (e–h) Interferometric phase (in radians).	53
3.4	Magnitude images for Scenarios 1–4 (columns 1–4, respectively). (a–d) Time-1 magnitude images. (e–h) Time-2 magnitude images.	54

3.5	Comparison of change detection results on simulated data Scenarios 1–4 (rows 1–4, respectively) with various coherence values γ . The first column has the true change maps with change pixels in black, and the remaining columns show the change detection results from the compared methods. The last column shows the results of the proposed inference approach with black indicating a high probability of change.	55
3.6	Receiver operating characteristic (ROC) curves for simulated data experiments. (a) Scenario 1 ($\gamma = 0.55$). (b) Scenario 2 ($\gamma = 0.65$). (c) Scenario 3 ($\gamma = 0.75$). (d) Scenario 4 ($\gamma = 0.85$).	57
3.7	Comparison of change detection results from Scenario 4 with γ values of 0.55, 0.65, 0.75, and 0.85 depicted in rows 1–4, respectively. The other simulation parameters are kept constant. The first column has the true change maps with change pixels in black, and the remaining columns show the change detection results from the compared methods. The last column shows the results of the proposed inference approach with black indicating a high probability of change.	59
3.8	Error rates of a sample coherence estimator plotted against window size. The error rates are computed using thresholds which minimize the error rates at each window size. The x-axis corresponds to the pixels per side of the square window used to compute the sample coherence. Also shown are the error rate from the NL-InSAR coherence estimates and the MAP change estimates from the proposed VEM approach. The underlying change map consists of randomly-placed square change regions, and the simulation assumes a coherence of $\gamma = 0.55$	61
3.9	Optical satellite images [1] of a region observed in the Gotcha change detection challenge problem dataset [2]. (a) The observed region in the measured data experiment is shown framed by a black box. (b) The hinged grate shown here is identified by the blue arrow in (a).	62
3.10	Magnitude images from the Gotcha dataset [2]. (a) Time-1 (FP1020). (b) Time-2 (FP1024).	63
3.11	Change detection results from the X-band Gotcha dataset. (a) Sample coherence with 3×3 window. (b) Sample coherence with 5×5 window. (c) NL-InSAR coherence estimate. (d) Proposed VEM inference approach with black indicating a high probability of change.	64

3.12	Optical image of the Troodos Mountains in Cyprus with a yellow overlay depicting the area impacted by a wildfire that began on July 3, 2021 [3,4].	65
3.13	Change detection results from C-band Copernicus Sentinel SAR data [5] collected over the region of Cyprus shown in Fig. 3.12. (a) Sample coherence with 3×3 window. (b) Sample coherence with 5×5 window. (c) NL-InSAR coherence estimate. (d) Proposed VEM inference approach with black indicating a high probability of change.	65
4.1	Proposed despeckling framework.	75
4.2	Comparison of despeckling algorithms using a 256×256 pixel simulated single-look PolSAR image [6]. The results are displayed in an RGB representation based on the Pauli decomposition. (a) Original data. (b) Boxcar. (c) Extended sigma. (d) IDAN. (e) Ground truth. (f) MuLoG-BM3D. (g) MuLoG-CNN. (h) MuLoG-CNN-P. (i) MS-CNN.	88
4.3	Whitened images from the simulated single-look PolSAR data. (a) Boxcar. (b) Extended sigma. (c) IDAN. (d) MuLoG-BM3D. (e) MuLoG-CNN. (f) MuLoG-CNN-P. (g) MS-CNN.	89
4.4	Co-polarization (top row) and cross-polarization (bottom row) signatures computed from a small region of a simulated single-look PolSAR image. The ellipticity and orientation angles of the transmitting antenna are given in degrees. (a,e) Ground truth. (b,f) Extended sigma. (c,g) MuLoG-CNN. (d,h) MS-CNN.	91
4.5	Single-look PolSAR image of 512×512 pixels from the Pi-SAR sensor used in despeckling comparison. The boxes indicate regions in which performance indexes are computed: white boxes for ENL and MOR, green for TCR, yellow for EPD-ROA, and blue for the whitened image assessments. The region within the red box is depicted in Fig. 4.6. . .	92
4.6	Detailed comparison of the despeckling algorithms using a 128×128 region cropped from Fig. 4.5. (a) Original data. (b) Boxcar. (c) Extended sigma. (d) IDAN. (e) MuLoG-BM3D. (f) MuLoG-CNN. (g) MuLoG-CNN-P. (h) MS-CNN.	93

4.7	Comparison of despeckling methods using a four-look PolSAR image of 200×200 pixels. The boxes in (a) indicate regions in which performance indexes are computed: white boxes for ENL and MOR, green for TCR, yellow for EPD-ROA, and blue for the whitened image assessments. (a) Original data. (b) Boxcar. (c) Extended sigma. (d) IDAN. (e) MuLoG-BM3D. (f) MuLoG-CNN. (g) MuLoG-CNN-P. (h) MS-CNN.	94
4.8	Whitened single-look PolSAR images from the region bound by the cyan box in Fig. 4.5. (a) Boxcar. (b) Extended sigma. (c) IDAN. (d) MuLoG-BM3D. (e) MuLoG-CNN. (f) MuLoG-CNN-P. (g) MS-CNN.	96
4.9	Estimated entropy from the four-look measured data. (a) Original data. (b) Extended sigma. (c) MuLoG-CNN. (d) MS-CNN.	97
4.10	Estimated alpha angle (in degrees) from the four-look measured data. (a) Original data. (b) Extended sigma. (c) MuLoG-CNN. (d) MS-CNN.	98
5.1	$\mathcal{P}_{\mathbf{V}}$ partitioned by six shifted copies of $\mathcal{P}_{\mathbf{C}}$ for the encoding matrix in (5.20).	111
5.2	(a) The value of \mathcal{L}_r is indicated by color over $\mathcal{P}_{\mathbf{C}}$. The wrapped hyperplanes in $\tilde{\mathcal{H}}$ are shown in red. The candidate solutions for (5.31) are marked with magenta squares, with the solution $\hat{\mathbf{r}}$ at the yellow cross. $\mathcal{P}_{\mathbf{C}}$ is partitioned into wrapped chambers, labeled 1–3. (b) \mathcal{L}_r as a surface plot.	115
5.3	(a) The value of \mathcal{L}_x is indicated by color over $\mathcal{P}_{\mathbf{V}}$. The wrapped hyperplanes in $\tilde{\mathcal{H}}_{\text{NC}}$ are shown in red. The candidate solutions for (5.2) are marked with magenta squares, with the solution $\hat{\mathbf{x}}$ at the yellow cross. $\mathcal{P}_{\mathbf{V}}$ is partitioned into wrapped chambers, labeled 1–6. (b) \mathcal{L}_x as a surface plot.	121
5.4	Example distribution, $f(\hat{\mathbf{x}} \mathbf{x}, \ell)$, for $m = 3$, $n = 2$	129
5.5	Mean computation time plotted against m for Algorithm 8 (MCRT), Algorithm 9 (NC), and the sphere decoding approach of [7] (SD). . .	132
5.6	DoA estimation performance against SNR with $T = 10$ snapshots from five-element planar array [8]. RMSE results for the MLE of (5.73) and the phase unwrapping MLE are shown along with the square root of the trace of the CRB for \mathbf{x} . The results are averaged over 10^6 trials. . . .	135

5.7	Velocity RMSE on $\mathcal{P}_{\mathbf{V}} \in \mathbb{R}^3$ for (a) the estimator in [9] and (b) the MLPUE. Velocity components and RMSE are shown in cm/s.	136
6.1	(a) For a uniform circular array of six sensors with λ spacing, illustration of the lattice Λ , a basis \mathbf{V} , a fundamental parallelogram, and the disc of all DoAs, \mathcal{B}_2 , translated periodically per the lattice. In this case, the translated discs intersect, and the array is aliased. (b) The lattice and shifted \mathcal{B}_2 for a four-sensor planar array that is alias-free, despite a minimum sensor separation of 0.855λ . (c) A second four-sensor planar array in which \mathcal{B} is the set of DoAs with $ \theta \leq 60^\circ$; the array is alias-free on \mathcal{B}	152
6.2	For a bi-tetrahedral (a) array of five sensors, (b) illustration of a fundamental parallelotope and the sphere of all DoAs, \mathcal{B}_3 , translated to the center of the parallelotope. Spheres translated to lattice points are merely tangent; the array is alias-free other than three pairs of ambiguous DoAs.	153
B.1	$\mathcal{P}_{\mathcal{C}}$ is partitioned into wrapped chambers, labeled 1–3. The points at which two wrapped hyperplanes in $\tilde{\mathcal{H}}$ (shown in red) intersect are marked with blue circles (denoted $\tilde{\mathcal{U}}$). The correspondence of each such point to a wrapped chamber (as described in Appendix B.2) is shown by an arrow.	168
B.2	The number of wrapped chambers for Algorithm 8 is plotted in red against m ; the red dash-dotted line is the upper bound on K from (B.28). The number of wrapped chambers for Algorithm (9) is in blue; the blue dash-dotted line is the upper bound on K_{NC} from (B.36). . .	175

Chapter 1: Introduction

This dissertation addresses two topics: inference from synthetic aperture radar (SAR) data and the use of lattice theory in array signal processing.

1.1 Inference Methods for Synthetic Aperture Radar

In the first portion of this dissertation, we propose a series of inference methods for multichannel SAR data. SAR is a microwave sensing technology capable of producing high-resolution images of terrain and targets. The airborne or spaceborne radar system traverses a flight path, transmitting and receiving a series of high-bandwidth radio frequency (RF) pulses, most often linear frequency modulated (LFM) chirps, towards a scene of interest. In spotlight mode SAR, the antenna is steered to illuminate a fixed patch of ground continuously. The backscattered pulses are mixed with a reference signal and low-pass filtered, resulting in a set of samples of the Fourier transform of the complex scene reflectivity confined to an annular region in the spatial frequency domain. The recorded pulses or phase history can then be processed into an image.

In this dissertation, we consider multichannel SAR imaging modalities, such as polarimetric SAR (PolSAR) or interferometric SAR (InSAR), in which data are collected via multiple channels. These data are considerably more informative than

single-channel collections; the inter-channel relationships are informative of physical scattering behavior [10, 11]. These multichannel radar systems have proven critical for monitoring changes on the earth’s surface, and numerous techniques employ these data for terrain classification or surface parameter estimation [12].

We use a tomographic formulation of SAR and a plane-wave assumption, permitting the expression of the sampled radar returns as

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{u}, \tag{1.1}$$

in which \mathbf{A} is a SAR imaging operator, the lexicographically-ordered 2D scene reflectivity samples are denoted by \mathbf{x} , and \mathbf{u} represents thermal noise.

Speckle in SAR imaging is a result of the coherent interaction of reflections from objects within a resolution cell. When illuminating a homogeneous medium that is rough on the scale of the radar’s central frequency, the reflectivity is modeled as a random variable with a complex normal distribution [13–15].

In Chapter 2 we give a method for multichannel coherent image recovery based on a hierarchical statistical model. This approach accommodates a reduced synthetic aperture but also suppresses speckle and thermal noise. The model presented here uses latent terrain class variables with an MRF prior to exploit both structure among coherent imaging channels and spatial correlation for improved image recovery.

In Chapter 3, we consider the problem of coherent change detection from repeat-pass SAR imagery. A Bayesian approach is formulated as an alternative to conventional window-based change detection statistics that entail losses to spatial resolution. The proposed approach assigns prior distributions to the unobserved model variables to exploit spatial structure both in the geophysical scattering qualities of the scene and among the scene disturbances that take place between the passes. Variational

expectation maximization is used to efficiently approximate the posterior distribution of the latent variables and the prior model hyperparameters. Experiments on simulated and measured interferometric SAR data pairs indicate the effectiveness of the proposed change detection method and highlight improvements over traditional window-based approaches.

In Chapter 4, we consider the problem of speckle reduction of multichannel SAR images. It is often a prerequisite to suppress speckle in some way before SAR data are considered suitable for analysis [16]. This is frequently realized by the application of a speckle filter, classically one which acts on the image locally [17]. We describe a polarimetric SAR despeckling approach that uses convolutional neural networks (CNNs). The approach operates in the matrix logarithm domain and uses a residual learning architecture that has been successful in natural image denoising. We have implemented this approach with a training set of spatially multilooked PolSAR images that are corrupted by a simulated speckle process. Further, the proposed approach provides improved despeckling over conventional filtering methods and operates with a comparably fast running time afforded by the use of CNNs.

1.2 Array Processing With Lattices

This dissertation next addresses the applications of lattices (regular arrangements of points in Euclidean space) to sensor array processing. In Chapter 5, we consider the problem of maximum likelihood (ML) parameter estimation in mixed integer linear models. That is, we consider a linear model with a measurement vector \mathbf{y} that is related to a real-valued unknown \mathbf{x} and an integer-valued unknown \mathbf{k} as

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{M}\mathbf{k} + \mathbf{u}, \tag{1.2}$$

where \mathbf{A} is an encoding matrix with full column rank and \mathbf{M} is an invertible matrix. This problem is relevant to many applications, including ranging [18–21], frequency estimation [22–24], interferometry [25], MIMO decoding [26–28], direction of arrival estimation [29], multichannel modulo sampling [30], and magnetic resonance imaging [9, 31].

In [32], a case of (1.2) is considered, in which there is a single real-valued variable and the encoding matrix meets a certain coprimeness condition. There, a polynomial-time ML solution method is given that leverages the classical Chinese remainder theorem. In Chapter 5, we provide two ML solution methods that extend polynomial-time computation to a wider range of problems. These extensions of the prior art account for both multivariate real-valued unknowns and general encoding matrices with full column rank. We then provide a solution method based on sphere decoding that applies generally to (1.2).

In Chapter 6, we focus on the use of lattice theory in the analysis and design of nonuniform sensor arrays in one, two, and three dimensions. Nonuniform array geometries are often used to extend the array aperture or to reduce mutual coupling among sensors. The work in this chapter is motivated by the use of nonuniform arrays in the direction of arrival (DoA) estimation problem. We give a necessary and sufficient condition for an array of isotropic sensor elements to be unambiguous for any specified set of DoAs. The set of unambiguous DoAs for a given array configuration is shown to be a parallelotope, a property that admits simple geometric interpretation and facilitates sensor array design. The result is an easily tested necessary and sufficient condition for arbitrary sets of DoAs, an alternative to a widely used sufficient

condition [33, Thm. 8]. We include additional properties in this chapter to aid in the design of nonuniform arrays that meet aliasing requirements.

Chapter 2: Speckle Suppression in Multichannel SAR Imaging

2.1 Introduction

An issue endemic to coherent imaging is the presence of speckle, a phenomenon which manifests as a granular quality in images and may obscure features of interest.¹ Speckle is the result of the coherent interaction of multiple reflections from within a resolution cell. It is common to account for the presence of speckle by modeling the scene reflectivity stochastically, canonically as a zero-mean Gaussian-distributed random variable [13, 14].

Also of concern for image formation and interpretation are gaps in the SAR coherent phase history, which may be caused by radio frequency jamming or interruptions due to hardware multiplexing; spectral gaps result in ringing artifacts which hinder image interpretation. The objective of this work is to estimate the multichannel scattering properties of a scene afflicted by speckle when the collection suffers from gaps in spatial frequency observations. We hypothesize that directly addressing speckle and ringing artifacts in a multichannel image formation procedure can produce results

¹This chapter is largely drawn from [34], coauthored with Lee C. Potter (© 2020 IEEE).

superior to the existing data flow of channel-by-channel imaging followed by speckle-reducing post-processing. To this end, we adopt a maximum a posteriori (MAP) estimation approach to recover the per-pixel multichannel SAR covariance with dual goals of speckle reduction and preservation of spatial resolution. Rather than placing a prior directly on the covariance, we employ a latent terrain class random vector on which a Markov random field (MRF) prior is placed. To accommodate texture within a terrain class, the per-pixel covariance within a class is assumed to be distributed as a complex inverse Wishart (CIW) matrix. To alleviate computational complexity, we use an expectation-maximization (EM) surrogate function for the MAP objective function and a probing method for estimating the block-diagonal of a matrix inverse.

2.1.1 Motivating Applications

The proposed approach is motivated by SAR imaging modalities in which data are collected via multiple channels, such as PolSAR and InSAR. Multichannel SAR data have been shown to be considerably more informative than single-channel collections because the relationships between channels capture information corresponding to physical scattering behavior [10, 11]. The information that these multichannel systems gather has proven critical for monitoring changes on the earth’s surface, and numerous techniques employ these data for terrain classification and surface parameter estimation [12].

A common multichannel SAR modality is PolSAR, in which diversity of the polarizations of the transmitted and received signals provides more information about the environment than would a single-channel sensor. A pair of orthogonal polarizations are used (e.g., horizontal and vertical polarizations) to emit and receive energy,

the permutations of which yield a reflectivity vector of four complex values for each pixel. The experimental results presented here use simulated and measured PolSAR data, although the proposed approach applies to other SAR modalities and, indeed, to other multichannel coherent imaging tasks.

2.1.2 Ameliorating Speckle by Post-Processing

The existing data flow is to first form single-channel SAR images, then post-process to reduce speckle before image interpretation [16]. The speckle suppression step is frequently realized by the application of a speckle filter, classically one which acts on the image locally [17]. The simplest such despeckling technique is a boxcar filter, which suppresses speckle effectively in homogeneous areas but degrades the resolution and the fine structural details of the image [16]. Further local filtering methods include the Lee [35] and Refined Lee [36] filters, which have, among other classical PolSAR speckle filtering techniques, been reviewed and analyzed by Foucher and López-Martínez [17].

In recent years, there has been a trend toward non-local patch-based speckle reduction based on the non-local means (NLM) filter [37]. Unlike the local filtering methods, the NLM ethos is based on the nonlocal averaging of pixels across the image. Applications of this principle to PolSAR despeckling have demonstrated robust performance [38–40].

There are also many variational PolSAR speckle reduction methods, in which the underlying idea is to suppress speckle by optimizing an energy functional consisting of a regularization term and a data fidelity term. Some such methods reduce complexity

by operating only on the diagonal elements of the covariance [41], ignoring the statistical correlations between channels. In [42], a scheme for multichannel SAR speckle reduction is proposed which maps multichannel SAR covariance images to real-valued images using a homomorphic transform, allowing for the use of general-purpose denoising algorithms. In [43], a MAP PolSAR despeckling approach is proposed which uses a matricial total variation (TV) regularization term to enforce smoothness, and this approach is modified in [44] with a nonlocal TV-based regularization term. A review of the recent developments in PolSAR speckle filtering can be found in [16].

2.1.3 Regularized Imaging

The speckle present in coherent imaging is central to the problem formulation in this work. In contrast, many regularized SAR imaging approaches eschew a speckle model and opt to estimate the scene reflectivity, rather than its variance. To combat artifacts due to gaps in the SAR phase history data, many of these techniques enforce sparsity on the magnitude of the reflectivity [45, 46]. For a review of this literature, see [47–49]. Another family of sparsity-based approaches is based on sparse Bayesian learning (SBL) [50–52], in which the reflectivity is modeled as a zero-mean circularly-symmetric Gaussian, and an independent Gamma hyperprior is placed on the variance of each reflectivity element [53].

The proposed approach can also be considered in the context of other coherent image recovery methods that estimate the second moment of the image. In [54], an image recovery framework is proposed for single-channel synthetic aperture LADAR; the approach adopts a statistical model for fully-developed speckle and uses an iterative EM optimization scheme that treats reflectivity as a nuisance parameter. The

variances of the reflectivity elements are found via MAP estimation with a generalized Gaussian MRF prior. In [55], this approach is broadened into a plug-and-play framework [56, 57] which allows for regularization by way of off-the-shelf Gaussian denoising algorithms. In both [54] and [55], the forward model operator is a scaled unitary matrix, a limitation which greatly reduces computational complexity.

Rather than treating reflectivity as a nuisance parameter in the MAP variance estimation as in [54], reflectivity and its variance are jointly estimated in a single-channel SAR image recovery approach in [58]. This approach uses a classical stochastic speckle model and a TV regularization term in a joint recovery approach which also accounts for the deleterious effects of arbitrary gaps in phase history.

The speckle-aware imaging approach proposed here draws motivation from [54, 55] where the variance of the scene reflectivity is estimated in the fully-developed speckle regime. Like [54, 55], our approach considers an image recovery task that jointly forms images and combats speckle. In contrast, the literature for speckle reduction customarily considers despeckling as an image post-processing problem [38, 42, 43]. Three novel components comprise the contribution of the proposed approach. First, we extend the Bayesian modeling and estimation in [54, 55] to the multi-variate case; rather than estimate a variance at each pixel, we estimate a positive semi-definite multichannel covariance matrix. In this way, informative relationships among the channels are exploited to combat both speckle and image side-lobes from data dropouts. Second, we present a novel algorithmic component that overcomes an intractable matrix inversion encountered when extending [54, 55] to a more general class of linear measurement models; in particular, a matrix probing scheme in the new framework allows for the correlation in the random speckle process induced by the

gaps in the phase history data. Finally, we incorporate into the proposed framework a terrain class-based MRF prior which models spatial dependencies in a latent terrain class label.

2.2 Imaging Model

2.2.1 Spotlight-Mode SAR Imaging

In spotlight-mode SAR, high-bandwidth pulses, most often linear frequency modulated (LFM) chirps, are transmitted from the radar system towards the scene of interest, and the reflected pulses recorded. The received signal is mixed with a reference sent through a low-pass filter, resulting in a set of what may be considered slices from the 2-D Fourier transform of the complex scene reflectivity [59].

The single-channel SAR data collection process can be approximated by a linear mapping [59]

$$\mathbf{y} = \mathbf{A}_m \mathbf{x} + \mathbf{w}, \quad (2.1)$$

where $\mathbf{x} \in \mathbb{C}^N$ is the scene reflectivity, $\mathbf{y} \in \mathbb{C}^M$ noisy linear measurements, and $\mathbf{w} \in \mathbb{C}^M$ complex zero-mean white Gaussian noise with covariance $\sigma_w^2 \mathbf{I}_M$; each corresponds to an image ordered lexicographically as a vector, with \mathbf{I}_M denoting an $M \times M$ identity matrix. Note that the vector \mathbf{w} in (2.1) models additive thermal noise present at the receiver and is not related to the speckle phenomenon. Assuming SAR to be a linear shift-invariant imaging system, the forward operator \mathbf{A}_m is characterized by the spatial frequency support of the system [60].

2.2.2 Multichannel SAR

For a SAR system with D channels, the reflectivity vectors at each pixel are

$$\mathbf{x}_i = [x_i^{(1)}, x_i^{(2)}, \dots, x_i^{(D)}]^\top, \quad i = 1, 2, \dots, N, \quad (2.2)$$

which compose the scene reflectivity vector $\mathbf{x} \in \mathbb{C}^{DN}$ as

$$\mathbf{x} = [\mathbf{x}_1^\top, \mathbf{x}_2^\top, \dots, \mathbf{x}_N^\top]^\top \quad (2.3)$$

where $(\cdot)^\top$ denotes transposition. The multichannel measurement and thermal noise vectors \mathbf{y} and \mathbf{w} are likewise formed by stacking their constituents.

A multichannel SAR linear system model relates the scene reflectivity to the observations as

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{w}, \quad (2.4)$$

in which the multichannel SAR forward operator $\mathbf{A} \in \mathbb{C}^{DM \times DN}$ is

$$\mathbf{A} = \mathbf{A}_m \otimes \mathbf{I}_D, \quad (2.5)$$

where \otimes is the Kronecker product operator.

2.2.3 Modeling Speckle in Multichannel SAR

Speckle in SAR imaging is a result of the coherent interaction of reflections from objects within a resolution cell. When illuminating a homogeneous medium that is rough on the scale of the radar's central frequency, the classical reflectivity model is [13, 14]:

$$p(\mathbf{x}_i \mid \boldsymbol{\Sigma}_i) = \mathcal{CN}(\mathbf{x}_i; \mathbf{0}, \boldsymbol{\Sigma}_i), \quad i = 1, 2, \dots, N \quad (2.6)$$

where $\mathcal{CN}(\mathbf{x}_i; \mathbf{0}, \cdot)$ denotes a complex zero-mean circularly-symmetric Gaussian distribution [15] defined by:

$$\mathcal{CN}(\mathbf{x}_i; \mathbf{0}, \mathbf{\Sigma}_i) = \frac{1}{\pi^D |\mathbf{\Sigma}_i|} \exp \{ -\mathbf{x}_i^H \mathbf{\Sigma}_i^{-1} \mathbf{x}_i \}, \quad (2.7)$$

in which $|\cdot|$ denotes a determinant. Note that this stochastic modeling of \mathbf{x}_i in (2.6) is to account for the speckle phenomenon in coherent imaging. Under this model, the scattering behavior within a pixel is characterized by the covariance $\mathbf{\Sigma}_i \in \mathbb{H}_+^D$, where \mathbb{H}_+^D is the set of $D \times D$ Hermitian positive-definite matrices. This is known as the fully-developed speckle model, and most PolSAR filtering techniques are derived from its assumption.

2.2.4 Markov Random Field Prior

The proposed model stipulates that each pixel belongs to one of ℓ terrain classes (e.g., grass or scrub), a construct similar to the non-homogeneous clutter model detailed in [61]. The terrain class labels are unknown, but the statistics of each class are assumed to be known *a priori*. The latent random vector $\mathbf{z} \in \mathcal{Z} = \{1, 2, \dots, \ell\}^N$ indicates the terrain class to which each pixel belongs. Given a pixel's terrain class, its covariance is conditionally distributed as a complex inverse Wishart (CIW) distribution [62], defined for $\mathbf{\Sigma} \in \mathbb{H}_+^D$ by

$$\text{CIW}(\mathbf{\Sigma}_i; \mathbf{G}, v) = \frac{|\mathbf{\Sigma}_i|^{-(v+D)} \exp \{ \text{tr}(-\mathbf{G} \mathbf{\Sigma}_i^{-1}) \} |\mathbf{G}|^v}{\pi^{D(D-1)/2} \prod_{i=1}^D \Gamma(v - i + 1)} \quad (2.8)$$

where $\text{tr}(\cdot)$ is the trace operator. The conditional distribution of $\mathbf{\Sigma}_i$ is given by

$$p(\mathbf{\Sigma}_i | z_i) = \begin{cases} \text{CIW}(\mathbf{\Sigma}_i; \mathbf{G}_1, v_1), & z_i = 1 \\ \vdots & \vdots \\ \text{CIW}(\mathbf{\Sigma}_i; \mathbf{G}_\ell, v_\ell), & z_i = \ell \end{cases}, \quad (2.9)$$

for $i = 1, 2, \dots, N$, where (\mathbf{G}_a, v_a) are the distribution parameters corresponding to class $a \in \{1, 2, \dots, \ell\}$. Each z_i is assigned a marginal prior distribution of $p(z_i = j) = \alpha_j$ for $i = 1, 2, \dots, N$, and $j = 1, 2, \dots, \ell$ where $\sum_j \alpha_j = 1$. A Markov random field (MRF) prior encodes spatial correlation among neighboring z_i . Specifically, we have, for $\gamma_n(\ell - 1) + \gamma_e = 1$,

$$p(z_i | z_j) = \begin{cases} \gamma_n, & z_i \neq z_j \\ \gamma_e, & z_i = z_j \end{cases}$$

as the pair-wise marginal probability of z_i conditional on any of its four neighbors $z_j \in \delta_i$ [63]. The total distribution of \mathbf{z} is given by

$$p(\mathbf{z}) = \frac{1}{Z} \prod_{i=1}^N p(z_i) \prod_{j \in \delta_i} p(z_i | z_j) \quad (2.10)$$

where Z is a normalizing term. A factor graph representing this probability model is shown in Fig. 2.1, where, for clarity, the MRF along the bottom of the graph is visualized as a chain. Square nodes are used to represent factors of the distribution and circular nodes denote unknown random variables.

2.3 MAP Estimation

The recovery of the per-pixel covariance from the observations is posed as a MAP estimation problem; i.e., we find

$$\hat{\Sigma} \in \arg \max_{\Sigma \in \Omega} p(\Sigma | \mathbf{y}) \quad (2.11)$$

where

$$\Sigma = \begin{bmatrix} \Sigma_1 & 0 & \cdots & 0 \\ 0 & \Sigma_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \Sigma_N \end{bmatrix}, \quad (2.12)$$

with a set of feasible values defined by

$$\Omega = \{\Sigma \mid \Sigma = \text{diag}_D(\Sigma) \text{ and } \Sigma \in \mathbb{H}_+^{DN}\} \quad (2.13)$$

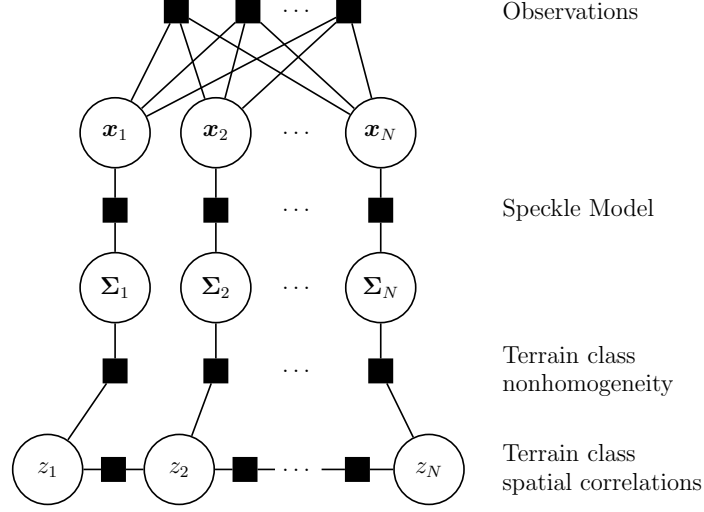


Figure 2.1: Factor graph depicting the conditional probabilities comprising the assumed signal model. The text at the right describes the physical characteristics captured in the model, depicted as factor nodes in the graph.

in which $\text{diag}_D(\mathbf{\Sigma})$ denotes a block-diagonal matrix of $D \times D$ blocks whose block-diagonal entries are the same as those of $\mathbf{\Sigma}$. From Bayes' theorem, the MAP problem is equivalent to

$$\hat{\mathbf{\Sigma}} \in \arg \min_{\mathbf{\Sigma} \in \Omega} \{ -\log p(\mathbf{y} \mid \mathbf{\Sigma}) - \log p(\mathbf{\Sigma}) \}, \quad (2.14)$$

where [53]

$$p(\mathbf{y} \mid \mathbf{\Sigma}) = \mathcal{CN}(\mathbf{0}, \mathbf{T}), \quad (2.15)$$

with

$$\mathbf{T} = \mathbf{A}\mathbf{\Sigma}\mathbf{A}^H + \sigma_w^2 \mathbf{I}_{DM}. \quad (2.16)$$

With (2.6), the MAP estimation optimization problem is

$$\hat{\mathbf{\Sigma}} \in \arg \min_{\mathbf{\Sigma} \in \Omega} \{ \log |\mathbf{T}| + \mathbf{y}^H \mathbf{T}^{-1} \mathbf{y} - \log p(\mathbf{\Sigma}) \}. \quad (2.17)$$

The objective function in (2.17) is non-convex and generally does not admit a solution in closed form [53]. As such, we use a surrogate for the objective function in (2.17) as is proposed in [54].

2.3.1 MAP Objective Surrogate Function

Treating reflectivity \mathbf{x} and terrain class \mathbf{z} as hidden variables, we form an EM surrogate for the MAP cost in (2.17) parameterized by the previous estimate Σ' and given by [54, 64]

$$\begin{aligned} Q(\Sigma; \Sigma') &= -\mathbb{E}_{\mathbf{x}, \mathbf{z} | \mathbf{y}, \Sigma'} [\log p(\mathbf{x}, \Sigma, \mathbf{z} | \mathbf{y})] \\ &= -\mathbb{E}_{\mathbf{x}, \mathbf{z} | \mathbf{y}, \Sigma'} [\log p(\mathbf{y} | \mathbf{x}) + \log p(\mathbf{x}, \Sigma, \mathbf{z})] + \text{const.}, \end{aligned}$$

where we have used the fact that $p(\mathbf{y} | \mathbf{x}, \Sigma, \mathbf{z}) = p(\mathbf{y} | \mathbf{x})$, and const. is used to collect terms which do not depend on Σ . We then have

$$\begin{aligned} Q(\Sigma; \Sigma') &= -\mathbb{E}_{\mathbf{x}, \mathbf{z} | \mathbf{y}, \Sigma'} [\log p(\mathbf{y} | \mathbf{x}) + \log p(\mathbf{x} | \Sigma) \\ &\quad + \log p(\Sigma | \mathbf{z}) + \log p(\mathbf{z})] + \text{const.} \\ &= -\mathbb{E}_{\mathbf{x}, \mathbf{z} | \mathbf{y}, \Sigma'} [\log p(\mathbf{x} | \Sigma) + \log p(\Sigma | \mathbf{z})] + \text{const.} \\ &= -\mathbb{E}_{\mathbf{x} | \mathbf{y}, \Sigma'} [\log p(\mathbf{x} | \Sigma)] - \mathbb{E}_{\mathbf{z} | \Sigma'} [\log p(\Sigma | \mathbf{z})] + \text{const.}, \end{aligned} \quad (2.18)$$

where we have used that $p(\mathbf{z} | \Sigma', \mathbf{y}, \mathbf{x}) = p(\mathbf{z} | \Sigma')$. The conditional posterior density $p(\mathbf{x} | \mathbf{y}, \Sigma')$ is a complex Gaussian-distributed random vector with mean and covariance [65]

$$\boldsymbol{\mu}_{\Sigma'} = \sigma_w^{-2} \mathbf{C}_{\Sigma'} \mathbf{A}^H \mathbf{y} \quad (2.19a)$$

$$\mathbf{C}_{\Sigma'} = [\sigma_w^{-2} \mathbf{A}^H \mathbf{A} + (\Sigma')^{-1}]^{-1}, \quad (2.19b)$$

respectively. With this, (2.18) can be simplified to

$$Q(\Sigma; \Sigma') = \log |\Sigma| + \text{tr}(\mathbf{S}_{\Sigma'} \Sigma^{-1}) - \mathbb{E}_{\mathbf{z}|\Sigma'} [\log p(\Sigma | \mathbf{z})] + \text{const.}, \quad (2.20)$$

where $\mathbf{S}_{\Sigma'} = \mathbf{C}_{\Sigma'} + \boldsymbol{\mu}_{\Sigma'} \boldsymbol{\mu}_{\Sigma'}^H$. Evaluation of $p(\mathbf{z} | \Sigma')$ over all possible realizations of \mathbf{z} in (2.20) is prohibitively costly. As such, we will approximate $p(\mathbf{z} | \Sigma')$ as a product of its marginal probabilities; that is, assume $p(\mathbf{z} | \Sigma') \approx \prod_{i=1}^N p(z_i | \Sigma')$. In this case, the surrogate in (2.20) is approximated as

$$Q(\Sigma; \Sigma') \approx \log |\Sigma| + \text{tr}(\mathbf{S}_{\Sigma'} \Sigma^{-1}) - \sum_{i=1}^N \sum_{k=1}^{\ell} p(z_i = k | \Sigma') \log p(\Sigma_i | z_i) + \text{const.} \quad (2.21)$$

The marginal distributions $p(z_i | \Sigma')$, for $i = 1, 2, \dots, N$, are approximated using the sum-product algorithm [65, 66]. The function in (2.21) decouples by pixel, and the i th block of its minimizer is

$$\hat{\Sigma}_i \triangleq \arg \min_{\Sigma_i \in \mathbb{H}_+^D} \log |\Sigma_i| + \text{tr}(\Sigma_i^{-1} (\mathbf{S}_{\Sigma'})_i) - \sum_{k=1}^{\ell} p(z_i = k | \Sigma') \log p(\Sigma_i | z_i) \quad (2.22)$$

$$= \frac{(\mathbf{S}_{\Sigma'})_i + \sum_{k=1}^{\ell} p(z_i = k | \Sigma') \mathbf{G}_k}{1 + \sum_{k=1}^{\ell} p(z_i = k | \Sigma') (v_k + D)}. \quad (2.23)$$

In (2.23), we have used the fact that the objective in (2.22) has a unique minimizer over \mathbb{H}_+^D [67].

2.3.2 Evaluating the Conditional Posterior Covariance

To evaluate (2.21), it is necessary to compute the conditional posterior covariance matrix $\mathbf{C}_{\Sigma'}$. Its direct, explicit evaluation is both memory-intensive and computationally prohibitive, involving as it does the inversion of a $DN \times DN$ matrix. For scaled unitary \mathbf{A}_m , $\mathbf{C}_{\Sigma'}$ is block-diagonal, making its inversion tractable for practical image sizes. This simplifying assumption, which is made in [54, 55], is met when

imaging sidelobes are not present. To accommodate more general imaging scenarios, first note that we need only compute the block-diagonal of $\mathbf{C}_{\Sigma'}$ because $\mathbf{C}_{\Sigma'}$ appears in (2.21) only via the trace of its product with the block-diagonal Σ^{-1} .

With this in mind, we consider computing the D-block-diagonal of the inverse of a Hermitian matrix, adapting an approach from Tang and Saad [68]. Their method exploits sparsity to ease the computational burden of computing the diagonal of the inverse of a symmetric and positive-definite matrix. In describing its adaptation, we adopt generic notation: let $\mathbf{B} \in \mathbb{H}_+^{DN}$ denote a known positive-definite Hermitian matrix from which we seek to approximate $\text{diag}_D(\mathbf{B}^{-1})$.

Consider \mathbf{B} to be composed of blocks as

$$\mathbf{B} = \begin{bmatrix} \mathbf{B}_{11} & \mathbf{B}_{12} & \cdots & \mathbf{B}_{1N} \\ \mathbf{B}_{21} & \mathbf{B}_{22} & \cdots & \mathbf{B}_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{B}_{N1} & \mathbf{B}_{N2} & \cdots & \mathbf{B}_{NN} \end{bmatrix} \quad (2.24)$$

where $\mathbf{B}_{ik} \in \mathbb{C}^{D \times D}$. This method relies on \mathbf{B}^{-1} having relatively few “non-trivial” blocks. This is to say that $\mathcal{S}_\epsilon(\mathbf{B}^{-1})$ is sparse, where the sparsified block-support operator $\mathcal{S}_\epsilon(\cdot) : \mathbb{C}^{DN \times DN} \rightarrow \{0, 1\}^{N \times N}$ is defined by

$$(\mathcal{S}_\epsilon(\mathbf{B}^{-1}))_{ik} = \begin{cases} 1, & \text{if } \|(\mathbf{B}^{-1})_{ik}\|_F^2 > \epsilon \\ 0, & \text{otherwise.} \end{cases} \quad (2.25)$$

for a small constant $\epsilon > 0$. The sparsified \mathbf{B}^{-1} , for which entries outside the sparsified block-support are set to zero, is denoted by $(\mathbf{B}^{-1})_\epsilon$.

To proceed, note the following relation:

$$\text{diag}_D(\mathbf{B}^{-1}) = \text{diag}_D(\mathbf{B}^{-1} \mathbf{V} \mathbf{V}^\top) \text{diag}_D^{-1}(\mathbf{V} \mathbf{V}^\top), \quad (2.26)$$

which holds trivially for orthogonal \mathbf{V} . In this estimation scheme, we exploit the approximately sparse structure of \mathbf{B}^{-1} to construct a tall, thin $\mathbf{V} \in \{0, 1\}^{DN \times DP}$ for

which (2.26) holds with $P \ll N$. In [68], \mathbf{V} is called a “probing” matrix. Here we decompose the probing matrix \mathbf{V} as

$$\mathbf{V} = [\mathbf{V}_1^\top \quad \mathbf{V}_2^\top \quad \cdots \quad \mathbf{V}_N^\top]^\top \quad (2.27)$$

where

$$\mathbf{V}_i = [\mathbf{V}_{i1} \quad \mathbf{V}_{i2} \quad \cdots \quad \mathbf{V}_{iP}], \quad i = 1, 2, \dots, N \quad (2.28)$$

and $\mathbf{V}_{ij} \in \{0, 1\}^{D \times D}$. The construction of the probing matrix rests on the following proposition, stated in [68] for $D = 1$:

Proposition 1. *Given Hermitian positive-definite $\mathbf{B} \in \mathbb{H}_+^{DN}$, let $\mathbf{V} \in \{0, 1\}^{DN \times DP}$ be full-rank with $P \leq N$, then*

$$\text{diag}_D(\mathbf{B}^{-1}) = \text{diag}_D(\mathbf{B}^{-1} \mathbf{V} \mathbf{V}^\top) \text{diag}_D(\mathbf{V} \mathbf{V}^\top) \quad (2.29)$$

holds if $\mathbf{V}_i \mathbf{V}_k^\top = \mathbf{0}$ for all (i, k) for which $i \neq k$ and $(\mathbf{B}^{-1})_{ik} \neq \mathbf{0}$.

A probing matrix which satisfies Proposition 1 can be formed by first constructing a graph in which each node corresponds to a pixel in the image, and the corresponding adjacency matrix is the sparse support of \mathbf{B}^{-1} . We are able to form a suitable \mathbf{V} by coloring the nodes of this graph. Graph coloring is known to be an NP-hard problem, and, as is common practice, we rely on a heuristic to find a graph coloring with few colors. A well-known greedy graph coloring algorithm is summarized in Algorithm 1 [68], and further details on graph coloring algorithms can be found in [69].

Algorithm 1 Greedy Graph Coloring Algorithm

Input: Graph corresponding to $N \times N$ matrix

Output: Colors corresponding to nodes of the graph

```
1: for  $i = 1$  to  $N$  do  
2:   Set  $\text{Color}(i) = 0$   
3: end for  
4: for  $i = 1$  to  $N$  do  
5:   Set  $\text{Color}(i) = \min\{k > 0 \mid k \neq \text{Color}(l) \ \forall l \in \text{Adjacent}(i)\}$   
6: end for  
7: return Color
```

As is done for the estimation of the diagonal of a matrix inverse in [68], a probing matrix can be formed from the colored nodes. Specifically, \mathbf{V} is constructed as

$$\mathbf{V}_{ik} = \begin{cases} \mathbf{I}_D, & \text{Color}(i) = k \\ \mathbf{0}, & \text{otherwise} \end{cases}. \quad (2.30)$$

In general the support of \mathbf{B}^{-1} is not known, and it must be estimated from the known \mathbf{B} . The support of the sparsified inverse can be estimated as

$$\mathcal{S}_\epsilon(\mathbf{B}^{-1}) \approx \mathcal{S}_\epsilon(\mathbf{B}^q) \quad (2.31)$$

where $q \in \mathbb{N}$ is small [68].

Next, the matrix $\mathbf{G} = \mathbf{B}^{-1}\mathbf{V}$ is computed by solving the linear systems

$$\mathbf{B}\mathbf{g}_i = \mathbf{v}_i, \quad i = 1, 2, \dots, DP, \quad (2.32)$$

from which we get $\mathbf{G} = [\mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_{DP}]$. The linear systems in (2.32) can be solved efficiently and in parallel by a Krylov subspace method such as the preconditioned conjugate gradient method. Finally, the desired output is found as $\text{diag}_D(\mathbf{B}^{-1}) \approx \text{diag}_D(\mathbf{G}\mathbf{V}^\top)$. The estimation procedure detailed above is summarized in Algorithm 2.

Algorithm 2 Matrix Inverse Block-Diagonal Estimation Probing Scheme

Input: $\mathbf{B} \in \mathbb{H}_+^{DN}$ **Output:** $\text{diag}_D((\mathbf{B}^{-1})_\epsilon)$

- 1: Color the graph associated with adjacency matrix $\mathcal{S}_\epsilon(\mathbf{B}^q)$ using Algorithm 1
 - 2: Construct \mathbf{V} per (2.30)
 - 3: **for** $i = 1$ **to** DP **do**
 - 4: Solve $\mathbf{B}\mathbf{g}_i = \mathbf{v}_i$ (e.g. with the preconditioned conjugate gradient method)
 - 5: **end for**
 - 6: Set $\mathbf{G} \leftarrow [\mathbf{g}_1, \dots, \mathbf{g}_{DP}]$
 - 7: Set $\text{diag}_D((\mathbf{B}^{-1})_\epsilon) \leftarrow \text{diag}_D(\mathbf{G}\mathbf{V}^\top)$
 - 8: **return** $\text{diag}_D((\mathbf{B}^{-1})_\epsilon)$
-

This method relies on coloring a graph corresponding to the sparse support of \mathbf{B}^{-1} with few colors. Physically, the P distinct colors used to label the graph partition the image into regions for which the speckle decorrelates, thereby averting the complexity of computing the full conditional posterior covariance. Rather than inverting the $DN \times DN$ matrix in (2.19b), the graph coloring with P colors allows us to compute the block-diagonal elements of the inverse by solving DP linear systems, where P is much smaller than N .

2.3.3 Summary of MAP Estimation Approach

Beginning with the matched filter estimate of reflectivity \mathbf{x} ,

$$\hat{\mathbf{x}} = \mathbf{A}^H \mathbf{y}, \quad (2.33)$$

an unbiased initial estimate of Σ_i is given by

$$\tilde{\Sigma}_i = \frac{\hat{\mathbf{x}}_i \hat{\mathbf{x}}_i^H}{\sum_k |(\mathbf{A}_m^H \mathbf{A}_m)_{ki}|^2} - \sigma_w^2 \mathbf{I}_D \sum_k |(\mathbf{A}_m)_{ki}|^2, \quad (2.34)$$

for $i = 1, 2, \dots, N$. We observe improved performance by applying a small diagonal loading to the unbiased estimate and re-scaling the off-diagonal elements as

$$(\tilde{\Sigma}_i)_{kj}^{(\text{new})} \leftarrow \frac{|\sum_l h_{i,l}(\tilde{\Sigma}_l)_{kj}|}{\sqrt{\sum_l h_{i,l}(\tilde{\Sigma}_l)_{kk} \sum_l h_{i,l}(\tilde{\Sigma}_l)_{jj}}} (\tilde{\Sigma}_i)_{kj} \quad (2.35)$$

where $h_{i,l}$ are the weights of a Gaussian kernel with a bandwidth of one pixel [42].

The iterative EM scheme for approximating a solution to (2.17) is summarized in Algorithm 3.

Algorithm 3 Iterative MAP Estimation Scheme

Input: $\mathbf{y}, \mathbf{A}, \sigma_w, N_s, \epsilon_s, \{(\mathbf{G}_i, v_i, \alpha_i)\}_{i=1}^\ell, \gamma$

Output: $\hat{\Sigma}$

- 1: Set $\hat{\mathbf{x}} \leftarrow \mathbf{A}^H \mathbf{y}$
 - 2: **for** $i = 1$ **to** N **do**
 - 3: Set $\tilde{\Sigma}_i \leftarrow \frac{\hat{\mathbf{x}}_i \hat{\mathbf{x}}_i^H - \sigma_w^2 \mathbf{I}_D \sum_k |(\mathbf{A}_m)_{ki}|^2}{\sum_k |(\mathbf{A}_m^H \mathbf{A}_m)_{ki}|^2}$
 - 4: Set $(\Sigma_i)_{kj}^{(0)} \leftarrow \frac{|\sum_l h_{il}(\tilde{\Sigma}_l)_{kj}|}{\sqrt{\sum_l h_{il}(\tilde{\Sigma}_l)_{kk} \sum_l h_{il}(\tilde{\Sigma}_l)_{jj}}} (\tilde{\Sigma}_i)_{kj}$
 - 5: Apply diagonal loading to $(\Sigma_i)^{(0)}$
 - 6: **end for**
 - 7: Set $k \leftarrow 0$
 - 8: **while** $k < N_s$ and $\epsilon < \epsilon_s$ **do**
 - 9: Compute $\boldsymbol{\mu}_{\Sigma'} = \sigma_w^{-2} \mathbf{C}_{\Sigma'} \mathbf{A}^H \mathbf{y}$
 - 10: Estimate $\text{diag}_D(\mathbf{C}_{\Sigma^{(k)}})$ with Algorithm 2
 - 11: Set $\text{diag}_D(\mathbf{S}_{\Sigma^{(k)}}) \leftarrow \text{diag}_D(\mathbf{C}_{\Sigma^{(k)}} + \boldsymbol{\mu}_{\Sigma^{(k)}} \boldsymbol{\mu}_{\Sigma^{(k)}}^H)$
 - 12: **for** $i = 1$ **to** N **do**
 - 13: Estimate $p(z_i | \Sigma^{(k)})$ using the sum-product algorithm
 - 14: $\Sigma_i^{(k+1)} \leftarrow \frac{(\mathbf{S}_{\Sigma^{(k)}})_i + \sum_{q=1}^\ell p(z_i=q | \Sigma^{(k)}) \mathbf{G}_q}{1 + \sum_{q=1}^\ell p(z_i=q | \Sigma^{(k)}) (v_q + D)}$
 - 15: **end for**
 - 16: Set $\epsilon \leftarrow \frac{\|\Sigma^{(k+1)} - \Sigma^{(k)}\|_F}{\|\Sigma^{(k)}\|_F}$
 - 17: Set $k \leftarrow k + 1$
 - 18: **end while**
 - 19: **return** $\Sigma^{(k)}$
-

The complexity of this iterative MAP estimation scheme in Algorithm 3 is dominated by the solving of linear systems in step 10. If we use the preconditioned

conjugate gradient (PCG) method and assume that the matrix vector product in (2.32) can be computed with NUFFT-based methods [70], step 10 requires

$$\mathcal{O}(\eta D^2 N P \log(DN)) \quad (2.36)$$

flops, where η is a constant which depends on the conditioning of $\mathbf{C}_{\Sigma^{(k)}}$ [71, 72]. The benefits of this method rely on P being much smaller than N , which is the case when $\mathbf{C}_{\Sigma^{(k)}}$ is approximately sparse and its corresponding graph requires relatively few colors. Computing $\mathbf{C}_{\Sigma^{(k)}}$ explicitly using Gauss-Jordan elimination requires $\mathcal{O}(D^3 N^3)$ flops and is generally non-feasible. If $P \ll N$, Algorithm 3 presents significant reductions in complexity, e.g. in the experiments that follow, we have that $P/N \approx 1.4 \times 10^{-3}$. To provide an indication of the computing burden in practice, note that an image chip in the measured data experiments of Section IV required 190.7 seconds of computation time in the MATLAB (R2018a) environment with an 3.20GHz Intel(R) Core(TM) i5-6500 CPU.

The proposed method relies on a terrain class-based prior to express prior beliefs about the structure of the scene. For such a prior it is necessary to have knowledge of the types of terrain present in the scene and their statistical properties. As such, this method requires training data.

In experiments with simulated data, we can define the statistical properties of terrain classes and ensure that the simulated data adheres to them. In application to measured data, it is necessary to estimate the statistics of the terrain classes through auxiliary data.

2.4 Experimental Results

This section provides an evaluation, using both synthetic and measured PolSAR data, of the proposed method and its performance compared to conventional processing chains. The conventional processing chains used for comparison involve applying two despeckling approaches to $\tilde{\Sigma}$, the biased empirical covariance matrix computed from the matched filter images formed channel by channel. The first, a simple yet commonplace boxcar filtering approach, is a convolutional filter with a small rectangular kernel (here a 3×3 window) of equal weights [73]. The other is MuLoG (MULTichannel LOGarithm with Gaussian denoising), a state-of-the-art multichannel SAR despeckling procedure that facilitates the use of Gaussian denoising algorithms on multichannel SAR data and has demonstrated impressive despeckling performance [42]. In each experiment, we use the default MuLoG parameters in the provided code with block-matching and 3D filtering (BM3D) [74] as the embedded Gaussian denoiser. The BM3D denoising algorithm has shown remarkable results in additive Gaussian noise removal from natural images, and performed notably well when embedded in the MuLoG framework in [42].

2.4.1 Evaluation of Algorithm Performance

The two properties we consider desirable in coherent image recovery are speckle reduction and preservation of spatial resolution. To evaluate the degree of speckle reduction we use the conventional equivalent number of looks (ENL), defined by [75]:

$$\text{ENL} = \frac{\mu^2}{\sigma^2}, \quad (2.37)$$

where the empirical mean μ and empirical standard deviation σ are computed within a homogeneous region. A higher ENL indicates greater speckle reduction. As ENL

is a single channel metric, we compute it on the intensity (the squared amplitude) image for each channel.

To measure edge preservation, we use edge-preservation degree based on the ratio of average (EPD-ROA), defined by

$$\text{EPD-ROA} = \frac{\sum_{i=1}^K |I_{d1}(i)/I_{d2}(i)|}{\sum_{i=1}^K |I_{s1}(i)/I_{s2}(i)|}, \quad (2.38)$$

in which I_{d1} and I_{d2} are adjacent intensity values of a despeckled image along a particular direction, and I_{s1} and I_{s2} are the corresponding adjacent intensity values of a reference image [76]. EPD-ROA values close to one indicate good edge preservation. As with ENL, we measure the EPD-ROA of each channel's intensity. The EPD-ROA is measured in two orientations, one considering the immediate horizontal neighbor and one the vertical.

Additionally, a loss metric is used to measure fidelity to the simulated data: the normalized root mean square error (NRMSE), defined by

$$\text{NRMSE}(\widehat{\Sigma}; \Sigma) = \frac{\|\widehat{\zeta}\widehat{\Sigma} - \Sigma\|_F}{\|\Sigma\|_F} \quad (2.39)$$

for estimate $\widehat{\Sigma}$, reference Σ , and least-squares scaling constant $\widehat{\zeta}$ found as

$$\widehat{\zeta} = \arg \min_{\zeta} \left\{ \|\zeta\widehat{\Sigma} - \Sigma\|_F \right\} \quad (2.40)$$

2.4.2 PolSAR data

Full PolSAR data comprise four complex channels corresponding to the pairs of transmission/reception modes. Assuming a basis of horizontal and linear polarizations, each pixel then has four reflectivity elements $x_i^{HH}, x_i^{HV}, x_i^{VH}, x_i^{VV}$, where H and V refer to the horizontal and vertical polarization states, respectively. The

first superscript character refers to the transmitting polarization and the second to the receiving polarization. It is common to assume a reciprocal medium in which $x_i^{HV} = x_i^{VH}$, meaning that each pixel can be represented by just three reflectivity elements [77]. The reflectivity at pixel i is then

$$\mathbf{x}_i = \begin{bmatrix} x_i^{(1)}, & x_i^{(2)}, & x_i^{(3)} \end{bmatrix}^T, \quad (2.41)$$

the elements of which correspond to the HH, HV, and VV channels, respectively. To visualize the experimental results, we use the span, a measure of total backscattering power defined for fully polarimetric data by [78]

$$\text{span}(\mathbf{\Sigma}_i) = (\mathbf{\Sigma}_i)_{11} + 2(\mathbf{\Sigma}_i)_{22} + (\mathbf{\Sigma}_i)_{33}. \quad (2.42)$$

2.4.3 Simulated Data Results

In this section, we generate simulated PolSAR data and use it to evaluate the performance of the proposed framework. In our simulations, we begin with focused SAR imagery and use a clustering method to segment the image into ℓ terrain classes, yielding a terrain class labeling vector $\mathbf{z} \in \{1, 2, \dots, \ell\}^N$. Within each class, the mean covariance is determined and used as the scale parameter \mathbf{G}_i for that class. From the labeling \mathbf{z} and the terrain class distribution parameters which have been determined from the data, a ground truth covariance can be had by drawing a realization from the CIW distribution corresponding to each pixel. In this manner, measured SAR data are used to generate realistic simulated data which adhere to the proposed terrain class model. From the ground truth covariance, the reflectivity can be drawn as a realization of the random vector

$$\mathbf{x}_i \sim \mathcal{CN}(\mathbf{x}_i; \mathbf{0}, \mathbf{\Sigma}_i), \quad i = 1, 2, \dots, N. \quad (2.43)$$

Note that the image reflectivity at pixel i is a random quantity by virtue of the speckle model, and per Goodman [13], a Gaussian distribution is assumed. The imaging process including additive thermal noise is then simulated according to (2.4), yielding a set of noisy linear measurements.

Dual-pol SAR Simulated Data Result

In a first experiment, dual-pol SAR data collected by the Sentinel-1 mission [79] are segmented into $\ell = 3$ terrain classes using k-means clustering, providing a ground truth terrain class labeling. Within each terrain class, the mean covariance is determined and assigned as the scale parameter \mathbf{G}_i for that class with a degree of freedom $v_i \in [4, 10]$. A ground truth covariance is then drawn as a realization of $\Sigma_i \sim \text{CIW}(\Sigma_i; \mathbf{G}_{z_i}, v_{z_i})$, and speckle is simulated by drawing a realization of $\mathbf{x}_i \sim \mathcal{CN}(\mathbf{x}_i; \mathbf{0}, \Sigma_i)$, for $i = 1, 2, \dots, N$. Finally, the linear observations are simulated according to (2.4) with thermal noise added to simulate a signal-to-noise ratio (SNR) of 16 dB. The synthetic aperture is reduced to simulate a scenario in which 14.1% of pulses are unobserved in the pattern $\{+9.0\%, -6.3\%, +43.8\%, -7.8\%, +33.2\%\}$, where $+$ denotes observations and $-$ gaps, an interruption pattern to be repeated in each subsequent experiment.

A matched filter estimate of \mathbf{x} is found as in (2.33), from which a biased estimate of Σ is determined as in (2.34, 2.35). This serves both as an initial estimate for the proposed iterative method and as an input to despeckling. In this section, we compare a conventional SAR processing chain (i.e. the despeckling an empirical covariance from conventionally-formed PolSAR imagery) to the proposed iterative MAP estimation scheme.

The ground truth span of our simulated dual-pol SAR image is depicted in Fig. 2.2a. The biased empirical covariance estimate from the linear measurements is shown in Fig. 2.2b; the result of its processing by boxcar despeckling is in Fig. 2.2c and by MuLog with BM3D in Fig. 2.2d.




Figure 2.2: Simulated SAR image estimation results: (a,f) the span of the ground truth covariance, (b,g) the single-look empirical covariance, (c,h) the boxcar-filtered estimate, (d,i) the MuLoG-filtered estimate, and (e,j) the MAP estimate with the top row involving dual-pol data and the bottom fully-polarimetric data.

The MAP estimate by the proposed method is found by using Algorithm 3. In our experiments we use equal marginal prior probabilities for each terrain class. For the conditional marginal probabilities of the terrain class variables, we use $\gamma_e/\gamma_n = 1.5$, a ratio which was found to work well and is used in each subsequent experiment. The preconditioned conjugate gradient method is used to solve the linear systems in Algorithm 2, and 10 iterations of the sum-product algorithm are used to estimate $p(z_i \mid \Sigma'_i)$, for $i = 1, 2, \dots, N$, in step 13 of Algorithm 3. The span of the resulting

Table 2.1: Recovery metrics from simulated fully polarimetric SAR data experiment.

Method	NRMSE	ENL	EPD-ROA (H)	EPD-ROA (V)
Boxcar	0.928	5.153	0.1305	0.1170
MuLoG	0.874	112.399	0.1663	0.1610
MAP	0.736	33.4412	0.687	0.6349

estimate is in Fig. 2.2e, and the recovery metrics for each method are shown in Table 2.1 where the ENL and EPD-ROA measurements have been averaged across channels.

Degraded spatial resolution is readily apparent in the boxcar-filtered result, paired with meager despeckling performance. The MuLoG-filtered results are spatially smooth and evince impressive despeckling performance that is both visually-apparent and supported by its ENL. The ENL measured on the MAP result is more than ten times that of the boxcar-processed image, indicating strong speckle suppression, but not to the degree supplied by MuLoG; yet, the image texture qualitatively appears to be better preserved by the MAP approach than by either filtering method. As the ground truth is not itself homogeneous, even within a terrain class, ENL alone renders an incomplete portrait of speckle suppression. Note as well that the MAP approach fares better in NRMSE than the filtering approaches, and in edge-preservation, it yields EPD-ROA values significantly closer to unity than either filtering approach in both the orientations measured.

Fully Polarimetric SAR Simulated Data Experiment

This simulation procedure is repeated using fully polarimetric SAR data from the Pi-SAR sensor [80]. Here two terrain classes are used with equal marginal prior

probabilities on each terrain class and $\gamma_e/\gamma_n = 1.5$. The span images of the results from the various methods from this fully polarimetric simulation are shown in Fig. 2.2, and the recovery metrics are listed in Table 2.1. The quantitative results in this case are similar to those from the dual-pol data: the MAP estimate provides speckle reduction with superior edge preservation and reduced reconstruction error.

2.4.4 Measured Data Results

The proposed approach is next tested on measured PolSAR data collected by the Advanced Detection Technology Sensor (ADTS), an airborne Ka band SAR sensor operated by MIT Lincoln Laboratory. Further details and sample imagery from the sensor can be found in [61]. Once again we analyze image recovery performance under a reduced-aperture regime, now in the absence of ground truth covariance for reference.

It is first necessary to estimate the statistics of the terrain classes which comprise the scenes used for this analysis in order to inform the prior model. To do so, we will use a corpus of PolSAR data to learn the distribution parameters for our prior model, which we will then use to test image recovery performance on other data of a similar terrain composition.

In [61], Novak et al. segment ADTS PolSAR images into four region classes and reported the mean covariance of each class. Four types of clutter regions were considered in this analysis: shadows, grass, mixed scrub, and trees. Each bin of the segmentation contained 2000 to 8000 pixels (equivalent to an area of several hundred square meters) [61]. We segment a corpus of twelve 128×128 pixel image chips using a maximum likelihood classifier in which we assume that the terrain classes

Table 2.2: Stockbridge ADTS data recovery metrics.

Method	ENL	EPD-ROA (H)	EPD-ROA (V)
Boxcar	2.581	0.711	0.782
MuLoG	7.207	0.669	0.746
MAP	70.923	0.842	0.883

are CIW-distributed with scale matrices corresponding to the empirical terrain class results reported in [61] and degrees of freedom $v_i = 4$, for $i \in \{1, 2, 3, 4\}$. Once this training data was segmented into the $\ell = 4$ terrain classes, we determine the maximum likelihood CIW parameters for each class. These distribution parameters are then used to inform the prior model used to test recovery performance.

As in the simulated data experiments, the matched filter estimate of the reflectivity is found as in (2.33), from which a biased estimate of the polarimetric covariance is found for each pixel as in (2.34, 2.35). Note that, unlike in the simulated data experiments, speckle is not simulated as it is assumed to be present in the data. The biased empirical covariance estimates for two sample images shown in Fig. 2.3b and Fig. 2.3g; the results of two speckle filtering methods, boxcar and MuLoG, are shown alongside them in Fig. 2.3. This procedure is applied to a set of twelve 128×128 images, the mean recovery metrics from which are reported in Table 2.2.

The proposed MAP inference method is carried out on the measured data with $\gamma_e/\gamma_n = 1.5$ and equal marginal prior probabilities on the terrain class. The recovered span images from the proposed inference method are shown in Fig. 2.3e and Fig. 2.3j for two representative image chips. Absent ground truth, we have plotted in Fig. 2.3a and Fig. 2.3f the span of the matched filter-based empirical covariances found as




Figure 2.3: Measured fully polarimetric Ka-band SAR data estimation results from two image chips: (a,f) the span of the full-data empirical covariance, (b,g) the single-look empirical covariance, (c,h) the boxcar-filtered estimate, (d,i) the MuLoG-filtered estimate, and (e,j) the MAP estimate. High-quality color graphics of reconstruction results have been made available as supplementary material.

in (2.34) using the uninterrupted measurements, to which spatial averaging has been applied using a 3×3 Gaussian kernel. These filtered full-data images are also used as references to compute the EPD-ROA metrics shown in Table 2.2.

The ENL metrics in Table 2.2, indicate that the proposed approach provides greater speckle suppression in homogeneous areas than the boxcar filter or MuLoG. Note that in the simulated data experiments, it was found that the MuLoG-filtered results had a higher ENL metric than the proposed method.

The EPD-ROA values are evidence of better edge preservation by the MAP approach than either of the filtering techniques. Considering the full-data span images to be imperfect references, the textures seem qualitatively to have been recovered more faithfully by the MAP approach.

2.4.5 Discussion

The proposed Bayesian method provides a principled balance between competing objectives when computing multichannel coherent images from under-sampled data. On one hand, we wish to maintain resolution, reduce imaging side-lobes, and preserve texture, and, on the other hand, we seek to suppress speckle and thermal noise. The proposed hierarchical model provides this balance, while exploiting correlations across channels. For example, the balance with respect to preserving texture arises from the CIW scale matrix: e.g., a larger determinant encourages richer texture at the expense of speckle reduction. In contrast, the existing processing chain performs despeckling that is agnostic to gaps in phase history data, to typical image texture, and to thermal noise power. Traditionally, speckle is ameliorated by averaging: using local sums, as in boxcar or Gaussian filtering, or by summing similar patches found throughout an image, as in BM3D. By tuning the filtering, such as increasing the filter support beyond a 3-by-3 neighborhood, any level of despeckling may be achieved, but at the expense of coarser image resolution and blurred image texture. The balancing of objectives via the Bayesian imaging procedure is evident in the numerical results presented in this section; for example, in the synthetic data experiment reported in Table 2.3, MuLoG provides more image smoothing, and hence a larger ENL, but at the expense of texture and resolution. The specific statistical model presented here uses a mixture model generated from terrain classes; for each class, a CIW scale matrix provides a prior, and the algorithm applies the mixture density without any explicit image segmentation required. The scale matrices are estimated from reference images and applied to mission images. Variations of the proposed framework might also be implemented with an MRF prior model [54], a TV regularization term [43,44],

Table 2.3: Recovery metrics from simulated dual-pol SAR data experiment.

Method	NRMSE	ENL	EPD-ROA (H)	EPD-ROA (V)
Boxcar	0.811	9.148	0.4958	0.535
MuLoG	0.752	504.272	0.457	0.491
MAP	0.682	125.940	0.914	0.925

or other Gibbs priors [81]. Further, an implicit prior could be used in a plug-and-play framework [55, 56], e.g. those in the form of neural network-based denoisers [82–84].

2.5 Conclusion

2.5.1 Summary

We have proposed a new method for multichannel coherent image recovery based on a hierarchical statistical model which not only accommodates a reduced synthetic aperture but also suppresses speckle and thermal noise. The model presented here uses latent terrain class variables with an MRF prior to exploit both structure among coherent imaging channels and spatial correlation for improved image recovery.

The algorithm to produce the MAP estimate based on this model uses a surrogate scheme from [54, 55], extending it for use with multiple channels in a reduced aperture regime by the novel application of a matrix inverse block-diagonal estimation scheme. This approach, tested on synthetic and measured PolSAR data, has been shown to suppress speckle and preserve spatial resolution in the recovery of multichannel coherent images.

2.5.2 Future Work

Future work might involve learning of the terrain classes from the noisy data, rather than requiring auxiliary data. The approach in this chapter assumes that the fully developed speckle model is followed throughout the scene. Instead, a mixture model might be used that accounts for areas in which man-made media are present that do not follow the speckle model.

Chapter 3: Coherent Change Detection in SAR Imaging

3.1 Introduction

An important application of SAR is the detection of changes to the complex scene reflectivity that occur between repeated collections of the same terrain.² In coherent change detection (CCD), both the magnitude and the phase components of SAR images are used to identify subtle changes in the scene, such as those caused by the movement of vehicles over grass [86], foot traffic [87], or land subsidence [88].

The complex correlation coefficient between a pair of SAR images (denoted time-1 and time-2) at a given pixel with index i is [89]

$$\gamma_i e^{j\phi_i} = \frac{\text{E} [y_i^{(1)} \overline{y_i^{(2)}}]}{\sqrt{\text{E} [|y_i^{(1)}|^2] \text{E} [|y_i^{(2)}|^2]}}, \quad (3.1)$$

where $y_i^{(k)} \in \mathbb{C}$ is the value of pixel i in the time- k image, $\text{E}[\cdot]$ denotes expectation, and $\overline{}$ denotes complex conjugation. The interferometric phase ϕ_i in (3.1) depends on the terrain topography as well as the displacement between the flight paths of the two passes. The magnitude of the correlation coefficient, denoted by γ_i and referred to as the coherence, is sensitive to the temporal changes in complex reflectivity that are of interest in CCD. Environmental effects can cause decorrelation, the degree of

²This chapter is largely drawn from [85], coauthored with Joshua N. Ash and Lee C. Potter (© 2023 IEEE).

which will depend on the time between passes and the types of scatterers present [86]. The coherence is also influenced by a number of other factors: thermal noise in the SAR system, the co-registration process of the images, the mismatch in acquisition geometries between the two passes, and the presence of volume scatterers in the scene [87, 90].

Change detection in SAR has relied largely on the computation of local statistics within sliding estimation windows, e.g., [91–94]. The coherence, for example, is locally estimated as

$$\hat{\gamma}_i = \frac{\left| \sum_j y_j^{(1)} \overline{y_j^{(2)}} \right|}{\sqrt{\sum_j |y_j^{(1)}|^2 \sum_j |y_j^{(2)}|^2}}, \quad (3.2)$$

in which each summation occurs over a small neighborhood or window of image samples [89, 91]. Man-made disturbances may result in small, isolated fluctuations in the scene coherence, for which a commensurately small estimation window is desirable. However, the use of a small estimation window to localize fine-scale scene changes leads to noisy coherence estimates and poor detection performance [86]. Other window-based CCD methods rely on likelihood ratio tests [93, 95] or involve Stokes parameters [96]. Rather than computing a statistic within a sliding rectangular window, locally-adaptive directional windows are used in [97] for despeckling and coherence estimation.

Various alternatives to window-based coherence estimation have been proposed. In [98], the two-dimensional discrete wavelet packet transform is used to adapt to the spatially nonhomogeneous character of the scene coherence. This approach allows for high resolution coherence estimation that is not biased by the interferometric phase. Nonlocal methods, based on the nonlocal means filter [37], have also been used to estimate both the coherence and the interferometric phase, e.g., the nonlocal

interferometric SAR (NL-InSAR) estimator [99]. Nonlocal methods exploit redundant patterns in the image, selecting pixels to combine in estimation based on the similarity of the patches that surround them. NL-InSAR and similar approaches [100, 101] estimate the coherence in a manner that may mitigate the biases and loss of resolution that are caused by scene heterogeneity in window-based approaches. However, these methods may be ill-suited to change detection scenarios with small, isolated regions of change.

Bayesian methods have been useful in the problem of joint SAR imaging and change detection from interrupted collections [63, 102]. These approaches incorporate prior information about the scene and the physical properties of change into Bayesian models which mitigate the impact of interruptions to the SAR data collections. In both [63, 102], Markov random field (MRF) priors are used to model the spatial sparsity and clustering tendencies of the scene changes. Another class of prior models called gamma MRFs was introduced to model the local dependencies among the variances of time-frequency coefficients of audio sources [103]. Gamma MRFs enforce positive correlations between neighboring variance elements and have the convenient property that they are conditionally conjugate to a Gaussian likelihood with unknown variance.

In this work we propose a Bayesian approach to SAR change detection that enables fine-resolution detection of temporal disturbances between pairs of co-registered SAR images. An MRF prior models the spatial dependencies of the scene changes, which are anticipated to appear in clusters. Gamma MRFs are used to model the spatial dependencies of the unknown mean backscatter powers of the image pair. These prior models are incorporated into a Bayesian framework from which the joint

posterior distribution of the latent variables is estimated using variational expectation maximization (EM) [104, 105], a framework which also allows for the iterative estimation of the model parameters. In addition, a 2-D polynomial regression based on the local phase differences is used to account for the spatially nonhomogeneous interferometric phase.

The proposed Bayesian approach assigns prior distributions to the latent model variables, exploiting spatial structure in the scene changes as well as in the scattering properties of the scene. This approach does not suffer from the loss of spatial resolution and fine-scale details as do conventional window-based SAR CCD methods. The proposed approach has the added benefit that it estimates the prior model hyperparameters from the data, rather than requiring manual tuning. In experiments using both simulated data and measured interferometric SAR images, this approach is shown to yield improved detection performance compared to the conventional window-based statistics used for SAR CCD.

3.2 Background

3.2.1 Variational Inference

Of primary interest in Bayesian inference is the estimation of the posterior distribution of the latent random vector \mathbf{x} given the observed random vector \mathbf{y} . The idea behind variational inference is to approximate the posterior with a distribution $q(\mathbf{x})$ that minimizes the Kullback-Leibler (KL) divergence between the two distributions.

The KL divergence between $q(\mathbf{x})$ and the posterior distribution is given by [104]

$$D_{\text{KL}}(q \parallel p) = \int q(\mathbf{x}) \log \frac{q(\mathbf{x})}{p(\mathbf{x} \mid \mathbf{y}; \Theta)} d\mathbf{x} \quad (3.3)$$

$$= \log p(\mathbf{y}; \Theta) + \int q(\mathbf{x}) \log \frac{q(\mathbf{x})}{p(\mathbf{x}, \mathbf{y}; \Theta)} d\mathbf{x} \quad (3.4)$$

$$= \log p(\mathbf{y}; \Theta) - \mathcal{L}(q; \Theta), \quad (3.5)$$

where Θ denotes a set of fixed parameters. Note that $\mathcal{L}(q; \Theta)$ is a lower bound on the log evidence, $\log p(\mathbf{y}; \Theta)$. Because the evidence $p(\mathbf{y}; \Theta)$ is constant with respect to the variational distribution q , maximizing $\mathcal{L}(q; \Theta)$ is equivalent to minimizing the KL divergence in (3.4). Optimization of $\mathcal{L}(q; \Theta)$ over the space of all candidates for $q(\mathbf{x})$ is generally infeasible, and the search is restricted to a family of distributions. This commonly involves a mean field assumption, in which a particular factorization for the variational distribution is postulated, e.g.,

$$q(\mathbf{x}) = \prod_{i=1} q_{x_i}(x_i). \quad (3.6)$$

By fixing the other factors, the q_{x_i} which maximizes $\mathcal{L}(q; \Theta)$ can be found as [65]

$$\log q_{x_i}(x_i) = \langle \log p(\mathbf{y}, \mathbf{x}; \Theta) \rangle_{q(\mathbf{x}_{-i})} + \text{const.} \quad (3.7)$$

where $\langle \cdot \rangle_{f(x)}$ denotes expectation under the probability distribution $f(x)$, and \mathbf{x}_{-i} denotes the variables other than x_i . A solution is found by cycling through the factors and updating them according to (3.7), using the current estimates for the other factors.

In variational EM, variational inference approximates the expectation step (E-step) of the EM algorithm [64], and like exact EM, the fixed model parameters are also estimated [105]. The maximization step (M-step) of variational EM for the

update of model parameter Θ_i is given by

$$\arg \max_{\Theta_i} \langle \log p(\mathbf{y}, \mathbf{x}; \Theta) \rangle_{q(\mathbf{x})}. \quad (3.8)$$

Note that unlike in exact EM, the variational M-step uses an approximate expectation to perform parameter optimization.

3.2.2 Gamma Markov Random Fields

Gamma MRFs were first used to model the dependencies among the variances of audio signal time-frequency elements [103, 106], making use of the fact that the signal energies often change slowly over both axes. These models have also been used for spectral unmixing problems [107, 108]. In this chapter, we use gamma MRFs to model the local dependencies of the mean backscatter powers of interferometric SAR image pairs in order to improve change detection performance.

A key feature of the gamma MRF defined in [103] is the introduction of auxiliary random variables which do not necessarily correspond to quantities of interest. Instead, these auxiliary variables are present to enforce positive correlations among the primary variables while preserving the conditional conjugacy of the prior model.

A gamma MRF models a joint probability distribution on $\mathbf{r} \in \mathbb{R}_+^{|\mathcal{V}|}$ and the auxiliary variables $\mathbf{z} \in \mathbb{R}_+^{|\mathcal{V}^z|}$ by defining a bipartite undirected graph between \mathbf{r} and \mathbf{z} , with \mathcal{V} and \mathcal{V}^z denoting sets of indices on which \mathbf{r} and \mathbf{z} are respectively defined.




Figure 3.1: The neighborhood structure of a 4-connected gamma MRF showing the vertical and horizontal neighbors for r_i . The auxiliary variables are shown in red.

The joint probability of \mathbf{r} and \mathbf{z} is expressed as [103]

$$p(\mathbf{r}, \mathbf{z}; \mathbf{a}) = \frac{1}{Z(\mathbf{a})} \pi(\mathbf{r}, \mathbf{z}; \mathbf{a}) \quad (3.9)$$

$$= \frac{1}{Z(\mathbf{a})} \prod_{i \in \mathcal{V}} \exp \left\{ - \left(1 + \sum_{j \in S_i^r} a_{i,j} \right) \log r_i \right\} \\ \times \prod_{j \in \mathcal{V}^z} \exp \left\{ \left(-1 + \sum_{i \in S_j^z} a_{i,j} \right) \log z_j - \sum_{i \in S_j^z} \frac{a_{i,j} z_j}{r_i} \right\}, \quad (3.10)$$

in which $\pi(\cdot; \mathbf{a})$ denotes the non-normalized probability density, $Z(\mathbf{a})$ is a normalizing factor or partition function, and S_i^r and S_j^z contain the indices of the neighboring elements of r_i and z_j , respectively. The hyperparameter $a_{i,j}$ determines the coupling strength between r_i and z_j for $i \in \mathcal{V}$ and $j \in S_i^r$.

The neighborhood structure of a 4-connected gamma MRF is illustrated in Fig. 3.1, in which the neighboring auxiliary nodes (in red) are shown for a given primary variable r_i . The auxiliary variables, denoted \mathbf{z}^h (horizontal) and \mathbf{z}^v (vertical), are present between the immediately adjoining elements of \mathbf{r} in the horizontal and vertical directions.

3.3 Coherent Change Detection

3.3.1 Bayesian Change Detection Model

We assume that the time-1 image of a repeat-pass interferometric SAR image pair conditionally follows a complex circularly-symmetric Gaussian distribution:

$$p(y_i^{(1)} | r_i^{(1)}) = \mathcal{CN}(y_i^{(1)}; 0, r_i^{(1)}), \quad i \in \mathcal{V}, \quad (3.11)$$

where \mathcal{V} is a set of pixel indices on a 2-D planar grid. The latent variable $r_i^{(1)}$ is referred to as the mean backscatter power of the time-1 image at pixel i . The mean backscatter power is a geophysical quantity that is dependent on the structural and electromagnetic properties of the scattering elements within the scene [86]. The model in (3.11) is common for distributed target environments where many statistically identical discrete scatterers occupy a given resolution cell [13, 109].

We next define the conditional distribution on the time-2 image as the Gaussian mixture

$$\begin{aligned} p(y_i^{(2)} | y_i^{(1)}, r_i^{(2)}, c_i; \phi_i, \gamma) = & \delta(c_i - 1) \mathcal{CN}(y_i^{(2)}; 0, r_i^{(2)}) \\ & + \delta(c_i + 1) \mathcal{CN}(y_i^{(2)}; \gamma e^{-j\phi_i} y_i^{(1)}, (1 - \gamma^2) r_i^{(2)}) \end{aligned} \quad (3.12)$$

where $\delta(\cdot)$ is the Kronecker delta function, ϕ_i is the unknown interferometric phase at pixel i , and $c_i \in \{-1, 1\}$ is a change indicator random variable. Here, $c_i = 1$ indicates a temporal change at pixel i , and $c_i = -1$ indicates the absence of change. The unknown parameter $\gamma \in (0, 1)$ determines the degree of change allowable within the no-change state. Note that (3.12) is a temporal change model similar to that in [94]; (3.12) differs by omitting thermal noise terms and allowing the mean backscatter power to differ between passes.

The interferometric phase ϕ is often assumed to be constant over the small neighborhoods that are conventionally used to compute change statistics. However, if ϕ is not constant within the estimation window, the sample coherence estimator will be biased [98]. For methods operating on a pixel-wise basis, the phase must be estimated, either by means of external digital elevation models (DEMs) or from the radar data. In this work, we estimate the spatially varying interferometric phase from the radar data by modeling the interferometric phase with a low-order two-dimensional polynomial.

To account for the spatial correlations between the neighboring change indicator variables, we define an MRF prior distribution for \mathbf{c} . In particular, this chapter focuses on the following prior distribution for the change indicator variables [110]

$$p(\mathbf{c} \mid \chi) = \frac{1}{Z(\chi)} \exp \left\{ \sum_{i \in \mathcal{V}} (\chi_0 c_i + \frac{\chi}{2} \sum_{j \in S_i^c} c_i c_j) \right\}, \quad (3.13)$$

in which χ_0 is assumed to be known *a priori* and S_i^c contains the neighboring elements of c_i for $i \in \mathcal{V}$. Values of $\chi_0 < 0$ enforce the notion that change pixels are less likely than no-change pixels. Larger values of the interaction parameter χ imply greater spatial correlation among the change indication variables, and when $\chi = 0$, the elements of \mathbf{c} are mutually independent. In this chapter, we use a 4-connected neighborhood structure for the change indicator variables.

Further, we assign gamma MRF priors to $\mathbf{r}^{(1)}$ and $\mathbf{r}^{(2)}$ that model the local dependencies of the mean backscatter power within the scene. This quantity may vary across the scene due to errors in the antenna pointing accuracy or spatial variations in the geophysical qualities of the scene. The structure of these spatial variations is meant to be captured by the gamma MRFs. The mean backscatter power of a given pixel may also differ between passes due to radiometric miscalibrations or scene

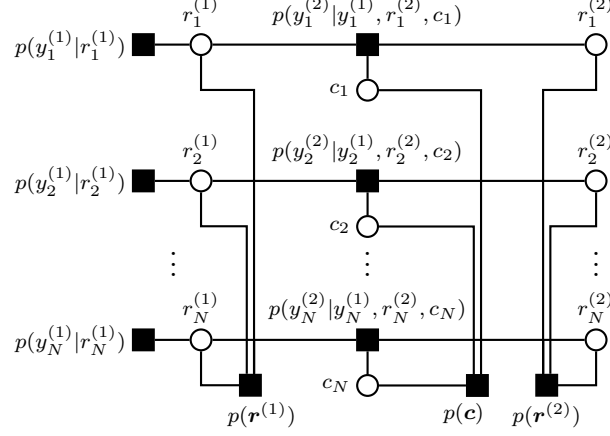


Figure 3.2: Factor graph representation of the proposed change detection model. The squares denote factor nodes and the circles variable nodes.

disturbances that occur between the collections [93]. We assume that the hyperparameters $a_{i,j}$ are uniform across and between both images and that a single hyperparameter a need be learned for the gamma MRFs. We use a 4-connected neighborhood structure for the gamma MRFs as depicted in Fig. 3.1. A factor graph representation of the proposed change detection model is shown in Fig. 3.2, in which square nodes denote factors and round nodes represent variables.

3.3.2 Inference

The joint posterior density of the latent variables is given by

$$\begin{aligned}
 p(\mathbf{r}^{(1)}, \mathbf{r}^{(2)}, \mathbf{z}^{(1)}, \mathbf{z}^{(2)}, \mathbf{c} \mid \mathbf{y}^{(1)}, \mathbf{y}^{(2)}; \Theta) &\propto p(\mathbf{c}; \chi) \\
 &\times p(\mathbf{r}^{(1)}, \mathbf{z}^{(1)}; a) p(\mathbf{r}^{(2)}, \mathbf{z}^{(2)}; a) \\
 &\times \prod_{i \in \mathcal{V}} p(y_i^{(1)} \mid r_i^{(1)}) p(y_i^{(2)} \mid y_i^{(1)}, r_i^{(2)}, c_i; \phi_i, \gamma),
 \end{aligned} \tag{3.14}$$

where $\Theta = \{\gamma, \chi, \phi, a\}$ is the set of model parameters, and \mathcal{V} is the set of pixel indices in the image. While exact computation of the posterior distribution in (3.14)

is computationally prohibitive, the marginal posteriors of the latent variables can be efficiently approximated by using variational EM on the factor graph of Fig. 3.2. In doing so, we use the following mean field approximation for the posterior:

$$q = \prod_{i \in \mathcal{V}} q_{r_i^{(1)}}(r_i^{(1)}) q_{r_i^{(2)}}(r_i^{(2)}) q_{c_i}(c_i) \prod_{j \in \mathcal{V}^z} q_{z_j^{(1)}}(z_j^{(1)}) q_{z_j^{(2)}}(z_j^{(2)}), \quad (3.15)$$

the factors of which are cyclically updated until convergence. Note that, due to the mean-field assumption, the computational complexity of a cycle of factor updates is linear in the number of pixels— $\mathcal{O}(|\mathcal{V}|)$. For brevity, the subscripts of the variational factors may be omitted when clarity is not impacted.

Next we describe the updates to the variational factors in (3.15). From (3.7) and (3.14), the factor associated with $r_i^{(1)}$ is updated as $q(r_i^{(1)}) = \mathcal{IG}(r_i^{(1)}; \alpha_{r,i}^{(1)}, \beta_{r,i}^{(1)})$ with

$$\alpha_{r,i}^{(1)} = 1 + a|S_i^r| \quad (3.16a)$$

$$\beta_{r,i}^{(1)} = |y_i^{(1)}|^2 + \sum_{j \in S_i^r} a \langle z_j^{(1)} \rangle, \quad (3.16b)$$

where $\mathcal{IG}(\cdot; \alpha, \beta)$ is the inverse gamma distribution's probability density function defined in Appendix A. The factor associated with $r_i^{(2)}$ is updated as $q(r_i^{(2)}) = \mathcal{IG}(r_i^{(2)}; \alpha_{r,i}^{(2)}, \beta_{r,i}^{(2)})$ with

$$\alpha_{r,i}^{(2)} = 1 + a|S_i^r| \quad (3.17a)$$

$$\beta_{r,i}^{(2)} = \frac{\langle \delta(c_i + 1) \rangle |y_i^{(2)} - \gamma e^{-j\phi_i} y_i^{(1)}|^2}{(1 - \gamma^2)} + \langle \delta(c_i - 1) \rangle |y_i^{(2)}|^2 + \sum_{j \in S_i^r} a \langle z_j^{(2)} \rangle. \quad (3.17b)$$

The time- k auxiliary variable factors for $k \in \{1, 2\}$ are given by

$$q(z_j^{(k)}) = \mathcal{G}(z_j^{(k)}; \alpha_{z,j}^{(k)}, \beta_{z,j}^{(k)}), \quad (3.18)$$

with

$$\alpha_{z,j}^{(k)} = a|S_j^z| \quad (3.19a)$$

$$\beta_{z,j}^{(k)} = \sum_{i \in S_j^z} \left\langle \frac{a}{r_i^{(k)}}, \right\rangle \quad (3.19b)$$

where $\mathcal{G}(\cdot; \alpha, \beta)$ is the gamma probability density function defined in Appendix A.

From (3.7) and (3.14), the factors associated with the change bits are updated using

$$\begin{aligned} \log q(c_i = -1) &= -\chi_0 - \log(1 - \gamma^2) \\ &- \left\langle \frac{|y_i^{(2)} - \gamma e^{-j\phi_i} y_i^{(1)}|^2}{r_i^{(2)}(1 - \gamma^2)} \right\rangle - \chi \sum_{j \in S_i^c} \langle c_j \rangle + \text{const.} \end{aligned} \quad (3.20a)$$

and

$$\log q(c_i = 1) = \chi_0 - \left\langle \frac{|y_i^{(2)}|^2}{r_i^{(2)}} \right\rangle + \chi \sum_{j \in S_i^c} \langle c_j \rangle + \text{const.} \quad (3.20b)$$

Further details regarding these updates are in Appendix A.

3.3.3 Parameter Estimation

Now we describe the procedures used to estimate the unknown model parameters. Two of the model parameters, γ and χ , are updated in a straightforward manner via the M-step in (3.8). The estimate for the interferometric phase ϕ is updated using a 2-D polynomial regression based on the local phase differences, rather than a conventional M-step update. The gamma MRF hyperparameter governing the local dependencies of the mean backscatter powers is estimated using contrastive divergence [111].

Gamma MRF Hyperparameter

We assume that a single hyperparameter, denoted by a , governs the prior models for the mean backscatter power of both the time-1 and time-2 images. This hyperparameter can then be learned from a single image as a prerequisite to the other parameter estimation or inference tasks. Because variational Bayesian methods have been shown to poorly estimate gamma MRF hyperparameters [112], we use a contrastive divergence approach. Contrastive divergence is an iterative approximate maximum likelihood method that uses Markov chain Monte Carlo (MCMC) samples to approximate the intractable log-partition function [111]. The approach described in Algorithm 4 is similar to one used in [103] for the estimation of the gamma MRF hyperparameters in audio source modeling, though, Algorithm 4 uses variational inference rather than MCMC samples to approximate the expectations in the contrastive divergence gradients. The variational inference procedure used to facilitate this approximation is described in Algorithm 5. This revised approach results in similar performance to [103] in our experiments, though it is faster as it requires fewer MCMC samples.

Interferometric Phase

We adopt a 2-D polynomial model for the local interferometric phase differences and, as in [113], find a vector of polynomial coefficients \mathbf{g} by solving a weighted least squares problem:

$$\hat{\mathbf{g}} = \arg \min_{\mathbf{g}} \left\| \mathbf{P} \begin{bmatrix} \phi_{\Delta 1}^T & \phi_{\Delta 2}^T \end{bmatrix}^T - \begin{bmatrix} \mathbf{B}_{\Delta 1}^T & \mathbf{B}_{\Delta 2}^T \end{bmatrix}^T \mathbf{g} \right\|_2^2,$$

where $\phi_{\Delta 1}$ and $\phi_{\Delta 2}$ are respectively the differences of the interferometric phase image along the first and second image dimensions; the matrices $\mathbf{B}_{\Delta 1}$ and $\mathbf{B}_{\Delta 2}$ are formed

Algorithm 4 (`cd_gmrf_normal`) Contrastive divergence on a gamma MRF model with Gaussian observations.

Input: Observations \mathbf{y} .

- 1: *Fixed parameters:* number of contrastive divergence steps N_{cd} , number of MCMC cycles N_{mc} , and step sizes $\boldsymbol{\lambda} \in \mathbb{R}_+^{N_{\text{cd}}}$
- 2: Initialize hyperparameter estimate \hat{a}
- 3: **for** $i \in \{1, \dots, N_{\text{cd}}\}$ **do**
- 4: $q \leftarrow \text{vi_gmrf_normal}(\mathbf{y}, a)$
- 5: Draw $\mathbf{r}^{(0)}, \mathbf{z}^{(0)}$ from q
- 6: **for** $k \in \{1, \dots, N_{\text{mc}}\}$ **do**
- 7: Draw $\mathbf{z}^{(k)}$ from $p(\mathbf{z} \mid \mathbf{r}^{(k-1)}; a)$
- 8: Draw $\mathbf{r}^{(k)}$ from $p(\mathbf{r} \mid \mathbf{y}, \mathbf{z}^{(k)}; a)$
- 9: **end for**
- 10: Draw y'_i from $\mathcal{CN}(y_i; 0, r_i^{(N_{\text{mc}})})$ for $i \in \mathcal{V}$
- 11: $q' \leftarrow \text{vi_gmrf_normal}(\mathbf{y}', a)$
- 12: $\hat{a} \leftarrow \hat{a} - \lambda_i \left(\left\langle \frac{\partial \log \pi(\mathbf{r}, \mathbf{z} | a)}{\partial a} \right\rangle_{q'} - \left\langle \frac{\partial \log \pi(\mathbf{r}, \mathbf{z} | a)}{\partial a} \right\rangle_q \right)$
- 13: **end for**

Output: Estimated hyperparameter \hat{a}

Algorithm 5 (`vi_gmrf_normal`) Variational inference on a gamma MRF model with complex Gaussian observations .

Input: Observations \mathbf{y} , gamma MRF hyperparameter a

- 1: Initialize q_r
- 2: **repeat**
- 3: $q_z \leftarrow \mathcal{G}(\cdot; \boldsymbol{\alpha}_z, \boldsymbol{\beta}_z)$ with $\alpha_{z,j} = a|S_j^z|$ and $\beta_{z,j} = a \sum_{i \in S_j^z} \langle r_i^{-1} \rangle$ for $j \in \mathcal{V}^z$
- 4: $q_r \leftarrow \mathcal{IG}(\cdot; \boldsymbol{\alpha}_r, \boldsymbol{\beta}_r)$ with $\alpha_{r,i} = 1 + |S_i^r|$ and $\beta_{r,i} = |y_i|^2 + a \sum_{j \in S_i^r} \langle z_j \rangle$ for $i \in \mathcal{V}$
- 5: **until** Terminated

Output: $q = q_z q_r$

by taking the corresponding differences in the rows of the polynomial design matrix \mathbf{B} , each column of which defines the spatial distribution of a polynomial term. The matrix \mathbf{P} is a row-pruning operator which excludes areas of temporal change from the computation of the polynomial coefficients. In particular, the pruning operation excludes elements with an estimated posterior probability of change above a threshold,

after which the element is considered unreliable. A threshold of 0.5 is used for the computed results, and performance was not found to be overly sensitive to the choice of threshold.

Once the polynomial coefficients are estimated, the method of least squares can be used to estimate the constant phase offset. The interferometric phase estimate is thus given by

$$\hat{\phi} = \mathbf{B}\hat{\mathbf{g}} + \hat{\phi}_0 \quad (3.21)$$

where $\hat{\phi}_0 \in [-\pi, \pi)$ is the constant phase offset estimate. The interferometric phase estimate is updated iteratively throughout the inference procedure along with the M-step parameter updates.

Coherence Parameter

Using (3.8), the updated estimate for γ is given by

$$\hat{\gamma} = \arg \max_{\gamma} \langle \log p(\mathbf{y}^{(2)} \mid \mathbf{y}^{(1)}, \mathbf{c}, \mathbf{r}^{(2)}; \phi, \gamma) \rangle \quad (3.22)$$

$$\begin{aligned} &= \arg \max_{\gamma} \left\{ \sum_{i \in \mathcal{V}} -\langle \delta(c_i + 1) \rangle \log(1 - \gamma^2) \right. \\ &\quad \left. - \langle \delta(c_i + 1) \rangle \left\langle \frac{\gamma^2 |y_i^{(1)}|^2 + |y_i^{(2)}|^2 - \gamma b_i}{r_i^{(2)}(1 - \gamma^2)} \right\rangle \right\}, \end{aligned} \quad (3.23)$$

in which

$$b_i = 2\text{Re}\{\bar{y}_i^{(2)} y_i^{(1)} e^{-j\phi_i}\}. \quad (3.24)$$

By setting to zero the derivative of the objective in (3.23) with respect to γ , we have a third-order polynomial,

$$\begin{aligned}
& - \left(\sum_{i \in \mathcal{V}} 2 \langle \delta(c_i + 1) \rangle \right) \gamma^3 + \left(\sum_{i \in \mathcal{V}} \left\langle \frac{b_i \delta(c_i + 1)}{r_i^{(2)}} \right\rangle \right) \gamma^2 \\
& + \left(\sum_{i \in \mathcal{V}} \langle \delta(c_i + 1) \rangle \left(2 - 2 \left\langle \frac{|y_i^{(1)}|^2 + |y_i^{(2)}|^2}{r_i^{(2)}} \right\rangle \right) \right) \gamma \\
& + \left(\sum_{i \in \mathcal{V}} \left\langle \frac{b_i \delta(c_i + 1)}{r_i^{(2)}} \right\rangle \right) = 0,
\end{aligned} \tag{3.25}$$

that may be used to find the optimal value for γ .

Change Indicator MRF Hyperparameter

Due to the difficulty of assessing the partition function in (3.13), given by

$$Z(\chi) = \sum_{\mathbf{c} \in \{-1, 1\}^{|\mathcal{V}|}} \exp \left\{ \sum_{i \in \mathcal{V}} \chi_0 c_i + \frac{\chi}{2} \sum_{j \in S_i^c} c_i c_j \right\}, \tag{3.26}$$

we use the following mean field approximation [114]:

$$Z_{\text{MF}}(\chi) = \prod_{i \in \mathcal{V}} \sum_{c_i \in \{-1, 1\}} \exp \left\{ \chi_0 c_i + \chi \sum_{j \in S_i^c} c_i \langle c_j \rangle \right\}. \tag{3.27}$$

Note that (3.27) is a latent variable version of the approximation used in the pseudo-likelihood of Besag [115], i.e., equivalent if $\langle c_j \rangle$ in (3.27) is replaced with c_j . With the mean field partition approximation and (3.8), the M-step update to χ is given by

$$\hat{\chi} = \arg \max_{\chi} \left\{ -\log Z_{\text{MF}}(\chi) + \sum_{i \in \mathcal{V}} \chi_0 \langle c_i \rangle + \chi \sum_{j \in S_i^c} \langle c_i \rangle \langle c_j \rangle \right\}, \tag{3.28}$$

which can be solved using Newton's method.

3.3.4 Summary

The proposed coherent change detection approach is summarized in Algorithm 6.

Algorithm 6 The proposed variational EM change detection algorithm

Input: Observations: $\mathbf{y}^{(1)}, \mathbf{y}^{(2)}$

- 1: Initialize $q_r, q_c, \hat{\phi}, \hat{\chi}, \hat{\gamma}$
- 2: $\hat{a} \leftarrow \text{cd_gmrf_normal}(\mathbf{y}^{(1)})$
- 3: **repeat**
- 4: **repeat**
- 5: Update $q_{\mathbf{z}^{(1)}}$ according to (3.19a) and (3.19b)
- 6: Update $q_{\mathbf{z}^{(2)}}$ according to (3.19a) and (3.19b)
- 7: Update $q_{\mathbf{r}^{(1)}}$ according to (3.16a) and (3.16b)
- 8: Update $q_{\mathbf{r}^{(2)}}$ according to (3.17a) and (3.17b)
- 9: Update q_c according to (3.20a) and (3.20b)
- 10: **until** Terminated
- 11: Update $\hat{\phi}$ according to (3.21)
- 12: Update $\hat{\gamma}$ according to (3.23)
- 13: Update $\hat{\chi}$ according to (3.28)
- 14: **until** Terminated

Output: Approximate change posterior distribution q_c

3.4 Experiments

In this section, we evaluate the proposed change detection approach using both simulated and measured interferometric SAR image pairs. The analysis includes comparisons to conventional window-based change detection as well as to coherence estimates from NL-InSAR [99], a more advanced approach that uses nonlocal techniques.

3.4.1 Simulated Data Experiments

Interferometric SAR data are first simulated according to four example scenarios that are intended to provide a variety of test conditions for the evaluation of change detection performance. The scenarios use various configurations of the underlying simulation parameters that may impact change detection: the interferometric phase,




Figure 3.3: Values used to simulate Scenarios 1–4 (columns 1–4, respectively). (a–d) Normalized mean backscatter power (in dB). (e–h) Interferometric phase (in radians).

mean backscatter power, coherence, and change maps (i.e., the arrangement of temporal disturbances in the scene).

To provide an empirical basis for the change maps used to simulate the data, we use ground truth temporal change maps derived from the Onera satellite change detection dataset [116]. The change maps derived from the Onera change detection dataset for the four scenarios appear in the first column of Fig. 3.5 as binary images with black pixels indicating a change between the two passes. Scenarios 1 and 3 use mean backscatter power images that have been realized by a Gibbs sampling procedure on gamma MRF models with $a = 20$ and $a = 16$, respectively. Scenarios 2 and 4 follow linear ramps that range in power from 0 dB to 30 dB. In each case the time-1 and time-2 images have the same mean backscatter power, the patterns for which are




Figure 3.4: Magnitude images for Scenarios 1–4 (columns 1–4, respectively). (a–d) Time-1 magnitude images. (e–h) Time-2 magnitude images.

shown in the first row of Fig. 3.3. The underlying interferometric phase data are two-dimensional polynomials which have been fit to interferogram phases obtained from image pairs in the Gotcha challenge problem dataset [2]. The phase images used to simulate each scenario are shown in the second row of Fig. 3.3. Interferometric SAR image pairs of size 256×256 pixels are then drawn according to the change model in (3.11) and (3.12) using γ values of 0.55, 0.65, 0.75, and 0.85 for Scenarios 1–4, respectively. The magnitudes of the resulting image pairs are shown in Fig. 3.4.

The change detection results from the four scenarios are shown in Fig. 3.5. The second and third columns of Fig. 3.5 show the sample coherence maps computed with 3×3 and 5×5 windows, respectively. The fourth column shows the coherence estimates from NL-InSAR [99], computed using the default parameters in the publicly




Figure 3.5: Comparison of change detection results on simulated data Scenarios 1–4 (rows 1–4, respectively) with various coherence values γ . The first column has the true change maps with change pixels in black, and the remaining columns show the change detection results from the compared methods. The last column shows the results of the proposed inference approach with black indicating a high probability of change.

available MATLAB code.³ Shown in the fifth column of Fig. 3.5, the results from the proposed method express the estimated posterior probability of change at a given pixel, with darker pixels indicating a higher probability of change. The results from the proposed method are computed according to Algorithm 6 with the variational

³<https://www.charles-deledalle.fr/pages/nlinsar.php>

factors $q_{\mathbf{r}^{(k)}}$ for $k \in \{1, 2\}$ initialized to the maximum likelihood estimates computed in 9×9 neighborhoods and a maximum of 35 iterations of variational EM. The change posterior estimates are initialized to 0.5 for $i \in \mathcal{V}$, and the interferometric phase parameters are initialized to the interferogram phases.

The receiver operating characteristic (ROC) curves for the four simulated data scenarios are shown in Fig. 3.6, in which the curves labeled “SC” correspond to the sample coherence computed with the specified window size and “VEM” refers to the proposed variational EM approach. The ROC curves show the proposed method to outperform the others with respect to reduced false alarm rates for a given probability of detection. The improved performance of the proposed approach is particularly apparent from the low-coherence simulations using Scenarios 1–2 (i.e., the first two rows in Fig. 3.5) in which many false alarms are visible in the coherence estimates. Qualitatively, the NL-InSAR estimates obscure the edges of the change regions and in some cases render small, isolated regions of change difficult to observe.

In these experiments, both NL-InSAR and the proposed approach are significantly slower than the classical window-based CCD methods, which averages 2.16 ms of computing time: NL-InSAR runs in 65.49 s and the proposed approach in 22.57 s. All experiments are computed in MATLAB R2020a with a 3.2 GHz Intel Core i5-6500 CPU.

As a further point of comparison, we compute the threshold values that, when applied to the coherence estimates, minimize the empirical error rate in the four simulated scenarios depicted in Fig. 3.5. The error rates for the coherence estimates thresholded in this manner are reported in Table 3.1. The error rates for the variational EM approach are computed using the maximum *a posteriori* (MAP) estimates

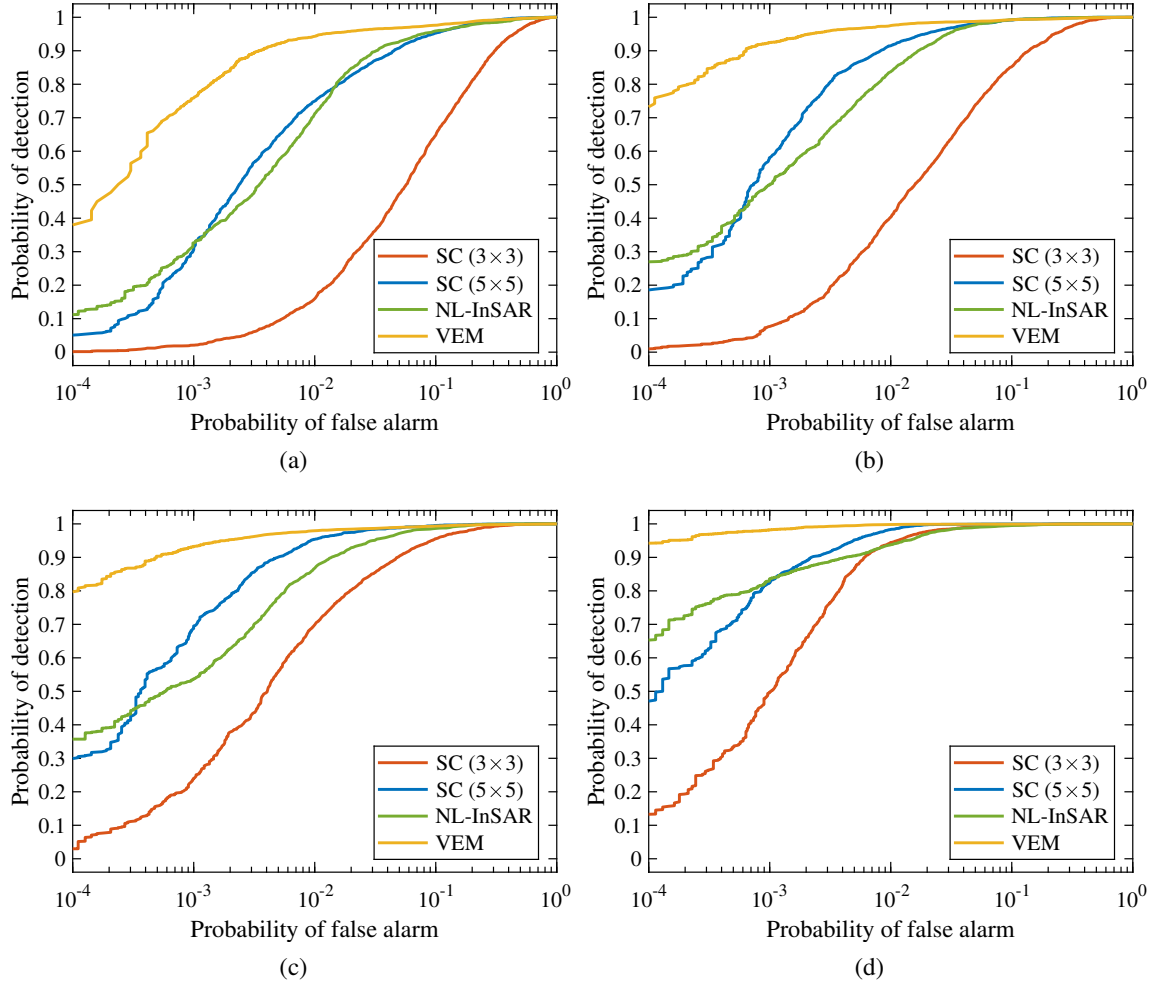


Figure 3.6: Receiver operating characteristic (ROC) curves for simulated data experiments. (a) Scenario 1 ($\gamma = 0.55$). (b) Scenario 2 ($\gamma = 0.65$). (c) Scenario 3 ($\gamma = 0.75$). (d) Scenario 4 ($\gamma = 0.85$).

of the change indicator variables. These results show the variational EM approach to improve upon the minimum error rates available from the coherence-based detectors.

Another series of experiments focuses on the impact of γ while keeping the other ground truth parameters constant. The change detection results computed using the compared methods on Scenario 4 with γ values of 0.55, 0.65, 0.75, and 0.85

Table 3.1: Error rates from the simulated data experiments using Scenarios 1–4. The error rates for the coherence estimates are computed using a threshold which minimizes the error rate.

	Scenario 1	Scenario 2	Scenario 3	Scenario 4
SC (3×3)	0.039	0.037	0.020	0.012
SC (5×5)	0.018	0.011	0.008	0.008
NL-InSAR	0.021	0.016	0.013	0.010
VEM	0.007	0.004	0.004	0.002

are shown in Fig. 3.7 with the minimum error rates reported in Table 3.2. Since the coherence quantifies the degree of pass-to-pass self-similarity between no-change pixels, we observe the expected result that all methods perform better for larger values of γ —with notably fewer numbers of false alarms. However, for all values of γ , the proposed VEM method appears to suppress substantially more false alarms while also faithfully representing fine details of the change regions. For example, with $\gamma = 0.55$, the “c”-shaped change region (identified by the red arrows in Fig. 3.7) is largely obscured by the window-based methods and NL-InSAR. However, the high contrast between the changed and unchanged regions of the VEM approach enables both fine localization of the change and identification of the change shape. Such high resolution change detection is important in forensic applications seeking to perform high resolution tracking or shape-based classification of a changed object (e.g., vehicle classification).

Next, the recovery of fine detail is evaluated by reusing the simulation parameters of Scenario 1 with the change map replaced by one consisting of square change regions with side lengths of ten pixels, placed randomly in the scene with at least a three pixel margin between them. The minimum error rates from the sample coherence using




Figure 3.7: Comparison of change detection results from Scenario 4 with γ values of 0.55, 0.65, 0.75, and 0.85 depicted in rows 1–4, respectively. The other simulation parameters are kept constant. The first column has the true change maps with change pixels in black, and the remaining columns show the change detection results from the compared methods. The last column shows the results of the proposed inference approach with black indicating a high probability of change.

various window sizes are plotted in Fig. 3.8 as the curve labeled “SC”. The x-axis corresponds to the number of pixels on each side of the square estimation window used to compute the sample coherence. The “U”-shaped error curve of the SC estimator highlights competing effects resulting from different window sizes. On the left of the figure, small window sizes produce poor coherence estimates, resulting in larger change

Table 3.2: Error rates from the simulated data experiments using Scenario 4 with various values for γ . The error rates for the coherence estimates are computed using a threshold which minimizes the error rate.

	$\gamma = 0.55$	$\gamma = 0.65$	$\gamma = 0.75$	$\gamma = 0.85$
SC (3×3)	0.062	0.047	0.027	0.012
SC (5×5)	0.023	0.013	0.009	0.008
NL-InSAR	0.021	0.017	0.015	0.010
VEM	0.011	0.005	0.003	0.002

detection error rates. Increasing the window size improves the coherence estimate; however, when the window becomes too large, it captures both change and non-change pixels. This degrades the coherence estimate and affects the ability to localize fine details, which contributes to the increased SC error rates seen for large window sizes on the right side of the figure. The error rate of the MAP estimate from the variational EM approach is shown to be lower than that of the sample coherence detectors using any of the window sizes considered in Fig. 3.8. This result demonstrates superior fidelity of the proposed approach to the structure of the underlying change maps used in these experiments without the need to select a window size.

3.4.2 Measured Data Experiments

Next we evaluate the performance of the proposed change detection algorithm using measured interferometric data from two SAR systems, one X-band and the other C-band. Unlike the simulated measurements, these data were produced under non-ideal conditions and lack precise pixel-level ground truth necessary for quantitative evaluation. Therefore, we limit the discussion to qualitative assessments of change detection performance.




Figure 3.8: Error rates of a sample coherence estimator plotted against window size. The error rates are computed using thresholds which minimize the error rates at each window size. The x-axis corresponds to the pixels per side of the square window used to compute the sample coherence. Also shown are the error rate from the NL-InSAR coherence estimates and the MAP change estimates from the proposed VEM approach. The underlying change map consists of randomly-placed square change regions, and the simulation assumes a coherence of $\gamma = 0.55$.

Gotcha Change Detection Data

We first evaluate the proposed change detection algorithm using a co-registered pair of measured HH-polarized SAR images from the X-band Gotcha change detection dataset [2]. The images have a sample spacing of 0.2 m, and they are from data that was collected on the same day and coherently aligned to a single reference. We use corresponding portions of the HH-polarization images FP1020 and FP1024 as the time-1 and time-2 images, respectively. The region observed in these sub-images is roughly that contained in the black rectangle in Fig. 3.9. The 256×200 pixel sub-images feature a paved path running along the right side of the scene and, to the



Figure 3.9: Optical satellite images [1] of a region observed in the Gotcha change detection challenge problem dataset [2]. (a) The observed region in the measured data experiment is shown framed by a black box. (b) The hinged grate shown here is identified by the blue arrow in (a).

left of it, a grassy area where foot traffic is known to have occurred in the interim between the two collections. The time-1 and time-2 magnitude images are shown in Fig. 3.10, from which the foot traffic is largely unobservable.

The change detection results from the compared methods are shown in Fig. 3.11. Each of the coherence estimates show the path to be a low coherence region, evidence of a temporal disturbance—and thus a false alarm, since actual changes were not present there. According to the proposed approach, the region occupied by the hinged metal grate, identified by the blue arrow in Fig. 3.9, is the only spot along the path that has a high probability of change. This hinged grate extends roughly the width




Figure 3.10: Magnitude images from the Gotcha dataset [2]. (a) Time-1 (FP1020). (b) Time-2 (FP1024).

of the path as shown in Fig. 3.9. The proposed method departs from the coherence estimates, in that it properly does not indicate change along the rest of the path. Each method shows clear evidence of the foot traffic in the grass, though the proposed VEM method results in a notably more contiguous change region about the track.

Sentinel-1 Change Detection Data

We also compare the change detection methods using a pair of SAR images formed with data from the C-band SAR systems of the Sentinel-1 constellation of satellites [5]. We use the Sentinel-1 images to detect the effects of a wildfire in the southern foothills of the Troodos Mountains in Cyprus that began on July 3, 2021.⁴ The images had a pixel spacing of 2.3 m in range and 14.0 m in azimuth before they were downsampled by a factor of two in range. The time-1 image in the pair was collected before the

⁴<https://emergency.copernicus.eu/mapping/list-of-components/EMSR515>




Figure 3.11: Change detection results from the X-band Gotcha dataset. (a) Sample coherence with 3×3 window. (b) Sample coherence with 5×5 window. (c) NL-InSAR coherence estimate. (d) Proposed VEM inference approach with black indicating a high probability of change.

fire, on June 30, 2021, and the time-2 image after, on July 6, 2021. The impacted area, as determined by optical image analysis [3], is shown in yellow on the map in Fig. 3.12.




Figure 3.12: Optical image of the Troodos Mountains in Cyprus with a yellow overlay depicting the area impacted by a wildfire that began on July 3, 2021 [3,4].




Figure 3.13: Change detection results from C-band Copernicus Sentinel SAR data [5] collected over the region of Cyprus shown in Fig. 3.12. (a) Sample coherence with 3×3 window. (b) Sample coherence with 5×5 window. (c) NL-InSAR coherence estimate. (d) Proposed VEM inference approach with black indicating a high probability of change.

The change detection results from the Cyprus Sentinel-1 data are shown in Fig. 3.13. The results from each method depict the extent of the burned area to be broadly in agreement with the estimate from the optical data. As seen in the earlier experiments, the proposed variational EM method exhibits greater contrast between the change and no-change regions. The results from NL-InSAR and the proposed approach both indicate small regions free of temporal change from within the burn region shown in Fig. 3.12, though, without precise ground truth knowledge of the temporal changes, it is difficult to speculate on the accuracy of these departures from the optical data.

3.5 Conclusion

3.5.1 Summary

In this chapter, we have introduced a Bayesian approach to SAR change detection that enables fine-resolution detection of temporal disturbances between pairs of co-registered SAR images. The proposed approach assigns prior distributions to the unobserved model variables that exploit spatial structure both in the scattering qualities of the scene and among the scene disturbances that occur between the passes. The proposed modeling advances the use of gamma MRFs for SAR change detection, using them to capture positive spatial correlations in the mean backscatter power elements. Further, a 2-D polynomial regression is used to account for the spatially nonhomogeneous interferometric phase.

The prior models are incorporated into a Bayesian framework, on which variational EM is used to approximate the posterior distribution of the latent variables and the prior model hyperparameters. In the Bayesian approach, the posterior probability of

change is directly inferred, rather than thresholding the coherence estimates obtained by local, or nonlocal, spatial averaging. To optimize the gamma MRF coupling hyperparameter, variational inference is used to efficiently approximate expectations in the contrastive divergence gradient.

The proposed approach does not suffer from losses of spatial resolution and fine-scale details as do conventional window-based CCD methods. In experiments using both simulated data and measured interferometric SAR image pairs, the new approach is shown to yield improved detection performance compared to the conventional window-based statistics used for SAR CCD.

3.5.2 Future Work

The approach of this chapter may be extended to form a joint imaging and change detection framework that is intended to mitigate the effects of interrupted data collections, as in [63]. A Bayesian approach similar to that proposed here could be used to compensate for missing data via prior distributions on the change variables, the scene reflectance, and the interferometric phase.

Chapter 4: Polarimetric SAR Despeckling

4.1 Introduction

Polarimetric SAR (PolSAR) systems use polarization diversity in the emitted and received electromagnetic waves to provide data that are considerably more informative than single-channel collections as the inter-channel relationships are informative of physical scattering behavior [16, 117, 118].⁵ Multichannel radar imaging modalities such as PolSAR have proven beneficial for tasks such as terrain classification [120], target recognition [10], and surface parameter estimation [121, 122].

The speckle phenomenon lends a noise-like, granular character to imagery which can impede interpretation and the completion of automated tasks [13]. The reduction of speckle, i.e. despeckling, improves both the visual quality of the image and the prospects for tasks such as segmentation and target recognition based on the data [36].

It is often necessary to suppress speckle before SAR data are considered suitable for analysis [16], a prerequisite realized by the application of a speckle filter, classically one which acts on the image locally [17]. Notable SAR despeckling algorithms have been proposed using local filtering [35, 76, 123], patch-based filtering [38, 39, 77, 124], region-growing [125], and variational methods [42–44].

⁵This chapter is largely drawn from [119], coauthored with Lee C. Potter (© 2022 IEEE).

Deep learning techniques that involve the use of hierarchical structures to extract feature representations from data have been responsible in recent years for state-of-the-art performance in numerous image processing tasks. Convolutional neural networks (CNN), flexible deep learning architectures composed of cascades of convolutions and other operations, have been used with great success in tasks such as super-resolution [126, 127], object detection [128, 129], and denoising [130, 131]. Since 2017, there has been a trend towards deep learning in single-channel SAR despeckling, e.g. [132–135]. CNNs have been applied to PolSAR data tasks [136] such as super-resolution [137, 138], land use classification [139, 140], and target recognition [141, 142]. Yet, there have been few CNN-based methods for PolSAR speckle reduction [143–146].

The need for large datasets in deep learning presents a hindrance to SAR despeckling applications because the speckle-free images which might be used in training are not observable. The issue has been managed in single-channel SAR despeckling by using optical images as clean intensity data [133, 134], a practice which does not have an obvious equivalent for PolSAR images. Temporal multilooking has also been used to approximate speckle-free intensity images [132, 147, 148]. However, the series of co-registered multi-temporal images used in these approaches may not be readily available in the quantity necessary for deep learning, particularly from fully polarimetric collections. Other approaches eschew the need for clean training data and instead learn to complete the despeckling task in a self-supervised manner using only noisy images [149–151].

PolSAR despeckling is further complicated by the nature of the data as images of complex-valued Hermitian positive definite polarimetric covariance matrices. Despeckling methods must respect the Hermitian positive definite constraint and manage

the complex-valued off-diagonal elements of the polarimetric covariance matrices. The phases of the off-diagonal elements hold valuable information, and the polarimetric characteristics that they capture ought to be preserved by any filtering process.

The approach presented in this chapter extends the single-channel SAR CNN-based speckle filtering in [132] to PolSAR despeckling, yielding a deep learning framework which estimates from PolSAR data the underlying polarimetric covariance matrices. With a residual learning approach [83, 152], a CNN recovers not the filtered image itself, but the speckle component which is to be removed from the corrupted image. Drawing from [42], a bijective transform enables real-valued CNN processing such that the network outputs represent valid covariance estimates. As training data, we use spatially-multilooked PolSAR images in lieu of unobservable ground truth covariance data and simulate speckle to form corrupted training images. We use both simulated and measured polarimetric SAR data to evaluate the proposed approach and introduce a suite of quantitative indexes to evaluate the competing objectives of speckle suppression and preservation of scene features and polarimetric information. The proposed approach markedly reduces speckle and faithfully preserves fine structural detail in comparison to state-of-the-art methods. A further benefit of the proposed approach is the fast running time due to the use of CNNs.

The remainder of the paper is organized as follows. Section 4.2 provides background information and a review of the relevant literature. Section 4.3 is a summary of the proposed despeckling method. Section 4.4 outlines an experimental procedure to compare the proposed method to existing methods. Section 4.5 presents and discusses the experimental results. Section 4.6 gives concluding remarks.

4.2 Background

4.2.1 Principles of Polarimetric SAR

The single-look scattering properties observed by a polarimetric SAR system at a given pixel can be described by a 2×2 complex scattering matrix given by

$$\mathbf{S} = \begin{bmatrix} S_{hh} & S_{hv} \\ S_{vh} & S_{vv} \end{bmatrix}, \quad (4.1)$$

where S_{hv} denotes the scattering element corresponding to the horizontal transmitting and vertical receiving polarizations. For a monostatic radar operating in a reciprocal medium, \mathbf{S} is symmetric and can be vectorized according to the lexicographic or Pauli basis as

$$\mathbf{s} = [S_{hh} \quad \sqrt{2}S_{hv} \quad S_{vv}]^T \quad (4.2)$$

$$\boldsymbol{\kappa} = \frac{1}{\sqrt{2}} [S_{hh} + S_{vv} \quad S_{hh} - S_{vv} \quad 2S_{hv}]^T, \quad (4.3)$$

respectively, wherein the superscript T denotes a matrix transpose. When the scattering medium is rough on the scale of the radar's central wavelength, it is common to model a scattering vector as following a complex zero-mean circularly-symmetric Gaussian distribution [13], given by

$$p(\mathbf{s} \mid \boldsymbol{\Sigma}) = \frac{1}{\pi^3 |\boldsymbol{\Sigma}|} \exp \{ -\mathbf{s}^H \boldsymbol{\Sigma}^{-1} \mathbf{s} \}, \quad (4.4)$$

in which $|\cdot|$ denotes a determinant. In (4.4), $\boldsymbol{\Sigma} \in \mathbb{H}_+^3$ is the covariance matrix $\boldsymbol{\Sigma} = \mathbb{E} [\mathbf{s} \mathbf{s}^H]$ with the superscript H denoting a complex conjugate transpose and \mathbb{H}_+^D denoting the set of $D \times D$ Hermitian positive definite matrices. The equivalent for the Pauli scattering vector is the coherency matrix, $\mathbf{T} = \mathbb{E} [\boldsymbol{\kappa} \boldsymbol{\kappa}^H]$. The stochastic scattering model in (4.4), often referred to as fully-developed speckle, accounts for the

speckle phenomenon in coherent imaging. When the model is valid, the covariance matrix $\mathbf{\Sigma}$ contains the important polarimetric information of the imaged surface.

Multilook PolSAR observations can be represented using the L-look sample covariance matrix given by

$$\mathbf{Z} = \frac{1}{L} \sum_{t=1}^L \mathbf{s}(t) \mathbf{s}^H(t) \quad (4.5)$$

where $\mathbf{s}(t)$ for $t = 1, \dots, L$ are independent realizations of the scattering vector. These multiple “looks” are most often taken from a small window centered upon the pixel of interest. The goal of PolSAR despeckling is to recover the unobserved covariance $\mathbf{\Sigma}$ from the sample covariance matrix \mathbf{Z} .

Under the fully-developed speckle model, the L-look sample covariance matrix follows a complex Wishart distribution, i.e. $\mathbf{Z} \sim \text{CW}(\mathbf{Z}; L^{-1}\mathbf{\Sigma}, L)$, where the complex Wishart density is defined for $\mathbf{X} \in \mathbb{H}_+^D$ as [153]

$$\text{CW}(\mathbf{X}; \mathbf{G}, v) = \frac{|\mathbf{X}|^{v-D} \exp\{-\text{tr}(\mathbf{G}^{-1}\mathbf{X})\}}{|\mathbf{G}|^v \tilde{\Gamma}_D(v)} \quad (4.6)$$

with scale matrix \mathbf{G} , degrees of freedom v , $\text{tr}(\cdot)$ denoting the trace function, and $\tilde{\Gamma}_D(\cdot)$ the complex multivariate gamma function.

4.2.2 Polarimetric SAR Data Notation

In describing the statistics of PolSAR, we have considered the behavior of a single pixel. We now adopt a notation to describe a covariance image wherein a D-channel $N_1 \times N_2$ pixel covariance image is represented by a fourth-order tensor $\mathbf{Z} \in \mathbb{C}^{D \times D \times N_1 \times N_2}$. We denote by $\mathbf{Z}_i \in \mathbb{H}_+^D$ the covariance matrix with pixel indices $i = (i_1, i_2)$, and denote by $\mathbf{Z}^{j_1 j_2} \in \mathbb{C}^{N_1 \times N_2}$ the spatial image corresponding to the (j_1, j_2) th entry of the covariance matrix.

4.2.3 Related Work

PolSAR despeckling has been an active area of research since the early 1990s. The early despeckling methods relied on local window filtering, in which each element of the despeckled image is estimated from the values of nearby pixels within a small window. The simplest and most common such technique is the boxcar filter, a convolutional filter with a small window of equally-weighted elements. In 1983, the sigma filter was developed by Lee [35], offering a simple and effective despeckling approach for single-channel SAR. The sigma filter operates by averaging nearby pixels, excluding those differing significantly in intensity. The improved sigma filter was developed to rectify deficiencies related to bias and blurring in its predecessor [154] and later was extended for use with PolSAR data [123].

Vasile *et al.* [125] developed the intensity-driven adaptive neighborhood (IDAN) filter, which uses a region-growing approach to group pixels belonging to the same statistical populations. Adaptive neighborhoods are formed with respect to the intensity information (the diagonal terms of the coherency matrix), and grouped pixels are given equal weight in the estimation of the coherency matrix at each pixel.

In recent years, patch-based methods have been used in PolSAR despeckling. Rather than selecting pixels with similar intensities, such methods compare the patches about pixels to determine relative importance. The nonlocal means (NLM) filter, developed by Buades *et al.* [37] for the denoising of natural images, is responsible for popularizing these methods, and patch-wise descendants of the NLM filter have proven effective in PolSAR despeckling [38, 39, 124].

Variational methods, which pose the estimation as an optimization problem, are also used for speckle reduction. Some variational despeckling methods, such as [41],

process only the diagonal elements of the coherency or covariance matrices, ignoring the information in the cross-channel correlations. Other works formulate an objective function which operates on the full covariance matrix, such as the variational model in [43] which introduces a matricial total variation (TV) prior term to encourage spatial smoothness. This approach is augmented in [44] with a nonlocal regularization term leading to improved despeckling performance. Like [43, 44], the proposed approach operates on the full polarimetric covariance matrix so as to consider the polarimetric information of the off-diagonal elements.

Another variational despeckling methodology, the multichannel logarithm with Gaussian denoising (MuLoG) framework [42], facilitates the use of single-channel denoisers designed to remove additive white Gaussian noise (AWGN). This is a considerable benefit as the bulk of work in image denoising has been dedicated to the removal of AWGN.

Considering the success of deep learning in natural image denoising [83, 130], its application to SAR despeckling is natural, e.g. [132–134, 147]. Among them, SAR-CNN in [132] adapts the denoising CNN (DnCNN) framework [83] to single-channel SAR data by way of a log transform. A residual learning strategy [152] is used that treats the output of the CNN as a residual image, i.e. the difference between the noisy observation and the latent clean image. Residual learning and batch normalization [155] have been shown to both accelerate training and improve denoising performance on natural images [83].

Yet there are comparatively few CNN-based despeckling methods for PolSAR. The MuLoG [42] framework allows for the straight-forward application of deep learning to multichannel SAR despeckling through the techniques developed for AWGN




Figure 4.1: Proposed despeckling framework.

denoising of natural images, e.g. [83, 130, 131]. Indeed, MuLoG has been studied with deep learning-based denoisers trained on both natural images [156] and measured SAR images [145]. A related study [147] compares the performance of SAR-CNN to MuLoG when paired with either natural image-based DnCNN denoisers or DnCNN denoisers trained on SAR images. In [143], an iterative approach to PolSAR despeckling is studied that uses a CNN architecture intended for image super-resolution. This approach is limited in that it only considers the diagonal elements of the covariance, ignoring the information of the off-diagonal elements. That shortcoming is addressed in a revised approach [144] which uses a CNN to derive speckle correction terms based on all covariance elements.

4.3 Proposed Despeckling Approach

In this section, we summarize a proposed CNN-based approach to polarimetric SAR despeckling that operates in the matrix logarithm domain. The proposed approach, referred to as “MS-CNN”, is depicted graphically in Fig. 4.1.

4.3.1 Approach Details

The PolSAR despeckling problem is complicated by the Hermitian positive definite constraint on $\mathbf{\Sigma}$, which for variational methods requires a potentially challenging optimization over the open cone of Hermitian positive definite matrices. This difficulty can be mitigated, as in [42], with a matrix logarithm transform, defined for $\mathbf{A} \in \mathbb{H}_+^D$ having eigendecomposition $\mathbf{A} = \mathbf{U} \text{diag}(\lambda_1, \dots, \lambda_D) \mathbf{U}^H$ by

$$\log \mathbf{A} = \mathbf{U} \text{diag}(\log \lambda_1, \dots, \log \lambda_D) \mathbf{U}^H. \quad (4.7)$$

Note that when applied to an image \mathbf{Z} , the $\log(\cdot)$ operator denotes pixel-wise application of the matrix logarithm to each \mathbf{Z}_i .

The matrix logarithm of a Hermitian positive definite covariance matrix is a Hermitian matrix. By filtering in the matrix log domain, we need only ensure that the output of the CNN remains Hermitian, which can be accomplished by operating on only the non-redundant elements of the matrix log-transformed covariance matrices. To do so, we define $\mathcal{V}(\cdot) : \mathbb{C}^{3 \times 3 \times N_1 \times N_2} \rightarrow \mathbb{R}^{9 \times N_1 \times N_2}$, a bijective operator which maps a complex fourth-order tensor to a nine channel real-valued image. The channels of the reparameterized image correspond to the non-redundant real and imaginary elements of a 3×3 Hermitian matrix.

As depicted in Fig. 4.1, the input is an L-look sample covariance image of $N_1 \times N_2$ pixels, \mathbf{Z} . At each pixel, \mathbf{Z}_i represents a corrupted measurement of the unobserved covariance $\mathbf{\Sigma}_i$. For $L < 3$, it is necessary to ensure \mathbf{Z}_i is full rank at each pixel, for which we adopt a technique used in [42] involving a minor rescaling of the off-diagonal elements. The covariance image undergoes a matrix log transformation at each pixel and is reparameterized using $\mathcal{V}(\cdot)$. The transformed image is then fed to

the CNN, and the output of the CNN is interpreted as a multichannel residual image. A despeckled estimate of the covariance image is given by subtracting the residual in the matrix log domain:

$$\hat{\Sigma} = \exp \{ \mathcal{V}^{-1} (\mathcal{V}(\log \mathbf{Z}) - f_{\boldsymbol{\theta}}(\mathcal{V}(\log \mathbf{Z}))) \}, \quad (4.8)$$

where $f_{\boldsymbol{\theta}}(\cdot)$ is a CNN using the parameter set $\boldsymbol{\theta}$. With the matrix logarithm and this reparameterization, we ensure that the despeckled covariance matrix estimates are Hermitian positive definite.

To train the network, we assume the availability of a suitable proxy for speckle-free ground truth covariance data. A training pair (\mathbf{x}, \mathbf{y}) is given by $\mathbf{x} = \mathcal{V}(\log \Sigma)$ and $\mathbf{y} = \mathcal{V}(\log \mathbf{Z})$, where \mathbf{Z} is an L-look sample covariance image drawn using the underlying covariance image Σ according to (4.4) and (4.5).

4.3.2 Network Architecture

The CNN architecture depicted within the dotted lines in Fig. 4.1 is that of [83], where the CNN is used to denoise natural images. The first layer is composed of 64 filters of size $3 \times 3 \times 9$. Next is a layer of rectified linear units (ReLUs) followed by 17 repetitions of convolutional layers with 64 filters of size $3 \times 3 \times 64$, batch normalization layers, and ReLUs. The final convolutional layer has nine filters of size $3 \times 3 \times 64$, leading to a nine channel output image. The use of 17 hidden layers was adopted here for the PolSAR despeckling problem; further expansion or contraction of the network failed to improve despeckling performance in our experiments.

We use the following loss function to form the deep learning objective

$$\ell(\boldsymbol{\theta}) = \sum_{m=1}^M \left\| \text{vec}(f_{\boldsymbol{\theta}}(\mathbf{y}^{(m)}) - (\mathbf{y}^{(m)} - \mathbf{x}^{(m)})) \right\|_1, \quad (4.9)$$

optimized with respect to the CNN parameter set θ , with the sum in (4.9) over M pairs from the training set. In our experiments, we found that an ℓ_1 loss function provided improved training performance over the ℓ_2 loss used with optical images in [83]; similar results were reported in [147] for single-channel SAR data.

4.3.3 Approach Summary

Thus, to extend deep learning to PolSAR despeckling, we adopt the residual network architecture used for natural images in [83] and single-channel SAR in [132]. In order to estimate the complex-valued covariance matrices, we employ the matrix logarithm and a reparameterization of the covariance matrix to an unconstrained set of real-valued parameters. Simulated speckle [13] is used to construct a corpus of training images. And, to better conform to the physics of the reparameterized variables, we use an ℓ_1 loss in training the network.

4.4 Experimental Procedure

In this section we detail experimental procedures to implement the algorithm described in Section 4.3 and compare its despeckling performance to existing methods. We also introduce a set of indexes for the quantitative assessment of polarimetric despeckling performance.

4.4.1 Training

In place of ground truth covariance information in the training set, we use PolSAR data from the L-band UAVSAR sensor [6, 157], images which have undergone a 36-look spatial averaging and equivalent downsampling. The noncoherent averaging coarsens the resolution in a compromise that reduces speckle so as to approximate the

unobservable ground truth covariance data. This method to generate clean covariance images relies on two assumptions which are likely to be violated at times: neighboring pixels combined by multilooking share statistical properties, and the neighboring observations are conditionally independent given the true covariance [158, 159]. This approach, while convenient and simple, may in its coarsening of image resolution diminish the value of the data for the most demanding of applications [151]. Further, the fundamental problem remains that we rely on an approximation to truly speckle-free images. We use a training set of 91490 such ground truth covariance images of size 40×40 that have been excised from 14 different collections. Ground truth images from two separate collections are later used in model validation.

The corrupted images used for training are L -look sample covariance images that are formed from clean images according to (4.4) and (4.5). For each number of looks, a separate CNN is trained using 60 epochs of stochastic gradient descent with weight decay of 10^{-3} , momentum of 0.9, batch size of 64, and a learning rate that begins at 0.1 and reduces to 10^{-4} during training. The network weights are initialized as in [160].

4.4.2 Conventional Despeckling Methods

We compare the proposed method to three well-established speckle filters: the simple boxcar filter [73], the extended sigma filter [123], and the IDAN filter [125]. The tuning parameters for these despeckling methods were selected as follows. For the boxcar filter, we use a 5×5 window. Two parameters are used to tune the performance of the extended sigma filter: the sigma value ξ , which controls the range in which pixels are considered similar, and a window size. In our experiments, we

use $\xi = 0.7$ for single-look data and $\xi = 0.9$ for four-look data, as is suggested in [123]. In either case, we use a filtering window of size 9×9 . The IDAN filter uses a tuning parameter which determines the maximum number of pixels included in a neighborhood, for which we select $N_{\max} = 50$ as in [125].

4.4.3 MuLoG Despeckling Framework

We also compare results from the MuLoG framework [42], which represents the state-of-the-art in multichannel SAR despeckling. The MuLoG framework poses a maximum *a posteriori* (MAP) optimization problem that is solved by the plug-and-play variant of the alternating direction method of multipliers (ADMM) [56], an approach which facilitates the use of single-channel AWGN denoising operators. In our experiments, we use the MATLAB implementation⁶ of MuLoG and adopt three choices for the denoising operator: the familiar block-matching 3D (BM3D) filter [74] used with MuLoG in [42], a DnCNN denoiser [83] trained on natural images, and a DnCNN denoiser trained on SAR images. We set the tuning parameter in MuLoG to $\beta = 1 + 2/L$ and use six iterations of ADMM, as suggested in [42].

For the DnCNN trained on natural images, we use a set of 14 pre-trained models from <https://github.com/cszn/DnCNN>; these have been trained for noise standard deviations of $\chi = \{\frac{10}{255}, \frac{15}{255}, \dots, \frac{75}{255}\}$ on sets of natural images [83]. The range of the training images is $[0, 1]$, and, due to the nonlinearity of DnCNN, we use a rescaling operator to match this range, defined by

$$\Psi_{v_0}(v) = \frac{v - \mathcal{R}_{0.3\%}(v_0)}{\mathcal{R}_{99.7\%}(v_0) - \mathcal{R}_{0.3\%}(v_0)}, \quad (4.10)$$

⁶<https://www.charles-deledalle.fr/pages/mulog.php>

where $\mathcal{R}_q(\cdot)$ returns the q^{th} quantile of its input, and we have used the quantile values suggested in [156]. We use as the denoising operator a function composition that for input \mathbf{v} is equal to

$$\Psi_{\mathbf{v}}^{-1}(\mathbf{g}_{\sigma^2}(\Psi_{\mathbf{v}}(\mathbf{v}))) \quad (4.11)$$

where \mathbf{g}_{σ^2} denotes a denoising operator for noise power σ^2 with σ selected according to the tuning parameter β as

$$\sigma = \arg \min_{\sigma \in \chi} \left| \frac{\beta^{-1/2}}{\mathcal{R}_{99.97\%}(\mathbf{v}) - \mathcal{R}_{0.3\%}(\mathbf{v})} - \sigma \right|. \quad (4.12)$$

We refer to MuLoG used with the pre-trained DnCNNs as “MuLoG-CNN.”

We also consider in the MuLoG framework a DnCNN trained with images from the training set used for the proposed approach. To form ground truth images appropriate for use with MuLoG, each patch is matrix log-transformed, decorrelated, then each channel is rescaled using (4.10), yielding nine single-channel images from each patch. Models are trained using this data for additive noise with standard deviations of $\{0.04, 0.08, \dots, 0.24\}$. We use 50 epochs of stochastic gradient descent with weight decay of 0.001, a momentum of 0.9, a batch size of 64, and a learning rate beginning at 0.1 and reducing to 10^{-4} during training. The denoisers trained on PolSAR data can then be applied in the same manner as those trained on natural images, i.e. as described in (4.11) and (4.12). We refer to MuLoG with the DnCNNs trained in this manner as “MuLoG-CNN-P.”

4.4.4 PolSAR Despeckling Performance Assessment

We describe in this section methods used to assess despeckling performance with consideration for speckle reduction, preservation of scene features, and preservation of polarimetric information. Assessment indexes are defined with respect to a covariance

image $\widehat{\Sigma}$ of $N_1 \times N_2$ pixels that has been estimated from an L-look sample covariance image \mathbf{Z} with reference to a true underlying covariance image Σ . The symbol η is used to denote the set of pixel locations in the image.

Loss Functions

We use various loss functions in the simulated data experiments to assess the fidelity of the estimated covariance image to the true underlying image and the preservation of polarimetric information. The first is the normalized mean square error (NMSE), given by

$$\text{NMSE} = 10 \log_{10} \left(\frac{\sum_{i \in \eta} \|\widehat{\Sigma}_i - \Sigma_i\|_F^2}{\sum_{i \in \eta} \|\Sigma_i\|_F^2} \right), \quad (4.13)$$

with $\|\cdot\|_F$ denoting a Frobenius norm. We also use a quadratic loss (QL) given by [67]

$$\text{QL} = \max_{i \in \eta} \|\widehat{\Sigma}_i \Sigma_i^{-1} - \mathbf{I}_3\|_2, \quad (4.14)$$

where $\|\cdot\|_2$ denotes a spectral norm. Note that both the NMSE and quadratic loss depend on all elements of the polarimetric covariance matrix. Additionally, the structural similarity (SSIM) index [161] is computed on the images corresponding to the log of the diagonal covariance matrix elements. The average SSIM of the images corresponding to the three intensity elements is recorded.

Whitened Covariance Image Evaluation

It is common in the despeckling literature to analyze the image formed by the element-wise ratio between the speckled single-channel image and the filtered image. The analysis of such ratio images is usually limited to an assessment of visible structure, the presence of which indicates an undue erosion of the underlying image

structure. For a multichannel analogue to the ratio image, we adopt a whitened image \mathbf{W} , formed as

$$\mathbf{W}_i = \hat{\Sigma}_i^{-1/2} \mathbf{Z}_i \hat{\Sigma}_i^{-1/2} \quad \text{for } i \in \eta. \quad (4.15)$$

For a scene with fully-developed speckle, a faithfully recovered covariance image results in \mathbf{W}_i that are independently and identically distributed as

$$\mathbf{W}_i \sim \text{CW}(\mathbf{W}_i; L^{-1} \mathbf{I}_3, L), \quad (4.16)$$

for $L \geq 3$. We introduce indexes regarding the whitened image which are intended to detect erosion of structure in the underlying image, providing an enrichment to the qualitative analysis of the ratio image.

The first is a Kolmogorov-Smirnov test computed on \mathbf{W}^{jj} for $j \in \{1, 2, 3\}$. For fully-developed speckle, each diagonal element of the whitened image ought to independently and identically follow a gamma distribution. The Kolmogorov-Smirnov test compares the empirical distribution function of a sample with a theorized cumulative distribution function, with a smaller Kolmogorov-Smirnov statistic (KSS) indicating greater correspondence.

Second, we use a modified Box-Pierce test [162], given by

$$\text{BP2} = |\eta| \left(-1 + \sum_{d_1=0}^K \sum_{d_2=0}^K \rho_v^2(d_1, d_2) \right), \quad (4.17)$$

to detect autocorrelations in the whitened image with $\rho_v(d_1, d_2)$ denoting the normalized sample autocorrelation function for image \mathbf{v} at lag (d_1, d_2) . When the fully developed speckle hypothesis is satisfied, the whitened image should be spatially decorrelated. The purpose of this test is to detect degradations of the spatial structure of the image from the despeckling process. In experiments, we use a maximum lag of $K = 5$.

In simulated data experiments, adherence to the speckle model is ensured by construction, and these whitened image metrics are valid everywhere. In the measured data experiments, we limit the quantitative analysis to a region of the image which appears to lack man-made media and is more likely to produce fully-developed speckle. In quantitative assessments, we report the KSS and BP2 results corresponding to \mathbf{W}^{11} .

Incoherent Decomposition Parameters

The entropy (H), anisotropy (An), and mean alpha angle ($\bar{\alpha}$) are parameters related to the eigenvalues and eigenvectors of the coherency matrix [163]. These parameters are sensitive to the type of scattering mechanism present within the resolution cell and are commonly used to compare PolSAR despeckling performance [17, 43, 125]. In the simulated data experiments, we report the relative absolute error between the true values and the estimates of these parameters as in [17, 43], using the median as an indicator of aggregate performance across the image.

Polarization Signatures

PolSAR data admit the synthesis of scattering coefficients corresponding to an arbitrary pair of transmit and receive polarizations. This allows the plotting of polarization signatures, 3D visualizations of the normalized scattering power density as a function of the polarizations of the transmitting and receiving antennas [164]. Because a visualization of all possible receive polarizations for each transmit polarization is difficult, it is common to plot either a co-polarization signature, in which the receive polarization is identical to the transmit polarization, or a cross-polarization

signature, in which the two are orthogonal [16, 17, 43]. We use both such plots in comparisons using simulated PolSAR data.

Measured Data Assessment

Next we list the performance indexes used exclusively on the measured PolSAR data, for which ground truth reference is unavailable. The first is the equivalent number of looks (ENL), which assesses suppression of speckle in homogeneous regions. We use a trace moment-based ENL estimator designed for PolSAR images [165]:

$$\text{ENL} = \frac{[\text{tr}(\langle \hat{\Sigma} \rangle)]^2}{\langle \text{tr}(\hat{\Sigma} \hat{\Sigma}) \rangle - \text{tr}(\langle \hat{\Sigma} \rangle \langle \hat{\Sigma} \rangle)}, \quad (4.18)$$

where $\langle \cdot \rangle$ denotes a spatial ensemble average.

We use the edge-preservation detection ratio of average (EPD-ROA) [76] to assess edge-preservation ability, defined for each j^{th} diagonal entry of the covariance matrix as

$$\text{EPD-ROA} = \frac{\sum_{i \in \eta} |\hat{\Sigma}_i^{jj} / \hat{\Sigma}_{i+d}^{jj}|}{\sum_{i \in \eta} |\mathbf{Z}_i^{jj} / \mathbf{Z}_{i+d}^{jj}|} \quad (4.19)$$

with lag $d = (d_1, d_2)$. EPD-ROA values closer to one indicate better edge preservation.

We use target-to-clutter ratio (TCR) to assess the preservation of point targets during despeckling [16]. The TCR is computed over a patch containing a point scatterer amongst homogeneous clutter, and it measures the absolute difference in decibels before and after despeckling of the intensity ratio between the strong scatterer and the surrounding area. Finally, we compute the mean of ratio (MOR), which is a ratio between the mean intensity in a homogeneous area before and after despeckling. A MOR close to one indicates good radiometric preservation. The TCR, EPD-ROA, and MOR results are reported as averages over the three intensity channels.

4.5 Results

In this section, we discuss experimental results computed using both simulated and measured polarimetric SAR data which are intended to compare the proposed method to existing despeckling methods.

4.5.1 Simulated Data Experiments

We first study the performance of the compared methods on a set of 18 ground truth covariance images, each formed by the same multilooking procedure used to generate training data. These data comprise a set of 256×256 images from collections which were not used for training. From the ground truth covariance images, we simulate speckle according to (4.4) and (4.5) using $L = 1$ and $L = 4$. For each number of looks, each clean image was used to generate 10 speckled images used in experimental trials. Experiments were computed in MATLAB R2020a with a 3.2 GHZ Intel Core i5-6500 CPU and a Nvidia GTX 1060 6GB GPU. Training and application of CNNs were carried out with the MatConvNet toolbox [166]. The MuLoG codes were run in parallel using four CPU cores.

In Fig. 4.2 the despeckled results computed from a simulated single-look ($L = 1$) image are displayed in an RGB representation based on the Pauli decomposition. The image depicting the ground truth is shown in the center panel along with the original data and the results from the compared methods. A visual analysis shows that the boxcar and IDAN filters reduce speckle in homogeneous regions yet fail to preserve fine structural detail, while the extended sigma filter appears to better preserve edges and point targets. The three MuLoG results are qualitatively similar to each other with the BM3D variant suffering from over-smoothing and poorer preservation of point

targets. Each of the CNN-based methods result in an apparently faithful restoration of the ground truth covariance image, with the proposed approach providing a better preservation of texture.

The proposed approach outperforms the compared methods on both the single-look and four-look data according to each of the loss index results reported in Table 4.1 and Table 4.2. Though operating with a running time similar to the extended sigma and IDAN filters, MS-CNN provides greater than 3 dB improvement in average NMSE over either in both the $L = 1$ and $L = 4$ simulations. The MuLoG variants provide comparable performance to MS-CNN in terms of the loss indexes but are more computationally intensive. The quantitative results regarding the incoherent decomposition parameters of entropy, anisotropy, and alpha angle are shown in Table 4.3. In these assessments of polarimetric information preservation, the proposed method is likewise shown to perform best among the compared methods.

Table 4.1: Quantitative despeckling assessment results from simulated single-look ($L = 1$) PolSAR data. The mean results over ten independent realizations of the speckle process from each of 18 ground truth covariance images of 256×256 pixels are shown with standard deviations in gray.

	NMSE (dB)	QL	SSIM
Original	0.92 \pm 0.7	33.5 \pm 17.6	0.574 \pm 0.10
Boxcar	-3.54 \pm 1.2	181.0 \pm 231.6	0.550 \pm 0.15
Ext. sigma	0.44 \pm 0.7	91.6 \pm 126.7	0.664 \pm 0.09
IDAN	-2.57 \pm 0.9	15.2 \pm 8.2	0.568 \pm 0.12
MuLoG-BM3D	-5.36 \pm 1.1	24.9 \pm 11.4	0.707 \pm 0.11
MuLoG-CNN	-5.23 \pm 1.1	27.2 \pm 12.2	0.722 \pm 0.10
MuLoG-CNN-P	-5.28 \pm 1.1	27.1 \pm 13.1	0.727 \pm 0.10
MS-CNN	-5.83 \pm 1.0	15.0 \pm 8.0	0.762 \pm 0.08




Figure 4.2: Comparison of despeckling algorithms using a 256×256 pixel simulated single-look PolSAR image [6]. The results are displayed in an RGB representation based on the Pauli decomposition. (a) Original data. (b) Boxcar. (c) Extended sigma. (d) IDAN. (e) Ground truth. (f) MuLoG-BM3D. (g) MuLoG-CNN. (h) MuLoG-CNN-P. (i) MS-CNN.




Figure 4.3: Whitened images from the simulated single-look PolSAR data. (a) Boxcar. (b) Extended sigma. (c) IDAN. (d) MuLoG-BM3D. (e) MuLoG-CNN. (f) MuLoG-CNN-P. (g) MS-CNN.

Table 4.2: Quantitative despeckling assessment results from simulated four-look ($L = 4$) PolSAR data. The mean results over ten independent realizations of the speckle process from each of 18 ground truth covariance images of 256×256 pixels are shown with standard deviations in gray.

	NMSE (dB)	QL	SSIM	Time (s)
Original	-4.85 \pm 0.6	14.3 \pm 7.6	0.820 \pm 0.06	-
Boxcar	-4.15 \pm 1.4	138.2 \pm 171.0	0.562 \pm 0.15	-
Ext. sigma	-4.54 \pm 0.7	29.9 \pm 29.9	0.844 \pm 0.05	0.4
IDAN	-4.39 \pm 1.1	12.9 \pm 4.7	0.701 \pm 0.09	0.4
MuLoG-BM3D	-7.39 \pm 1.0	12.2 \pm 6.9	0.864 \pm 0.05	24.4
MuLoG-CNN	-7.17 \pm 1.1	12.0 \pm 7.0	0.873 \pm 0.05	10.4
MuLoG-CNN-P	-7.44 \pm 0.9	12.9 \pm 9.7	0.870 \pm 0.05	10.9
MS-CNN	-8.28 \pm 0.7	11.7 \pm 8.1	0.876 \pm 0.05	0.2

Fig. 4.3 shows the \mathbf{W}^{11} element of the whitened images produced via (4.15). Structure is visible in the boxcar, extended sigma, and IDAN whitened images, indicating a loss of spatial detail from the despeckling process. The whitened images from

Table 4.3: Quantitative polarimetric information preservation assessment results from simulated PolSAR data. The mean results over ten independent realizations of the speckle process from each of 18 ground truth covariance images of 256×256 pixels are shown.

	Single-look					Four-look				
	KSS	BP2	H	An	$\bar{\alpha}$	KSS	BP2	H	An	$\bar{\alpha}$
Boxcar	0.074	1581.0	0.097	0.855	0.113	0.186	3506.8	0.081	0.709	0.092
Ext. sigma	0.305	717.2	0.342	1.100	0.185	0.349	399.2	0.153	1.031	0.122
IDAN	0.163	1009.1	0.089	0.887	0.097	0.150	1500.3	0.073	0.757	0.082
MuLoG-BM3D	0.035	270.1	0.082	0.777	0.081	0.036	282.9	0.073	0.734	0.076
MuLoG-CNN	0.026	148.2	0.085	0.775	0.082	0.067	1779.5	0.074	0.724	0.080
MuLoG-CNN-P	0.026	123.7	0.083	0.772	0.082	0.073	168.5	0.074	0.732	0.080
MS-CNN	0.035	268.7	0.073	0.676	0.077	0.074	625.6	0.061	0.602	0.065

MuLoG show very little visible structure, while apparent structure is visible within the road at the bottom of the MS-CNN image. The visual analysis is supported by the modified Box-Pierce test results in Table 4.3, which indicate significant autocorrelation in the whitened images from the boxcar, extended sigma, and IDAN methods in the $L = 1$ case. MuLoG-CNN-P exhibits the least autocorrelation according to the Box-Pierce test for any of the methods in both the $L = 1$ and $L = 4$ experiments. The Kolmogorov-Smirnov statistics in Table 4.3 show the MuLoG variants to adhere best to the theorized statistical properties of the whitened image, followed by MS-CNN.

In Fig. 4.4, co-polarization and cross-polarization signatures computed from the average covariance matrix within the 7×6 pixel region identified by the red arrow in Fig. 4.2 are shown for a subset of the compared despeckling methods. It is notable that the MS-CNN signatures, both co-polarization and cross-polarization, are in closer agreement with the respective ground truth signatures than those of the other methods shown, indicating preservation of polarimetric information. Of the methods

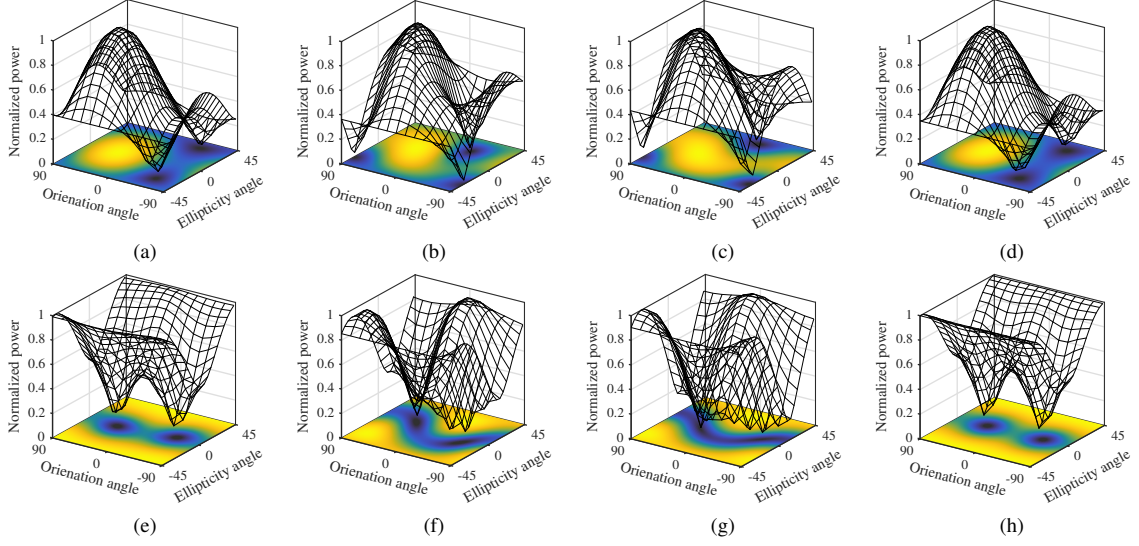


Figure 4.4: Co-polarization (top row) and cross-polarization (bottom row) signatures computed from a small region of a simulated single-look PolSAR image. The ellipticity and orientation angles of the transmitting antenna are given in degrees. (a,e) Ground truth. (b,f) Extended sigma. (c,g) MuLoG-CNN. (d,h) MS-CNN.

not displayed, IDAN produced signatures distinctly different from the ground truth, and the other MuLoG variants produced results similar to those from MuLoG-CNN.

4.5.2 Measured Data Experiments

We next compare the despeckling approaches using measured PolSAR images from the Pi-SAR sensor. Pi-SAR, an L-band airborne PolSAR system with a 3 m resolution in both azimuth and slant range, was developed by the former Communications Research Laboratory and NASDA (now JAXA) of Japan [80, 167]. A single-look 512×512 pixel PolSAR image from data collected by Pi-SAR over Tsukuba, Japan is shown in Fig. 4.5 with boxes indicating regions in which specific performance indexes are computed: white boxes for ENL and MOR, green for TCR, yellow for EPD-ROA, and blue for the whitened image assessments. To allow a detailed comparison, a




Figure 4.5: Single-look PolSAR image of 512×512 pixels from the Pi-SAR sensor used in despeckling comparison. The boxes indicate regions in which performance indexes are computed: white boxes for ENL and MOR, green for TCR, yellow for EPD-ROA, and blue for the whitened image assessments. The region within the red box is depicted in Fig. 4.6.

128×128 pixel region cropped from within the red box in Fig. 4.5 is shown in Fig. 4.6 along with the despeckled images from that region. Despeckling results from a four-look 200×200 pixel Pi-SAR image of Niigata City, Japan are shown in Fig. 4.7. The quantitative results computed from the two images are shown in Table 4.4 and Table 4.5.

From the cropped images in Fig. 4.6, it is clear that the MuLoG variants provide significant speckle suppression in homogeneous regions. MS-CNN provides a less smooth image with a greater preservation of texture within the vegetated regions. The ENL results show the CNN-based MuLoG variants to provide the greatest speckle




Figure 4.6: Detailed comparison of the despeckling algorithms using a 128×128 region cropped from Fig. 4.5. (a) Original data. (b) Boxcar. (c) Extended sigma. (d) IDAN. (e) MuLoG-BM3D. (f) MuLoG-CNN. (g) MuLoG-CNN-P. (h) MS-CNN.

suppression on the single-look data, followed by MS-CNN, which reduces speckle significantly more than the conventional methods. IDAN returns the highest ENL on the four-look data followed by the proposed method. Lacking a clean reference image, it is difficult to judge the correct balance between speckle suppression and textural preservation.

Notably, the extended sigma filter appears to preserve point targets more than the other methods, while the IDAN filter obscures point scatterers present in the original image. The MuLoG variants tend to smear point scatterers more than either the extended sigma filter or MS-CNN. The “EPD-HD” and “EPD-VD” columns in Table 4.4 and Table 4.5 show the EPD-ROA computed using immediate horizontally and vertically neighboring pixels, respectively. The EPD-ROA results show the extended sigma filter and MS-CNN to best preserve edges in both the single-look and




Figure 4.7: Comparison of despeckling methods using a four-look PolSAR image of 200×200 pixels. The boxes in (a) indicate regions in which performance indexes are computed: white boxes for ENL and MOR, green for TCR, yellow for EPD-ROA, and blue for the whitened image assessments. (a) Original data. (b) Boxcar. (c) Extended sigma. (d) IDAN. (e) MuLoG-BM3D. (f) MuLoG-CNN. (g) MuLoG-CNN-P. (h) MS-CNN.

four-look images, and the TCR and MOR indexes show largely mixed results. The Kolmogorov-Smirnov statistics in Table 4.4 and Table 4.5 show MS-CNN to underperform all methods but IDAN, indicating substandard fit to the theorized distribution. The extended sigma filter and MuLoG-BM3D both show a lower modified Box-Pierce statistic than the proposed method.

The \mathbf{W}^{11} element of the whitened images formed from the region bound by the cyan box in the single-look image in Fig. 4.5 are shown in Fig. 4.8. Spatial structure is visible in the whitened images from each method, most notably from IDAN. The whitened images from the MuLoG variants appear similar to that of the proposed method, with each having small constant areas. Note that the simulated data

Table 4.4: Quantitative despeckling assessment results from the measured PolSAR data ($L = 1$).

	ENL	EPD-HD	EPD-VD	TCR	MOR	KSS	BP2
Original	2.3	-	-	-	-	-	-
Boxcar	15.0	0.33	0.30	6.32	1.01	0.067	554.7
IDAN	35.3	0.33	0.30	6.21	0.94	0.120	727.9
Ext. sigma	3.9	0.41	0.40	0.71	1.00	0.313	346.0
MuLoG-BM3D	61.6	0.34	0.30	2.84	1.03	0.068	422.8
MuLoG-CNN	110.1	0.34	0.29	1.99	1.03	0.071	621.2
MuLoG-CNN-P	170.5	0.33	0.29	2.02	1.00	0.096	707.5
MS-CNN	89.3	0.35	0.31	1.97	1.00	0.100	480.3

Table 4.5: Quantitative despeckling assessment results from the measured PolSAR data ($L = 4$).

	ENL	EPD-HD	EPD-VD	TCR	MOR	KSS	BP2
Original	3.9	-	-	-	-	-	-
Boxcar	9.1	0.64	0.54	7.62	1.12	0.084	217.3
IDAN	51.5	0.71	0.55	5.68	0.94	0.095	159.4
Ext. sigma	4.9	0.90	0.82	0.11	0.99	0.358	162.6
MuLoG-BM3D	22.3	0.77	0.64	1.33	0.94	0.061	122.4
MuLoG-CNN	23.6	0.78	0.64	0.41	0.94	0.090	107.1
MuLoG-CNN-P	33.8	0.75	0.61	0.98	0.94	0.057	163.8
MS-CNN	38.1	0.79	0.67	2.07	0.89	0.087	152.0




Figure 4.8: Whitened single-look PolSAR images from the region bound by the cyan box in Fig. 4.5. (a) Boxcar. (b) Extended sigma. (c) IDAN. (d) MuLoG-BM3D. (e) MuLoG-CNN. (f) MuLoG-CNN-P. (g) MS-CNN.

experiments contain only fully-developed speckle by construction, and thus spatial correlations in the whitened images indicated an undue erosion of polarimetric information by the despeckling process. For measured data, this is not necessarily the case, e.g., with the presence of man-made media.

Fig. 4.9 shows the estimated entropy of the four-look image from a subset of the considered despeckling methods. The entropy image from the extended sigma filter has a granular character similar to that of the original data. Compared to the original image and the extended sigma filter result, both MuLoG and MS-CNN yield high entropy values within the vegetated regions, indicating a more random scattering process. The MuLoG and MS-CNN results are quite similar with MuLoG appearing smoother and MS-CNN better preserving small low entropy regions. The estimated alpha angle images in Fig. 4.10 show an extended sigma filter result that again appears similar to the original data. The MuLoG image is much smoother with apparent structure over the man-made media. The result from the proposed




Figure 4.9: Estimated entropy from the four-look measured data. (a) Original data. (b) Extended sigma. (c) MuLoG-CNN. (d) MS-CNN.

method is also smooth compared to the original data with low-angle regions that are less contiguous than the MuLoG image. A satisfying interpretation of these results is difficult due to the lack of ground truth. Note that the MuLoG variants not shown were found to produce results similar to MuLoG-CNN in estimates of both entropy and mean alpha angle.

4.6 Conclusion

4.6.1 Summary

In this chapter, we have proposed a CNN-based polarimetric SAR despeckling approach that operates in the matrix logarithm domain and uses a residual learning architecture that has been successful in natural image denoising. We have implemented this approach with a training set of spatially multilooked PolSAR images that are corrupted by a simulated speckle process.




Figure 4.10: Estimated alpha angle (in degrees) from the four-look measured data. (a) Original data. (b) Extended sigma. (c) MuLoG-CNN. (d) MS-CNN.

In experimental results on both simulated and measured PolSAR data, the proposed method has demonstrated improved scene feature preservation and comparable despeckling performance compared to state-of-the-art methods. Further, the proposed approach provides improved despeckling over conventional filtering methods and operates with a comparably fast running time afforded by the use of CNNs.

4.6.2 Future Work

A continuation of the work in this chapter might entail application to other multi-channel SAR imaging modalities. Alternative methods of procuring unspeckled SAR data for training purposes might also be considered, e.g., using temporally multi-looked images. Another avenue for future work is the use of complex-valued CNNs

in a similar approach; accounting for the relationship between the real and imaginary components of the covariance matrix in this way could improve despeckling performance and allow for training with fewer samples.

Chapter 5: Maximum Likelihood Estimation in Mixed Integer Linear Models

5.1 Introduction

In many sensing and communication applications, an unknown quantity is encoded in the phase of a signal. Because phase is unambiguous only on $[0, 2\pi)$, a measurement may experience an integer number of phase wraps. Let \mathbb{R} , \mathbb{Z} , and \mathbb{Q} denote the fields of real numbers, integers, and rational numbers, respectively. To estimate the quantity of interest, we consider a linear model with an observation vector $\mathbf{y} \in \mathbb{R}^m$ that is related to parameter vectors $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{k} \in \mathbb{Z}^m$ as

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{M}\mathbf{k} + \mathbf{u}. \quad (5.1)$$

The encoding matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ has full column rank and no all-zero row, $\mathbf{M} \in \mathbb{R}^{m \times m}$ is nonsingular, and the zero-mean noise $\mathbf{u} \in \mathbb{R}^m$ is assumed to be Gaussian with covariance $\mathbf{\Sigma} \in \mathbb{R}^{m \times m}$. The model in (5.1) arises in applications such as MIMO decoding [26–28], frequency estimation [22–24], ranging [18–21], interferometry [25], direction of arrival estimation [29], multichannel modulo sampling [30], and magnetic resonance imaging [9, 31, 168].

The maximum likelihood (ML) estimate of \mathbf{x} and \mathbf{k} in (5.1) is given by

$$\{\hat{\mathbf{x}}, \hat{\mathbf{k}}\} = \arg \min_{\mathbf{x} \in \mathbb{R}^n, \mathbf{k} \in \mathbb{Z}^m} \mathcal{L}_{x,k}(\mathbf{x}; \mathbf{y}), \quad (5.2)$$

where

$$\mathcal{L}_{x,k}(\mathbf{x}, \mathbf{k}; \mathbf{y}) = \|\Sigma^{-1/2}(\mathbf{y} - \mathbf{A}\mathbf{x} - \mathbf{M}\mathbf{k})\|^2. \quad (5.3)$$

For general \mathbf{A} , \mathbf{M} , and covariance Σ , (5.2) is an NP-hard mixed-integer quadratic program; for $\mathbf{M}^{-1}\mathbf{A} \in \mathbb{Q}^{m \times n}$, (5.2) is equivalent to a *closest lattice point problem* (CPP) [169, 170]. As such, it can be solved with Voronoi cell algorithms [171] or methods known as sphere decoders [172, 173]. But these methods have non-polynomial worst-case computational complexity [174], and the computation time depends on the random noise realization. Polynomial-time, and even linear-time, algorithms for the CPP problem have long been known for certain lattices [175–177]. In [32], a special case of (5.1) is considered, for which (5.2) can be solved in polynomial time by leveraging the classical Chinese remainder theorem (CRT). The case in [32] has that $n = 1$, Σ and \mathbf{M} are diagonal, and $\mathbf{M}^{-1}\mathbf{A}$ meets a certain coprime restriction. In some applications, this coprime restriction limits potential data acquisition strategies, and relaxation of the restriction has been shown to enable design of \mathbf{A} with improved estimation error variance [21, 178, 179].

In this chapter⁷, we first survey the existing literature, then provide background on the CRT and lattices. We present two algorithms for the problem in (5.2) that, for fixed n , have worst-case computational complexity that is polynomial in m . The

⁷This chapter is largely drawn from [7] (© 2023 IEEE) and [180], coauthored with Shen Zhao and Lee C. Potter. Manuscript [180] was supported in part by FA238523C0002 and has been granted public release, PA# AFRL-2024-1975. Distribution Statement A. Approved for public release: distribution is unlimited. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Air Force Research Laboratory or the U.S. Government.

two polynomial-time methods require that $\mathbf{\Sigma}$ and \mathbf{M} be diagonal, as in [32]. First, we describe a solution method that allows for a multivariate unknown $\mathbf{x} \in \mathbb{R}^n$, but the method applies to a restricted class of rational \mathbf{A} matrices. The solution, similar to its existing counterpart for $n = 1$ [32], is found in two parts: one part real-valued and the other integer-valued. The real part is found by constructing and evaluating a short list of candidate solutions. The integer part is found through a multivariate extension of the CRT [181, 182]. The second method lifts the restriction on \mathbf{A} and allows arbitrary rational \mathbf{A} with full column rank; this method works by simply evaluating a list of candidate solutions. We then provide a method that applies more generally to the problem in (5.2). This approach relies on sphere decoding and requires only that $\mathbf{M}^{-1}\mathbf{A}$ be rational. Unlike the other methods in this chapter, this more general method does not require that the covariance matrix $\mathbf{\Sigma}$ be diagonal, thus accommodating correlated noise. Application of the theory is illustrated in computed examples. Most proofs are relegated to appendices.

5.2 Prior Art

For $n = 1$, methods based on the CRT have been used to address special cases of (5.1) [32, 183–185]. Alternatively, parameter estimation involving (5.1) can be addressed by solving a closest lattice point problem [169, 170, 179]. Though, to use conventional algorithms such as sphere decoding, it is necessary to find a basis for a lattice generated by a rank-deficient matrix. Li et al. [20] study the $n = 1$ case of (5.1) with $\mathbf{M} = \mathbf{I}$, $\mathbf{\Sigma} = \mathbf{I}$, and $\mathbf{A} = [\Gamma_1^{-1}, \Gamma_2^{-1}, \dots, \Gamma_m^{-1}]^T$, where the elements of $\mathbf{\Gamma}$ are pairwise coprime integers (wavelengths, in ranging); explicit construction of a lattice basis is given. The case considered in [20] is broadened in [21] to allow for

arbitrary integer entries in $\mathbf{\Gamma}$. The lattice basis is explicitly derived in terms of the entries of \mathbf{A} , with a construction that relies on unimodular extension of \mathbf{A} . Though, for $n > 1$, the unimodular extension does not, in general, exist [186].

The CRT-based estimator proposed in [32] applies to the $n = 1$ case of (5.1) with diagonal $\mathbf{\Sigma}$ and \mathbf{M} such that $\mathbf{M}^{-1}\mathbf{A} = \alpha[\Gamma_1^{-1}, \Gamma_2^{-1}, \dots, \Gamma_m^{-1}]^T$, for some $\alpha \in \mathbb{R}$, with $\mathbf{\Gamma} \in \mathbb{Z}^m$ having pairwise coprime entries. Thus, the method in [32] accommodates non-white, uncorrelated noise. In [187], the $n = 1$ case of (5.1) is considered with arbitrary covariance matrix $\mathbf{\Sigma}$ and $\mathbf{M}^{-1}\mathbf{A} = [\Gamma_1^{-1}, \Gamma_2^{-1}, \dots, \Gamma_m^{-1}]^T$, with $\mathbf{\Gamma} \in \mathbb{Z}^m$ having pairwise coprime entries. Hence, [187] accommodates correlated noise. In [188, 189], (5.1) is studied for modulo sampling in the noiseless case with $\mathbf{M} = \mathbf{I}$, sparse structure in \mathbf{x} , and a restriction on \mathbf{A} . In [190], a method from [177] is modified to solve the special case of (5.2) with $n = 1$, diagonal $\mathbf{\Sigma}$, and $\mathbf{M}^{-1}\mathbf{A}$ not subject to a coprime restriction.

The general problem in (5.1) for $n > 1$ is widely addressed for MIMO communications. There, the intrinsic rank deficiency is overcome by regularization or, equivalently, maximum *a posteriori* probability (MAP) estimation [19, 26–28]. The problem in (5.2) is equivalent to the CPP [7]

$$\arg \min_{\mathbf{k} \in \mathbb{Z}^m} \|\mathbf{B}\mathbf{k} - \mathbf{B}\mathbf{y}\|^2, \quad (5.4)$$

for $\mathbf{B}^T\mathbf{B} = \mathbf{M}^T\mathbf{\Sigma}^{-1}\mathbf{M} - \mathbf{E}\mathbf{E}^T$ and $\mathbf{E} = \mathbf{M}^T\mathbf{\Sigma}^{-1}\mathbf{A}(\mathbf{A}^T\mathbf{\Sigma}^{-1}\mathbf{A})^{-1/2}$. The CPP solution method in [177] requires \mathbf{B} to be nonsingular, and its complexity depends on the eigenstructure of $\mathbf{B}^T\mathbf{B}$. However, \mathbf{B} is inherently singular, and thus, [177] does not apply to our problem.

5.3 Background

5.3.1 Lattices

For $\mathbf{Q} \in \mathbb{R}^{c \times d}$ with full column rank, the set

$$\Lambda_{\mathbf{Q}} = \{\mathbf{Q}\mathbf{k} \mid \mathbf{k} \in \mathbb{Z}^d\} \quad (5.5)$$

is the c -dimensional *lattice* generated by the *lattice basis* \mathbf{Q} [191]. The integer d is the rank of the lattice. All lattice points $\mathbf{x} \in \Lambda_{\mathbf{Q}}$ admit a unique representation, $\mathbf{x} = \mathbf{Q}\mathbf{k}$, for $\mathbf{k} \in \mathbb{Z}^d$ [192]. If \mathbf{Q} is invertible ($c = d$), $\Lambda_{\mathbf{Q}}$ is called a *full-rank* lattice. The set

$$\mathcal{P}_{\mathbf{Q}} = \{\mathbf{Q}\boldsymbol{\lambda} \mid \boldsymbol{\lambda} \in [0, 1]^n\} \quad (5.6)$$

is called the *fundamental parallelotope* [193] of $\Lambda_{\mathbf{Q}}$ for \mathbf{Q} . A fundamental parallelotope of $\Lambda_{\mathbf{Q}}$ contains no lattice points of $\Lambda_{\mathbf{Q}}$ except for its vertices.

A square matrix $\mathbf{U} \in \mathbb{Z}^{c \times c}$ is called *unimodular* if it has integer entries and $\det(\mathbf{U}) = \pm 1$, where $\det(\cdot)$ denotes the determinant of a matrix. The following are equivalent [194, Thm. 4.3].

- (i) \mathbf{U} is unimodular.
- (ii) \mathbf{U}^{-1} is unimodular.
- (iii) The lattice generated by the columns of \mathbf{U} is \mathbb{Z}^c .
- (iv) The Hermite normal form of \mathbf{U} is \mathbf{I}_c .
- (v) \mathbf{U} can be transformed into \mathbf{I}_c by elementary column operations.

A lattice basis \mathbf{Q} is not unique for a given lattice and, hence, neither is the corresponding fundamental parallelotope $\mathcal{P}_{\mathbf{Q}}$. We can express any lattice basis for $\Lambda_{\mathbf{Q}}$ as $\mathbf{Q}\mathbf{U}$, for some unimodular matrix \mathbf{U} , and $\Lambda_{\mathbf{Q}} = \Lambda_{\mathbf{Q}\mathbf{U}}$, for any unimodular \mathbf{U} .

The volume of a parallelotope, denoted by $\det \Lambda_Q \triangleq \text{vol}(\mathcal{P}_Q) = \sqrt{\det(Q^T Q)}$, is the same for all lattice bases. If $c = d$, the operation that translates $\mathbf{x} \in \mathbb{R}^d$ onto a fundamental parallelotope \mathcal{P}_Q is

$$\langle \mathbf{x} \rangle_Q \triangleq \mathbf{x} - Q \lfloor Q^{-1} \mathbf{x} \rfloor, \quad (5.7)$$

where $\lfloor \cdot \rfloor$ returns the greatest integer less than or equal to its input, acting element-wise when applied to a vector.

5.3.2 Hermite Normal Form

Definition 2. An integer matrix $\mathbf{H} \in \mathbb{Z}^{c \times d}$ is said to be in Hermite normal form (HNF) if it has entries that satisfy the following:

$$\begin{aligned} h_{i,i} &> 0, \quad \forall i \\ h_{i,j} &= 0, \quad \forall i > j \\ 0 &\leq h_{i,j} < h_{i,i}, \quad \forall i < j. \end{aligned} \quad (5.8)$$

A matrix \mathbf{H} in HNF is upper triangular and has full column rank; it can be expressed as $\mathbf{H} = [\widetilde{\mathbf{H}}^T, \mathbf{0}_{d,(c-d)}]^T$, where $\widetilde{\mathbf{H}} \in \mathbb{Z}^{d \times d}$ is a nonsingular, upper triangular matrix.⁸ Every full column rank integer matrix $\mathbf{T} \in \mathbb{Z}^{c \times d}$ has unique decomposition $\mathbf{H} = \mathbf{U}\mathbf{T}$, with unimodular \mathbf{U} and \mathbf{H} in HNF [194, Cor. 4.2a]. The HNF and the corresponding unimodular transformation can be computed in polynomial time [191, p. 149].

5.3.3 Chinese Remainder Theorems

The classical CRT for a single unknown variable is as follows [195, 196].

⁸The HNF is sometimes defined to be lower triangular or according to various alternative restrictions on the matrix entries.

Theorem 3 (Chinese remainder theorem). *Let $\mathbf{b} \in \mathbb{Z}^m$. Let $\mathbf{c} \in \mathbb{Z}^m$ have pairwise coprime entries. The system of linear congruences given by*

$$\mathbf{1}_m x \equiv \mathbf{b} \pmod{\mathbf{c}} \quad (5.9)$$

has integer solutions, and any two solutions are congruent modulo $\prod_{i=1}^m c_i$.

There are solution methods to the univariate Chinese remaindering problem with complexity $O(m^2)$ [195, p. 23].

The CRT has been extended to accommodate multivariate unknowns [181, 182]. The following is a direct implication of the multivariate CRT in [182]:

Theorem 4 (Multivariate CRT). *Let $\mathbf{W} \in \mathbb{Z}^{m \times n}$. Let $\mathbf{c} \in \mathbb{Z}^m$ have pairwise coprime entries. The system of linear congruences given by*

$$\mathbf{W}\mathbf{x} \equiv \mathbf{b} \pmod{\mathbf{c}} \quad (5.10)$$

has integer solutions for all $\mathbf{b} \in \mathbb{Z}^m$ iff

$$\gcd([g_{i,1}, \dots, g_{i,n}, c_i]^\top) = 1, \quad (5.11)$$

for each $i \in \{1, \dots, m\}$, where $\gcd(\cdot)$ is the positive column-wise greatest common divisor of an integer matrix input.

Lemma 5. *Let $\mathbf{W} \in \mathbb{Z}^{m \times n}$; $\mathbf{c} \in \mathbb{Z}^m$; $\mathbf{b} \in \mathbb{Z}^m$; and $\widetilde{\mathbf{W}} = [\mathbf{W}, \text{diag}(\mathbf{c})]^\top \in \mathbb{Z}^{(m+n) \times m}$. Let $\mathbf{H} = [\widetilde{\mathbf{H}}^\top, \mathbf{0}_{m,n}]^\top$ be the HNF of $\widetilde{\mathbf{W}}$. The system of congruences*

$$\mathbf{W}\mathbf{x} \equiv \mathbf{b} \pmod{\mathbf{c}} \quad (5.12)$$

has integer solutions iff $\widetilde{\mathbf{H}}^{-\top} \mathbf{b} \in \mathbb{Z}^m$.

Proof. The system of linear congruences in (5.12) has integer solutions iff there are $\mathbf{k} \in \mathbb{Z}^m$ and $\mathbf{x} \in \mathbb{Z}^n$ such that $\mathbf{W}\mathbf{x} = \mathbf{b} - \mathbf{c} \circ \mathbf{k}$ or, equivalently, if there is $\tilde{\mathbf{k}} \in \mathbb{Z}^{m+n}$ such that $\tilde{\mathbf{W}}^\top \tilde{\mathbf{k}} = \mathbf{b}$. Here, \circ denotes the element-wise product. From [194, Cor. 5.3b], $\tilde{\mathbf{W}}^\top \tilde{\mathbf{k}} = \mathbf{b}$ has an integer solution iff $\tilde{\mathbf{H}}^{-\top} \mathbf{b} \in \mathbb{Z}^m$. \square

The following multivariate CRT is more general than Theorem 4 as it does not assume that the moduli are pairwise coprime:

Theorem 6 (Generalized multivariate CRT). *Let $\mathbf{W} \in \mathbb{Z}^{m \times n}$; $\mathbf{c} \in \mathbb{Z}^m$; and $\tilde{\mathbf{W}} = [\mathbf{W}, \text{diag}(\mathbf{c})]^\top \in \mathbb{Z}^{(m+n) \times m}$. Let $\mathbf{H} = [\tilde{\mathbf{H}}^\top, \mathbf{0}_{m,n}]^\top$ be the HNF of $\tilde{\mathbf{W}}$. The system of congruences*

$$\mathbf{W}\mathbf{x} \equiv \mathbf{b} \pmod{\mathbf{c}} \quad (5.13)$$

has integer solutions for all $\mathbf{b} \in \mathbb{Z}^m$ iff $\tilde{\mathbf{H}} = \mathbf{I}_m$.

Proof. From Lemma 5, (5.13) has integer solutions for all $\mathbf{b} \in \mathbb{Z}^m$ iff $\tilde{\mathbf{H}}^{-\top} \mathbf{b} \in \mathbb{Z}^m$ for all $\mathbf{b} \in \mathbb{Z}^m$, which is true iff $\tilde{\mathbf{H}}^{-\top} \in \mathbb{Z}^{m \times m}$. As $\tilde{\mathbf{H}}^\top \in \mathbb{Z}^{m \times m}$, $\tilde{\mathbf{H}}^{-\top} \in \mathbb{Z}^{m \times m}$ iff $\tilde{\mathbf{H}}^\top$, and hence $\tilde{\mathbf{H}}$, is unimodular. Because $\tilde{\mathbf{H}}$ is upper-triangular, its determinant is the product of its diagonal entries. Thus, if $\tilde{\mathbf{H}}$ is unimodular, its diagonal entries are ones. By the definition of the HNF, the off-diagonal entries of $\tilde{\mathbf{H}}$ are zeros. Therefore, if $\tilde{\mathbf{H}}$ is unimodular, then $\tilde{\mathbf{H}} = \mathbf{I}_m$. The converse is also true. Therefore, $\tilde{\mathbf{H}}$ is unimodular iff $\tilde{\mathbf{H}} = \mathbf{I}_m$. \square

A method is given in Algorithm 7 for computing an integer solution to a multivariate system of congruences.

Algorithm 7 Solution to a system of linear congruences with a multivariate unknown.

Input: $\mathbf{b} \in \mathbb{Z}^m$, $\mathbf{c} \in \mathbb{Z}^m$, and $\mathbf{W} \in \mathbb{Z}^{m \times n}$, such that $\mathbf{W}\mathbf{x} \equiv \mathbf{b} \pmod{\mathbf{c}}$ has integer solutions.

Output: Solution $\mathbf{x} \in \mathbb{Z}^n$ such that $\mathbf{W}\mathbf{x} \equiv \mathbf{b} \pmod{\mathbf{c}}$.

1: Let $\widetilde{\mathbf{W}} = [\mathbf{W}, \text{diag}(\mathbf{c})]^\top$

2: By Theorem 6, $\mathbf{H} = [\mathbf{I}_m, \mathbf{0}_{m,n}]^\top$ is the HNF of $\widetilde{\mathbf{W}}$, with $\mathbf{H} = \mathbf{U}\widetilde{\mathbf{W}}$ for unimodular $\mathbf{U} \in \mathbb{Z}^{(m+n) \times (m+n)}$. Let $\widetilde{\mathbf{U}} \in \mathbb{Z}^{m \times n}$ be the first m rows of the first n columns of \mathbf{U} .

3: $\mathbf{x} \leftarrow \widetilde{\mathbf{U}}^\top \mathbf{b}$

5.3.4 Negative Log-Likelihood Function

The negative log-likelihood function for \mathbf{x} and \mathbf{k} in (5.1) is, up to an additive constant, proportional to $\mathcal{L}_{x,k}$ in (5.3). Note that $\mathcal{L}_{x,k}$ is periodic in \mathbf{x} and \mathbf{k} , with

$$\mathcal{L}_{x,k}(\mathbf{x}, \mathbf{k}; \mathbf{y}) = \mathcal{L}_{x,k}(\mathbf{x} + \boldsymbol{\delta}, \mathbf{k} - \mathbf{M}^{-1}\mathbf{A}\boldsymbol{\delta}; \mathbf{y}), \quad (5.14)$$

for any $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{k} \in \mathbb{Z}^m$, and $\boldsymbol{\delta} \in \{\mathbf{x} \mid \mathbf{M}^{-1}\mathbf{A}\mathbf{x} \in \mathbb{Z}^m\}$. If $\mathbf{M}^{-1}\mathbf{A}$ is rational, then

$$\{\mathbf{x} \mid \mathbf{M}^{-1}\mathbf{A}\mathbf{x} \in \mathbb{Z}^m\} = \Lambda_{\mathbf{V}}, \quad (5.15)$$

for some invertible lattice basis $\mathbf{V} \in \mathbb{R}^{n \times n}$. The basis \mathbf{V} can be formed as in [7]. First, let $\mathbf{M}^{-1}\mathbf{A} = \mathbf{Q}_{\text{num}} \oslash \mathbf{Q}_{\text{den}}$, for $\mathbf{Q}_{\text{num}}, \mathbf{Q}_{\text{den}} \in \mathbb{Z}^{m \times n}$, and let $\mathbf{D} = \text{diag}(\text{lcm}(\mathbf{Q}_{\text{den}})) \in \mathbb{Z}^{n \times n}$, where \oslash denotes element-wise division, and $\text{lcm}(\cdot)$ gives the column-wise positive least common multiple of an integer-valued matrix input, with zeros ignored. Let

$$\mathbf{T} = \mathbf{M}^{-1}\mathbf{A}\mathbf{D} \in \mathbb{Z}^{m \times n}, \quad (5.16)$$

and let $\mathbf{H} = [\widetilde{\mathbf{H}}^\top, \mathbf{0}_{n,(m-n)}]^\top$ be the HNF of \mathbf{T} , with

$$\mathbf{H} = \mathbf{U}\mathbf{T}, \quad (5.17)$$

for some unimodular $\mathbf{U} \in \mathbb{Z}^{m \times m}$. From (5.15) and (5.16),

$$\begin{aligned} \mathbf{x} \in \Lambda_{\mathbf{V}} &\Leftrightarrow \mathbf{A}\mathbf{x} \in \mathbb{Z}^m \Leftrightarrow \mathbf{T}\mathbf{D}^{-1}\mathbf{x} \in \mathbb{Z}^m \\ &\Leftrightarrow \mathbf{U}^{-1}\mathbf{H}\mathbf{D}^{-1}\mathbf{x} \in \mathbb{Z}^m \Leftrightarrow \mathbf{H}\mathbf{D}^{-1}\mathbf{x} \in \mathbb{Z}^m \\ &\Leftrightarrow \mathbf{x} \in \mathbf{D}\widetilde{\mathbf{H}}^{-1}\mathbb{Z}^n, \end{aligned}$$

where we have used that \mathbf{U} is bijective on \mathbb{Z}^m . Thus,

$$\Lambda_{\mathbf{V}} = \{\mathbf{D}\widetilde{\mathbf{H}}^{-1}\boldsymbol{\lambda} \mid \boldsymbol{\lambda} \in \mathbb{Z}^n\} \quad (5.18)$$

is a full-rank lattice with a basis $\mathbf{V} = \mathbf{D}\widetilde{\mathbf{H}}^{-1}$. Therefore, for any $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{k} \in \mathbb{Z}^m$,

$$\mathcal{L}_{x,k}(\mathbf{x}, \mathbf{k}; \mathbf{y}) = \mathcal{L}_{x,k}(\mathbf{x} + \mathbf{V}\boldsymbol{\lambda}, \mathbf{k} - \mathbf{M}^{-1}\mathbf{A}\mathbf{V}\boldsymbol{\lambda}; \mathbf{y}), \quad (5.19)$$

for any $\boldsymbol{\lambda} \in \mathbb{Z}^n$. The parallelotope $\mathcal{P}_{\mathbf{V}}$ defines a region in which the solution to (5.2) is unique, and we can therefore limit our search to this region.

5.4 ML Estimator for Restricted Encoding Matrices

In this section we provide an ML estimation procedure for a special case of (5.2), in which $n \geq 1$, the covariance matrix $\boldsymbol{\Sigma}$ and \mathbf{M} are diagonal, and $\mathbf{M}^{-1}\mathbf{A}$ is a rational matrix that meets a certain condition. For simplicity in this section, we will assume, without loss of generality, that $\mathbf{M} = \mathbf{I}$. The restriction on \mathbf{A} is analogous to the coprime moduli restriction used in [32]. Hence in this section, we extend the algorithm in [32] to provide polynomial-time (for fixed n) ML estimation for multivariate unknowns.

5.4.1 Preliminaries

In [32], it is assumed that $\mathbf{A} = \alpha[\Gamma_1^{-1}, \Gamma_2^{-1}, \dots, \Gamma_m^{-1}]^T$, for some $\alpha \in \mathbb{R}$, with $\boldsymbol{\Gamma} \in \mathbb{Z}^m$ having pairwise coprime entries. The solution for x in (5.2) is constructed

as a sum of two terms: The first is an estimate of a “common remainder,” that is, a real number within an interval of length $|C|$, where $C = \alpha^{-1} \gcd(\mathbf{\Gamma})$. The common remainder is found by a search over a list of only m candidate values. The second term is an integer multiple of C , which can be solved for efficiently by the use of the univariate CRT.

For $n > 1$, the interval in which the common remainder is found becomes the n -dimensional parallelotope \mathcal{P}_C corresponding to a lattice basis $\mathbf{C} \in \mathbb{R}^{n \times n}$. The basis is given by $\mathbf{C} = \mathbf{V} \text{diag}^{-1}(\mathbf{z})$, with $\mathbf{z} = (\text{lcm}(\mathbf{A}\mathbf{V}))^\top \in \mathbb{Z}^n$, and \mathbf{V} found from \mathbf{A} as in (5.18). Thus, $\Lambda_{\mathbf{V}}$ is a sub-lattice of $\Lambda_{\mathbf{C}}$, and \mathcal{P}_C provides a tiling of $\mathcal{P}_{\mathbf{V}}$, i.e., translated copies of \mathcal{P}_C can be arranged within $\mathcal{P}_{\mathbf{V}}$ such that $\mathcal{P}_{\mathbf{V}}$ is covered with no overlap between the copies. This tiling is illustrated, for the encoding matrix

$$\mathbf{A} = \begin{bmatrix} 1/2 & 0 & 1 \\ -1/6 & 1/3 & 2/3 \end{bmatrix}^\top, \quad (5.20)$$

in Fig. 5.1; in this case $\mathbf{V} = [2, 1; 0, 3]$ and $\mathbf{C} = [1, 1/3; 0, 1]$. Thus, $\mathbf{z} = [2, 3]^\top$.

Let $\boldsymbol{\beta} \in \mathbb{Z}^{m \times n}$ with

$$\beta_{j,k} = \begin{cases} ((\mathbf{a}_j \mathbf{C})_k)^{-1}, & (\mathbf{a}_j \mathbf{C})_k \neq 0 \\ 0, & \text{otherwise} \end{cases}, \quad (5.21)$$

where $\mathbf{a}_j \in \mathbb{Q}^{1 \times n}$ denotes the j th row of \mathbf{A} . Also let $\boldsymbol{\omega} \in \mathbb{Z}^m$ denote the row-wise positive least common multiple of $\boldsymbol{\beta}$, i.e.,

$$\boldsymbol{\omega} = (\text{lcm}(\boldsymbol{\beta}^\top))^\top. \quad (5.22)$$

For example with $n = 1$, $m = 3$, and $\mathbf{A} = [\frac{1}{2}, \frac{1}{4}, \frac{1}{6}]^\top$, we have $\mathbf{V} = 12$, $\mathbf{C} = 2$, $\mathbf{z} = 1$, and $\boldsymbol{\beta} = [1, 2, 3]^\top$. Note that for $n = 1$, $\omega_j = |\beta_j|$, $j \in \{1, \dots, m\}$.

The following condition is, for $n > 1$, analogous to the assumption of pairwise coprime moduli, used for $n = 1$ in [32]:

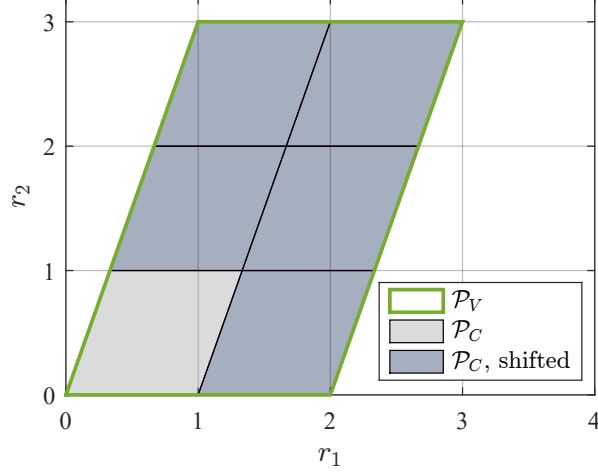


Figure 5.1: \mathcal{P}_V partitioned by six shifted copies of \mathcal{P}_C for the encoding matrix in (5.20).

Condition 7. For any $\mathbf{b} \in \mathbb{Z}^m$, there is $\boldsymbol{\gamma} \in \mathbb{Z}^n$ such that

$$\boldsymbol{\omega} \circ (\mathbf{AC})\boldsymbol{\gamma} \equiv \mathbf{b} \pmod{\boldsymbol{\omega}}. \quad (5.23)$$

The following proposition is directly implied by Theorem 6:

Proposition 8. Let $\widetilde{\mathbf{W}} = [\boldsymbol{\omega} \circ (\mathbf{AC}), \text{diag}(\boldsymbol{\omega})]^\top \in \mathbb{Z}^{(m+n) \times m}$ and $\mathbf{H} = [\widetilde{\mathbf{H}}^\top, \mathbf{0}_{m,n}]^\top$ be the HNF of $\widetilde{\mathbf{W}}$. Condition 7 is equivalent to $\widetilde{\mathbf{H}} = \mathbf{I}_m$.

The solution method in this section requires that \mathbf{A} meets Condition 7. We will show that this restriction leads to a fast ML estimation procedure.

A solution to (5.2) for \mathbf{x} is constructed as

$$\widehat{\mathbf{x}} = \widehat{\mathbf{r}} + \mathbf{C}\widehat{\boldsymbol{\gamma}} \in \mathcal{P}_V, \quad (5.24)$$

with $\widehat{\mathbf{r}} \in \mathcal{P}_C$ and $\widehat{\boldsymbol{\gamma}} \in \langle \mathbb{Z}^n \rangle_{\mathbf{C}^{-1}\mathbf{V}}$. Such an approach is appealing because the integer-valued unknown $\boldsymbol{\gamma}$ can be estimated efficiently thanks to the multivariable CRT; the

real-valued unknown \mathbf{r} can, as in [32], be found by searching a relatively short list of candidate solutions.

5.4.2 Estimating the Integer Remainders

Recall that $\mathbf{C}^{-1}\mathbf{V} = \text{diag}(\mathbf{z})$, with $\mathbf{z} \in \mathbb{Z}^n$. According to the formulation in (5.24), we consider an equivalent problem to (5.2), given by:

$$\{\hat{\gamma}, \hat{\mathbf{r}}, \hat{\mathbf{k}}\} = \arg \min_{\gamma \in \langle \mathbb{Z}^n \rangle_{\mathbf{C}^{-1}\mathbf{V}}, \mathbf{r} \in \mathcal{P}_{\mathbf{C}}, \mathbf{k} \in \mathbb{Z}^m} \mathcal{L}_{\gamma, \mathbf{r}, \mathbf{k}}(\gamma, \mathbf{r}; \mathbf{y}), \quad (5.25)$$

where

$$\mathcal{L}_{\gamma, \mathbf{r}, \mathbf{k}}(\gamma, \mathbf{r}, \mathbf{k}; \mathbf{y}) = \left\| \Sigma^{-1/2} (\mathbf{y} - \mathbf{A}(\mathbf{r} + \mathbf{C}\gamma) - \mathbf{k}) \right\|^2.$$

For any \mathbf{q}, \mathbf{k} , and γ such that

$$\omega \circ (\mathbf{A}\mathbf{C}\gamma + \mathbf{k}) = \mathbf{q}, \quad (5.26)$$

we have that $\mathcal{L}_{q, \mathbf{r}}(\mathbf{q}, \mathbf{r}; \mathbf{y}) = \mathcal{L}_{\gamma, \mathbf{r}, \mathbf{k}}(\gamma, \mathbf{r}, \mathbf{k}; \mathbf{y})$, where

$$\mathcal{L}_{q, \mathbf{r}}(\mathbf{q}, \mathbf{r}; \mathbf{y}) = \left\| \Sigma^{-1/2} (\mathbf{y} - \mathbf{A}\mathbf{r} - \mathbf{q} \oslash \omega) \right\|^2. \quad (5.27)$$

The bijection between $\{\mathbf{q} \in \mathbb{Z}^m\}$ and $\{\gamma \in \langle \mathbb{Z}^n \rangle_{\mathbf{C}^{-1}\mathbf{V}}, \mathbf{k} \in \mathbb{Z}^m\}$ shown by Theorem 27 suggests that, rather than finding $\hat{\gamma} \in \langle \mathbb{Z}^n \rangle_{\mathbf{C}^{-1}\mathbf{V}}$ and $\hat{\mathbf{k}} \in \mathbb{Z}^m$ directly, we find $\hat{\mathbf{q}} \in \mathbb{Z}^m$ such that (5.26) is satisfied. This leads to yet another equivalent problem, given by

$$\{\hat{\mathbf{q}}, \hat{\mathbf{r}}\} = \arg \min_{\mathbf{q} \in \mathbb{Z}^m, \mathbf{r} \in \mathcal{P}_{\mathbf{C}}} \mathcal{L}_{q, \mathbf{r}}(\mathbf{q}, \mathbf{r}; \mathbf{y}). \quad (5.28)$$

For any $\mathbf{r} \in \mathcal{P}_{\mathbf{C}}$, $\hat{\mathbf{q}}$ can be found by solving a CPP:

$$\hat{\mathbf{q}} = \arg \min_{\mathbf{q} \in \mathbb{Z}^m} \left\| \Sigma^{-1/2} (\mathbf{y} - \mathbf{A}\mathbf{r} - \text{diag}^{-1}(\omega)\mathbf{q}) \right\|^2. \quad (5.29)$$

Because $\Sigma^{-1/2} \text{diag}^{-1}(\omega)$ has orthogonal columns (as it is diagonal), $\hat{\mathbf{q}}$ can be found via Babai's rounding procedure [197]. That is [198, pg. 379],

$$\hat{\mathbf{q}} = \lceil \omega \circ (\mathbf{y} - \mathbf{A}\mathbf{r}) \rceil, \quad (5.30)$$

where $\lceil \mathbf{x} \rceil$ denotes element-wise rounding of \mathbf{x} to the nearest integers, with $\lceil \frac{1}{2} \rceil = 1$.

5.4.3 Estimating the Real-Valued Common Remainder

After inserting $\hat{\mathbf{q}}$ into (5.28), we consider

$$\begin{aligned} \hat{\mathbf{r}} &= \arg \min_{\mathbf{r} \in \mathcal{P}_C} \mathcal{L}_{q,r}(\hat{\mathbf{q}}, \mathbf{r}; \mathbf{y}) \\ &= \arg \min_{\mathbf{r} \in \mathcal{P}_C} \left\| \Sigma^{-1/2} \mathbf{d}_{\mathbf{1}_m \odot \omega}(\mathbf{y} - \mathbf{A}\mathbf{r}) \right\|^2 \\ &= \arg \min_{\mathbf{r} \in \mathcal{P}_C} \mathcal{L}_r(\mathbf{r}; \mathbf{y}), \end{aligned} \quad (5.31)$$

where $\mathbf{1}_m$ is a length- m column vector of ones, and

$$\mathbf{d}_{\mathbf{z}}(\mathbf{x}) \triangleq \mathbf{x} - \mathbf{z} \circ \lceil \mathbf{x} \odot \mathbf{z} \rceil, \quad (5.32)$$

for $\mathbf{z} \in \mathbb{R}^m$, with $\mathbf{0} \prec \mathbf{z}$. Here, \prec and \preceq denote element-wise inequalities. The operator $\mathbf{d}_{\mathbf{z}}$ has the property that

$$-\mathbf{z}/2 \preceq \mathbf{d}_{\mathbf{z}}(\mathbf{x}) \prec \mathbf{z}/2 \quad (5.33)$$

If $-\mathbf{z}/2 \preceq \mathbf{x} \prec \mathbf{z}/2$, then $\mathbf{d}_{\mathbf{z}}(\mathbf{x}) = \mathbf{x}$.

Define \mathcal{H} as

$$\mathcal{H} = \{ \mathbf{r} \in \mathbb{R}^n \mid \exists j \in \{1, \dots, m\} \text{ s.t. } \mathbf{d}_{\omega_j^{-1}}(y_j - \mathbf{a}_j \mathbf{r}) = -\frac{1}{2} \omega_j^{-1} \} \quad (5.34)$$

$$= \{ \mathbf{r} \in \mathbb{R}^n \mid \exists j \in \{1, \dots, m\} \text{ s.t. } \omega_j \mathbf{a}_j \mathbf{r} - \langle \omega_j y_j + \frac{1}{2} \rangle_1 \in \mathbb{Z} \}. \quad (5.35)$$

Note that \mathcal{H} is a union of countably infinitely many vector hyperplanes (subspaces with codimension one) in \mathbb{R}^n . Due to the rounding operation in \mathbf{d} , \mathcal{L}_r is non-differentiable for $\mathbf{r} \in \mathcal{H}$. The hyperplanes in \mathcal{H} partition \mathbb{R}^n into *chambers*, convex regions within which $\lceil \boldsymbol{\omega} \circ (\mathbf{y} - \mathbf{A}\mathbf{r}) \rceil$ is constant.⁹ A chamber is thus uniquely identified by $\boldsymbol{\ell} \in \Phi \subset \mathbb{Z}^m$, according to

$$\Psi(\boldsymbol{\ell}) = \{\mathbf{r} \in \mathbb{R}^n \mid \lceil \boldsymbol{\omega} \circ (\mathbf{y} - \mathbf{A}\mathbf{r}) \rceil = \boldsymbol{\ell}\}, \quad (5.36)$$

where $\Psi : \Phi \mapsto \mathbb{R}^n$, and Φ is the set of integer vectors that correspond to chambers in \mathbb{R}^n formed by \mathcal{H} . The closure of any chamber is a polytope (the convex hull of finitely many points).

The function \mathcal{L}_r in (5.31) is continuous for $\mathbf{r} \in \mathbb{R}^n$ and differentiable for $\mathbf{r} \in \mathbb{R}^n \setminus \mathcal{H}$. For any $\boldsymbol{\ell} \in \Phi$, the rounded portion of \mathcal{L}_r is constant within $\mathbf{r} \in \Psi(\boldsymbol{\ell})$. Thus, \mathcal{L}_r is quadratic over $\Psi(\boldsymbol{\ell})$, and \mathcal{L}_r is a piecewise quadratic function. In Fig. 5.2, \mathcal{L}_r for $\mathbf{r} \in \mathcal{P}_C$ is shown, with value indicated by color, for a measurement from (5.1) with \mathbf{A} from (5.20).

Let $\tilde{\mathcal{H}} = \mathcal{P}_C \cap \mathcal{H}$. The set $\tilde{\mathcal{H}}$ is the intersection of \mathcal{P}_C with a union of finitely many hyperplanes, and it can be considered to partition \mathcal{P}_C into “wrapped chambers,” i.e., chambers that are wrapped about the parallelotope \mathcal{P}_C . The wrapped chamber corresponding to $\boldsymbol{\ell} \in \tilde{\Phi}$ is

$$\begin{aligned} \tilde{\Psi}(\boldsymbol{\ell}) &= \{\langle \mathbf{r} \rangle_C \mid \mathbf{r} \in \Psi(\boldsymbol{\ell})\} \\ &= \{\mathbf{r} \in \mathcal{P}_C \mid \exists \boldsymbol{\lambda} \in \mathbb{Z}^n \text{ s.t. } \lceil \boldsymbol{\omega} \circ (\mathbf{y} - \mathbf{A}\mathbf{r}) \rceil = \boldsymbol{\ell} + \boldsymbol{\omega} \circ (\mathbf{A}\mathbf{C}\boldsymbol{\lambda})\}. \end{aligned} \quad (5.37)$$

⁹The term “chamber” has elsewhere been used interchangeably with the more common term “cell” to refer to an intersection of half-planes in the context of hyperplane arrangements [199, 200]. Here, we use “chamber” to refer to a similar concept, with the distinction that the cells of hyperplane arrangements are open, whereas our chambers are neither open nor closed.

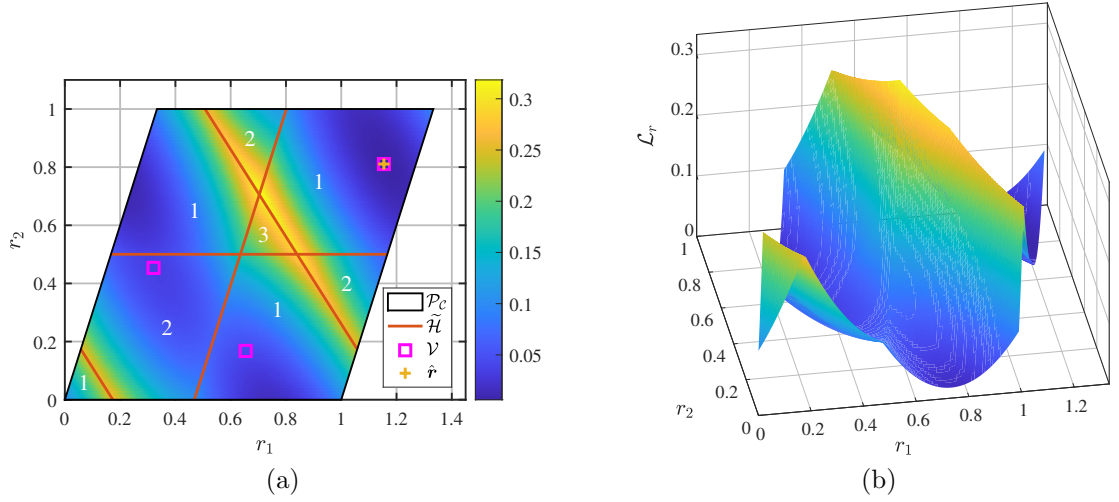


Figure 5.2: (a) The value of \mathcal{L}_r is indicated by color over \mathcal{P}_C . The wrapped hyperplanes in $\tilde{\mathcal{H}}$ are shown in red. The candidate solutions for (5.31) are marked with magenta squares, with the solution $\hat{\mathbf{r}}$ at the yellow cross. \mathcal{P}_C is partitioned into wrapped chambers, labeled 1–3. (b) \mathcal{L}_r as a surface plot.

For any $\ell \in \Phi$ and $\lambda \in \mathbb{Z}^n$, $\tilde{\Psi}(\ell) = \tilde{\Psi}(\ell + \omega \circ (\mathbf{A}\mathbf{C}\lambda))$, as shown by Proposition 22 (Appendix B.1). In Fig. 5.2, $\tilde{\mathcal{H}}$ is depicted by red line segments, and \mathcal{P}_C is shown to be divided by $\tilde{\mathcal{H}}$ into three wrapped chambers, labeled with white text.

Theorem 9. *For some $\ell \in \Phi$, $\hat{\mathbf{r}} \in \text{int } \tilde{\Psi}(\ell)$.*

Proof. We have that $\hat{\mathbf{r}} \in \mathcal{P}_C$, and

$$\mathcal{P}_C = \left(\bigcup_{\ell \in \Phi} \text{int } \tilde{\Psi}(\ell) \right) \cup \tilde{\mathcal{H}}. \quad (5.38)$$

For no $\mathbf{r} \in \mathcal{H}$ does \mathcal{L}_r have a local minimum at \mathbf{r} , as shown by Lemma 24 (Appendix B.1); the statement and proof follow those given for $n = 1$ in [32, Thm. 2]. As $\tilde{\mathcal{H}} \subset \mathcal{H}$, we have that $\hat{\mathbf{r}} \notin \tilde{\mathcal{H}}$. Therefore, $\hat{\mathbf{r}} \in \text{int } \tilde{\Psi}(\ell)$, for some $\ell \in \Phi$. \square

According to Theorem 9, $\hat{\mathbf{r}} \in \text{int } \tilde{\Psi}(\ell)$, for some $\ell \in \Phi$, where int denotes the interior of a set. Now assume, for any $\hat{\ell} \in \Phi$, that $\hat{\mathbf{r}} \in \tilde{\Psi}(\hat{\ell})$. From (5.37), for any

$\mathbf{r} \in \tilde{\Psi}(\hat{\ell})$, that

$$\mathcal{L}_r(\mathbf{r}; \mathbf{y}) = \left\| \Sigma^{-1/2} \left(\mathbf{y} - \mathbf{A}\mathbf{r} - \hat{\ell} \odot \boldsymbol{\omega} - \mathbf{A}\mathbf{C}\boldsymbol{\lambda} \right) \right\|^2, \quad (5.39)$$

for some $\boldsymbol{\lambda} \in \mathbb{Z}^n$. Therefore, $\hat{\mathbf{r}}$ can be found by generalized least squares as

$$\begin{aligned} \hat{\mathbf{r}} &= (\Sigma^{-1/2} \mathbf{A})^\dagger \Sigma^{-1/2} (\mathbf{y} - \hat{\ell} \odot \boldsymbol{\omega} - \mathbf{A}\mathbf{C}\boldsymbol{\lambda}) \\ &= \langle (\Sigma^{-1/2} \mathbf{A})^\dagger \Sigma^{-1/2} (\mathbf{y} - \hat{\ell} \odot \boldsymbol{\omega}) \rangle_C \\ &= \langle (\Sigma^{-1/2} \mathbf{A})^\dagger \Sigma^{-1/2} (\mathbf{y} - \lceil \boldsymbol{\omega} \circ (\mathbf{y} - \mathbf{A}\mathbf{r}) \rceil \odot \boldsymbol{\omega}) \rangle_C, \end{aligned} \quad (5.40)$$

for any $\mathbf{r} \in \tilde{\Psi}(\hat{\ell})$. Because $\hat{\ell}$ is unknown (i.e., we do not know which wrapped cell contains $\hat{\mathbf{r}}$), it is useful to find a set χ containing one point from each wrapped chamber (see Appendix B.3 for details). We then form \mathcal{V} , a set of candidate solutions to (5.31), as

$$\mathcal{V} = \left\{ \langle (\Sigma^{-1/2} \mathbf{A})^\dagger \Sigma^{-1/2} (\mathbf{y} - \lceil \boldsymbol{\omega} \circ (\mathbf{y} - \mathbf{A}\mathbf{r}) \rceil \odot \boldsymbol{\omega}) \rangle_C \mid \mathbf{r} \in \chi \right\}. \quad (5.41)$$

For the example in Fig. 5.2, there are three wrapped chambers and, thus, three elements in \mathcal{V} , which are marked with magenta squares; the solution to (5.2) is shown with a yellow cross. The solution to (5.31) is

$$\hat{\mathbf{r}} = \arg \min_{\mathbf{r} \in \mathcal{V}} \mathcal{L}_r(\mathbf{r}; \mathbf{y}), \quad (5.42)$$

and it can be found by evaluating \mathcal{L}_r for each element of the candidate set \mathcal{V} .

5.4.4 Constructing the ML Solution

By inserting $\hat{\mathbf{r}}$ into (5.30), we get

$$\hat{\mathbf{q}} = \lceil \boldsymbol{\omega} \circ (\mathbf{y} - \mathbf{A}\hat{\mathbf{r}}) \rceil. \quad (5.43)$$

Algorithm 8 ML solution for (5.1) with Condition 7 satisfied.

Input: Measurements $\mathbf{y} \in \mathbb{R}^m$; diagonal covariance $\mathbf{\Sigma}$; and full column rank $\mathbf{A} \in \mathbb{Q}^{m \times n}$ satisfying Condition 7.

Output: ML solution $\{\hat{\mathbf{x}}, \hat{\mathbf{k}}\}$.

- 1: Let χ be a set containing a point in $\tilde{\Psi}(\ell)$ for each $\ell \in \Phi$
 - 2: Let \mathcal{V} be the set of candidate solutions in (5.41)
 - 3: $\hat{\mathbf{r}} \leftarrow \arg \min_{\mathbf{r} \in \mathcal{V}} \mathcal{L}_r(\mathbf{r}; \mathbf{y})$
 - 4: $\hat{\mathbf{q}} \leftarrow \lceil \boldsymbol{\omega} \circ (\mathbf{y} - \mathbf{A}\hat{\mathbf{r}}) \rceil$
 - 5: Use Algorithm 7 to find $\hat{\boldsymbol{\gamma}}$ as the solution for $\boldsymbol{\gamma}$ to $\boldsymbol{\omega} \circ (\mathbf{A}\mathbf{C}\boldsymbol{\gamma}) \equiv \hat{\mathbf{q}} \pmod{\boldsymbol{\omega}}$
 - 6: $\hat{\mathbf{x}} \leftarrow \langle \mathbf{C}\hat{\boldsymbol{\gamma}} + \hat{\mathbf{r}} \rangle_{\mathbf{v}}$
 - 7: $\hat{\mathbf{k}} \leftarrow \lceil \mathbf{y} - \mathbf{A}\hat{\mathbf{x}} \rceil$
-

We then find $\hat{\boldsymbol{\gamma}} \in \langle \mathbb{Z}^n \rangle_{\mathbf{C}^{-1}\mathbf{v}}$ as the solution to the following system of congruences:

$$\boldsymbol{\omega} \circ (\mathbf{A}\mathbf{C})\boldsymbol{\gamma} \equiv \hat{\mathbf{q}} \pmod{\boldsymbol{\omega}}, \quad (5.44)$$

using Algorithm 7; the congruences are known to have a solution by Theorem 6.

Finally, the solution to (5.2) is formed as

$$\hat{\mathbf{x}} = \langle \mathbf{C}\hat{\boldsymbol{\gamma}} + \hat{\mathbf{r}} \rangle_{\mathbf{v}} \quad (5.45)$$

$$\hat{\mathbf{k}} = \lceil \mathbf{y} - \mathbf{A}\hat{\mathbf{x}} \rceil. \quad (5.46)$$

The proposed method to solve (5.2) for this case is summarized in Algorithm 8.

5.4.5 Computational Complexity

The ML solution method in Algorithm 8 involves two primary tasks. The first is to solve a system of linear congruence equations in multiple integer variables, which can be done in polynomial time [194, Cor. 5.3b]. The solution method given in Algorithm 7 computes a HNF, which can be reused for subsequent measurements from the same \mathbf{A} . The second task is to build a set of candidate solutions for (5.31) and evaluate them. The number of candidate solutions is given by $|\mathcal{V}| = |\{\tilde{\Psi}(\ell)\}_{\ell \in \Phi}|$.

For fixed \mathbf{V} , \mathbf{C} , and n , we get that $|\mathcal{V}|$ is $O(m^n)$, as shown in Appendix B.4. Here again, much of the computation to form the set of candidate solutions can be reused for inference on subsequent measurements. Thus, for fixed \mathbf{V} , \mathbf{C} , and n , both primary tasks can be completed in time that is polynomial in m .

5.4.6 Comment on the $n = 1$ Case

The problem treated in this section is simpler when $n = 1$, which is the scenario considered in [32]. With $n = 1$, $\omega_j = |\beta_j| = |(Ca_j)^{-1}|$, for $j \in \{1, \dots, m\}$, and $C^{-1}\mathbf{V} = \text{lcm}(\boldsymbol{\beta})$. The fundamental parallelotope \mathcal{P}_V is simply an interval with length $|\mathbf{V}|$. A univariate CRT method [195, p. 23] can be used rather than Algorithm 7 in step 5 of Algorithm 8. The construction of χ is particularly simple in this case, as discussed in Appendix B.3. The set χ , and hence \mathcal{V} , has cardinality m . Therefore, with $n = 1$, just m candidates need be considered to estimate the real-valued unknown, as in [32].

5.5 ML Estimator for General Encoding Matrices

In this section, we remove the restriction on \mathbf{A} given by Condition 7 and provide an ML estimator for $n \geq 1$ that constructs a short list of candidate points in \mathbb{R}^n provably containing the ML solution. That is, we consider the case of (5.2) in which $n \geq 1$, the covariance matrix $\boldsymbol{\Sigma}$ and \mathbf{M} are diagonal, and $\mathbf{A} \in \mathbb{Q}^{m \times n}$ is an arbitrary matrix with full column rank. As in Section 5.4, we assume, without loss of generality that $\mathbf{M} = \mathbf{I}_m$. The development parallels Section 5.4; however, the multivariate CRT is no longer necessarily applicable, typically resulting in a larger list of candidate solutions.

According to Proposition 36, for any \mathbf{x} , the Babai rounding method [197] gives the optimal \mathbf{k} in (5.2) as $\widehat{\mathbf{k}} = \lceil \mathbf{y} - \mathbf{A}\mathbf{x} \rceil$. Inserting this into $\mathcal{L}_{x,k}$ gives us

$$\mathcal{L}_{x,k}(\mathbf{x}, \widehat{\mathbf{k}}; \mathbf{y}) = \mathcal{L}_x(\mathbf{x}; \mathbf{y}) = \left\| \Sigma^{-1/2} \mathbf{d}_1(\mathbf{y} - \mathbf{A}\mathbf{x}) \right\|^2, \quad (5.47)$$

with \mathbf{d} defined in (5.32). We define \mathcal{H}_{NC} as the set of points at which \mathcal{L}_x is non-differentiable:

$$\mathcal{H}_{\text{NC}} = \left\{ \mathbf{x} \in \mathbb{R}^n \mid \exists j \in \{1, \dots, m\} \text{ s.t. } \mathbf{a}_j \mathbf{x} - \langle y_j + \tfrac{1}{2} \rangle_1 \in \mathbb{Z} \right\}, \quad (5.48)$$

using the subscript _{NC} to distinguish from analogous quantities in Section 5.4. The hyperplanes in \mathcal{H}_{NC} partition \mathbb{R}^n into chambers, within which $\lceil \mathbf{y} - \mathbf{A}\mathbf{x} \rceil$ is constant. A chamber is thus uniquely identified by $\ell \in \Phi_{\text{NC}}$, according to

$$\Psi_{\text{NC}}(\ell) = \{ \mathbf{x} \in \mathbb{R}^n \mid \lceil \mathbf{y} - \mathbf{A}\mathbf{x} \rceil = \ell \}, \quad (5.49)$$

where $\Psi_{\text{NC}} : \Phi_{\text{NC}} \mapsto \mathbb{R}^n$, and $\Phi_{\text{NC}} \subset \mathbb{Z}^m$ is the set of integer vectors that correspond to chambers in \mathbb{R}^n formed by \mathcal{H}_{NC} . Let $\tilde{\mathcal{H}}_{\text{NC}} = \mathcal{P}_{\mathbf{V}} \cap \mathcal{H}_{\text{NC}}$; i.e., $\tilde{\mathcal{H}}_{\text{NC}}$ is the intersection of $\mathcal{P}_{\mathbf{V}}$ with a union of finitely many hyperplanes, and it can be considered to partition $\mathcal{P}_{\mathbf{V}}$ into wrapped chambers. Now, similarly to (5.37), we define $\tilde{\Psi}_{\text{NC}}(\ell)$ as

$$\begin{aligned} \tilde{\Psi}_{\text{NC}}(\ell) &= \{ \langle \mathbf{x} \rangle_{\mathbf{V}} \mid \mathbf{x} \in \Psi_{\text{NC}}(\ell) \} \\ &= \{ \mathbf{x} \in \mathcal{P}_{\mathbf{V}} \mid \exists \boldsymbol{\lambda} \in \mathbb{Z}^n \text{ s.t. } \lceil \mathbf{y} - \mathbf{A}\mathbf{x} \rceil = \ell + \mathbf{A}\mathbf{V}\boldsymbol{\lambda} \}. \end{aligned} \quad (5.50)$$

Note that (5.50) uses the lattice basis \mathbf{V} , rather than basis \mathbf{C} used in (5.37). By an analogous statement to Theorem 9, we have that $\widehat{\mathbf{x}} \in \text{int } \tilde{\Psi}_{\text{NC}}(\ell)$, for some $\ell \in \Phi_{\text{NC}}$.

Assume, for any $\widehat{\ell} \in \Phi_{\text{NC}}$, that $\widehat{\mathbf{x}} \in \tilde{\Psi}_{\text{NC}}(\widehat{\ell})$. From (5.50) we have, for any $\mathbf{x} \in \tilde{\Psi}_{\text{NC}}(\widehat{\ell})$, that

$$\mathcal{L}_x(\mathbf{x}; \mathbf{y}) = \left\| \Sigma^{-1/2} (\mathbf{y} - \mathbf{A}\mathbf{x} - \widehat{\ell} - \mathbf{A}\mathbf{V}\boldsymbol{\lambda}) \right\|^2, \quad (5.51)$$

for some $\boldsymbol{\lambda} \in \mathbb{Z}^n$. Therefore, $\widehat{\boldsymbol{x}}$ can be found by generalized least squares as

$$\begin{aligned}\widehat{\boldsymbol{x}} &= (\boldsymbol{\Sigma}^{-1/2} \mathbf{A})^\dagger \boldsymbol{\Sigma}^{-1/2} (\mathbf{y} - \widehat{\boldsymbol{\ell}} - \mathbf{A} \mathbf{V} \boldsymbol{\lambda}) \\ &= \langle (\boldsymbol{\Sigma}^{-1/2} \mathbf{A})^\dagger \boldsymbol{\Sigma}^{-1/2} (\mathbf{y} - \widehat{\boldsymbol{\ell}}) \rangle_{\mathbf{V}} \\ &= \langle (\boldsymbol{\Sigma}^{-1/2} \mathbf{A})^\dagger \boldsymbol{\Sigma}^{-1/2} (\mathbf{y} - \lceil \mathbf{y} - \mathbf{A} \mathbf{x} \rceil) \rangle_{\mathbf{V}},\end{aligned}\tag{5.52}$$

for any $\boldsymbol{x} \in \widetilde{\Psi}_{\text{NC}}(\widehat{\boldsymbol{\ell}})$. We again form a set of candidate solutions to (5.2), as

$$\mathcal{V}_{\text{NC}} = \{ \langle (\boldsymbol{\Sigma}^{-1/2} \mathbf{A})^\dagger \boldsymbol{\Sigma}^{-1/2} (\mathbf{y} - \lceil \mathbf{y} - \mathbf{A} \mathbf{x} \rceil) \rangle_{\mathbf{V}} \mid \boldsymbol{x} \in \chi_{\text{NC}} \},\tag{5.53}$$

where the set χ_{NC} is constructed so as to contain a point in each wrapped chamber.

The set χ_{NC} can be constructed analogously to χ in Section 5.4. The solution to (5.2) for \boldsymbol{x} is

$$\widehat{\boldsymbol{x}} = \arg \min_{\boldsymbol{x} \in \mathcal{V}_{\text{NC}}} \mathcal{L}_x(\boldsymbol{x}; \mathbf{y}),\tag{5.54}$$

and it can be found by evaluating \mathcal{L}_x for each element of the candidate set \mathcal{V}_{NC} . The solution for \mathbf{k} is then given by $\widehat{\mathbf{k}} = \lceil \mathbf{y} - \mathbf{A} \widehat{\boldsymbol{x}} \rceil$.

In Fig. 5.3, \mathcal{L}_x for $\boldsymbol{x} \in \mathcal{P}_{\mathbf{V}}$ is shown, with value indicated by color, with \mathbf{A} from (5.20) for the measurement as Fig. 5.2. In this case, the more general method given in this section yields six wrapped chambers, and hence six candidate solutions, whereas the approach in Section 5.4 must evaluate only three candidate solutions.

The proposed method to solve (5.2) for this case is summarized in Algorithm 9. Note that this method is more straightforward than Algorithm 8, in the sense that it involves only a search over a candidate set and does not require solving a system of congruence equations. The cardinality of the candidate set is again polynomial in m , as shown in Appendix B.5. However, $|\mathcal{V}_{\text{NC}}|$ is typically larger than $|\mathcal{V}|$, the cardinality of the candidate set constructed in Section 5.4.

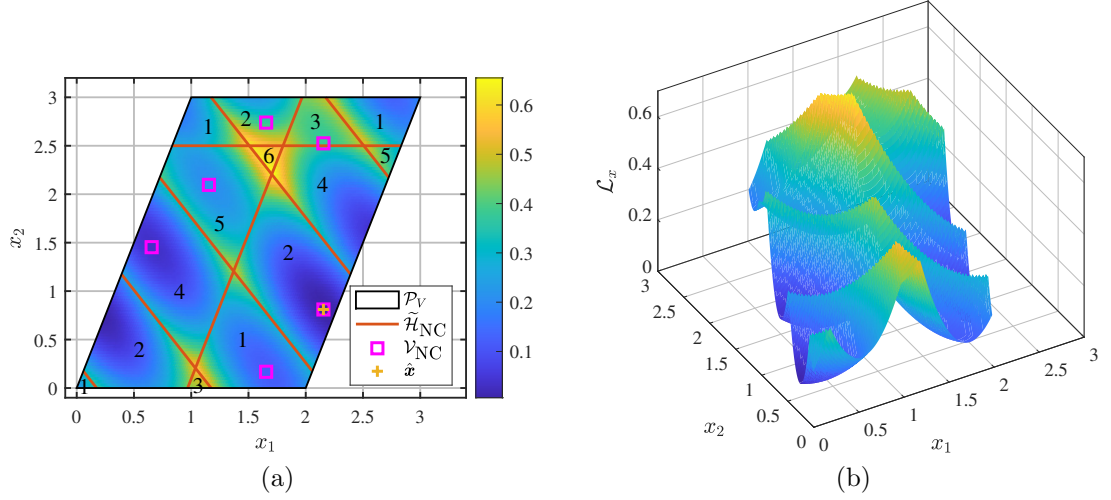


Figure 5.3: (a) The value of \mathcal{L}_x is indicated by color over \mathcal{P}_V . The wrapped hyperplanes in $\tilde{\mathcal{H}}_{\text{NC}}$ are shown in red. The candidate solutions for (5.2) are marked with magenta squares, with the solution $\hat{\mathbf{x}}$ at the yellow cross. \mathcal{P}_V is partitioned into wrapped chambers, labeled 1–6. (b) \mathcal{L}_x as a surface plot.

Algorithm 9 ML solution for (5.1) with unrestricted \mathbf{A} .

Input: Measurements $\mathbf{y} \in \mathbb{R}^m$; diagonal covariance Σ ; and $\mathbf{A} \in \mathbb{Q}^{m \times n}$.

Output: ML solution $\{\hat{\mathbf{x}}, \hat{\mathbf{k}}\}$.

- 1: Let χ_{NC} be a set containing a point in $\tilde{\Psi}_{\text{NC}}(\ell)$ for each $\ell \in \Phi_{\text{NC}}$
 - 2: Let \mathcal{V}_{NC} be the set of candidate solutions defined in (5.53)
 - 3: $\hat{\mathbf{x}} \leftarrow \arg \min_{\mathbf{x} \in \mathcal{V}_{\text{NC}}} \mathcal{L}_x(\mathbf{x}; \mathbf{y})$
 - 4: $\hat{\mathbf{k}} \leftarrow \lceil \mathbf{y} - \mathbf{A}\hat{\mathbf{x}} \rceil$
-

As in Section 5.4, the problem considered in this section is simpler when $n = 1$. In that case, $C^{-1}V = \text{lcm}(\beta)$, and $|\chi_{\text{NC}}| = |\mathcal{V}_{\text{NC}}| \leq mV/C$, as shown in Appendix B.5. A benefit of both Algorithm 8 and Algorithm 9 is that much of the computation can be done ahead of time and reused on arbitrarily many sets of measurements when the same \mathbf{A} is used repeatedly.

5.6 ML Estimator for General Mixed Integer Linear Models

The methods presented in Section 5.4 and Section 5.5 require that the covariance Σ be diagonal, and thus, that the noise in (5.2) be uncorrelated. Though, in certain applications the noise is correlated, such as when the measurements correspond to phase differences appearing in sufficient statistics for multi-snapshot, multi-channel data [29]. In this section we provide an ML solution method for (5.2) that applies more generally than the methods presented in Section 5.4 and Section 5.5.

Some applications considered in this chapter share three features: (1) no prior knowledge of the distribution of the parameters, rendering MAP estimation inappropriate or subject to extensive tuning of a regularization parameter; (2) correlated noise; (3) multivariate unknowns, $\mathbf{x} \in \mathbb{R}^n$. In this section, we provide an ML solution method for the problem in (5.1) that simultaneously accommodates these three characteristics.

For this method, we require only that $\mathbf{M}^{-1}\mathbf{A}$ be rational. We also assume in this section, without loss of generality, that the measurements have been whitened; thus, $\Sigma = \mathbf{I}_m$. A MATLAB implementation of this approach is available.¹⁰

5.6.1 Estimation Task

For any \mathbf{k} , minimization of (5.3) over \mathbf{x} yields

$$\mathbf{x}_k = \mathbf{A}^\dagger (\mathbf{y} - \mathbf{M}\mathbf{k}), \quad (5.55)$$

¹⁰<https://github.com/tuckerda/Mixed-Integer-Linear-Models-MLE>

where \dagger denotes the Moore-Penrose pseudo-inverse. Minimizing $\mathcal{L}(\mathbf{x}_k, \mathbf{k}; \mathbf{y})$ with respect to the integer vector \mathbf{k} gives us

$$\mathbf{k}^\star = \arg \min_{\mathbf{k} \in \mathbb{Z}^m} \mathcal{L}(\mathbf{x}_k, \mathbf{k}; \mathbf{y}) = \arg \min_{\mathbf{k} \in \mathbb{Z}^m} \|\tilde{\mathbf{y}} - \mathbf{B}\mathbf{k}\|^2, \quad (5.56)$$

an integer least squares problem (ILS) with $\mathbf{B} = (\mathbf{I}_m - \mathbf{A}\mathbf{A}^\dagger)\mathbf{M}$ and $\tilde{\mathbf{y}} = \mathbf{B}\mathbf{M}^{-1}\mathbf{y}$. Note that the range of \mathbf{B} is the nullspace of \mathbf{A}^\top . For $m = n$, we get that $\mathbf{B} = \mathbf{0}$, and any $\mathbf{k} \in \mathbb{Z}^m$ minimizes (5.56).

The estimation task is then to find the point in the lattice $\Lambda_B = \{\mathbf{B}\mathbf{k} \mid \mathbf{k} \in \mathbb{Z}^m\}$ that is closest to the transformed noisy observation $\tilde{\mathbf{y}}$. This is a closest lattice point problem, which is NP-hard [201], though the expected complexity may be a polynomial [174]. Note that the lattice-generating matrix \mathbf{B} has nullity n and is therefore not a lattice basis. To use conventional algorithms such as sphere decoding [169, 170] to solve (5.56), it is necessary to first find a basis for Λ_B .

5.6.2 Lattice Basis Construction

To find a basis for Λ_B , we use that $\mathbf{B}\mathbf{T} = \mathbf{0}$. Partition \mathbf{U}^{-1} from (5.17) into $\tilde{\mathbf{G}} \in \mathbb{Z}^{m \times n}$ and $\mathbf{G}_\perp \in \mathbb{Z}^{m \times (m-n)}$ such that

$$\mathbf{U}^{-1} = [\tilde{\mathbf{G}}, \mathbf{G}_\perp]. \quad (5.57)$$

Because only the first n rows of \mathbf{H} are non-zero, we have from (5.17) that $\tilde{\mathbf{G}} = \mathbf{T}\tilde{\mathbf{H}}^{-1}$.

Using that unimodular \mathbf{U}^{-1} is bijective on \mathbb{Z}^m , we have

$$\begin{aligned} \Lambda_B &= \mathbf{B}\mathbb{Z}^m = \mathbf{B}\mathbf{U}^{-1}\mathbb{Z}^m \\ &= \mathbf{B}\tilde{\mathbf{G}}\mathbb{Z}^n + \mathbf{B}\mathbf{G}_\perp\mathbb{Z}^{m-n} \\ &= \mathbf{B}\mathbf{T}\mathbf{H}_n^{-1}\mathbb{Z}^n + \mathbf{B}\mathbf{G}_\perp\mathbb{Z}^{m-n}. \end{aligned} \quad (5.58)$$

Since $\mathbf{B}\mathbf{T} = \mathbf{0}$, this gives us

$$\Lambda_B = \mathbf{B}\mathbf{G}_\perp \mathbb{Z}^{m-n}, \quad (5.59)$$

and therefore, $\tilde{\mathbf{B}} = \mathbf{B}\mathbf{G}_\perp$ is a basis for Λ_B .

An alternative to the basis construction in (5.59) is to apply the modified LLL algorithm [202, 203] to \mathbf{B} , as in [187]. Unlike the original LLL algorithm [204], the modified LLL algorithm accepts as input a linearly dependent set of vectors. The modified LLL algorithm operates by detecting linear dependencies among a set of vectors, and this process can be plagued by rounding errors [202]. We find that this method is prone to numerical difficulties even for small m [205].

5.6.3 Lattice Basis Reduction

In order to solve (5.56), we first compute the LLL-QRZ factorization of $\tilde{\mathbf{B}}$ as

$$\tilde{\mathbf{B}} = \mathbf{Q}[\mathbf{R}^\top, \mathbf{0}_{m-n,n}]^\top \mathbf{Z}^{-1} = \mathbf{Q}_1 \mathbf{R} \mathbf{Z}^{-1} \quad (5.60)$$

where $\mathbf{Q} = [\mathbf{Q}_1, \mathbf{Q}_2] \in \mathbb{R}^{m \times m}$ is orthonormal, $\mathbf{Z}^{m \times m}$ is unimodular, and $\mathbf{R}^{(m-n) \times (m-n)}$ is upper-triangular [206]. The matrix \mathbf{R} satisfies the size-reduction and Lovász conditions [204], given respectively by

$$|r_{i,j}| \leq |r_{i,i}|/2, \quad 1 \leq i < j \leq m-n \quad (5.61)$$

$$\delta r_{i-1,i-1}^2 \leq r_{i-1,i}^2 + r_{i,i}^2, \quad 1 < i \leq m-n, \quad (5.62)$$

with $\delta \in (1/4, 1]$. We use this factorization to pose a new closest lattice point problem that is equivalent to (5.56), given by

$$\mathbf{z}^* = \arg \min_{\mathbf{z} \in \mathbb{Z}^{(m-n)}} \|\mathbf{Q}_1^\top \tilde{\mathbf{y}} - \mathbf{R}\mathbf{z}\|^2. \quad (5.63)$$

The closest lattice point problem in (5.63) can be solved using a search algorithm such as sphere decoding. For our examples, we use a depth-first sphere decoding algorithm [169, 170, 207]. After a solution has been found for (5.63), we get that $\mathbf{k}^\star = \mathbf{G}_\perp \mathbf{Z} \mathbf{z}^\star$ and $\mathbf{x}_{\mathbf{k}^\star} = \mathbf{A}^\dagger (\mathbf{y} - \mathbf{M} \mathbf{k}^\star)$. Thus, for any optimal $\{\mathbf{k}^\star, \mathbf{x}_{\mathbf{k}^\star}\}$ in (5.3), there exists an equivalent solution $\{\hat{\mathbf{k}}, \hat{\mathbf{x}}\}$ with $\hat{\mathbf{x}} \in \mathcal{P}_V$, given by

$$\hat{\mathbf{x}} = \langle \mathbf{x}_{\mathbf{k}^\star} \rangle_{\mathcal{P}} \quad (5.64)$$

and $\hat{\mathbf{k}} = \mathbf{k}^\star - \mathbf{M}^{-1} \mathbf{A} (\langle \hat{\mathbf{x}} \rangle_{\mathcal{P}} - \mathbf{x}_{\mathbf{k}^\star})$. From (5.55), (5.63), and (5.64),

$$\hat{\mathbf{x}} = \langle \mathbf{A}^\dagger (\mathbf{y} - \mathbf{M} \mathbf{G}_\perp \mathbf{Z} \mathbf{z}^\star) \rangle_{\mathcal{P}}. \quad (5.65)$$

5.6.4 Babai Point

Simple rounding of the non-integer least squares solution to (5.63) can, in some cases, provide the ML solution ℓ^\star , thereby foregoing the complexity of sphere decoding. Define

$$\ell_B = \lceil \tilde{\mathbf{B}}^\dagger \tilde{\mathbf{y}} \rceil \quad (5.66)$$

where $\lceil \cdot \rceil$ denotes element-wise rounding to the nearest integer. The solution ℓ_B is referred to as the *Babai point* [197] or the *zero-forcing equalization point* [174]. Define d_{\min} as the length of the shortest non-zero lattice point in Λ_B [208]. As noted in [19, 178], for example, if

$$\|\tilde{\mathbf{y}} - \tilde{\mathbf{B}} \ell_B\| \leq d_{\min}/2, \quad (5.67)$$

then $\ell_B = \ell^\star$. For applications in which the same $\tilde{\mathbf{B}}$ is used with many noisy measurements [209], then $\tilde{\mathbf{B}}^\dagger$ may be pre-computed and (5.67) tested before resorting to sphere decoding for integer solution in (5.63), thereby providing accelerated computation, on average. Further, if Λ_B admits an orthogonal basis $\tilde{\mathbf{B}} \mathbf{U}$, for some unimodular

\mathbf{U} , then the $(m - n)$ -dimensional problem in (5.63) decouples for that basis. In that case, $\lceil (\tilde{\mathbf{B}}\mathbf{U})^\dagger \tilde{\mathbf{y}} \rceil$ is likewise optimal, as shown by Proposition 36 (Appendix B.6).

Other instances of lattices that admit polynomial-time solution of (5.63) are given in [24, 176, 190, 210].

5.6.5 Estimation Procedure

An ML parameter estimation procedure for the model in (5.1) is given in Algorithm 10. Note that steps 1–4, 6, and 7 can be computed beforehand and reused to form estimates from different sets of noisy measurements that share the same \mathbf{A} and \mathbf{M} . This reuse of precomputed values can greatly reduce computation times when many measurements have the same noise covariance and sensor configuration.

We also seek to improve the average computational cost by first computing the Babai point as in (5.66) and determining whether it is optimal as in (5.67). If the Babai point is optimal, then sphere decoding can be foregone. Otherwise we solve a closest lattice point problem using a sphere decoding algorithm [169, 170, 207]. In the case that $m = n$, we have that any $\mathbf{k} \in \mathbb{Z}^m$ is optimal in (5.56), and we can give \mathbf{k}^* an arbitrary value.

5.6.6 Estimator Distribution

We next consider the distribution of the estimator, which in turn enables optimized design of \mathbf{A} [31, 178, 211]. The Voronoi cell of Λ_B , denoted $\text{Vor}(\Lambda_B)$, is the closure of the subset in \mathbb{R}^m containing all points closer to the origin than to any other lattice point. The lattice Λ_B lies in the nullspace of \mathbf{A}^\top , and $\text{Vor}(\Lambda_B)$ is therefore unbounded in the range of \mathbf{A} . The integer parameter estimate $\hat{\mathbf{k}}$ is considered *correct* if $\mathbf{B}\hat{\mathbf{k}} = \mathbf{B}\mathbf{k}$ [178] or, equivalently, if $\mathbf{B}\mathbf{M}^{-1}\mathbf{u} \in \text{Vor}(\Lambda_B)$ [19]. From [178],

Algorithm 10 ML parameter estimation for linear model in (5.1).

Input: Measurements $\mathbf{y} \in \mathbb{R}^m$; full-column rank $\mathbf{A} \in \mathbb{R}^{m \times n}$; and invertible $\mathbf{M} \in \mathbb{R}^{m \times m}$, such that $\mathbf{M}^{-1}\mathbf{A} \in \mathbb{Q}^{m \times n}$.

Output: ML parameter estimates $\hat{\mathbf{x}}$ and $\hat{\mathbf{k}}$

- 1: Compute \mathbf{V} , a lattice basis for Λ , as in (5.18), (5.16), (5.17)
 - 2: Compute \mathbf{G}_\perp as in (5.57)
 - 3: $\mathbf{B} \leftarrow (\mathbf{I}_m - \mathbf{A}\mathbf{A}^\dagger)\mathbf{M}$
 - 4: $\tilde{\mathbf{B}} \leftarrow \mathbf{B}\mathbf{G}_\perp$
 - 5: **if** $m > n$ **then**
 - 6: Compute the LLL-QRZ factorization of $\tilde{\mathbf{B}}$, with \mathbf{Q}_1 and \mathbf{Z} as given in (5.60).
 - 7: $d_{\min} \leftarrow \min_{\mathbf{z} \in \mathbb{Z}^{m-n}, \mathbf{z} \neq \mathbf{0}} \|\mathbf{R}\mathbf{z}\|$
 - 8: $\tilde{\mathbf{y}} \leftarrow \mathbf{B}\mathbf{M}^{-1}\mathbf{y}$
 - 9: $\mathbf{z}_B \leftarrow \lceil \mathbf{R}^{-1}\mathbf{Q}_1^\top \tilde{\mathbf{y}} \rceil$
 - 10: **if** $\|\mathbf{Q}_1^\top \tilde{\mathbf{y}} - \mathbf{R}\mathbf{z}_B\| \leq \frac{1}{2}d_{\min}$ **then**
 - 11: $\mathbf{z}^* \leftarrow \mathbf{z}_B$
 - 12: **else**
 - 13: $\mathbf{z}^* \leftarrow \arg \min_{\mathbf{z} \in \mathbb{Z}^{(m-n)}} \|\mathbf{Q}_1^\top \tilde{\mathbf{y}} - \mathbf{R}\mathbf{z}\|^2$
 - 14: **end if**
 - 15: $\mathbf{k}^* \leftarrow \mathbf{G}_\perp \mathbf{Z} \mathbf{z}^*$
 - 16: **else**
 - 17: $\mathbf{k}^* \leftarrow \mathbf{0}$
 - 18: **end if**
 - 19: $\mathbf{x}_{\mathbf{k}^*} \leftarrow \mathbf{A}^\dagger (\mathbf{y} - \mathbf{M}^{-1}\mathbf{k}^*)$
 - 20: $\hat{\mathbf{x}} \leftarrow \langle \mathbf{x}_{\mathbf{k}^*} \rangle_{\mathcal{P}}$
 - 21: $\hat{\mathbf{k}} \leftarrow \mathbf{k}^* - \mathbf{M}^{-1}\mathbf{A}(\hat{\mathbf{x}} - \mathbf{x}_{\mathbf{k}^*})$.
-

$\mathbf{B}\mathbf{M}^{-1}\mathbf{u} \in \text{Vor}(\Lambda_B)$ if and only if $\mathbf{B}\mathbf{M}^{-1}\mathbf{u} + \mathbf{A}\mathbf{s} \in \text{Vor}(\Lambda_B)$ for all $\mathbf{s} \in \mathbb{R}^n$. Since $\mathbf{u} = \mathbf{B}\mathbf{M}^{-1}\mathbf{u} + \mathbf{A}\mathbf{s}$ for some \mathbf{s} ,

$$\mathbf{B}\mathbf{M}^{-1}\mathbf{u} \in \text{Vor}(\Lambda_B) \Leftrightarrow \mathbf{u} \in \text{Vor}(\Lambda_B). \quad (5.68)$$

That is, the probability that $\hat{\mathbf{k}}$ is correct is equal to the probability that \mathbf{u} is in $\text{Vor}(\Lambda_B)$. This probability, denoted P_c , is upper bounded by the probability that \mathbf{u} is within a sphere of volume $\det \Lambda_B$ and lower bounded by the probability that \mathbf{u} is in a sphere of radius $d_{\min}/2$ [19]. Thus,

$$F_{r,0}(d_{\min}^2/4) \leq P_c \leq F_{r,0}((\det \Lambda_B/\alpha)^{2/r}) \quad (5.69)$$

where $F_{r,0}$ is the central χ^2 cumulative distribution function (cdf) with $r = m - n$ degrees of freedom, $\alpha = \pi^{r/2}/\Gamma(1 + r/2)$, and Γ is the Gamma function.

Conditioned on $\mathbf{k}^* = \mathbf{G}_\perp \boldsymbol{\ell}^*$, the estimate $\mathbf{x}_{\mathbf{k}^*}$ (without translation onto \mathcal{P}) in (5.55) is linear in \mathbf{y} and is therefore Gaussian with variance $\boldsymbol{\Sigma}_{\text{ML}} = (\mathbf{A}^\top \mathbf{A})^{-1}$. Let $\mathcal{N}_{\mathcal{P}}(\mu, \Sigma)$ denote a Gaussian distribution wrapped onto \mathcal{P} . Given the true \mathbf{x} and $\boldsymbol{\ell}$, by total probability the distribution of the estimate $\hat{\mathbf{x}}$ is then a mixture of wrapped Gaussians,

$$f(\hat{\mathbf{x}} | \mathbf{x}, \boldsymbol{\ell}) = \sum_{\boldsymbol{\ell}' \in \mathbb{Z}^r} w_{\boldsymbol{\ell}'} \mathcal{N}_{\mathcal{P}}(\mathbf{x} + \mathbf{A}^\dagger \mathbf{M} \mathbf{G}_\perp (\boldsymbol{\ell} - \boldsymbol{\ell}'), \boldsymbol{\Sigma}_{\text{ML}}). \quad (5.70)$$

Each weight $w_{\boldsymbol{\ell}'}$ in the mixture is the integration of the probability density of $\tilde{\mathbf{y}}$ over a translated Voronoi cell, $\{\text{Vor}(\Lambda_B) + \tilde{\mathbf{B}}(\boldsymbol{\ell}' - \boldsymbol{\ell})\}$. We can numerically approximate this distribution by limiting the sum to the top several closest lattice points and approximating $w_{\boldsymbol{\ell}'}$, as in (5.69), using a non-central χ^2 cdf with non-centrality parameter $\|\tilde{\mathbf{B}}(\boldsymbol{\ell}' - \boldsymbol{\ell})\|^2$. For the non-central case, the right-hand side term in (5.69) provides only an approximation when $\boldsymbol{\ell}' \neq \boldsymbol{\ell}$, while the left-hand side remains a valid lower-bound. To illustrate, Fig. 5.4 displays the distribution of $f(\hat{\mathbf{x}} | \mathbf{x}, \boldsymbol{\ell})$ for whitened matrices $\mathbf{M} = \boldsymbol{\Sigma}^{-1/2}$,

$$\mathbf{A} = \frac{\boldsymbol{\Sigma}^{-1/2}}{6} \begin{bmatrix} -2 & 7 \\ 2 & 2 \\ 4 & 1 \end{bmatrix}, \quad \boldsymbol{\Sigma} = \frac{1}{100} \begin{bmatrix} 4 & 1 & -1 \\ 1 & 4 & 1 \\ -1 & 1 & 4 \end{bmatrix}.$$

Here, the probability of correctly detecting the wrapping integers is only $P_c \approx 0.6988$, with upper and lower bounds in (5.69) equal because $r = 1$. The most probable six components are visible via the colormap used to display the Gaussian mixture on the fundamental parallelotope.




Figure 5.4: Example distribution, $f(\hat{\mathbf{x}} | \mathbf{x}, \boldsymbol{\ell})$, for $m = 3$, $n = 2$.

5.6.7 Direction of Arrival Estimation

Consider a planar array of L isotropic sensors that is used to estimate the direction of arrival (DoA) of a narrowband signal $s(t) \in \mathbb{C}$ from a single emitter located in the far field of the array. The demodulated array output vector is given by

$$\mathbf{b}(t) = \mathbf{a}(\mathbf{x})s(t) + \boldsymbol{\varepsilon}(t), \quad t = 1, 2, \dots, T, \quad (5.71)$$

where $\mathbf{a}(\mathbf{x})$ is the array steering vector and $\boldsymbol{\varepsilon}(t)$ is zero-mean circular Gaussian noise that is uncorrelated across snapshots indexed by t . The array steering vector is given by

$$\mathbf{a}(\mathbf{x}) = [e^{i\pi \mathbf{p}_1^T \mathbf{x}} \quad e^{i\pi \mathbf{p}_2^T \mathbf{x}} \quad \dots \quad e^{i\pi \mathbf{p}_L^T \mathbf{x}}]^T, \quad (5.72)$$

The position vectors $\mathbf{p}_j \in \mathbb{R}^n$, for $j \in \{1, 2, \dots, L\}$, give the sensor positions in units of one-half wavelength. For linear arrays, $n = 1$ and $x = \sin \theta$, where θ is the azimuth angle of the source. For planar arrays, $n = 2$ and $\mathbf{x} = [\cos \theta \cos \phi, \sin \theta \cos \phi]^T$, with ϕ denoting the elevation angle to the source. For general arrays, $n = 3$ and $\mathbf{x} = [\cos \theta \cos \phi, \sin \theta \cos \phi, \sin \phi]^T$.

The sample covariance, given by $\hat{\mathbf{R}} = \frac{1}{T} \sum_{t=1}^T \mathbf{b}(t)\mathbf{b}^H(t)$, is a sufficient statistic of the data [29], and the MLE for \mathbf{x} in (5.71) is given by [22, p. 288]

$$\arg \max_{\mathbf{x} \in \mathcal{B}} \mathbf{a}^H(\mathbf{x}) \hat{\mathbf{R}} \mathbf{a}(\mathbf{x}), \quad (5.73)$$

where \mathcal{B} is the set of feasible values for \mathbf{x} . For $n = 1$, the set $\mathcal{B}_1 = \{\mathbf{x} \mid \theta \in [-\frac{\pi}{2}, \frac{\pi}{2}]\}$ is the unit ball in \mathbb{R} , $\|x\| \leq 1$. For a planar array, $n = 2$ and the set $\mathcal{B}_2 = \{\mathbf{x} \mid \theta \in [0, 2\pi), \phi \in [0, \frac{\pi}{2}]\}$ of all DoAs is the unit ball in \mathbb{R}^2 , $\|\mathbf{x}\| \leq 1$. These DoAs correspond to source directions on the upper hemisphere. If, for the planar array, the source is constrained to lie in the plane of the sensors, then \mathcal{B}'_2 is the unit circle in the plane. For a 3D array, $n = 3$ and the set $\mathcal{B}_3 = \{\mathbf{x} \mid \theta \in [0, 2\pi), \phi \in [-\frac{\pi}{2}, \frac{\pi}{2}]\}$ of all DoAs is a sphere in \mathbb{R}^3 , $\{\mathbf{x} \mid \|\boldsymbol{\omega}\| = 1\}$.

The evaluation of the MLE in (5.73) often involves a costly grid search over the non-convex cost surface. This grid search can be accelerated by the use of a multi-resolution grid search as in [212], though this involves some tuning of the succession of grid spacings over which the search will be conducted. The grid search will also necessarily introduce a bias.

Phase unwrapping estimators [23, 24, 31] offer a fast grid-free alternative. A phase unwrapping estimator uses only the phase information to estimate the direction of arrival. As such, this option is viable only when estimating the DoA of a single source [213–215].

Let \mathbf{v} denote the vector of $\binom{L}{2}$ lower-triangular entries of $\hat{\mathbf{R}}$ in column-major order. That is,

$$\mathbf{v} = [\hat{r}_{2,1}, \hat{r}_{3,1}, \dots, \hat{r}_{L,L-1}]^T \quad (5.74)$$

where $\widehat{r}_{i,j}$ is the $(i,j)^{\text{th}}$ entry of $\widehat{\mathbf{R}}$. The whitened measurement vector \mathbf{y} in (5.1) is then given by

$$\mathbf{y} = \widetilde{\Sigma}^{-1/2}(2\pi)^{-1}\angle\mathbf{v}, \quad (5.75)$$

where \angle returns the phase of a complex quantity, and $\widetilde{\Sigma} = \text{cov}((2\pi)^{-1}\angle\mathbf{v})$. The covariance matrix $\widetilde{\Sigma}$ can be approximated as in [31]; see Appendix C for details. The encoding matrix \mathbf{A} is given by

$$\mathbf{A} = \frac{1}{2}\widetilde{\Sigma}^{-1/2}[(\mathbf{p}_2 - \mathbf{p}_1), (\mathbf{p}_3 - \mathbf{p}_1), \dots, (\mathbf{p}_L - \mathbf{p}_{L-1})]^\top \quad (5.76)$$

and $\mathbf{M} = \widetilde{\Sigma}$. This construction allows the use of Algorithm 10 to estimate DoA from only the phase data of the sample covariance matrix. This approach offers a fast, grid-free alternative to the MLE in (5.73).

The phase unwrapping estimator conducts a depth-first search to find a solution to (5.63), though the solution may not correspond to a direction in the feasible set \mathcal{B} [216]. A non-feasible $\widehat{\omega}$ is projected to the feasible set as

$$\widehat{\omega}_{\text{MLPUE}} = \arg \min_{\omega \in \mathcal{B}} \min_{\mathbf{q} \in \mathbb{Z}^n} \|\widehat{\mathbf{x}} - \omega - \mathbf{V}\mathbf{q}\|. \quad (5.77)$$

More generally, multiple solutions may be accepted from the sphere decoding step, only to keep the projected solution maximizing the likelihood.

5.7 Examples

In this section we provide computed examples for the three ML solution methods given in the chapter.

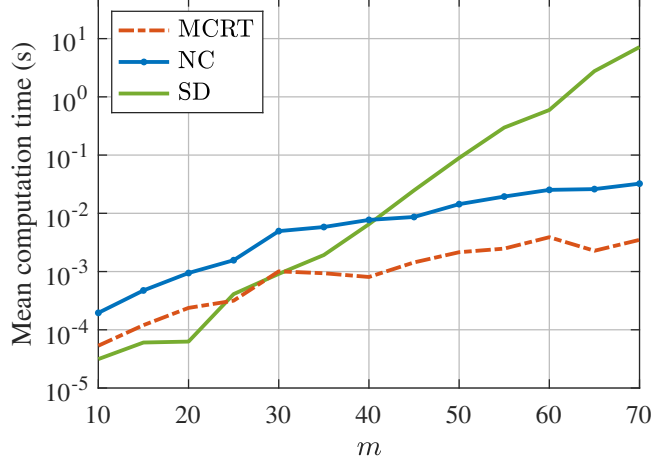


Figure 5.5: Mean computation time plotted against m for Algorithm 8 (MCRT), Algorithm 9 (NC), and the sphere decoding approach of [7] (SD).

5.7.1 Computation Time Comparison

First, we compare the computation times of the methods. We form a series of encoding matrices \mathbf{A} with $n = 2$ and $m \in [10, 70]$. Each such \mathbf{A} satisfies Condition 7 of Section 5.4 and has $\mathbf{V} = \mathbf{I}_2$ and $\mathbf{C} = \text{diag}([1/5, 1/3]^\top)$. The noise has covariance $\Sigma = 10\mathbf{I}_m$. In Fig. 5.5, the average times to compute the ML estimate via Algorithm 8, (labeled “MCRT”), Algorithm 9 (labeled “NC”), and the Algorithm 10 (labeled “SD”) are plotted. The results are from 100 random draws of noise and parameters \mathbf{x} and \mathbf{k} . The algorithms have identical estimation error; here, we compare only the computation time. Note that each algorithm uses a precomputation step that is not included in the timing results. This step need only be done once for a given \mathbf{A} and Σ . For reasonably small m , the sphere decoding approach is faster than either polynomial-time method proposed here. As m becomes large, the mean computation time of the sphere decoding approach increases rapidly [174]; indeed,

for no fixed signal-to-noise ratio (SNR) is there a polynomial upper bound on the expected complexity of sphere decoding that holds for all m [217]; the rate of the exponential function dictating average complexity depends on SNR. Compared to the polynomial-time algorithms, the sphere decoding approach also exhibits greater variance in computing time, which can result in undesirable decoding and complications in hardware implementation [218]. Nevertheless, the sphere decoding approach remains an efficient option for many practical problems involving (5.1).

5.7.2 Encoding Matrices in Application

Table 5.1 refers to encoding matrices from the literature and provides context relevant to the two polynomial time algorithms given in this chapter. For sensor arrays, \mathbf{A} corresponds to pairwise differences of the element positions. Table 5.1 shows that there are encoding matrices in use that meet Condition 7 and those that do not. For the \mathbf{A} matrix in [219], $K = K_{\text{NC}}$. Thus, despite meeting Condition 7, Algorithm 8 provides no computational benefit, and it is advisable to use the simpler Algorithm 9. Though, for the encoding matrix made by the wavelengths denoted “set E” in [178], the computational benefit is significant, with $K_{\text{NC}} > 585K$.

5.7.3 Direction of Arrival Estimation With Planar Array

As an example for the estimation approach in Section 5.6, we simulate array outputs from the cross-shaped five-element planar array configuration from [8], with positions given by $[0, 0]^T$, $[-4, 0]^T$, $[0, 4]^T$, $[0, -5]^T$, and $[5, 0]^T$, in units of one-half wavelength. The MLE in (5.73) is evaluated using a multi-resolution grid search over \mathbf{x} [212]. We use an initial grid search with a grid spacing in \mathbf{x} of 10^{-2} , followed

Table 5.1: Algorithms and problem sizes for examples of the linear model (5.1) appearing in application.

Reference	m	n	Condition 7	Alg.	K	K_{NC}
[216]	10	3	false	Alg. 9	-	264
[219]	10	3	true	Alg. 9	200	200
[220]	10	1	false	Alg. 9	-	34
[221]	10	2	false	Alg. 9	-	1143
[178, set E]	5	1	true	Alg. 8	5	2927
[178, set F]	5	1	false	Alg. 9	-	2311
Appendix B.5	16	3	true	Alg. 8	1087	8953

by a second search over the grid cells adjacent to the optimal point in the first search, with a spacing of 10^{-3} .

We compare the MLE in (5.73), which uses the full sample covariance data, amplitude and phase, to an approximate MLE that uses phase unwrapping estimation. The phase unwrapping estimator uses only the phase information of the sample covariance. Phase unwrapping estimation allows us to use Algorithm 10. To do so, we make the approximation that the phase noise follows a zero-mean Gaussian distribution with the covariance matrix determined as in [31]; see Appendix C for details.

The root mean square error (RMSE) results of the two estimators from $T = 10$ snapshots are plotted in Fig. 5.6 against signal-to-noise ratio (SNR) for a true angle of arrival of $\theta = \pi/2$, $\phi = \pi/4$. The results are computed from 10^6 trials at each SNR. The approximate MLE based on phase unwrapping (labeled MLPUE in Fig. 5.6, for “ML phase unwrapping estimator”) matches the RMSE performance of the MLE. The square root of the trace (denoted $\text{tr}(\cdot)$ in Fig. 5.6) of the Cramér-Rao bound (CRB) [222] of \mathbf{x} in (5.71) is also shown. Both estimators approach this lower bound

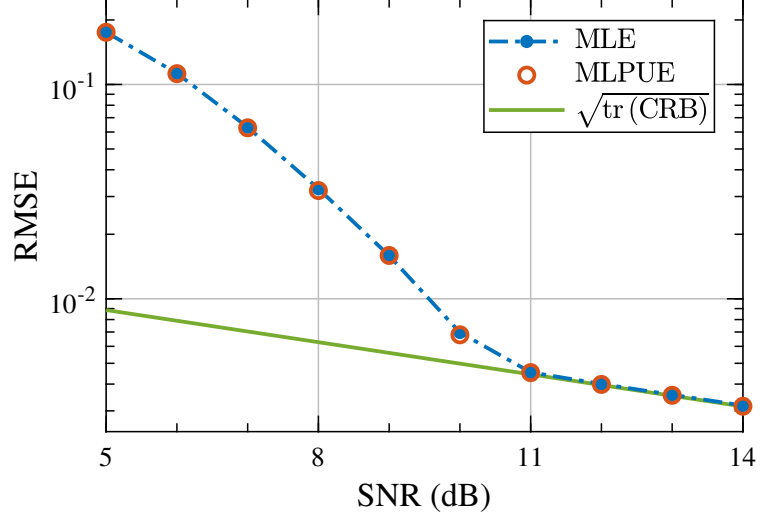


Figure 5.6: DoA estimation performance against SNR with $T = 10$ snapshots from five-element planar array [8]. RMSE results for the MLE of (5.73) and the phase unwrapping MLE are shown along with the square root of the trace of the CRB for \mathbf{x} . The results are averaged over 10^6 trials.

for SNR values exceeding a threshold [223] of roughly 11 dB. In this example, the MLPUE approach is on average over 85 times faster than the multi-resolution grid search.

5.7.4 Magnetic Resonance Flow Imaging

We next demonstrate the use of the estimation approach in Section 5.6 in phase-contrast magnetic resonance imaging (PC-MRI). PC-MRI encodes the velocity of spins into the image voxel phase through the first moment, $\mathbf{p} \in \mathbb{R}^3$, of the time varying gradient field:

$$b_k(t) = s(t)e^{i(\phi_0 + \gamma \mathbf{p}_k^T \mathbf{x})} + \varepsilon_k(t), \quad (5.78)$$

where $k = 1, 2, \dots, L$ indexes encodings, ϕ_0 is an unknown background phase, γ is the gyromagnetic ratio, $\mathbf{x} \in \mathbb{R}^3$ is velocity, and $\varepsilon_k(t)$ is zero-mean circular Gaussian

noise that is uncorrelated across T coils indexed by t . Proceeding with the sample covariance matrix as in Section 5.7.3, \mathbf{y} is formed from the complex measurements as in (5.75). Each entry in \mathbf{y} is a scaled phase difference from two encodings, averaged over coils, thereby canceling the unknown background phase.

Pelc et al. [9] introduced 3-directional flow encoding and a velocity estimator.

Phase differences from balanced four-point encoding yield

$$\mathbf{A}^T = \frac{\gamma\rho}{\pi} \begin{bmatrix} 1 & 1 & 0 & 0 & -1 & -1 \\ 1 & 0 & 1 & -1 & 0 & 1 \\ 0 & 1 & 1 & 1 & 1 & 0 \end{bmatrix}, \quad (5.79)$$

where ρ is the absolute value of the first moment in each spatial direction. Phase differences are then pre-processed to conveniently yield a decoupled set of three equations in $n = 3$ unknown velocity components. Specifically, the pre-processing is given by a linear operator $\mathbf{P} \in \mathbb{R}^{3 \times 6}$, given by

$$\mathbf{P} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & -1 \\ 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 \end{bmatrix}. \quad (5.80)$$

This preprocessing yields $\mathbf{PA} = \frac{2\gamma\rho}{\pi}\mathbf{I}$ and a least squares estimate, $\mathbf{x}_{\text{LS}} = \frac{\pi}{2\gamma\rho}\mathbf{P}\mathbf{y}$.




Figure 5.7: Velocity RMSE on $\mathcal{P}_{\mathbf{V}} \in \mathbb{R}^3$ for (a) the estimator in [9] and (b) the MLPUE. Velocity components and RMSE are shown in cm/s.

Fig. 5.7 shows RMSE versus true \mathbf{x} across one choice of the fundamental parallelopete $\mathcal{P}_{\mathbf{V}}$. The example is computed for $\text{var}(\varepsilon_k(t)) = 1$, $s(t) = 10$, $T = 8$ coils, and $\gamma\rho = \frac{\pi}{100}$ s/cm, with 10^3 trials at each \mathbf{x} . The estimator \mathbf{x}_{LS} is aliased on half of \mathcal{P} , as shown in Fig. 5.7a, whereas the MLPUE allows un-aliased estimation across most of \mathcal{P} , illustrated in Fig. 5.7b. With the MLPUE, wrap-around error is encountered only when the true speed is within approximately 2 cm/s of a boundary. With the four-point balanced encoding, the error covariance matrices for both \mathbf{x}_{LS} and the MLPUE are a scaled identity, signifying i.i.d. errors; but, the processing in [9] results in a 95% confidence sphere with 3.7% larger volume, because it ignores noise correlation in the phase differences. Finally, the sphere decoding for the MLPUE is, on average, over 500 times faster than grid-search estimation in (5.73) using 100 grid points per dimension [224].

5.8 Summary

In this chapter, we have proposed three methods for the ML estimation task in (5.2). Generally, the problem can be solved using the sphere decoding approach given in Section 5.6. The computation time of this approach is dependent on the noise realization, and the worst-case computational complexity is not polynomial in m , the length of the measurement vector. In [32], an ML solution method for (5.2) is given for the special case of a single unknown variable ($n = 1$), diagonal noise covariance, and a restriction on the encoding matrix \mathbf{A} . This method has worst-case polynomial-time complexity. In this chapter, we have presented estimation algorithms that extend polynomial-time computation to a wider range of problems with diagonal covariance. The extensions presented here account for both multivariate unknowns

($n > 1$) and general encoding matrices with full column rank. The method given in Section 5.4 use a multivariate CRT to further reduce algorithm complexity.

5.9 Future Work

In this chapter, we have described two ML estimation methods for the problem in (5.2) that have polynomial complexity in m , the number of measurements. Both of these methods require that the measurement noise be uncorrelated. We have also described a sphere decoding approach that accommodates correlated noise, though the worst-case computational complexity is non-polynomial in m . It remains to find a polynomial-time algorithm that accommodates correlated noise. A structure is present for the special case of (5.2) with measurements corresponding to phase differences and the covariance matrix Σ found according to the phase noise approximation described in Appendix C. This structure may admit an efficient estimation procedure.

Another avenue of continued work is the leveraging of the proposed techniques for MAP estimation. In this chapter, we have considered estimation from a single set of measurements. Though, often, as in phase-contrast MRI, a parameter is estimated for each voxel in an image. The per-voxel estimation proposed here could be a component in an estimation method that leverages spatial correlations to improve estimation performance, as in [31, 225].

It is often necessary to estimate the DoA of multiple targets simultaneously, and a notable limitation of the methods in this chapter is that they apply only to situations with a single target. Future research could extend these methods to apply to multiple targets. Such an extension would also allow for the simultaneous estimation

of multiple frequency components in frequency estimation problems and the heights of multiple scatterers within a resolution cell in interferometric SAR.

Chapter 6: Aliasing in Nonuniform Arrays

6.1 Introduction

A sensor array is said to be “alias-free” on some set \mathcal{B} of directions of arrival (DoAs) if array steering vectors are distinct for all DoAs in \mathcal{B} . In this chapter¹¹, a necessary and sufficient condition is presented for an arbitrary sensor array to be alias-free. The results apply to linear, planar, and 3D arrays; they hold for any predefined subset of DoAs for which the array must be unambiguous. The results extend existing literature to provide simple geometrical interpretation for use in array design.

In [33], Godara and Cantoni give necessary and sufficient conditions for alias-free arrays in one, two, or three dimensions. For L sensors, the test entails a combinatorial search over L -tuples of integers. For planar and 3D arrays, [33] also provides easily tested, but restrictive, sufficient conditions, requiring a smallest inter-sensor spacing not exceeding $\lambda/2$, where λ is the wavelength of the signal. The sufficient conditions have been widely used in design [226–228]. A slightly less restrictive sufficient condition for planar arrays is given in [229, Thm. 2].

Motivated by array design, alternative sufficient conditions have been devised for linear arrays. To avoid aliasing over the full range of azimuth angles, uniform linear

¹¹This chapter is largely drawn from [216], coauthored with Shen Zhao, Rizwan Ahmad, and Lee C. Potter (© 2022 IEEE).

arrays require an inter-sensor spacing that is no more than $\lambda/2$ [230]. It may be prohibitive to have two sensors within $\lambda/2$, due either to physical size limitations or mutual coupling. Nonuniform linear arrays allow for alias-free operation despite a minimum inter-sensor spacing that exceeds $\lambda/2$ [8, 221, 231, 232]. A difference *co-array* is a virtual array with positions given by the pairwise differences of the physical sensor locations. Particular attention has been given to array geometries for which the co-array has uniform sampling and is fully filled, such as minimum redundancy arrays [233], nested arrays [234], and co-prime arrays [235]. Thinned arrays refer to the special case of a uniform array from which some elements have been removed. For a thinned linear array of L sensors with $\lambda/2$ spacing, a necessary and sufficient condition for no aliasing is given in [236]: for reference sensor at 0 and positions $n_i\lambda/2$, the integers $\{n_2, n_3, \dots, n_L\}$ must have a greatest common divisor (gcd) of 1. Specialized to a nonuniform linear array of $L = 3$ sensors at positions $\{0, q_1, q_2\}$ with $0 < q_1 < q_2$, [237] posits the sufficient condition that $q_2/q_1 = m/\ell$ for co-prime integers m, ℓ , with $|\ell| > \lfloor 2q_1/\lambda \rfloor$ or $|m| > \lfloor 2q_2/\lambda \rfloor$. Likewise for $L = 3$, [213, 238] establish the sufficient condition that the co-array contains a pair of elements not more than $\lambda/2$ apart.

A task very similar to array design arises in phase-contrast magnetic resonance imaging [31, 168, 225, 239, 240]. There, the velocity of spins at a voxel is encoded into image phase differences. Encoded images are acquired by adopting a time-varying gradient field, and the first moments of these waveforms are analogous to sensor positions. For velocity in one or several dimensions, unambiguous performance is required up to some maximum speed, such as 300 cm/s.

6.2 Signal Model

Consider data collected from an array with L sensors receiving a narrow-band signal from a single emitter located in the far field. The demodulated complex-valued $L \times 1$ array output vector is given by

$$\mathbf{b}(t) = \mathbf{a}(\mathbf{x})s(t) + \boldsymbol{\varepsilon}(t), \quad t = 1, 2, \dots, T, \quad (6.1)$$

where $s(t) \in \mathbb{C}$ is the baseband complex amplitude of the impinging wavefront, T is the number of snapshots, $\mathbf{a}(\mathbf{x}) \in \mathbb{C}^{L \times 1}$ is the array steering vector, $\boldsymbol{\varepsilon}(t) \in \mathbb{C}^{L \times 1}$ is circularly symmetric Gaussian noise with zero mean and covariance $\sigma^2 \mathbf{I}_L$, and $\boldsymbol{\varepsilon}(t)$ is uncorrelated across snapshots.

For an array of isotropic sensors in n dimensions, the array steering vector is given by

$$\mathbf{a}(\mathbf{x}) = [e^{i\pi \mathbf{p}_1^T \mathbf{x}} \quad e^{i\pi \mathbf{p}_2^T \mathbf{x}} \quad \dots \quad e^{i\pi \mathbf{p}_L^T \mathbf{x}}]^T, \quad (6.2)$$

where $i = \sqrt{-1}$ and the vector $\mathbf{x} \in \mathbb{R}^d$ corresponds to the DoA of the impinging signal. The position vectors $\mathbf{p}_j \in \mathbb{R}^n$, for $j \in \{1, 2, \dots, L\}$, give the sensor positions in units of one-half wavelength. For linear arrays, $n = 1$ and $x = \sin \theta$, where θ is the azimuth angle of the source. For planar arrays, $n = 2$ and $\mathbf{x} = [\cos \theta \cos \phi, \sin \theta \cos \phi]^T$, with ϕ denoting the elevation angle to the source. For general arrays, $n = 3$ and $\mathbf{x} = [\cos \theta \cos \phi, \sin \theta \cos \phi, \sin \phi]^T$.

6.3 Alias-free

For an array to be alias-free, the steering vectors in (6.2) must be distinct for distinct DoAs. Using the first sensor as a phase reference, the $(L-1)$ phase differences

are $\pi [\mathbf{p}_k - \mathbf{p}_1]^\top \mathbf{x}$, for $k \in \{2, 3, \dots, L\}$. We define the $(L - 1)$ -by- n encoding matrix

$$\mathbf{A} = \frac{1}{2} [(\mathbf{p}_2 - \mathbf{p}_1), (\mathbf{p}_3 - \mathbf{p}_1), \dots, (\mathbf{p}_L - \mathbf{p}_1)]^\top, \quad (6.3)$$

and assume that $\text{rank}(\mathbf{A}) = n$ [33, Thm. 4]. Note that the encoding matrix \mathbf{A} is determined entirely by the geometry of the array. Next, define \mathcal{B} as the set of arrival directions. For $n = 1$, the set $\mathcal{B}_1 = \{\mathbf{x} \mid \theta \in [-\frac{\pi}{2}, \frac{\pi}{2}]\}$ is the unit ball in \mathbb{R} , $\|\mathbf{x}\| \leq 1$. For a planar array, $n = 2$ and the set $\mathcal{B}_2 = \{\mathbf{x} \mid \theta \in [0, 2\pi), \phi \in [0, \frac{\pi}{2}]\}$ of all DoAs is the unit ball in \mathbb{R}^2 , $\|\mathbf{x}\| \leq 1$. These DoAs correspond to source directions on the upper hemisphere. If, for the planar array, the source is constrained to lie in the plane of the sensors, then \mathcal{B}'_2 is the unit circle in the plane. For a 3D array, $n = 3$ and the set $\mathcal{B}_3 = \{\mathbf{x} \mid \theta \in [0, 2\pi), \phi \in [-\frac{\pi}{2}, \frac{\pi}{2}]\}$ of all DoAs is a sphere in \mathbb{R}^3 , $\{\mathbf{x} \mid \|\mathbf{x}\| = 1\}$. These DoAs correspond to source directions on the sphere. Note that \mathcal{B} may alternatively be defined as some subset of \mathcal{B}_n in n dimensions, for example, with prior knowledge restricting the range of possible DoAs [241].

Definition 10 (Alias-free on \mathcal{B}). *An array with the encoding matrix \mathbf{A} is alias-free on \mathcal{B} if, for any $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{B}$,*

$$\mathbf{A}(\mathbf{x}_2 - \mathbf{x}_1) \in \mathbb{Z}^{L-1} \Rightarrow \mathbf{x}_2 = \mathbf{x}_1.$$

When considering the set of all directions of arrival, $\mathcal{B} = \mathcal{B}_n$ and the array is simply said to be “alias-free.”

Thus, we consider the system of linear congruence equations

$$\mathbf{A}\mathbf{x} \equiv \mathbf{0}_{L-1} \pmod{\mathbf{1}_{L-1}}, \quad (6.4)$$

which can be reformulated as

$$\mathbf{A}\mathbf{x} = \mathbf{k}, \quad (6.5)$$

where $\mathbf{k} \in \mathbb{Z}^{L-1}$ is called a vector of *wrapping integers*. The set of solutions to (6.4) for an array with encoding matrix \mathbf{A} is

$$\Omega_{\mathbf{A}} = \{\mathbf{x} \in \mathbb{R}^n \mid \mathbf{A}\mathbf{x} \in \mathbb{Z}^{L-1}\}. \quad (6.6)$$

If (6.4) admits any nonzero solution, then $\Omega_{\mathbf{A}} = \Lambda_{\mathbf{V}}$; that is, the solutions form a lattice $\Lambda_{\mathbf{V}} \subset \mathbb{R}^n$ [181, 242, 243], for some lattice basis \mathbf{V} . The solution set $\Omega_{\mathbf{A}}$ is thus either $\{\mathbf{0}\}$ or a lattice; in either case the solution set is closed under addition. Note that if $\Omega_{\mathbf{A}} = \{\mathbf{0}\}$, then the array is said to be alias-free on all \mathbb{R}^n .

We have that $\Omega_{\mathbf{A}} = \Lambda_{\mathbf{V}}$ if and only if, for all $\mathbf{x} \in \mathbb{R}^n$,

$$\mathbf{A}\mathbf{x} \in \mathbb{Z}^{L-1} \Leftrightarrow \mathbf{x} \in \Lambda_{\mathbf{V}}. \quad (6.7)$$

Definition 11. A matrix $\mathbf{Z} \in \mathbb{Z}^{c \times d}$ with $c \geq d$ is called *primitive* if, for all $\mathbf{q} \in \mathbb{Q}^d$ [244],

$$\mathbf{Z}\mathbf{q} \in \mathbb{Z}^c \Rightarrow \mathbf{q} \in \mathbb{Z}^d. \quad (6.8)$$

A column vector is *primitive* if its elements have a greatest common divisor of unity [245]. Any c -by- d primitive matrix can be extended to a c -by- c unimodular matrix [186, 244].

Lemma 12. A matrix $\mathbf{Z} \in \mathbb{Z}^{c \times d}$ with $c \geq d$ is primitive if and only if $\mathbf{U}\mathbf{Z}$ is primitive for any unimodular $\mathbf{U} \in \mathbb{Z}^{c \times c}$.

Proof. For any unimodular $\mathbf{U} \in \mathbb{Z}^{c \times c}$ and any $\mathbf{q} \in \mathbb{Q}^d$,

$$\mathbf{U}\mathbf{Z}\mathbf{q} \in \mathbb{Z}^c \Leftrightarrow \mathbf{Z}\mathbf{q} \in \mathbf{U}^{-1}\mathbb{Z}^c \Leftrightarrow \mathbf{Z}\mathbf{q} \in \mathbb{Z}^c.$$

From Definition 11, \mathbf{Z} is primitive if and only if $\mathbf{U}\mathbf{Z}$ is primitive. □

Lemma 13. For invertible $\mathbf{V} \in \mathbb{R}^{n \times n}$, if $\mathbf{Z} = \mathbf{A}\mathbf{V}$ is primitive, then $\Omega_{\mathbf{A}} = \Lambda_{\mathbf{V}}$.

Proof. If $\mathbf{Z} = \mathbf{A}\mathbf{V}$ is primitive, then for all $\mathbf{x} \in \mathbb{R}^n$,

$$\begin{aligned} \mathbf{A}\mathbf{x} \in \mathbb{Z}^{L-1} &\Leftrightarrow \mathbf{Z}\mathbf{V}^{-1}\mathbf{x} \in \mathbb{Z}^{L-1} \Leftrightarrow \mathbf{V}^{-1}\mathbf{x} \in \mathbb{Z}^n \\ &\Leftrightarrow \mathbf{x} \in \mathbf{V}\mathbb{Z}^n \Leftrightarrow \mathbf{x} \in \Lambda_{\mathbf{V}}. \end{aligned} \quad (6.9)$$

For the second equivalence, sufficiency is given by Definition 11, and necessity is due to \mathbf{Z} being integer-valued. From (6.7), $\Omega_{\mathbf{A}} = \Lambda_{\mathbf{V}}$. \square

Proposition 14. *For invertible $\mathbf{V} \in \mathbb{R}^{n \times n}$, $\Omega_{\mathbf{A}} = \Lambda_{\mathbf{V}}$ if and only if $\mathbf{A} = \mathbf{Z}\mathbf{V}^{-1}$, for some primitive matrix $\mathbf{Z} \in \mathbb{Z}^{(L-1) \times n}$.*

Proof. Necessity: This is given by Lemma 13.

Sufficiency: Assume that $\Omega_{\mathbf{A}} = \Lambda_{\mathbf{V}}$. From (6.7), for all $\mathbf{x} \in \mathbb{R}^n$,

$$\mathbf{A}\mathbf{x} \in \mathbb{Z}^{L-1} \Leftrightarrow \mathbf{x} \in \Lambda_{\mathbf{V}} \Leftrightarrow \mathbf{x} \in \mathbf{V}\mathbb{Z}^n. \quad (6.10)$$

Thus, $\mathbf{A}\mathbf{V}\mathbb{Z}^n \subset \mathbb{Z}^{L-1}$, which is true if and only if $\mathbf{A}\mathbf{V} \in \mathbb{Z}^{(L-1) \times n}$.

From (6.10), for any $\mathbf{q} \in \mathbb{Q}^n$,

$$\mathbf{A}\mathbf{V}\mathbf{q} \in \mathbb{Z}^{L-1} \Leftrightarrow \mathbf{V}\mathbf{q} \in \mathbf{V}\mathbb{Z}^n \Leftrightarrow \mathbf{q} \in \mathbb{Z}^n. \quad (6.11)$$

Therefore, from Definition 11, $\mathbf{Z} = \mathbf{A}\mathbf{V}$ is primitive. \square

Corollary 15. *For any invertible $\mathbf{V} \in \mathbb{R}^{n \times n}$, $\Omega_{\mathbf{A}} = \Lambda_{\mathbf{V}}$ if and only if $\Omega_{\mathbf{U}\mathbf{A}} = \Lambda_{\mathbf{V}}$, for any unimodular $\mathbf{U} \in \mathbb{Z}^{(L-1) \times (L-1)}$.*

Proof. From Lemma 12, $\mathbf{A}\mathbf{V}$ is primitive if and only if $\mathbf{U}\mathbf{A}\mathbf{V}$ is primitive. Thus, according to Lemma 13, $\Omega_{\mathbf{A}} = \Lambda_{\mathbf{V}}$ if and only if $\Omega_{\mathbf{U}\mathbf{A}} = \Lambda_{\mathbf{V}}$. \square

Corollary 16. *For any invertible $\mathbf{V} \in \mathbb{R}^{n \times n}$, $\Omega_{\mathbf{A}} = \Lambda_{\mathbf{V}}$ if and only if $\Omega_{\mathbf{A}\mathbf{S}} = \Lambda_{\mathbf{S}^{-1}\mathbf{V}}$, for any invertible $\mathbf{S} \in \mathbb{R}^{n \times n}$.*

Proof. From Proposition 14, if $\Omega_{\mathbf{A}} = \Lambda_{\mathbf{V}}$, then $\mathbf{A} = \mathbf{Z}\mathbf{V}^{-1}$, for some primitive $\mathbf{Z}^{(L-1) \times n}$. This is true if and only if $\mathbf{A}\mathbf{S} = \mathbf{Z}\mathbf{V}^{-1}\mathbf{S}$, for any invertible $\mathbf{S} \in \mathbb{R}^{n \times n}$, which is true if and only if $\Omega_{\mathbf{A}\mathbf{S}} = \Lambda_{\mathbf{S}^{-1}\mathbf{V}}$. \square

For linear arrays, the lattice basis becomes a scalar $v \in \mathbb{R}$. From Proposition 14, we have that a 1D array has $\Omega_{\mathbf{A}} = \Lambda_v$ if and only if the sensor displacements in (6.3) are setwise coprime integer multiples of a common real-valued factor $1/v$.

If \mathbf{A} is rational, $\Omega_{\mathbf{A}} = \Lambda_{\mathbf{V}}$ is a full-rank lattice with invertible basis $\mathbf{V} \in \mathbb{R}^{n \times n}$; the basis \mathbf{V} can be formed as in (5.16), (5.17), (5.18), (5.18) in Section 5.3.4 [7], with $\mathbf{M} = \mathbf{I}_{L-1}$. Therefore, a sufficient condition for $\Omega_{\mathbf{A}}$ to be a full-rank lattice is for \mathbf{A} to be rational, i.e., for the sensors to have been placed on a Cartesian grid, such as might naturally happen in fabrication. The parallelotope $\mathcal{P}_{\mathbf{V}}$ defines a region in which the solution to (6.4) is unique.

For a given DoA \mathbf{x}_0 , let $\phi_0 = 2\pi(\mathbf{A}\mathbf{x}_0 - \mathbf{k}_0) = \angle \mathbf{a}(\mathbf{x}_0)$ be the phases of the steering vector; \mathbf{k}_0 is the associated vector of wrapping integers. The least-squares solution $\hat{\mathbf{x}} = \left\langle \mathbf{A}^\dagger \left(\frac{\phi_0}{2\pi} + \hat{\mathbf{k}} \right) \right\rangle_{\mathbf{V}}$ for $\hat{\mathbf{k}} \neq \mathbf{k}_0$ corresponds to a beampattern side-lobe. Here, $\langle \cdot \rangle_{\mathbf{V}}$ refers to the operation defined in (5.7), by which a vector is translated to the parallelotope $\mathcal{P}_{\mathbf{V}}$.

6.4 Necessary and Sufficient Condition

For the remainder of this chapter, we assume that $\Omega_{\mathbf{A}} = \Lambda_{\mathbf{V}}$, for some invertible $\mathbf{V} \in \mathbb{R}^{n \times n}$. From Definition 10, we immediately have:

Proposition 17. *An array is alias-free on \mathcal{B} if and only if translations $\mathcal{B} + \{\mathbf{u}_i\}$ and $\mathcal{B} + \{\mathbf{u}_j\}$ are disjoint for all $\mathbf{u}_i \neq \mathbf{u}_j \in \Lambda_{\mathbf{V}}$.*

Specializing to the case of the full range of DoAs, defined as \mathcal{B}_n above, yields the following corollary.

Corollary 18. *An array is alias-free if and only if*

$$\min_{\mathbf{u} \in \Lambda_{\mathbf{V}} \setminus \{\mathbf{0}\}} \|\mathbf{u}\| > 2, \quad (6.12)$$

i.e., the smallest norm of a non-zero lattice point in $\Lambda_{\mathbf{V}}$ is greater than 2.

Proof. Recall \mathcal{B}_n as defined in 6.3: unit balls in \mathbb{R} and \mathbb{R}^2 or the sphere in \mathbb{R}^3 . Then, $\mathcal{B} + \{\mathbf{u}_i\}$ and $\mathcal{B} + \{\mathbf{u}_j\}$ have an empty intersection for all $\mathbf{u}_i \neq \mathbf{u}_j \in \Lambda_{\mathbf{V}}$ if and only if $\|\mathbf{u}_j - \mathbf{u}_i\| > 2$ for all $\mathbf{u}_i \neq \mathbf{u}_j \in \Lambda_{\mathbf{V}}$. As lattices are closed under addition, the array is alias-free if and only if $\|\mathbf{u}\| > 2$, for all $\mathbf{u} \in \Lambda_{\mathbf{V}} \setminus \mathbf{0}$. \square

Godara and Cantoni [33, Thm. 6] provide an alternative necessary and sufficient test for alias-free on the full set of DoAs: the image of the unit ball under \mathbf{A} contains no integer point other than the origin. The test is applied by combinatorial search over $\mathbf{k} \in \mathbb{Z}^{L-1}$, with $\|\mathbf{k}\| < \|\mathbf{A}\|_2$. In contrast, construction of the lattice basis \mathbf{V} provides for an easily tested necessary and sufficient condition for arbitrary \mathcal{B} . Specialization to the set of all possible directions of arrival in one or several dimensions is considered below.

6.4.1 Linear Arrays

For $n = 1$, the fundamental parallelotope $\mathcal{P}_{\mathbf{v}} \subset \mathbb{R}$ reduces to an interval and is unique, up to a sign. By the Chinese remainder theorem and Corollary 18, a linear array is alias-free if and only if

$$\text{LCM}(\mathbf{1}_{(L-1)} \odot \mathbf{A}) > 2, \quad (6.13)$$

where $\text{LCM}(\cdot)$ denotes the least common multiple, which is the smallest positive real number that is an integer multiple of all input numbers. For example,

$$\text{LCM}\left(\frac{\pi}{3}, \frac{\pi}{2}, \frac{2\pi}{3}\right) = 2\pi.$$

Equation (6.13) is an equivalent and alternative test to [33, Thm.11]: there exists sensor separation less than or equal to $\lambda/2$, or for minimum spacing $q\alpha > \lambda/2$ with integer q , there exists a third element such that its distance from any one of the two elements is not an integer multiple of α . The simplicity of (6.13) aids in design.

6.4.2 Nonlinear Arrays

For planar or 3D arrays, the lattice basis, and hence the corresponding fundamental parallelootope, are not unique. Equation (6.12) provides a simple necessary and sufficient test, in contrast to the widely used sufficient condition [33, Thm.8]. Note that evaluating (6.12) is equivalent to solving a shortest lattice vector problem in n dimensions.

6.4.3 Invariances and Scaling

From (6.6), we see that $\Lambda_{\mathbf{V}}$ is invariant to translations of the point set $\{\mathbf{p}_1, \dots, \mathbf{p}_L\}$. According to Corollary 15, $\Lambda_{\mathbf{V}}$ is invariant to unimodular transformation of \mathbf{A} . Corollary 16 describes the effect of an invertible transformation of \mathbf{A} on the lattice of solutions to (6.4); this property can be used to examine the effects of dilation or contraction of a point set in array design. We also have the following property showing that an alias-free array remains alias-free following orthonormal transformation (including rotation or reflection) of its point set.

Proposition 19. *If the array with encoding matrix \mathbf{A} is alias-free, then the array with encoding matrix \mathbf{AB}^\top is alias-free, for any n -by- n matrix \mathbf{B} with orthonormal columns.*

Proof. Assume that $\Omega_{\mathbf{A}} = \Lambda_{\mathbf{V}}$ and that the array with encoding matrix \mathbf{A} is alias-free. By Corollary 16, $\Omega_{\mathbf{AB}^\top} = \Lambda_{\mathbf{BV}}$ (note that $\mathbf{B}^{-1} = \mathbf{B}^\top$). From Corollary 18,

$$\begin{aligned} \min_{\mathbf{u} \in \Lambda_{\mathbf{V}} \setminus \{\mathbf{0}\}} \|\mathbf{u}\| > 2 &\Leftrightarrow \min_{\boldsymbol{\lambda} \in \mathbb{Z}^n \setminus \{\mathbf{0}\}} \|\mathbf{V}\boldsymbol{\lambda}\| > 2 \\ &\Leftrightarrow \min_{\boldsymbol{\lambda} \in \mathbb{Z}^n \setminus \{\mathbf{0}\}} \|\mathbf{BV}\boldsymbol{\lambda}\| > 2 \Leftrightarrow \min_{\mathbf{u} \in \Lambda_{\mathbf{BV}} \setminus \{\mathbf{0}\}} \|\mathbf{u}\| > 2. \end{aligned}$$

Therefore, the array with encoding matrix \mathbf{AB}^\top is alias-free. \square

6.5 Examples

6.5.1 Linear Array

In [221] a four-element linear array is proposed for meteor trajectory estimation with element positions given, in units of one-half wavelength, by $\{0, 4, 9, 40\}$. The array, an extension of an interferometric system in [8], is alias-free for the full range of azimuth angles. The application requires a minimum inter-sensor spacing much larger than $\lambda/2$ to limit the effects of mutual coupling. As an example, we design an alternative array configuration to minimize the Cramér-Rao lower bound (CRB) on estimation error variance subject to the constraints that the minimum inter-sensor spacing and peak side-lobe level meet those of the array in [221].

The source signal $\{s(t)\}$ follows a zero-mean circularly symmetric Gaussian distribution with power ρ^2 and is independently and identically distributed across snapshots, indexed by t [22, 230]. Under this model with T snapshots and L sensors, the

CRB is given by (e.g., [222, 223])

$$\text{CRB} = \frac{L \cdot \text{snr} + 1}{2\pi^2 T L^2 M \text{snr}^2} \quad (6.14)$$

where snr denotes the signal-to-noise ratio ρ^2/σ^2 and

$$M = \frac{1}{L} \sum_{k=1}^L \left(p_k - \frac{1}{L} \sum_{n=1}^L p_n \right)^2 \quad (6.15)$$

is the moment of inertia.

For a 20 dB signal-to-noise ratio as in [8], the array in [221] has a CRB of 5.09×10^{-7} and a peak side-lobe level of 0.369 dB. Assuming that $|\theta| < 80^\circ$ [221], we follow a design procedure similar to that in [31] to obtain the element locations $0.986 \cdot \{0, 5, 41, 78\}$. This new array configuration has a peak side-lobe level of 0.369 dB, a reduced CRB of 1.25×10^{-7} , and from (6.13), an unambiguous interval for $\sin \theta$ of 1.97.

Next, we interpret (6.13) for an arbitrary three-element linear array with rational positions. Without loss of generality, let $p_1 = 0$ and $(p_2 - p_1) \leq (p_3 - p_2)$. Select co-prime integers q and ℓ such that $\frac{p_3 - p_1}{p_3 - p_2} = \frac{q}{\ell} \in (1, 2]$. Then, from (6.13) with $\mathcal{B} = (-1, 1)$, the array is alias-free on \mathcal{B} if and only if

$$\frac{q - \ell}{(p_2 - p_1)} \geq 1. \quad (6.16)$$

In [213, 214, 238], the co-array is constrained to have a pair of positions not more than $\lambda/2$ apart; from (6.16), this sufficient condition is not necessary, allowing additional design flexibility. If $[-1, 1]$ is a strict subset of the unambiguous interval $\mathcal{P}_{\mathbf{V}}$, then the two end-fire DoAs, $\pm \sin(\pi/2)$, are disambiguated.

6.5.2 Planar Arrays

Consider a uniform circular array with $L = 6$ and one-wavelength spacing [246–248]. The lattice has basis $\mathbf{V} = \text{diag}(\{2, 2\sqrt{3}/3\})$. Corollary 18 is violated, because a vertex in $\mathcal{P}_{\mathbf{V}}$, given by $[0, \frac{2}{3}\sqrt{3}]^T$, has norm $\frac{2}{3}\sqrt{3} \approx 1.1547 < 2$. This is evident in Fig. 6.1a, where the boundary of \mathcal{B}_2 is shown translated to each of the several lattice points displayed, and the translated unit balls overlap. For a regular hexagonal array, the largest spacing at which the unit balls become tangent is $\frac{\sqrt{3}}{3}\lambda \approx 0.5774$ wavelength. The minimum spacing $\frac{\sqrt{3}}{3}\lambda > \frac{1}{2}\lambda$ violates the widely used sufficient condition [33, Thm. 9].

Fig. 6.1b depicts the lattice and shifted unit balls for a four-element array with positions $[0, 0]^T$, $[\pm 1.71, 0]^T$, and $[0.26, 2.11]^T$. The array is alias-free on \mathcal{B}_2 , although no fundamental parallelogram contains a unit ball and the smallest inter-sensor spacing is 0.855λ .

Fig. 6.1c shows the lattice and shifted \mathcal{B} for a four-element array with a reduced set of DoAs given by $\mathcal{B} = \{\mathbf{x} \mid \theta \in [-\pi/3, \pi/3], \phi \in [0, \pi/2]\}$. Note that the array is not alias-free for the full set of DoAs \mathcal{B}_2 ; the minimum separation between the parallelogram vertices is 1.335, and the condition of Corollary 18 is not satisfied. However, for the subset \mathcal{B} , the condition in Proposition 17 is satisfied, and the array is alias-free on \mathcal{B} . The lack of aliasing is illustrated by the disjointedness of the hourglass-shaped translations of \mathcal{B} in Fig. 6.1c. The array is designed from the desired parallelogram via $\mathbf{A} = \mathbf{Z}\mathbf{V}^{-1}$ with primitive matrix \mathbf{Z} yielding positions $[0, 0]^T$, $[2.67, 0.42]^T$, $[1.78, 2.48]^T$, and $[0.89, 4.54]^T$; the minimum separation between sensors is 1.122λ . This example shows how Proposition 17 may be used to design arrays that avoid aliasing over a subset of DoAs, as is sought in [221].

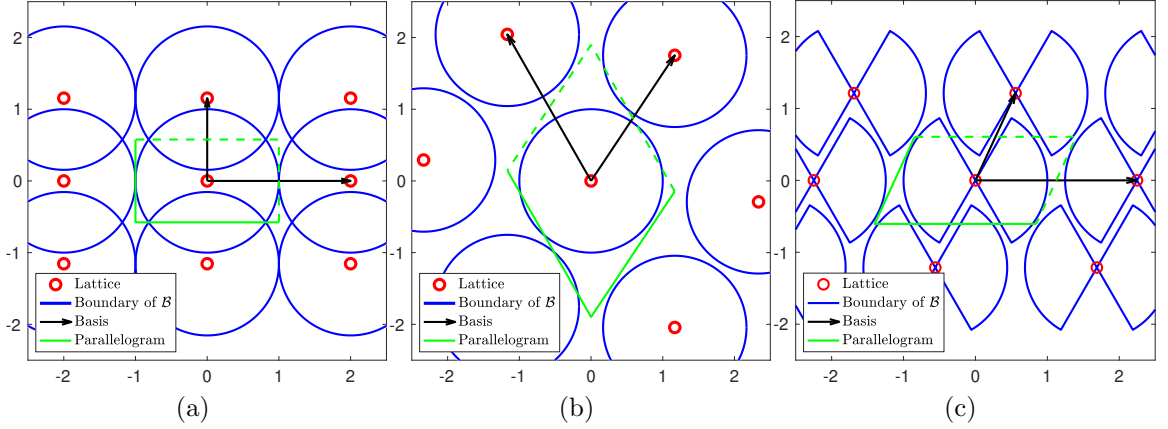


Figure 6.1: (a) For a uniform circular array of six sensors with λ spacing, illustration of the lattice Λ , a basis \mathbf{V} , a fundamental parallelogram, and the disc of all DoAs, \mathcal{B}_2 , translated periodically per the lattice. In this case, the translated discs intersect, and the array is aliased. (b) The lattice and shifted \mathcal{B}_2 for a four-sensor planar array that is alias-free, despite a minimum sensor separation of 0.855λ . (c) A second four-sensor planar array in which \mathcal{B} is the set of DoAs with $|\theta| \leq 60^\circ$; the array is alias-free on \mathcal{B} .

6.5.3 3D Array

A bi-tetrahedron of five sensors [219, 249, 250] is shown in Fig. 6.2a with positions $[0, 0, \pm 1]^\top$, $[2, 0, 0,]^\top$, and $[-1, \pm 1, 0,]^\top$ in units of one-half wavelength. The positions correspond to the vertices of an irregular polyhedron and are chosen so that there exists a cubic fundamental parallelotope circumscribing a sphere of diameter 2. The minimum distance between sensors is $\sqrt{3}\lambda/2$. This is 73% too large to satisfy the commonly used sufficient condition [33, Thm. 8], which requires that the minimum distance between sensors not exceed $\lambda/2$. The tangent points on each face correspond to the three pairs (az, el) of ambiguous DoAs, $(0, \pm 90^\circ)$, $(\pm 90^\circ, 0)$, and $(90^\circ \pm 90^\circ, 0)$; any small contraction of the array geometry removes these points of tangency. The

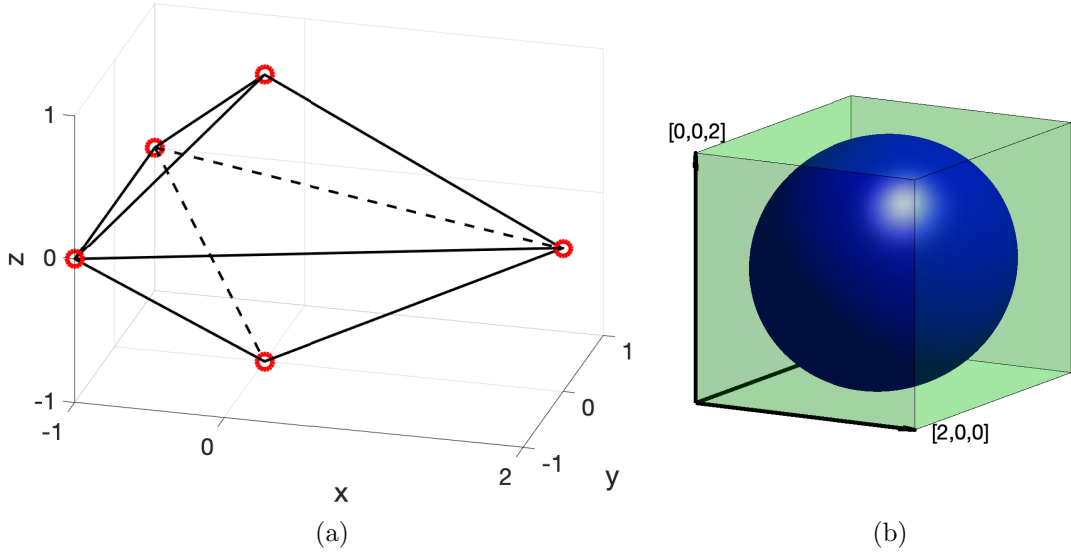


Figure 6.2: For a bi-tetrahedral (a) array of five sensors, (b) illustration of a fundamental parallelotope and the sphere of all DoAs, \mathcal{B}_3 , translated to the center of the parallelotope. Spheres translated to lattice points are merely tangent; the array is alias-free other than three pairs of ambiguous DoAs.

array is not isotropic: the CRB for x_1 is three times smaller than for x_2 and x_3 . The unambiguous set for 3D arrays is underestimated in some prior literature, e.g., [9, 219].

6.6 Summary

In this chapter, the results in [33] have been extended to provide a simply tested necessary and sufficient condition characterizing array geometries that provide unambiguous steering vectors. First, Proposition 17 considers an arbitrary set, \mathcal{B} , of candidate DoAs, with application illustrated in Fig. 6.1c. In contrast, [33] only considers the case in which \mathcal{B} corresponds to all possible DoAs. Second, the test in (6.12) provides a simple necessary and sufficient condition for alias-free array geometries in

2D and 3D, offering a less restrictive alternative to a widely used sufficient condition [33, Thms. 8, 9]. The test in (6.12) is easily applied, yet it offers increased flexibility for array geometry design over the sufficient condition in [33]. Third, the direct test for 1D arrays in (6.13) is simpler than [33, Thm. 11] and therefore advantageous in a design procedure. The new results facilitate the design of nonuniform arrays for widely separated sensors or increased aperture without an increase in the number of sensors. The unambiguous set of spatial frequencies was shown to be a parallelotope, admitting both intuitive geometric interpretation and application of the test for any subset of DoAs.

6.7 Future Work

The properties given in this chapter may allow for efficient 2D and 3D sensor array design and analysis techniques similar to those offered for 1D arrays in [31, 178]. In particular, Lemma 13 allows for an array design process that begins with the specification of a desired lattice basis \mathbf{V} . Next, a set of candidate primitive matrices can be formed. For each \mathbf{Z} in the set, the encoding matrix \mathbf{A} is given by $\mathbf{A} = \mathbf{Z}\mathbf{V}^{-1}$, and from \mathbf{A} , the sensor locations can be determined. The result is a set of sensor configurations that can be analyzed with respect to the CRB and the probability of an outlier [223]. This approach may lead to an efficient design process, in which only the sensor configurations with a desirable unambiguous region will be considered. Corollary 16 can also be used to scale candidate array configurations.

Note that the signal model in (6.2) uses the traditional assumption of isotropic sensor elements. In [251], ambiguity analysis with arbitrary complex-valued gain

functions per sensor was considered. Future research might address these alternative scenarios.

Appendix A: Variational Inference Details for Chapter 3

The gamma distribution's probability density is denoted by

$$\mathcal{G}(v; \alpha, \beta) = \frac{\beta^\alpha v^{\alpha-1}}{\Gamma(\alpha)} \exp\{-\beta v\}, \quad (\text{A.1})$$

with shape parameter $\alpha > 0$ and rate parameter $\beta > 0$. The inverse gamma distribution's probability density is denoted by

$$\mathcal{IG}(v; \alpha, \beta) = \frac{\beta^\alpha v^{-(\alpha+1)}}{\Gamma(\alpha)} \exp\{-\beta v^{-1}\} \quad (\text{A.2})$$

with shape parameter $\alpha > 0$ and scale parameter $\beta > 0$. From (3.7) and (3.14), the variational factor associated with $r_i^{(1)}$ is updated as

$$\begin{aligned} \log q(r_i^{(1)}) &= \log p(y_i^{(1)} | r_i^{(1)}) + \langle \log p(r_i^{(1)} | \mathbf{z}^{(1)}) \rangle_{q(\mathbf{z}^{(1)})} \\ &\quad + \text{const.} \\ &= -\log r_i^{(1)} - \frac{|y_i^{(1)}|^2}{r_i^{(1)}} - (1 + a|S_i^r|) \log r_i^{(1)} \\ &\quad - \sum_{j \in S_i^r} \frac{a \langle z_j^{(1)} \rangle}{r_i^{(1)}} + \text{const.}, \end{aligned} \quad (\text{A.3})$$

giving $q(r_i^{(1)}) = \mathcal{IG}(r_i^{(1)}; \alpha_{r,i}^{(1)}, \beta_{r,i}^{(1)})$, with

$$\alpha_{r,i}^{(1)} = 1 + a|S_i^r| \quad (\text{A.4a})$$

$$\beta_{r,i}^{(1)} = |y_i^{(1)}|^2 + \sum_{j \in S_i^r} a \langle z_j^{(1)} \rangle. \quad (\text{A.4b})$$

The factor associated with $r_i^{(2)}$ is updated as

$$\begin{aligned} \log q(r_i^{(2)}) &= \langle \log p(y_i^{(2)} \mid y_i^{(1)}, r_i^{(2)}, c_i) \rangle_{q(\mathbf{c})} \\ &\quad + \langle \log p(r_i^{(2)} \mid \mathbf{z}^{(2)}) \rangle_{q(\mathbf{z}^{(2)})} + \text{const.} \end{aligned} \quad (\text{A.5})$$

$$\begin{aligned} &= -\langle \delta(c_i + 1) \rangle \left[\frac{|y_i^{(2)} - \gamma e^{-j\phi_i} y_i^{(1)}|^2}{r_i^{(2)}(1 - \gamma^2)} \right] \\ &\quad - \langle \delta(c_i - 1) \rangle \left[\frac{|y_i^{(2)}|^2}{r_i^{(2)}} \right] - (1 + a|S_i^r|) \log r_i^{(2)} \\ &\quad - \log r_i^{(2)} - \sum_{j \in S_i^r} \frac{a \langle z_j^{(2)} \rangle}{r_i^{(2)}} + \text{const.}, \end{aligned} \quad (\text{A.6})$$

giving $q(r_i^{(2)}) = \mathcal{IG}(r_i^{(2)}; \alpha_{r,i}^{(2)}, \beta_{r,i}^{(2)})$, with

$$\alpha_{r,i}^{(2)} = 1 + a|S_i^r| \quad (\text{A.7a})$$

$$\begin{aligned} \beta_{r,i}^{(2)} &= \frac{\langle \delta(c_i + 1) \rangle |y_i^{(2)} - \gamma e^{-j\phi_i} y_i^{(1)}|^2}{(1 - \gamma^2)} \\ &\quad + \langle \delta(c_i - 1) \rangle |y_i^{(2)}|^2 + \sum_{j \in S_i^r} a \langle z_j^{(2)} \rangle. \end{aligned} \quad (\text{A.7b})$$

The time- k auxiliary variable factors for $k \in \{1, 2\}$ are updated according to

$$\begin{aligned} \log q(z_j^{(k)}) &= \langle \log p(z_j^{(k)} \mid \mathbf{r}^{(k)}) \rangle_{q(\mathbf{r}^{(k)})} + \text{const.} \\ &= (-1 + a|S_j^z|) \log z_j^{(k)} \\ &\quad - \sum_{i \in S_j^z} \left\langle \frac{a}{r_i^{(k)}} \right\rangle z_j^{(k)} + \text{const.} \end{aligned} \quad (\text{A.8})$$

This implies that $q(z_j^{(k)}) = \mathcal{G}(z_j^{(k)}; \alpha_{z,j}^{(k)}, \beta_{z,j}^{(k)})$, with

$$\alpha_{z,j}^{(k)} = a|S_j^z| \quad (\text{A.9a})$$

$$\beta_{z,j}^{(k)} = \sum_{i \in S_j^z} \left\langle \frac{a}{r_i^{(k)}} \right\rangle. \quad (\text{A.9b})$$

Appendix B: Proofs for Chapter 5

B.1 Proofs for Section 5.4

This appendix has statements that are relevant to the ML solution method in Section 5.4, in which Condition 7 is assumed to be met.

Proposition 20. *For any $\mathbf{r} \in \mathcal{H}$ and $\boldsymbol{\lambda} \in \mathbb{Z}^n$, $\mathbf{r} + \mathbf{C}\boldsymbol{\lambda} \in \mathcal{H}$.*

Proof. From (5.35), $\omega_j \mathbf{a}_j \mathbf{r} - \langle \omega_j y_j + \frac{1}{2} \rangle_1 \in \mathbb{Z}$, for some $j \in \{1, \dots, m\}$. Then for any $\boldsymbol{\lambda} \in \mathbb{Z}^n$,

$$\omega_j \mathbf{a}_j (\mathbf{r} + \mathbf{C}\boldsymbol{\lambda}) - \langle \omega_j y_j + \frac{1}{2} \rangle_1 \in \mathbb{Z}, \quad (\text{B.1})$$

as $\omega_j \mathbf{a}_j \mathbf{C}\boldsymbol{\lambda} \in \mathbb{Z}$. Therefore, $\mathbf{r} + \mathbf{C}\boldsymbol{\lambda} \in \mathcal{H}$. □

Proposition 21. *If $\mathbf{r} \in \mathcal{H}$, then $\langle \mathbf{r} \rangle_{\mathbf{C}} \in \tilde{\mathcal{H}}$.*

Proof. From (5.7), $\langle \mathbf{r} \rangle_{\mathbf{C}} = \mathbf{r} - \mathbf{C} \lfloor \mathbf{C}^{-1} \mathbf{r} \rfloor$. Thus from Proposition 20, $\langle \mathbf{r} \rangle_{\mathbf{C}} \in \mathcal{H}$, and since $\langle \mathbf{r} \rangle_{\mathbf{C}} \in \mathcal{P}_{\mathbf{C}}$, $\langle \mathbf{r} \rangle_{\mathbf{C}} \in \tilde{\mathcal{H}}$. □

Proposition 22. *If $\tilde{\mathbf{r}} \in \tilde{\Psi}(\boldsymbol{\ell})$, then $\tilde{\mathbf{r}} \in \tilde{\Psi}(\boldsymbol{\ell} + \boldsymbol{\omega} \circ (\mathbf{A}\mathbf{C}\boldsymbol{\lambda}))$, for any $\boldsymbol{\lambda} \in \mathbb{Z}^n$.*

Proof. From (5.37), if $\tilde{\mathbf{r}} \in \tilde{\Psi}(\ell)$, then there is $\mathbf{r} \in \Psi(\ell)$ such that $\langle \mathbf{r} \rangle_C = \tilde{\mathbf{r}}$. From (5.36), $\lceil \boldsymbol{\omega} \circ (\mathbf{y} - \mathbf{A}\mathbf{r}) \rceil = \ell$, and therefore, for any $\boldsymbol{\lambda} \in \mathbb{Z}^n$,

$$\begin{aligned} \lceil \boldsymbol{\omega} \circ (\mathbf{y} - \mathbf{A}(\mathbf{r} - \mathbf{C}\boldsymbol{\lambda} + \mathbf{C}\boldsymbol{\lambda})) \rceil &= \ell \\ \lceil \boldsymbol{\omega} \circ (\mathbf{y} - \mathbf{A}(\mathbf{r} - \mathbf{C}\boldsymbol{\lambda})) \rceil &= \ell + \boldsymbol{\omega} \circ (\mathbf{A}\mathbf{C}\boldsymbol{\lambda}) \end{aligned} \quad (\text{B.2})$$

Thus, $\mathbf{r} - \mathbf{C}\boldsymbol{\lambda} \in \Psi(\ell + \boldsymbol{\omega} \circ (\mathbf{A}\mathbf{C}\boldsymbol{\lambda}))$. Since $\langle \mathbf{r} - \mathbf{C}\boldsymbol{\lambda} \rangle_C = \tilde{\mathbf{r}}$, $\tilde{\mathbf{r}} \in \tilde{\Psi}(\ell + \boldsymbol{\omega} \circ (\mathbf{A}\mathbf{C}\boldsymbol{\lambda}))$, according to (5.37). \square

Proposition 23. *For any $\boldsymbol{\lambda} \in \mathbb{Z}^n$, $\mathcal{L}_r(\mathbf{r}; \mathbf{y}) = \mathcal{L}_r(\mathbf{r} + \mathbf{C}\boldsymbol{\lambda}; \mathbf{y})$*

Proof.

$$\begin{aligned} \mathcal{L}_r(\mathbf{r} + \mathbf{C}\boldsymbol{\lambda}; \mathbf{y}) &= \left\| \boldsymbol{\Sigma}^{-1/2} \mathbf{d}_{\mathbf{1}_m \otimes \boldsymbol{\omega}}(\mathbf{y} - \mathbf{A}(\mathbf{r} + \mathbf{C}\boldsymbol{\lambda})) \right\|^2 \\ &= \left\| \boldsymbol{\Sigma}^{-1/2} (\mathbf{y} - \mathbf{A}\mathbf{r} - \mathbf{A}\mathbf{C}\boldsymbol{\lambda} \right. \\ &\quad \left. - \lceil \boldsymbol{\omega} \circ (\mathbf{y} - \mathbf{A}\mathbf{r} - \mathbf{A}\mathbf{C}\boldsymbol{\lambda}) \rceil \otimes \boldsymbol{\omega}) \right\|^2 \\ &= \mathcal{L}_r(\mathbf{r}; \mathbf{y}) \end{aligned}$$

\square

Lemma 24. *At no element of \mathcal{H} does \mathcal{L}_r have a local minimum.*

Proof. From (5.31),

$$\mathcal{L}_r(\mathbf{r}; \mathbf{y}) = \sum_{i=1}^m \frac{1}{\sigma_{ii}^2} \mathbf{d}_{\omega_i^{-1}}^2(y_i - \mathbf{a}_i \mathbf{r}).$$

For any $\mathbf{r}_0 \in \mathcal{H}$, $\mathbf{d}_{\omega_j^{-1}}(y_j - \mathbf{a}_j \mathbf{r}_0) = -\frac{1}{2}\omega_j^{-1}$, for some $j \in \{1, \dots, m\}$, and therefore,

$$y_j = \mathbf{a}_j \mathbf{r}_0 + \omega_j^{-1} \lceil \omega_j(y_j - \mathbf{a}_j \mathbf{r}_0) \rceil - \frac{1}{2}\omega_j^{-1}. \quad (\text{B.3})$$

Let $\mathbf{u} \in \mathbb{R}^n$ be an arbitrary unit vector. The right directional derivative of \mathcal{L}_r with respect to \mathbf{u} is

$$\begin{aligned}
\partial_{\mathbf{u}^+} \mathcal{L}_r(\mathbf{r}_0; \mathbf{y}) &= \lim_{h \rightarrow 0^+} \frac{\mathcal{L}_r(\mathbf{r}_0 + h\mathbf{u}) - \mathcal{L}_r(\mathbf{r}_0)}{h} \\
&= \lim_{h \rightarrow 0^+} \sum_{i=1}^m \frac{\sigma_{ii}^{-2}}{h} \left(d_{\omega_i^{-1}}^2(y_i - \mathbf{a}_i(\mathbf{r}_0 + h\mathbf{u})) \right. \\
&\quad \left. - d_{\omega_i^{-1}}^2(y_i - \mathbf{a}_i\mathbf{r}_0) \right) \\
&= \lim_{h \rightarrow 0^+} \frac{\sigma_{jj}^{-2}}{h} \left(d_{\omega_j^{-1}}^2(y_j - \mathbf{a}_j(\mathbf{r}_0 + h\mathbf{u})) \right. \\
&\quad \left. - d_{\omega_j^{-1}}^2(y_j - \mathbf{a}_j\mathbf{r}_0) \right) \\
&= \lim_{h \rightarrow 0^+} \frac{\sigma_{jj}^{-2}}{h} \left((\omega_j^{-1} \lceil \omega_j(y_j - \mathbf{a}_j\mathbf{r}_0) \rceil \right. \\
&\quad \left. - \omega_j^{-1} \lceil \omega_j(y_j - \mathbf{a}_j(\mathbf{r}_0 + h\mathbf{u})) \rceil \right. \\
&\quad \left. - \frac{1}{2} \omega_j^{-1} - h\mathbf{a}_j\mathbf{u} \right)^2 - \frac{1}{4} \omega_j^{-2} \Big). \tag{B.4}
\end{aligned}$$

If $\mathbf{a}_j\mathbf{u} > 0$, we have

$$\lceil \omega_j(y_j - \mathbf{a}_j(\mathbf{r}_0 + h\mathbf{u})) \rceil - \lceil \omega_j(y_j - \mathbf{a}_j\mathbf{r}_0) \rceil = -1, \tag{B.5}$$

in which we have used that $h > 0$. Using (B.5) in (B.4), we get that

$$\partial_{\mathbf{u}^+} \mathcal{L}_r(\mathbf{r}_0; \mathbf{y}) = -(\sigma_{jj}^2 \omega_j)^{-1} \mathbf{a}_j\mathbf{u}, \tag{B.6}$$

and thus $\partial_{\mathbf{u}^+} \mathcal{L}_r(\mathbf{r}_0; \mathbf{y}) < 0$, when $\mathbf{a}_j\mathbf{u} > 0$. (Note that $\sigma_{jj}^2 \omega_j > 0$.) Alternatively, if $\mathbf{a}_j\mathbf{u} \leq 0$,

$$\lceil \omega_j(y_j - \mathbf{a}_j(\mathbf{r}_0 + h\mathbf{u})) \rceil - \lceil \omega_j(y_j - \mathbf{a}_j\mathbf{r}_0) \rceil = 0, \tag{B.7}$$

and from (B.4), we have

$$\partial_{\mathbf{u}^+} \mathcal{L}_r(\mathbf{r}_0; \mathbf{y}) = (\sigma_{jj}^2 \omega_j)^{-1} \mathbf{a}_j\mathbf{u}. \tag{B.8}$$

With $\mathbf{a}_j\mathbf{u} < 0$, we also have that $\partial_{\mathbf{u}^+} \mathcal{L}_r(\mathbf{r}_0; \mathbf{y}) < 0$. Thus, $\partial_{\mathbf{u}^+} \mathcal{L}_r(\mathbf{r}_0; \mathbf{y}) < 0$. Following a similar procedure for the left derivative, we find that $\partial_{\mathbf{u}^-} \mathcal{L}_r(\mathbf{r}_0; \mathbf{y}) < 0$. Thus, \mathcal{L}_r

does not have a local minimum at \mathbf{r}_0 , and we conclude that \mathcal{L}_r has a local minimum at no point in \mathcal{H} . \square

Lemma 25. *If $\boldsymbol{\omega}$ has pairwise coprime entries, then Condition 7 is satisfied.*

Proof. From (5.22), $\omega_i = \text{lcm}([\beta_{i,1}, \dots, \beta_{i,n}]^\top)$, for each $i = 1, \dots, m$. So, there is $h_{i,j} \in \mathbb{Z}$, for each $j \in \{1, \dots, n\}$, such that

$$h_{i,j} = \begin{cases} \frac{\omega_i}{\beta_{i,j}}, & \beta_{i,j} \neq 0 \\ 0, & \text{else,} \end{cases}$$

with $\gcd([h_{i,1}, \dots, h_{i,n}]^\top) = 1$. Thus, from (5.21),

$$\gcd([\omega_i \mathbf{a}_i \mathbf{C}, \omega_i]^\top) = \gcd([h_{i,1}, \dots, h_{i,n}, \omega_i]^\top) = 1.$$

According to Theorem 4, there are integer solutions to

$$\boldsymbol{\omega} \circ (\mathbf{A}\mathbf{C})\boldsymbol{\lambda} \equiv \mathbf{q} \pmod{\boldsymbol{\omega}}, \quad (\text{B.9})$$

for any $\mathbf{q} \in \mathbb{Z}^m$. Hence, Condition 7 is satisfied. \square

Lemma 25 shows that $\boldsymbol{\omega}$ having pairwise coprime entries is sufficient to satisfy Condition 7, but it is not necessary. As an example, consider the encoding matrix

$$\mathbf{A} = \begin{bmatrix} 4 & 0 & 1 \\ 2 & 1 & 2 \end{bmatrix}^\top, \quad (\text{B.10})$$

which has $\boldsymbol{\omega} = [1, 2, 4]^\top$. The entries of $\boldsymbol{\omega}$ are not pairwise coprime, yet \mathbf{A} satisfies Condition 7.

Proposition 26. *With $n = 1$, Condition 7 is satisfied iff $\boldsymbol{\omega}$ has pairwise coprime entries.*

Proof. (\Leftarrow): Given by Lemma 25.

(\Rightarrow): Let $\widetilde{\mathbf{W}} = [\boldsymbol{\omega} \circ (\mathbf{A}\mathbf{C}), \text{diag}(\boldsymbol{\omega})]^\top \in \mathbb{Z}^{(m+n) \times m}$, and let $\mathbf{H} = [\widetilde{\mathbf{H}}^\top, \mathbf{0}_{m,n}]^\top$ be the HNF of $\widetilde{\mathbf{W}}$. From Proposition 8, Condition 7 is equivalent to $\widetilde{\mathbf{H}} = \mathbf{I}_m$. As $\widetilde{\mathbf{W}} = \mathbf{U}^{-1}\mathbf{H}$, Condition 7 is equivalent to the possibility of extending $\widetilde{\mathbf{W}}$ to an $(m+n) \times (m+n)$ unimodular matrix. According to [186, Lemma 2], this is possible iff the full-rank minors of $\widetilde{\mathbf{W}}$ are setwise coprime. For $n = 1$, we have that $\omega_j a_j C = \pm 1$, for $j \in \{1, 2, \dots, m\}$. Therefore, each element of the first row of $\widetilde{\mathbf{W}}$ is ± 1 , and the absolute values of the full-rank minors of $\widetilde{\mathbf{W}}$ are given by

$$\mathcal{M} = \left\{ \prod_{i=1}^m \omega_i \right\} \cup \left\{ \prod_{i \neq j} \omega_i \right\}_{j=1}^m. \quad (\text{B.11})$$

For example, if $\mathbf{A} = [1/2, 1/3, 1/5]^\top$, then $C = 1$ and $\boldsymbol{\omega} = [2, 3, 5]^\top$. Thus,

$$\widetilde{\mathbf{W}} = \begin{bmatrix} 1 & 1 & 1 \\ 2 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 5 \end{bmatrix}, \quad (\text{B.12})$$

and $\mathcal{M} = \{30, 15, 10, 6\}$.

Suppose that \mathcal{M} is setwise coprime and that $\boldsymbol{\omega}$ does not have pairwise coprime entries. Then $\omega_k = \alpha \omega_\ell$ for some $k \neq \ell$ and $\alpha \in \mathbb{Z}$. Thus, the factor ω_ℓ is present in each element in \mathcal{M} . This is a contradiction, as we have assumed that \mathcal{M} is setwise coprime. Thus, if the full-rank minors of $\widetilde{\mathbf{W}}$ are setwise coprime, then $\boldsymbol{\omega}$ has pairwise coprime entries. \square

Theorem 27. *For encoding matrix \mathbf{A} satisfying Condition 7, the function $f : \{\langle \mathbb{Z}^n \rangle_{\mathbf{C}^{-1}\mathbf{V}} \times \mathbb{Z}^m\} \mapsto \mathbb{Z}^m$ defined by $f(\boldsymbol{\gamma}, \mathbf{k}) = \boldsymbol{\omega} \circ (\mathbf{A}\mathbf{C}\boldsymbol{\gamma} + \mathbf{k})$ is a bijection.*

Proof. Onto: As \mathbf{A} meets Condition 7, for any $\mathbf{q} \in \mathbb{Z}^m$, there is $\boldsymbol{\lambda} \in \mathbb{Z}^n$ such that

$$\boldsymbol{\omega} \circ (\mathbf{A}\mathbf{C})\boldsymbol{\lambda} = \mathbf{q} \pmod{\boldsymbol{\omega}}. \quad (\text{B.13})$$

Let $\gamma = \langle \lambda \rangle_{C^{-1}\mathbf{V}}$. It follows from (B.13) and (5.7) that

$$\begin{aligned}\omega \circ (\mathbf{AC})(\gamma + C^{-1}\mathbf{V}[V^{-1}\mathbf{C}\lambda]) &\equiv \mathbf{q} \pmod{\omega} \\ \Leftrightarrow \omega \circ (\mathbf{AC}\gamma) &\equiv \mathbf{q} \pmod{\omega},\end{aligned}$$

where we have used that $\mathbf{AV} \in \mathbb{Z}^{m \times n}$. Therefore, $f(\gamma, \mathbf{k}) = \omega \circ (\mathbf{AC}\gamma + \mathbf{k}) = \mathbf{q}$ for some $\mathbf{k} \in \mathbb{Z}^m$.

One-to-one: For any $\gamma_1, \gamma_2 \in \langle \mathbb{Z}^n \rangle_{C^{-1}\mathbf{V}}$ and $\mathbf{k}_1, \mathbf{k}_2 \in \mathbb{Z}^m$, assume that $f(\gamma_1, \mathbf{k}_1) = f(\gamma_2, \mathbf{k}_2)$. Then,

$$\begin{aligned}\omega \circ (\mathbf{AC}\gamma_1 + \mathbf{k}_1) &= \omega \circ (\mathbf{AC}\gamma_2 + \mathbf{k}_2) \\ \Rightarrow \mathbf{AC}(\gamma_2 - \gamma_1) &\in \mathbb{Z}^m \Leftrightarrow \mathbf{C}(\gamma_2 - \gamma_1) \in \Lambda_{\mathbf{V}} \\ \Leftrightarrow \gamma_2 - \gamma_1 &\in C^{-1}\mathbf{V}\mathbb{Z}^n \Leftrightarrow \gamma_1 = \gamma_2 \Leftrightarrow \mathbf{k}_1 = \mathbf{k}_2.\end{aligned}\tag{B.14}$$

□

Proposition 28. *If Condition 7 is satisfied for an encoding matrix \mathbf{A} , then it is satisfied for the encoding matrix $\tilde{\mathbf{A}} = \mathbf{AS}$, where \mathbf{S} is any n -by- n invertible matrix.*

Proof. If \mathbf{V} is the basis for $\{\mathbf{x} \mid \mathbf{Ax} \in \mathbb{Z}^m\}$, then according to Corollary 16, $\mathbf{S}^{-1}\mathbf{V}$ is the basis for $\{\mathbf{x} \mid \tilde{\mathbf{A}}\mathbf{x} \in \mathbb{Z}^m\}$. If \mathbf{C} is a basis for a super-lattice of $\Lambda_{\mathbf{V}}$ (as defined in Section 5.4.1), then $\tilde{\mathbf{C}} = \mathbf{S}^{-1}\mathbf{C}$ is a basis for a super-lattice of $\Lambda_{\tilde{\mathbf{V}}}$. Continuing with this notation, we get that $\tilde{\beta} = \beta$ and that $\tilde{\omega} = \omega$. Thus, for any $\mathbf{b} \in \mathbb{Z}^m$,

$$\tilde{\omega} \circ (\tilde{\mathbf{A}}\tilde{\mathbf{C}})\gamma \equiv \mathbf{b} \pmod{\tilde{\omega}} \Leftrightarrow \omega \circ (\mathbf{AC})\gamma \equiv \mathbf{b} \pmod{\omega}.$$

□

B.2 Correspondence Between Hyperplane Intersections and Chambers

Let \mathcal{U} denote the set of points at which n hyperplanes in \mathcal{H} intersect, and let $\tilde{\mathcal{U}} = \mathcal{P}_C \cap \mathcal{U} = \langle \mathcal{U} \rangle_C$. Recall that \mathcal{H} , defined in (5.35), comprises hyperplanes (subspaces of codimension one) in \mathbb{R}^n defined by the normal vectors $\{\mathbf{a}_j^\top\}_{j \in \{1, \dots, m\}}$, where $\mathbf{a}_j \in \mathbb{Q}^{1 \times n}$ is the j th row of \mathbf{A} . Because \mathbf{A} is assumed to have full column rank, we have that \mathcal{U} is nonempty. Further, we assume that there is no intersection of more than n hyperplanes in \mathcal{H} , which is true almost surely.

Let $\text{cl } \Psi(\ell)$ denote the polytope given by the closure of $\Psi(\ell)$. Define a directed graph G , in which each node is a point in \mathcal{U} . Let any pair of nodes in G be connected by an edge if both points occupy the same edge in $\text{cl } \Psi(\ell)$, for any $\ell \in \Phi$. The direction of the edges will be determined by a linear objective function on the vertices.

Let $G(\Psi(\ell))$ be the sub-graph of G having only the nodes that are vertices of $\text{cl } \Psi(\ell)$. The graph $G(\Psi(\ell))$ is thus a *polytopal digraph*. Given a polytope and a linear objective function on the vertices of the polytope, a polytopal digraph is a directed graph with nodes and edges corresponding to the polytope vertices and extremal edges, with each edge directed toward the vertex with the higher objective function value [252, 253]. The linear objective function we use is $p(\mathbf{v}) = \mathbf{d}^\top \mathbf{v}$, where $\mathbf{d} \in \mathbb{R}^n$ is an arbitrary unit vector such that $\mathbf{d}^\top \mathbf{a}_j \neq 0$, for all $j \in \{1, \dots, m\}$. This ensures that \mathbf{d} is orthogonal to no hyperplane in \mathcal{H} . Thus, the edge between any pair of connected nodes \mathbf{v} and $\boldsymbol{\eta}$ is directed towards \mathbf{v} if $\mathbf{d}^\top \mathbf{v} > \mathbf{d}^\top \boldsymbol{\eta}$, and towards $\boldsymbol{\eta}$ otherwise. Also, let $G(\mathbf{v})$ denote a sub-graph of G containing only the node \mathbf{v} and the nodes connected to \mathbf{v} .

Proposition 29. *Let $\mathbf{v} \in \mathcal{U}$. For any node $\boldsymbol{\eta}$ connected to \mathbf{v} , $\mathbf{d}^\top \mathbf{v} \neq \mathbf{d}^\top \boldsymbol{\eta}$.*

Proof. Both \mathbf{v} and $\boldsymbol{\eta}$ are in one of the hyperplanes that constitute \mathcal{H} . So, $\mathbf{v} - \boldsymbol{\eta}$ is also in that hyperplane. We assume that \mathbf{d} is orthogonal to no hyperplane in \mathcal{H} , and thus, $\mathbf{d}^\top(\mathbf{v} - \boldsymbol{\eta}) \neq 0$. \square

Lemma 30. *For any $\mathbf{v} \in \mathcal{U}$, the remaining nodes in $G(\mathbf{v})$ can be partitioned into pairs such that, within each pair, the edge between \mathbf{v} and one neighboring node directs towards \mathbf{v} , and the edge between \mathbf{v} and the other neighboring node directs away from \mathbf{v} .*

Proof. Any vertex $\mathbf{v} \in \mathcal{U}$ is an intersection of n hyperplanes, and there are $2n$ nodes in G that are connected to \mathbf{v} . The points in \mathcal{U} appear along lines formed by the intersection of $n - 1$ hyperplanes. Hence, for any point connected to \mathbf{v} in G , there will be a corresponding point (also connected to \mathbf{v}) in the opposing direction. Each pair of opposing points is collinear with \mathbf{v} , and \mathbf{v} is located between the pair. Let $\{\boldsymbol{\eta}_1, \boldsymbol{\eta}_2\}$ be any such pair. We have that $\boldsymbol{\eta}_2 = \mathbf{v} - \alpha(\boldsymbol{\eta}_1 - \mathbf{v})$ for some $\alpha > 0$. (This expresses that $\boldsymbol{\eta}_1$ is in the opposite direction from \mathbf{v} as $\boldsymbol{\eta}_2$.) The edge between \mathbf{v} and $\boldsymbol{\eta}_1$ directs to $\boldsymbol{\eta}_1$ iff

$$\begin{aligned} \mathbf{d}^\top(\boldsymbol{\eta}_1 - \mathbf{v}) > 0 &\Leftrightarrow \mathbf{d}^\top(-\alpha(\boldsymbol{\eta}_1 - \mathbf{v})) < 0 \\ &\Leftrightarrow \mathbf{d}^\top(\boldsymbol{\eta}_2 - \mathbf{v}) < 0, \end{aligned}$$

which is true iff the edge between \mathbf{v} and $\boldsymbol{\eta}_2$ directs to \mathbf{v} . Hence, one of $\{\boldsymbol{\eta}_1, \boldsymbol{\eta}_2\}$ has an edge directed towards \mathbf{v} , and one has an edge directed away from \mathbf{v} . \square

Lemma 31. *For any $\mathbf{v} \in \mathcal{U}$, there is only one chamber, denoted $\Psi(\boldsymbol{\ell})$ for some $\boldsymbol{\ell} \in \Phi$, for which \mathbf{v} is a sink (source) in $G(\Psi(\boldsymbol{\ell}))$. (Note that $\mathbf{v} \in \text{cl } \Psi(\boldsymbol{\ell})$.)*

Proof. Any of the 2^n chambers that have \mathbf{v} in its closure can be uniquely identified with a set consisting of one node from each of the n pairs described by Lemma 30.

One of each such pair has an edge that directs towards (away from) \mathbf{v} . Therefore, for only one of the 2^n chambers are all edges between \mathbf{v} and vertices in the chamber's closure directed towards (away from) \mathbf{v} . Thus, for this chamber, denoted $\Psi(\ell)$, \mathbf{v} is a sink (source) in $G(\Psi(\ell))$. \square

Theorem 32 ([253, Theorem 19.4.1]). *A polytopal digraph has one source and one sink.*

Theorem 33. *There is a bijection between \mathcal{U} and $\{\Psi(\ell)\}_{\ell \in \Phi}$.*

Proof. Let $\pi : \mathcal{U} \mapsto \{\Psi(\ell)\}_{\ell \in \Phi}$ be a function that maps $\mathbf{v} \in \mathcal{U}$ to the (unique, from Lemma 31) chamber $\Psi(\ell)$ for which \mathbf{v} is a source in $G(\Psi(\ell))$.

Onto: According to Theorem 32, one vertex of $\text{cl } \Psi(\ell)$ is a source in $G(\Psi(\ell))$. For any $\ell \in \Phi$, the source in $G(\Psi(\ell))$ is given by $\pi(\Psi(\ell))$. Therefore, π is onto.

One-to-one: Assume that $\pi(\mathbf{v}) = \pi(\boldsymbol{\eta}) = \Psi(\ell)$, for some $\ell \in \Phi$. Thus, both \mathbf{v} and $\boldsymbol{\eta}$ are sources in $G(\Psi(\ell))$. Because $G(\Psi(\ell))$ is a polytopal digraph with a unique source, $\mathbf{v} = \boldsymbol{\eta}$. Therefore, π is one-to-one. \square

Lemma 34. *For any $\mathbf{v} \in \mathcal{U}$ and any $\boldsymbol{\lambda} \in \mathbb{Z}^n$, $\pi(\mathbf{v} + \mathbf{C}\boldsymbol{\lambda}) = \pi(\mathbf{v}) + \{\mathbf{C}\boldsymbol{\lambda}\}$, a Minkowski sum.*

Proof. The chamber $\pi(\mathbf{v})$ is formed by hyperplanes in \mathcal{H} . Due to the periodicity in \mathcal{H} , $\pi(\mathbf{v} + \mathbf{C}\boldsymbol{\lambda})$ is a translation of $\pi(\mathbf{v})$ by $\mathbf{C}\boldsymbol{\lambda}$. \square

Theorem 35. *There is a bijection between $\tilde{\mathcal{U}}$ and $\{\tilde{\Psi}(\ell)\}_{\ell \in \Phi}$.*

Proof. Define a mapping $\tilde{\pi} : \tilde{\mathcal{U}} \mapsto \{\tilde{\Psi}(\ell)\}_{\ell \in \Phi}$, which is given by $\tilde{\pi}(\tilde{\mathbf{v}}) = \langle \pi(\tilde{\mathbf{v}}) \rangle_{\mathbf{C}}$.

Onto: According to Theorem 33: for any $\ell \in \Phi$ there is $\mathbf{v} \in \mathcal{U}$ such that $\pi(\mathbf{v}) = \Psi(\ell)$. Then $\pi(\langle \mathbf{v} \rangle_C + \mathbf{C}[\mathbf{C}^{-1}\mathbf{v}]) = \Psi(\ell)$, as $\mathbf{v} = \langle \mathbf{v} \rangle_C + \mathbf{C}[\mathbf{C}^{-1}\mathbf{v}]$. From Lemma 34, we have that $\pi(\langle \mathbf{v} \rangle_C) + \{\mathbf{C}[\mathbf{C}^{-1}\mathbf{v}]\} = \Psi(\ell)$. Thus,

$$\begin{aligned}\langle \pi(\langle \mathbf{v} \rangle_C) \rangle_C &= \langle \Psi(\ell) \rangle_C \\ \tilde{\pi}(\langle \mathbf{v} \rangle_C) &= \tilde{\Psi}(\ell).\end{aligned}\tag{B.15}$$

Therefore, $\tilde{\pi}$ is onto.

One-to-one: Assume, for any $\mathbf{v}, \boldsymbol{\eta} \in \mathcal{U}$ that $\tilde{\pi}(\mathbf{v}) = \tilde{\pi}(\boldsymbol{\eta})$, and therefore that $\langle \pi(\mathbf{v}) \rangle_C = \langle \pi(\boldsymbol{\eta}) \rangle_C$. Then, $\pi(\mathbf{v}) = \pi(\boldsymbol{\eta}) + \mathbf{C}\boldsymbol{\lambda}$, for some $\boldsymbol{\lambda} \in \mathbb{Z}^n$. From Lemma 34, $\pi(\mathbf{v}) = \pi(\boldsymbol{\eta} + \mathbf{C}\boldsymbol{\lambda})$, and from Theorem 33, $\mathbf{v} = \boldsymbol{\eta} + \mathbf{C}\boldsymbol{\lambda}$. Since $\mathbf{v}, \boldsymbol{\eta} \in \mathcal{P}_C$, this gives us $\mathbf{v} = \boldsymbol{\eta}$. Thus, π is one-to-one.

Therefore, $\tilde{\pi}$ is a bijection between $\tilde{\mathcal{U}}$ and $\{\tilde{\Psi}(\ell)\}_{\ell \in \Phi}$.

□

The correspondence between $\tilde{\mathcal{U}}$ and $\{\tilde{\Psi}(\ell)\}_{\ell \in \Phi}$ is shown in Fig. B.1. For each of the three elements in $\tilde{\mathcal{U}}$ (marked with blue circles), the wrapped chamber given by $\tilde{\pi}$ is indicated by an arrow.

B.3 Constructing χ

B.3.1 Constructing χ with $n \geq 1$

Algorithm 8 in Section 5.4 requires a set χ that contains a point in $\tilde{\Psi}(\ell)$ for each $\ell \in \Phi$. According to Theorem 35 in Appendix B.2, we can construct χ by forming a point in $\tilde{\pi}(\mathbf{v})$, for each $\mathbf{v} \in \tilde{\mathcal{U}}$.

Any $\mathbf{v} \in \tilde{\mathcal{U}}$ is the intersection of n hyperplanes defined by $\{\mathbf{a}_{j_i}^\top\}_{i \in \{1, \dots, n\}}$, for some $\mathbf{j} \in \{1, \dots, m\}^n$. Let $\{\boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_n\}$ be the n vertices in $\text{cl } \pi(\mathbf{v})$ that are connected to \mathbf{v}




Figure B.1: \mathcal{P}_C is partitioned into wrapped chambers, labeled 1–3. The points at which two wrapped hyperplanes in $\tilde{\mathcal{H}}$ (shown in red) intersect are marked with blue circles (denoted $\tilde{\mathcal{U}}$). The correspondence of each such point to a wrapped chamber (as described in Appendix B.2) is shown by an arrow.

in G . (Recall that the closure of $\pi(\mathbf{v})$ is a polytope.) Consider an n -simplex defined by the $n + 1$ extremal points $\{\mathbf{v}, \boldsymbol{\eta}_1, \dots, \boldsymbol{\eta}_n\}$. This simplex is a subset of $\text{cl } \pi(\mathbf{v})$.

We find a point in the interior of this simplex as the convex combination

$$\phi(\mathbf{v}) = \alpha_0 \mathbf{v} + \sum_{i=1}^n \alpha_i \boldsymbol{\eta}_i, \quad (\text{B.16})$$

where $\alpha_0 + \sum_{i=1}^n \alpha_i = 1$ and $\alpha_i \in (0, 1)$, for $i \in \{0, \dots, n\}$. First note that

$$\boldsymbol{\eta}_i = \mathbf{v} + (\mathbf{I}_n - \mathbf{P}_i) \mathbf{d} \gamma_i \quad (\text{B.17})$$

for some $\gamma_i > 0$, where \mathbf{P}_i is the orthogonal projection onto $\text{span}(\{\mathbf{a}_{j_q}^\top\}_{q \in \{1, \dots, n\} \setminus i})$.

That is, $(\mathbf{I}_n - \mathbf{P}_i)$ is a projection to the line that contains both \mathbf{v} and $\boldsymbol{\eta}_i$.

From (B.16) and (B.17), we have

$$\phi(\mathbf{v}) = \alpha_0 \mathbf{v} + \sum_{i=1}^n \alpha_i (\mathbf{v} + (\mathbf{I}_n - \mathbf{P}_i) \mathbf{d} \gamma_i) \quad (\text{B.18})$$

$$= \mathbf{v} + \sum_{i=1}^n \alpha_i (\mathbf{I}_n - \mathbf{P}_i) \mathbf{d} \gamma_i. \quad (\text{B.19})$$

We choose the weights such that $\alpha_i = \frac{\epsilon}{\gamma_i}$, for each $i \in \{1, \dots, n\}$, for some $\epsilon > 0$. Letting $\alpha_0 = 1 - \sum_{i=1}^n \frac{\epsilon}{\gamma_i}$, we have that $\sum_{i=0}^n \alpha_i = 1$. For a convex combination, it is necessary that $\sum_{i=1}^n \alpha_i = \sum_{i=1}^n \frac{\epsilon}{\gamma_i} < 1$, and thus, that $\epsilon < (\sum_{i=1}^n \gamma_i^{-1})^{-1}$. Recalling that \mathbf{d} is a unit vector, we have from (B.17) that

$$\begin{aligned} \mathbf{d}^\top (\boldsymbol{\eta}_i - \mathbf{v}) &= \mathbf{d}^\top (\mathbf{I}_n - \mathbf{P}_i) \mathbf{d} \gamma_i \\ \mathbf{d}^\top (\boldsymbol{\eta}_i - \mathbf{v}) &= \|(\mathbf{I}_n - \mathbf{P}_i) \mathbf{d}\|^2 \gamma_i \\ \mathbf{d}^\top (\boldsymbol{\eta}_i - \mathbf{v}) &\leq \gamma_i. \end{aligned}$$

Therefore, it is sufficient that

$$\epsilon < \left(\sum_{i=1}^n (\mathbf{d}^\top (\boldsymbol{\eta}_i - \mathbf{v}))^{-1} \right)^{-1} \quad (\text{B.20})$$

to ensure that $\sum_{i=1}^n \frac{\epsilon}{\gamma_i} < 1$. We have a point,

$$\phi(\mathbf{v}) = \mathbf{v} + \epsilon \sum_{i=1}^n (\mathbf{I}_n - \mathbf{P}_i) \mathbf{d}, \quad (\text{B.21})$$

in the interior of the simplex and, thus, in $\pi(\mathbf{v})$. Then $\langle \phi(\mathbf{v}) \rangle_C \in \tilde{\pi}(\mathbf{v})$. The set χ is formed as $\chi = \{\langle \phi(\mathbf{v}) \rangle_C \mid \mathbf{v} \in \tilde{\mathcal{U}}\}$.

B.3.2 Comment on the $n = 1$ Case

For $n = 1$, the construction of χ is greatly simplified. In this case the unit vector is merely $d = \pm 1$ (we assume $d = 1$, without loss of generality), and $(\mathbf{I}_n - \mathbf{P}_i) = 1$. We also have that $\tilde{\mathcal{U}} = \tilde{\mathcal{H}}$. Thus, each point in χ is given by $\phi(v) = v + \epsilon$, for some

$v \in \tilde{\mathcal{H}}$, where ϵ is chosen such that $\epsilon < |\eta - v|$, for any $v \neq \eta \in \tilde{\mathcal{H}}$. Therefore, $\chi = \{\langle v + \epsilon \rangle_C \mid v \in \tilde{\mathcal{H}}\}$.

Note that, for $n = 1$, we have that $\mathcal{H} = \{a_j^{-1}y_j + C(\lambda + 1/2) \mid \lambda \in \mathbb{Z}\}$, and

$$\tilde{\mathcal{H}} = \{\langle a_j^{-1}y_j + C/2 \rangle_C \mid j \in \{1, \dots, m\}\}. \quad (\text{B.22})$$

Therefore, χ has a cardinality of m .

B.4 Bounds on Problem Size for Restricted Encoding Matrices

This appendix provides bounds on quantities related to the size of the problem considered in Section 5.4, in which we assume that Condition 7 is met by \mathbf{A} .

B.4.1 Upper Bound on N

From (5.6), for any $\mathbf{r} \in \mathcal{P}_C$ and some $\boldsymbol{\lambda} \in [0, 1]^n$,

$$\begin{aligned} \omega_j \mathbf{a}_j \mathbf{r} &= \omega_j \mathbf{a}_j \mathbf{C} \boldsymbol{\lambda} \\ &\in \omega_j \mathbf{a}_j \mathbf{C} [0, 1]^n \\ &= [\min(\mathbf{0}_n^\top, \omega_j \mathbf{a}_j \mathbf{C}) \mathbf{1}_n, \max(\mathbf{0}_n^\top, \omega_j \mathbf{a}_j \mathbf{C}) \mathbf{1}_n) \\ &= [t_j^{\min}, t_j^{\max}), \end{aligned} \quad (\text{B.23})$$

with $t_j^{\min}, t_j^{\max} \in \mathbb{Z}$, for each $j \in \{1, \dots, m\}$ and \min and \max denoting element-wise minimum and maximum operations, respectively.

From (5.35) and (B.23), any $\mathbf{r} \in \tilde{\mathcal{H}}$ satisfies the following linear equality:

$$\omega_j \mathbf{a}_j \mathbf{r} = \langle \omega_j y_j + \tfrac{1}{2} \rangle_1 + \ell, \quad (\text{B.24})$$

for some $j \in \{1, \dots, m\}$ and $\ell \in \{t_j^{\min}, \dots, t_j^{\max} - 1\}$. Note that $\{t_j^{\min}, \dots, t_j^{\max} - 1\}$ is a set of integers with cardinality $N^j = \sum_{k=1}^n |\omega_j (\mathbf{a}_j \mathbf{C})_k|$. Thus, (B.24) describes a

total of $N = \sum_{j=1}^m N^j$ hyperplanes, and $\tilde{\mathcal{H}} \subset \mathcal{H}_{\text{sub}}$, where

$$\mathcal{H}^{\text{sub}} = \{\mathbf{r} \in \mathbb{R}^n \mid \exists j \in \{1, \dots, m\} \text{ s.t.} \\ \omega_j \mathbf{a}_j \mathbf{r} - \langle \omega_j \mathbf{y}_j + \frac{1}{2} \rangle_1 \in \{t_j^{\min}, \dots, t_j^{\max} - 1\}\}, \quad (\text{B.25})$$

i.e., \mathcal{H}^{sub} is a union of the N hyperplanes in \mathcal{H} that intersect with \mathcal{P}_C . Thus, $\tilde{\mathcal{H}} = \mathcal{P}_C \cap \mathcal{H} = \mathcal{P}_C \cap \mathcal{H}^{\text{sub}}$.

Recall that $\boldsymbol{\beta} \in \mathbb{Z}^{m \times n}$, with

$$\beta_{j,k} = \begin{cases} 1/(\mathbf{a}_j \mathbf{C})_k, & (\mathbf{a}_j \mathbf{C})_k \neq 0 \\ 0, & \text{otherwise} \end{cases},$$

and that $\boldsymbol{\beta}_j \in \mathbb{Z}^{1 \times n}$ is the j th row of $\boldsymbol{\beta}$. Thus, if $\beta_{j,i} \neq 0$, $\beta_{j,i} = \frac{z_i}{(\mathbf{a}_j \mathbf{V})_i}$. For each $j \in \{1, \dots, m\}$,

$$\begin{aligned} N^j &= \sum_{k=1}^n |\text{lcm}(\boldsymbol{\beta}_j^\top)(\mathbf{a}_j \mathbf{C})_k| \\ &\leq \sum_{k=1}^n \left| \left(\prod_{\{i \mid \beta_{j,i} \neq 0\}} \beta_{j,i} \right) (\mathbf{a}_j \mathbf{C})_k \right| \\ &= \sum_{k=1}^n \prod_{\{i \neq k \mid \beta_{j,i} \neq 0\}} |\beta_{j,i}| \leq \prod_{\{i \mid \beta_{j,i} \neq 0\}} n |\beta_{j,i}| \\ &\leq \prod_{\{i \mid \beta_{j,i} \neq 0\}} n \left| \frac{z_i}{(\mathbf{a}_j \mathbf{V})_i} \right| \leq n \prod_{i=1}^n z_i, \end{aligned} \quad (\text{B.26})$$

in which we have used that $\mathbf{a}_j \mathbf{V} \in \mathbb{Z}^{1 \times n}$. From (B.26), an upper bound on N is given by

$$N \leq mn \left(\prod_{i=1}^n z_i \right) = mn \det(\mathbf{C}^{-1} \mathbf{V}). \quad (\text{B.27})$$

Thus, for fixed \mathbf{C} , \mathbf{V} , and n , N is $O(m)$.

B.4.2 Bound on Number of Wrapped Chambers

Let $\mathcal{U} \subset \mathcal{H}$ denote the set of points at which n hyperplanes in \mathcal{H} intersect, and let $\mathcal{U}^{\text{sub}} \subset \mathcal{U}$ denote the set of points at which n hyperplanes in \mathcal{H}^{sub} intersect. $|\mathcal{U}^{\text{sub}}|$

is maximized when none of the N hyperplanes that it comprises are parallel to one another, in which case, $|\mathcal{U}^{\text{sub}}| = \binom{N}{n}$. From (B.27), an upper bound on $|\mathcal{U}^{\text{sub}}|$ is given by

$$|\mathcal{U}^{\text{sub}}| \leq \binom{N}{n} \leq (m \det(\mathbf{C}^{-1} \mathbf{V}))^n n^{n-1}. \quad (\text{B.28})$$

Thus, for fixed \mathbf{C} , \mathbf{V} , and n , $|\mathcal{U}^{\text{sub}}|$ is $O(m^n)$.

In general, some hyperplanes in \mathcal{H}^{sub} will be parallel to one another, and this admits a refinement of the upper bound on $|\mathcal{U}^{\text{sub}}|$. Letting τ denote the set of n -combinations of $\{1, \dots, m\}$, an upper bound on $|\mathcal{U}^{\text{sub}}|$ is given by

$$|\mathcal{U}^{\text{sub}}| \leq \sum_{\mathcal{I} \in \tau} \prod_{j \in \mathcal{I}} N^j. \quad (\text{B.29})$$

We have that $\tilde{\mathcal{U}} = \mathcal{P}_{\mathbf{C}} \cap \mathcal{U} = \mathcal{P}_{\mathbf{C}} \cap \mathcal{U}^{\text{sub}}$ and, therefore, that $|\tilde{\mathcal{U}}| \leq |\mathcal{U}^{\text{sub}}|$.

Let $K = |\{\tilde{\Psi}(\ell) \mid \ell \in \Phi\}|$ denote the number of wrapped chambers. We have from Theorem 35 that there is a bijection between the set of wrapped chambers and $\tilde{\mathcal{U}}$. Thus,

$$K = |\tilde{\mathcal{U}}| \leq |\mathcal{U}^{\text{sub}}|, \quad (\text{B.30})$$

and the upper bounds on $|\mathcal{U}^{\text{sub}}|$ in (B.28) and (B.29) apply to the number of wrapped chambers as well. Thus, K is $O(m^n)$.

B.5 Bounds on Problem Size for General Encoding Matrices

This appendix provides bounds on quantities related to the size of the problem considered in Section 5.5, in which the restriction on \mathbf{A} used in Section 5.4 is lifted.

B.5.1 Bound on N_{NC}

For $\mathbf{x} \in \mathcal{P}_{\mathbf{V}}$, we have, for some $\boldsymbol{\lambda} \in [0, 1)^n$, that

$$\begin{aligned} \mathbf{a}_j \mathbf{x} &= \mathbf{a}_j \mathbf{V} \boldsymbol{\lambda} \\ &\in \mathbf{a}_j \mathbf{V} [0, 1)^n \\ &= [\min(\mathbf{0}_n^\top, \mathbf{a}_j \mathbf{V}) \mathbf{1}_n, \max(\mathbf{0}_n^\top, \mathbf{a}_j \mathbf{V}) \mathbf{1}_n) \\ &= [t_j^{\min}, t_j^{\max}) \end{aligned} \tag{B.31}$$

with $t_j^{\min}, t_j^{\max} \in \mathbb{Z}$, for each $j \in \{1, \dots, m\}$.

From (5.48) and (B.31), any $\mathbf{x} \in \tilde{\mathcal{H}}_{\text{NC}}$ satisfies

$$\mathbf{a}_j \mathbf{x} = \langle y_j + \tfrac{1}{2} \rangle_1 + \ell, \tag{B.32}$$

for some $j \in \{1, \dots, m\}$ and $\ell \in \{t_j^{\min}, \dots, t_j^{\max} - 1\}$. Note that $\{t_j^{\min}, \dots, t_j^{\max} - 1\}$ is a set of integers with cardinality $N_{\text{NC}}^j = \sum_{k=1}^n |(\mathbf{a}_j \mathbf{V})_k|$. Thus (B.32) describes a total of $N_{\text{NC}} = \sum_{j=1}^m N_{\text{NC}}^j$ hyperplanes. We therefore have that $\tilde{\mathcal{H}}_{\text{NC}} \subset \mathcal{H}_{\text{NC}}^{\text{sub}}$, where

$$\begin{aligned} \mathcal{H}_{\text{NC}}^{\text{sub}} &= \{ \mathbf{x} \in \mathbb{R}^n \mid \exists j \in \{1, \dots, m\} \text{ s.t.} \\ &\quad \mathbf{a}_j \mathbf{x} - \langle y_j + \tfrac{1}{2} \rangle_1 \in \{t_j^{\min}, \dots, t_j^{\max} - 1\} \}, \end{aligned} \tag{B.33}$$

i.e., $\mathcal{H}_{\text{NC}}^{\text{sub}}$ is a union of the N_{NC} hyperplanes in \mathcal{H}_{NC} that intersect with $\mathcal{P}_{\mathbf{V}}$. Therefore,

$$\tilde{\mathcal{H}}_{\text{NC}} = \mathcal{P}_{\mathbf{V}} \cap \mathcal{H}_{\text{NC}} = \mathcal{P}_{\mathbf{V}} \cap \mathcal{H}_{\text{NC}}^{\text{sub}}.$$

Recalling that $\mathbf{z} = (\text{lcm}(\mathbf{A}\mathbf{V}))^\top$, we have an upper bound on N_{NC}^j given by $N_{\text{NC}}^j \leq \sum_{k=1}^n z_k$. Therefore,

$$N_{\text{NC}} \leq m \mathbf{z}^\top \mathbf{1}_n = m \text{tr}(\mathbf{C}^{-1} \mathbf{V}). \tag{B.34}$$

B.5.2 Comparison Between N and N_{NC}

Like N_{NC} , N in Appendix B.4 determines the complexity of an ML solution method. Both quantities give the number of hyperplanes that intersect a parallelo-
tope. The method in Section 5.4, to which N is relevant, has a restrictive condition
on the encoding matrix \mathbf{A} . It may be surprising that, even when Condition 7 is met,
it is possible that $N > N_{\text{NC}}$. This is the case, for example, when

$$\mathbf{A} = \begin{bmatrix} 3 & 0 & 1 & 3 \\ 4 & 4 & 1 & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix}^{\text{T}}, \quad (\text{B.35})$$

with $\mathbf{V} = \mathbf{I}_3$ and $\mathbf{C} = \text{diag}([\frac{1}{3}, \frac{1}{4}, 1]^{\text{T}})$. In this case, $N = 26$ and $N_{\text{NC}} = 20$. Note
that for $n = 1$, $N = m$, and thus, $N \leq N_{\text{NC}}$.

B.5.3 Bound on the Number of Wrapped Chambers

Following Appendix B.4, upper bounds on $|\mathcal{U}^{\text{sub}}|$ are given by

$$|\mathcal{U}^{\text{sub}}| \leq \binom{N_{\text{NC}}}{n} \leq (m \text{tr}(\mathbf{C}^{-1}\mathbf{V}))^n n^{-1}. \quad (\text{B.36})$$

$$|\mathcal{U}^{\text{sub}}| \leq \sum_{\mathcal{I} \in \tau} \prod_{j \in \mathcal{I}} N_{\text{NC}}^j. \quad (\text{B.37})$$

We have that $\tilde{\mathcal{U}} = \mathcal{P}_{\mathbf{V}} \cap \mathcal{U} = \mathcal{P}_{\mathbf{V}} \cap \mathcal{U}^{\text{sub}}$ and, therefore, that $|\tilde{\mathcal{U}}| \leq |\mathcal{U}^{\text{sub}}|$. Let
 $K_{\text{NC}} = |\{\tilde{\Psi}_{\text{NC}}(\ell) \mid \ell \in \Phi\}|$ denote the number of wrapped chambers. By a statement
analogous to Theorem 35 in Appendix B.2, it can be shown that there is a bijection
between the set of wrapped chambers and $\tilde{\mathcal{U}}$. Thus,

$$K_{\text{NC}} = |\tilde{\mathcal{U}}| \leq |\mathcal{U}^{\text{sub}}|, \quad (\text{B.38})$$

and the upper bounds on $|\mathcal{U}^{\text{sub}}|$ in (B.36) and (B.37) apply to the number of wrapped
chambers as well. Therefore, K is of order $O(m^n)$ (for fixed n , \mathbf{C} , and \mathbf{V}).




Figure B.2: The number of wrapped chambers for Algorithm 8 is plotted in red against m ; the red dash-dotted line is the upper bound on K from (B.28). The number of wrapped chambers for Algorithm (9) is in blue; the blue dash-dotted line is the upper bound on K_{NC} from (B.36).

B.5.4 Comparison Between K and K_{NC}

Here we compare the number of wrapped chambers for the algorithms in Section 5.4 (denoted K) and Section 5.5 (denoted K_{NC}). The number of wrapped chambers for the two methods are shown against m in Fig. B.2. The encoding matrices are such that, as m increases, \mathbf{V} and \mathbf{C} are constant. The encoding matrices are made to satisfy Condition 7 from Section 5.4. In Fig. B.2, K and K_{NC} are plotted against $m = 5, \dots, 20$. The upper bound on K from (B.28) is shown as well as the upper bound on K_{NC} from (B.36). The plot shows the computational benefit that Algorithm 8 may have over Algorithm 9 when Condition 7 is met. For example, with $m = 20$, $K_{\text{NC}} > 8K$.

B.6 Proofs for Section 5.6

Proposition 36. *For any $\mathbf{t} \in \mathbb{R}^m$ and $\mathbf{B} \in \mathbb{R}^{m \times n}$ having orthogonal (though, not necessarily orthonormal) columns,*

$$\widehat{\mathbf{k}} = \arg \min_{\mathbf{k} \in \mathbb{Z}^n} \|\mathbf{t} - \mathbf{B}\mathbf{k}\| = \lceil \mathbf{B}^\dagger \mathbf{t} \rceil. \quad (\text{B.39})$$

That is, the Babai rounding method [197] solves a closest lattice point problem when the lattice basis has orthogonal columns.

Proof.

$$\begin{aligned} \|\mathbf{t} - \mathbf{B}\mathbf{k}\|^2 &= \|(\mathbf{B}\mathbf{B}^\dagger + \mathbf{P}_B^\perp)\mathbf{t} - \mathbf{B}\mathbf{k}\|^2 \\ &= \|\mathbf{P}_B^\perp \mathbf{t}\|^2 + \|\mathbf{B}\mathbf{B}^\dagger \mathbf{t}\|_2^2 + \|\mathbf{B}\mathbf{k}\|^2 - 2\mathbf{t}^\top \mathbf{B}\mathbf{k} \\ &= \|\mathbf{P}_B^\perp \mathbf{t}\|^2 + \sum_{i=1}^n \|\mathbf{b}_i\|^2 ((\mathbf{B}^\dagger \mathbf{t})_i - k_i)^2, \end{aligned}$$

where $\mathbf{B}\mathbf{B}^\dagger$ is the orthogonal projection onto the range of \mathbf{B} , \mathbf{P}_B^\perp is the orthogonal projection onto the kernel of \mathbf{B} , and \mathbf{b}_i is the i th column of \mathbf{B} . Since the objective is separable in \mathbf{k} , it is clear that $\widehat{k}_i = \lceil (\mathbf{B}^\dagger \mathbf{t})_i \rceil$, and $\widehat{\mathbf{k}} = \lceil \mathbf{B}^\dagger \mathbf{t} \rceil$. \square

Appendix C: Phase Noise Model

In [31], the noise in the phase entries of a sample covariance matrix from multi-snapshot data is modeled by jointly Gaussian random variables. In this appendix, we rederive this approximation and express the resulting approximate covariance matrix. Following the signal model in Section 5.6.7 and Section 6.2, consider an array of L isotropic sensors that is impinged upon by a narrowband signal $s(t) \in \mathbb{C}$ from a single emitter located in the far field of the array. The demodulated array output vector is given by

$$\mathbf{b}(t) = \mathbf{a}(\mathbf{x})s(t) + \boldsymbol{\varepsilon}(t), \quad t = 1, \dots, T, \quad (\text{C.1})$$

where $\mathbf{a}(\mathbf{x})$ is the array steering vector for the parameter $\mathbf{x} \in \mathbb{R}^n$, and $\boldsymbol{\varepsilon}(t)$ is zero-mean circular Gaussian noise that is uncorrelated across snapshots indexed by t . The array steering vector is given by

$$\mathbf{a}(\mathbf{x}) = [e^{i\pi \mathbf{p}_1^T \mathbf{x}} \quad e^{i\pi \mathbf{p}_2^T \mathbf{x}} \quad \dots \quad e^{i\pi \mathbf{p}_L^T \mathbf{x}}]^T, \quad (\text{C.2})$$

where $i = \sqrt{-1}$. The position vectors $\mathbf{p}_j \in \mathbb{R}^n$, for $j \in \{1, 2, \dots, L\}$, give the sensor positions in units of one-half wavelength.

The sample correlation matrix of \mathbf{b} is $\hat{\mathbf{R}} = \frac{1}{T} \sum_{t=1}^T \mathbf{b}(t)\mathbf{b}^H(t)$. Let \mathbf{v} denote the vector of $\binom{L}{2}$ lower-triangular entries of $\hat{\mathbf{R}}$ in column-major order. That is,

$$\mathbf{v} = [\hat{r}_{2,1}, \hat{r}_{3,1}, \dots, \hat{r}_{L,L-1}]^T \quad (\text{C.3})$$

where $\widehat{r}_{i,j}$ is the (i,j) th entry of $\widehat{\mathbf{R}}$.

Under the conditional (deterministic) signal model [222], the distribution of $\widehat{r}_{i,j}$ depends on the fixed, unknown signal $\{s(t)\}$. Because we do not observe this signal directly, we involve the noisy array outputs in the following approximation:

$$|s(t)|^2 \approx \frac{1}{L} \mathbf{b}^H(t) \mathbf{b}(t). \quad (\text{C.4})$$

Using (C.4), we have

$$\mathbb{E} [\widehat{r}_{i,j}] \approx \widehat{\rho} a_i(\mathbf{x}) \bar{a}_j(\mathbf{x}), \quad i \neq j \quad (\text{C.5})$$

$$\text{var} (\widehat{r}_{i,j}) \approx T^{-1} (2\widehat{\rho}\sigma^2 + \sigma^4), \quad i \neq j \quad (\text{C.6})$$

$$\text{cov} (\widehat{r}_{i,j}, \widehat{r}_{k,j}) \approx T^{-1} a_i(\mathbf{x}) \bar{a}_k(\mathbf{x}) \widehat{\rho} \sigma^2, \quad i \neq j, k \neq j, i \neq k \quad (\text{C.7})$$

$$\text{cov} (\widehat{r}_{i,j}, \widehat{r}_{i,k}) \approx T^{-1} \bar{a}_j(\mathbf{x}) a_k(\mathbf{x}) \widehat{\rho} \sigma^2, \quad i \neq j, i \neq k, j \neq k \quad (\text{C.8})$$

$$\text{cov} (\widehat{r}_{i,j}, \widehat{r}_{k,\ell}) = 0, \quad i \neq j, k \neq \ell, i \neq k, j \neq \ell, \quad (\text{C.9})$$

where $-$ denotes complex conjugation, and

$$\widehat{\rho} \triangleq \frac{1}{LT} \sum_{t=1}^T \mathbf{b}^H(t) \mathbf{b}(t). \quad (\text{C.10})$$

In the above, we have used the conditional signal model for $\{s(t)\}$. If the signal follows the unconditional model, then $\widehat{\rho} = T^{-1} \sum_{t=1}^T \mathbb{E} [|s(t)|^2]$ may be used, rather than estimating the power of $\{s(t)\}$ from the data.

The complex logarithm of $z \in \mathbb{C} \setminus \{0\}$ is given by $\text{Log } z \triangleq \log |z| + i\angle z$. Let

$$\begin{aligned} \mathbf{f}: \{\mathbb{C} \setminus 0\}^m &\rightarrow [-\tfrac{1}{2}, \tfrac{1}{2})^m \\ \mathbf{z} &\mapsto \frac{i}{4\pi} (\text{Log } \bar{\mathbf{z}} - \text{Log } \mathbf{z}). \end{aligned} \quad (\text{C.11})$$

Thus, $\mathbf{f}(\mathbf{z}) = (2\pi)^{-1} \angle \mathbf{z} \in [-\frac{1}{2}, \frac{1}{2})^m$ for any $\mathbf{z} \in \{\mathbb{C} \setminus 0\}^m$.

Now, let $\mathbf{y} = \mathbf{f}(\mathbf{v}) \in [-\frac{1}{2}, \frac{1}{2})^{(\frac{L}{2})}$. We use nonlinear error propagation to approximate the covariance matrix of \mathbf{y} , i.e., we find the covariance of a first-order Taylor

approximation of \mathbf{y} . The linear approximation of \mathbf{y} around its expected value is

$$\mathbf{y} \approx \tilde{\mathbf{y}} = \mathbf{f}(\mathbb{E}[\mathbf{v}]) + \mathbf{J}_{\mathbf{f}}(\mathbb{E}[\mathbf{v}])(\mathbf{v} - \mathbb{E}[\mathbf{v}]) + \mathbf{J}_{\mathbf{f}}^c(\mathbb{E}[\mathbf{v}])(\bar{\mathbf{v}} - \mathbb{E}[\bar{\mathbf{v}}]) \quad (\text{C.12})$$

$$= \mathbf{f}(\mathbb{E}[\mathbf{v}]) + 2 \operatorname{Re} \{ \mathbf{J}_{\mathbf{f}}(\mathbb{E}[\mathbf{v}])(\mathbf{v} - \mathbb{E}[\mathbf{v}]) \}, \quad (\text{C.13})$$

in which $\mathbf{J}_{\mathbf{f}}(\mathbf{z}) \in \mathbb{R}^{m \times m}$ is the Jacobian of \mathbf{f} evaluated at \mathbf{z} and $\mathbf{J}_{\mathbf{f}}^c(\mathbf{z}) \in \mathbb{R}^{m \times m}$ is the conjugate Jacobian of \mathbf{f} [254]. That is, $(\mathbf{J}_{\mathbf{f}}(\mathbf{z}))_{j,k} = \frac{\partial f_j(\mathbf{z})}{\partial v_k}$ and $(\mathbf{J}_{\mathbf{f}}^c(\mathbf{z}))_{j,k} = \frac{\partial f_j(\mathbf{z})}{\partial \bar{v}_k}$, for $j, k \in \{1, \dots, m\}$. Note that $\mathbf{J}_{\mathbf{f}}(\mathbf{z}) = \overline{\mathbf{J}_{\mathbf{f}}^c(\mathbf{z})}$, as \mathbf{f} is real-valued [254]. The Jacobian of \mathbf{f} evaluated at $\mathbf{z} \in \{\mathbb{C} \setminus 0\}^m$ is

$$\mathbf{J}_{\mathbf{f}}(\mathbf{z}) = \operatorname{diag} \left(-\frac{i}{4\pi} [z_1^{-1}, z_2^{-1}, \dots, z_m^{-1}]^T \right). \quad (\text{C.14})$$

The covariance of $\tilde{\mathbf{y}}$ is

$$\operatorname{cov}(\tilde{\mathbf{y}}) = \operatorname{cov}(2 \operatorname{Re} \{ \mathbf{J}_{\mathbf{f}}(\mathbb{E}[\mathbf{v}])(\mathbf{v} - \mathbb{E}[\mathbf{v}]) \}) \quad (\text{C.15})$$

$$= 2 \mathbf{J}_{\mathbf{f}}(\mathbb{E}[\mathbf{v}]) \operatorname{cov}(\mathbf{v})(\mathbf{J}_{\mathbf{f}}(\mathbb{E}[\mathbf{v}]))^H. \quad (\text{C.16})$$

Thus,

$$\begin{aligned} \operatorname{var}(\tilde{y}_{i,j}) &= (2\pi)^{-2} \frac{\operatorname{var}(\hat{r}_{i,j})}{2|\mathbb{E}[\hat{r}_{i,j}]|^2} = (2\pi)^{-2} \frac{1 + 2\frac{\hat{\rho}}{\sigma^2}}{2T\left(\frac{\hat{\rho}}{\sigma^2}\right)^2}, \quad i > j \\ \operatorname{cov}(\tilde{y}_{i,j}, \tilde{y}_{k,j}) &= (2\pi)^{-2} \frac{\operatorname{cov}(\hat{r}_{i,j}, \hat{r}_{k,j})}{2\mathbb{E}[\hat{r}_{i,j}]\mathbb{E}[\hat{r}_{k,j}]} = (2\pi)^{-2} \frac{1}{2T\frac{\hat{\rho}}{\sigma^2}}, \quad i > j, k > j, i \neq k \\ \operatorname{cov}(\tilde{y}_{i,j}, \tilde{y}_{i,k}) &= (2\pi)^{-2} \frac{1}{2T\frac{\hat{\rho}}{\sigma^2}}, \quad i > j, i > k, j \neq k \\ \operatorname{cov}(\tilde{y}_{i,j}, \tilde{y}_{j,k}) &= -\operatorname{cov}(\tilde{y}_{i,j}, \tilde{y}_{k,j}) = -(2\pi)^{-2} \frac{1}{2T\frac{\hat{\rho}}{\sigma^2}}, \quad i > j > k \\ \operatorname{cov}(\tilde{y}_{i,j}, \tilde{y}_{k,i}) &= -(2\pi)^{-2} \frac{1}{2T\frac{\hat{\rho}}{\sigma^2}}, \quad k > i > j \\ \operatorname{cov}(\tilde{y}_{i,j}, \tilde{y}_{k,\ell}) &= 0, \quad i > j, k > \ell, \{i, j\} \cap \{k, \ell\} = \emptyset. \end{aligned} \quad (\text{C.17})$$

Equivalently, for any $i > j$, $k > \ell$, we have that

$$\text{cov}(\tilde{y}_{i,j}, \tilde{y}_{k,\ell}) = \begin{cases} (2\pi)^{-2} \frac{1+2\frac{\hat{\rho}}{\sigma^2}}{2T(\frac{\hat{\rho}}{\sigma^2})^2}, & i = k, j = \ell \\ (2\pi)^{-2} \frac{1}{2T\frac{\hat{\rho}}{\sigma^2}}, & i \neq k, j = \ell, \\ (2\pi)^{-2} \frac{1}{2T\frac{\hat{\rho}}{\sigma^2}}, & i = k, j \neq \ell, \\ -(2\pi)^{-2} \frac{1}{2T\frac{\hat{\rho}}{\sigma^2}}, & j = k \\ -(2\pi)^{-2} \frac{1}{2T\frac{\hat{\rho}}{\sigma^2}}, & i = \ell \\ 0, & \{i, j\} \cap \{k, \ell\} = \emptyset \end{cases} \quad (\text{C.18})$$

Hence, the covariance matrix approximation can be expressed as

$$\text{cov}(\tilde{\mathbf{y}}) = (2\pi)^{-2} \frac{1}{2T\frac{\hat{\rho}}{\sigma^2}} \left(\mathbf{\Phi} \mathbf{\Phi}^\top + \left(\frac{\hat{\rho}}{\sigma^2} \right)^{-1} \mathbf{I} \right), \quad (\text{C.19})$$

where $\mathbf{\Phi} \in \{-1, 0, 1\}^{\binom{L}{2} \times L}$ is a matrix with rank $L-1$ that corresponds to the ordering of elements in \mathbf{v} . When $L = 4$, for example, $\mathbf{\Phi}$ is given by

$$\mathbf{\Phi} = \begin{bmatrix} -1 & 1 & 0 & 0 \\ -1 & 0 & 1 & 0 \\ -1 & 0 & 0 & 1 \\ 0 & -1 & 1 & 0 \\ 0 & -1 & 0 & 1 \\ 0 & 0 & -1 & 1 \end{bmatrix}. \quad (\text{C.20})$$

Bibliography

- [1] Google, “Wright-Patterson Air Force Base Area B,” 2021, Accessed Sept. 16, 2021. [Online]. Available: <https://www.google.com/maps>
- [2] S. M. Scarborough, L. Gorham, M. J. Minardi, U. K. Majumder, M. G. Judge, L. Moore, L. Novak, S. Jaroszewski, L. Spoldi, and A. Pieramico, “A challenge problem for SAR change detection and data compression,” in *Proc. SPIE*, vol. 7699, 2010, pp. 287–291.
- [3] Directorate Space, Security and Migration, European Commission Joint Research Centre, “Copernicus Emergency Management Service, EMSR515 Melini: Grading product, version 1, release 1,” Jul. 2021, Accessed May 12, 2022. [Online]. Available: <https://emergency.copernicus.eu>
- [4] Google Earth version 9.163.0.0, “Troodos Mountains, Cyprus,” May 2020, Accessed June 2, 2022. [Online]. Available: <https://earth.google.com>
- [5] Copernicus, “Sentinel-1 data, processed by ESA,” Jul. 2021, Retrieved from ASF DAAC May 12, 2022. [Online]. Available: <https://asf.alaska.edu/>
- [6] NASA, “Dataset: UAVSAR,” Retrieved from ASF DAAC August 18, 2020, 2020. [Online]. Available: <https://asf.alaska.edu/>
- [7] D. Tucker, S. Zhao, and L. C. Potter, “Maximum likelihood estimation in mixed integer linear models,” *IEEE Signal Process. Lett.*, vol. 30, pp. 1557–1561, 2023.
- [8] J. Jones, A. R. Webster, and W. K. Hocking, “An improved interferometer design for use with meteor radars,” *Radio Sci.*, vol. 33, no. 1, pp. 55–65, 1998.
- [9] N. J. Pelc, M. A. Bernstein, A. Shimakawa, and G. H. Glover, “Encoding strategies for three-direction phase-contrast MR imaging of flow,” *J. Magn. Reson. Imaging*, vol. 1, no. 4, pp. 405–413, 1991.
- [10] L. M. Novak, S. D. Halversen, G. Owirka, and M. Hiett, “Effects of polarization and resolution on SAR ATR,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 33, no. 1, pp. 102–116, Jan. 1997.

- [11] E. Ertin and L. C. Potter, "Polarimetric classification of scattering centers using M-ary Bayesian decision rules," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 36, no. 3, pp. 738–749, Jul. 2000.
- [12] C. López-Martínez and X. Fabregas, "Polarimetric SAR speckle noise model," *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 10, pp. 2232–2242, Oct. 2003.
- [13] J. W. Goodman, "Some fundamental properties of speckle," *J. Opt. Soc. Am.*, vol. 66, no. 11, pp. 1145–1150, Nov. 1976.
- [14] J.-S. Lee, K. W. Hoppel, S. A. Mango, and A. R. Miller, "Intensity and phase statistics of multilook polarimetric and interferometric SAR imagery," *IEEE Trans. Geosci. Remote Sens.*, vol. 32, no. 5, pp. 1017–1028, 1994.
- [15] R. G. Gallagher, *Circularly-Symmetric Gaussian Random Vectors*. Unpublished, Jan. 2008. [Online]. Available: <https://pdfs.semanticscholar.org/815e/637fafa233b44e2067f6ce345734012bd844.pdf>
- [16] X. Ma, P. Wu, Y. Wu, and H. Shen, "A review on recent developments in fully polarimetric SAR image despeckling," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 11, no. 3, pp. 743–758, Mar. 2018.
- [17] S. Foucher and C. López-Martínez, "Analysis, evaluation, and comparison of polarimetric SAR speckle filtering techniques," *IEEE Trans. Image Process.*, vol. 23, no. 4, pp. 1751–1764, Apr. 2014.
- [18] P. Teunissen, "The least-squares ambiguity decorrelation adjustment: A method for fast GPS integer ambiguity estimation," *J. Geod.*, vol. 70, pp. 65–82, 11 1995.
- [19] A. Hassibi and S. Boyd, "Integer parameter estimation in linear models with applications to GPS," *IEEE Trans. Signal Process.*, vol. 46, no. 11, pp. 2938–2952, 1998.
- [20] W. Li, X. Wang, X. Wang, and B. Moran, "Distance estimation using wrapped phase measurements in noise," *IEEE Trans. Signal Process.*, vol. 61, no. 7, pp. 1676–1688, 2013.
- [21] A. Akhlaq, R. G. McKilliam, and R. Subramanian, "Basis construction for range estimation by phase unwrapping," *IEEE Signal Process. Lett.*, vol. 22, no. 11, pp. 2152–2156, 2015.
- [22] P. Stoica and R. Moses, *Spectral Analysis of Signals*. Upper Saddle River, NJ, USA: Prentice-Hall, 2005.

- [23] S. Tretter, “Estimating the frequency of a noisy sinusoid by linear regression,” *IEEE Trans. Inf. Theory*, vol. 31, no. 6, pp. 832–835, 1985.
- [24] R. G. McWilliam, B. G. Quinn, I. V. L. Clarkson, and B. Moran, “Frequency estimation by phase unwrapping,” *IEEE Trans. Signal Process.*, vol. 58, no. 6, pp. 2953–2963, 2010.
- [25] W. Xu, E.-C. Chang, L. K. Kwok, H. Lim, W. Cheng, and A. Heng, “Phase-unwrapping of SAR interferogram with multi-frequency or multi-baseline,” in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, Sep. 1994, pp. 730–732.
- [26] M. O. Damen, H. El Gamal, and G. Caire, “On maximum-likelihood detection and the search for the closest lattice point,” *IEEE Trans. Inf. Theory*, vol. 49, no. 10, pp. 2389–2402, 2003.
- [27] H. El Gamal, G. Caire, and M. O. Damen, “Lattice coding and decoding achieve the optimal diversity-multiplexing tradeoff of MIMO channels,” *IEEE Trans. Inf. Theory*, vol. 50, no. 6, pp. 968–985, 2004.
- [28] A. D. Murugan, H. El Gamal, M. O. Damen, and G. Caire, “A unified framework for tree search decoding: rediscovering the sequential decoder,” *IEEE Trans. Inf. Theory*, vol. 52, no. 3, pp. 933–953, 2006.
- [29] B. Ottersten, P. Stoica, and R. Roy, “Covariance matching estimation techniques for array signal processing applications,” *Digital Signal Processing*, vol. 8, no. 3, pp. 185–210, 1998.
- [30] L. GAN and H. Liu, “High dynamic range sensing using multi-channel modulo samplers,” in *2020 IEEE 11th Sensor Array and Multichannel Signal Processing Workshop (SAM)*, 2020, pp. 1–5.
- [31] S. Zhao, R. Ahmad, and L. C. Potter, “Venc design and velocity estimation for phase contrast MRI,” *IEEE Trans. Med. Imag.*, vol. 41, no. 12, pp. 3712–3724, 2022.
- [32] W. Wang, X. Li, W. Wang, and X.-G. Xia, “Maximum likelihood estimation based robust Chinese remainder theorem for real numbers and its fast algorithm,” *IEEE Trans. Signal Process.*, vol. 63, no. 13, pp. 3317–3331, 2015.
- [33] L. C. Godara and A. Cantoni, “Uniqueness and linear independence of steering vectors in array space,” *J. Acoust. Soc. Am.*, vol. 70, no. 2, pp. 467–475, Aug. 1981.
- [34] D. Tucker and L. C. Potter, “Speckle suppression in multi-channel coherent imaging: A tractable Bayesian approach,” *IEEE Trans. Comput. Imag.*, vol. 6, pp. 1429–1439, 2020.

- [35] J.-S. Lee, “A simple speckle smoothing algorithm for synthetic aperture radar images,” *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-13, no. 1, pp. 85–89, 1983.
- [36] J.-S. Lee, M. R. Grunes, and G. de Grandi, “Polarimetric SAR speckle filtering and its implication for classification,” *IEEE Trans. Geosci. Remote Sens.*, vol. 37, no. 5, pp. 2363–2373, 1999.
- [37] A. Buades, B. Coll, and J.-M. Morel, “A non-local algorithm for image denoising,” in *Proc. IEEE Comput. Soc. Conf. Comput. Vision Pattern Recognit.*, vol. 2, 2005, pp. 60–65.
- [38] C.-A. Deledalle, L. Denis, F. Tupin, A. Reigber, and M. Jäger, “NL-SAR: A unified nonlocal framework for resolution-preserving (Pol)(In)SAR denoising,” *IEEE Trans. Geosci. Remote Sens.*, vol. 53, no. 4, pp. 2021–2038, 2015.
- [39] H. Zhong, J. Zhang, and G. Liu, “Robust polarimetric SAR despeckling based on nonlocal means and distributed Lee filter,” *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 7, pp. 4198–4210, Jul. 2014.
- [40] G. Liu and H. Zhong, “Nonlocal means filter for polarimetric SAR data despeckling based on discriminative similarity measure,” *IEEE Geosci. Remote Sens. Lett.*, vol. 11, no. 2, pp. 514–518, Feb. 2014.
- [41] H. Li and Z. Qin, “PolSAR image speckle reduction algorithm based on TV-PDE,” in *Proc. SPIE*, vol. 8002, 2011, pp. 178 – 184.
- [42] C.-A. Deledalle, L. Denis, S. Tabti, and F. Tupin, “MuLoG, or how to apply Gaussian denoisers to multi-channel SAR speckle reduction?” *IEEE Trans. Image Process.*, vol. 26, no. 9, pp. 4389–4403, 2017.
- [43] X. Nie, H. Qiao, and B. Zhang, “A variational model for PolSAR data speckle reduction based on the Wishart distribution,” *IEEE Trans. Image Process.*, vol. 24, no. 4, pp. 1209–1222, Apr. 2015.
- [44] X. Nie, H. Qiao, B. Zhang, and X. Huang, “A nonlocal TV-based variational method for PolSAR data speckle reduction,” *IEEE Trans. Image Process.*, vol. 25, no. 6, pp. 2620–2634, Jun. 2016.
- [45] M. Çetin and W. C. Karl, “Feature-enhanced synthetic aperture radar image formation based on nonquadratic regularization,” *IEEE Trans. Image Process.*, vol. 10, no. 4, pp. 623–631, Apr. 2001.
- [46] H. E. Güven, A. Güngör, and M. Çetin, “An augmented Lagrangian method for complex-valued compressed SAR imaging,” *IEEE Trans. Comput. Imag.*, vol. 2, no. 3, pp. 235–250, Sep. 2016.

- [47] M. Çetin, I. Stojanović, N. O. Önhon, K. Varshney, S. Samadi, W. C. Karl, and A. S. Willsky, “Sparsity-driven synthetic aperture radar imaging: Reconstruction, autofocusing, moving targets, and compressed sensing,” *IEEE Signal Process. Mag.*, vol. 31, no. 4, pp. 27–40, Jul. 2014.
- [48] L. C. Potter, E. Ertin, J. T. Parker, and M. Çetin, “Sparsity and compressed sensing in radar imaging,” *Proceedings of the IEEE*, vol. 98, no. 6, pp. 1006–1020, Jun. 2010.
- [49] N. A. Goodman and L. C. Potter, “Pitfalls and possibilities of radar compressive sensing,” *Appl. Opt.*, vol. 54, no. 8, pp. C1–C13, Mar. 2015.
- [50] S. Zhang, Y. Liu, and X. Li, “Sparse aperture InSAR imaging via sequential multiple sparse Bayesian learning,” in *Sensors*, 2017.
- [51] G. Xu, M. Xing, L. Zhang, Y. Liu, and Y. Li, “Bayesian inverse synthetic aperture radar imaging,” *IEEE Geosci. Remote Sens. Lett.*, vol. 8, no. 6, pp. 1150–1154, Nov. 2011.
- [52] H. Liu, B. Jiu, H. Liu, and Z. Bao, “Superresolution ISAR imaging based on sparse Bayesian learning,” *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 8, pp. 5005–5013, Aug. 2014.
- [53] M. E. Tipping, “Sparse Bayesian learning and the relevance vector machine,” *J. Mach. Learn. Res.*, vol. 1, pp. 211–244, Sep. 2001. [Online]. Available: <https://doi.org/10.1162/15324430152748236>
- [54] C. J. Pellizzari, R. Trahan, H. Zhou, S. Williams, S. E. Williams, B. Nemati, M. Shao, and C. A. Bouman, “Synthetic aperture LADAR: A model-based approach,” *IEEE Trans. Comput. Imag.*, vol. 3, no. 4, pp. 901–916, Dec. 2017.
- [55] —, “Optically-coherent image formation and denoising using plug and play inversion framework,” *Applied Optics*, vol. 56, no. 16, pp. 4735–4744, Jun. 2017.
- [56] S. V. Venkatakrisnan, C. A. Bouman, and B. Wohlberg, “Plug-and-play priors for model based reconstruction,” in *2013 IEEE Global Conf. Signal Inf. Process.*, Dec. 2013, pp. 945–948.
- [57] S. H. Chan, X. Wang, and O. A. Elgendy, “Plug-and-play ADMM for image restoration: Fixed-point convergence and applications,” *IEEE Trans. Comput. Imag.*, vol. 3, no. 1, pp. 84–98, Mar. 2017.
- [58] X. Dong and Y. Zhang, “SAR image reconstruction from undersampled raw data using maximum a posteriori estimation,” *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 8, no. 4, pp. 1651–1664, Apr. 2015.

- [59] D. C. Munson Jr, J. Dennis O'Brien, and K. Jenkins, "A tomographic formulation of spotlight-mode synthetic aperture radar," *Proceedings of the IEEE*, vol. 17, pp. 917–925, Sep. 1983.
- [60] T. J. Kragh and A. A. Kharbouch, "Monotonic iterative algorithms for SAR image restoration," in *2006 IEEE Int. Conf. Image Process.*, Oct. 2006, pp. 645–648.
- [61] L. M. Novak, M. C. Burl, R. D. Chaney, and G. J. Owirka, "Optimal processing of polarimetric synthetic-aperture radar imagery," *The Lincoln Laboratory Journal*, pp. 273–290, 1990.
- [62] J. A. Tague and C. I. Caldwell, "Expectations of useful complex Wishart forms," *Multidimensional Systems and Signal Processing*, vol. 5, no. 3, pp. 263–279, Jul. 1994.
- [63] J. N. Ash, "Joint imaging and change detection for robust exploitation in interrupted SAR environments," in *Proc. SPIE*, vol. 8746, 2013, pp. 144–152.
- [64] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. R. Stat. Soc. Series B Stat. Methodol.*, vol. 39, no. 1, pp. 1–38, 1977.
- [65] C. M. Bishop, *Pattern Recognition and Machine Learning*. New York, NY, USA: Springer, 2006.
- [66] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 498–519, Feb. 2001.
- [67] P. Zwiernik, C. Uhler, and D. Richards, "Maximum likelihood estimation for linear Gaussian covariance models," *J. R. Stat. Soc. Series B Stat. Methodol.*, vol. 79, no. 4, pp. 1269–1292, 2016.
- [68] J. M. Tang and Y. Saad, "A probing method for computing the diagonal of a matrix inverse," *Numerical Linear Algebra with Applications*, vol. 19, no. 3, pp. 485–501, Apr. 2010.
- [69] M. Kubale, *Graph Colorings*. Providence, RI, USA: American Mathematical Society, 2004.
- [70] N. Nguyen and Q. H. Liu, "The regular Fourier matrices and nonuniform fast Fourier transforms," *SIAM Journal on Scientific Computing*, vol. 21, no. 1, pp. 283–293, 1999.

- [71] R. Ahmad, C. D. Austin, and L. C. Potter, "Toeplitz embedding for fast iterative regularized imaging," in *SPIE Algorithms for Synthetic Aperture Radar Imagery*, vol. 8051, 2011.
- [72] J. R. Shewchuk, "An introduction to the conjugate gradient method without the agonizing pain," Carnegie Mellon University, Pittsburgh, PA, USA, Tech. Rep., 1994.
- [73] G. Farage, S. Foucher, and G. B. Béné, "Comparison of PolSAR speckle filtering techniques," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, Jul. 2006, pp. 1760–1763.
- [74] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.
- [75] Hua Xie, L. E. Pierce, and F. T. Ulaby, "SAR speckle reduction using wavelet denoising and Markov random field modeling," *IEEE Trans. Geosci. Remote Sens.*, vol. 40, no. 10, pp. 2196–2212, Oct. 2002.
- [76] H. Feng, B. Hou, and M. Gong, "SAR image despeckling based on local homogeneous-region segmentation by using pixel-relativity measurement," *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 7, pp. 2724–2737, Jul. 2011.
- [77] L. Torres, S. J. S. Sant'Anna, C. da Costa Freitas, and A. C. Frery, "Speckle reduction in polarimetric sar imagery with stochastic distances and nonlocal means," *Pattern Recognit.*, vol. 47, no. 1, pp. 141–157, Jan. 2014.
- [78] Y. Yamaguchi, T. Nishikawa, M. Sengoku, and W. M. Boerner, "2-D polarimetric imaging by an FM-CW radar," in *Proc. Proc. IEEE Antennas Propag. Soc. Int. Symp. URSI Nat. Radio Sci. Meeting*, vol. 3, Jun. 1994, pp. 1998–2001.
- [79] M. Berger, J. Moreno, J. A. Johannessen, P. F. Levelt, and R. F. Hanssen, "ESA's sentinel missions in support of earth system science," *Remote Sens. Environ.*, vol. 120, pp. 84–90, 2012, the Sentinel Missions - New Opportunities for Science.
- [80] Y. Yamaguchi, K. Kimura, H. Yamada, S. Uratsuka, and W.-M. Boerner, "L-band polarimetric AIR/Pi-SAR images around Niigata city," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, vol. 1, Jun. 2002, pp. 423–425.
- [81] C. Bouman and K. Sauer, "A generalized gaussian image model for edge-preserving map estimation," *IEEE Trans. Image Process.*, vol. 2, no. 3, pp. 296–310, July 1993.

- [82] L. Denis, C. Deledalle, and F. Tupin, “From patches to deep learning: Combining self-similarity and neural networks for sar image despeckling,” in *IGARSS 2019 - 2019 IEEE International Geoscience and Remote Sensing Symposium*, July 2019, pp. 5113–5116.
- [83] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, “Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising,” *IEEE Trans. Image Process.*, vol. 26, no. 7, pp. 3142–3155, Jul. 2017.
- [84] C. J. Pellizzari, M. F. Spencer, and C. A. Bouman, “Coherent-image reconstruction using convolutional neural networks,” in *Imaging and Applied Optics*, 2019, p. MTu4D.4.
- [85] D. Tucker, J. N. Ash, and L. C. Potter, “SAR coherent change detection with variational expectation maximization,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 59, no. 3, pp. 2163–2175, 2023.
- [86] M. Preiss and N. J. S. Stacy, “Coherent change detection: Theoretical description and experimental results,” Defence Sci. Technol. Org., Tech. Rep. DSTO-TR-1851, 2006.
- [87] D. A. Yocky and B. F. Johnson, “Repeat-pass dual-antenna synthetic aperture radar interferometric change-detection post-processing,” *Photogrammetric Eng. and Remote Sens.*, pp. 425–429, 1998.
- [88] T. Strozzi, U. Wegmuller, C. L. Werner, A. Wiesmann, and V. Spreckels, “JERS SAR interferometry for land subsidence monitoring,” *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 7, pp. 1702–1708, 2003.
- [89] R. Touzi, A. Lopes, J. Bruniquel, and P. W. Vachon, “Coherence estimation for SAR imagery,” *IEEE Trans. Geosci. Remote Sens.*, vol. 37, no. 1, pp. 135–149, 1999.
- [90] F. Gatelli, A. Monti Guarnieri, F. Parizzi, P. Pasquali, C. Prati, and F. Rocca, “The wavenumber shift in SAR interferometry,” *IEEE Trans. Geosci. Remote Sens.*, vol. 32, no. 4, pp. 855–865, 1994.
- [91] C. V. Jakowatz, D. E. Wahl, P. H. Eichel, D. C. Ghiglia, and P. A. Thompson, *Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach*. New York, NY, USA: Springer, 1996.
- [92] M. Cha, R. D. Phillips, P. J. Wolfe, and C. D. Richmond, “Two-stage change detection for synthetic aperture radar,” *IEEE Trans. Geosci. Remote Sens.*, vol. 53, no. 12, pp. 6547–6560, 2015.

- [93] M. Preiss, D. A. Gray, and N. J. S. Stacy, “Detecting scene changes using synthetic aperture radar interferometry,” *IEEE Trans. Geosci. Remote Sens.*, vol. 44, no. 8, pp. 2041–2054, 2006.
- [94] D. E. Wahl, D. A. Yocky, C. V. Jakowatz, and K. M. Simonson, “A new maximum-likelihood change estimator for two-pass SAR coherent change detection,” *IEEE Trans. Geosci. Remote Sens.*, vol. 54, no. 4, pp. 2460–2469, 2016.
- [95] M. Newey, G. Benitz, and S. Kogon, “A generalized likelihood ratio test for SAR CCD,” in *Proc. Conf. Rec. 46th Asilomar Conf. Signals, Syst., Comput.*, 2012, pp. 1727–1730.
- [96] R. Sabry, “A new coherency formalism for change detection and phenomenology in SAR imagery: A field approach,” *IEEE Trans. Geosci. Remote Sens.*, vol. 6, no. 3, pp. 458–462, 2009.
- [97] J.-S. Lee, S. R. Cloude, K. P. Papathanassiou, M. R. Grunes, and I. H. Woodhouse, “Speckle filtering and coherence estimation of polarimetric SAR interferometry data for forest applications,” *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 10, pp. 2254–2263, 2003.
- [98] C. López-Martínez, X. Fabregas, and E. Pottier, “Wavelet transform-based interferometric SAR coherence estimator,” *IEEE Signal Process. Lett.*, vol. 12, no. 12, pp. 831–834, 2005.
- [99] C.-A. Deledalle, L. Denis, and F. Tupin, “NL-InSAR: Nonlocal interferogram estimation,” *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 4, pp. 1441–1452, 2011.
- [100] G. Ferraioli, C.-A. Deledalle, L. Denis, and F. Tupin, “PARISAR: Patch-based estimation and regularized inversion for multibaseline SAR interferometry,” *IEEE Trans. Geosci. Remote Sens.*, vol. 56, no. 3, pp. 1626–1636, 2018.
- [101] F. Sica, D. Cozzolino, X. X. Zhu, L. Verdoliva, and G. Poggi, “InSAR-BM3D: A nonlocal filter for SAR interferometric phase restoration,” *IEEE Trans. Geosci. Remote Sens.*, vol. 56, no. 6, pp. 3456–3467, 2018.
- [102] Y. Yang, X. Cong, K. Long, Y. Luo, W. Xie, and Q. Wan, “MRF model-based joint interrupted SAR imaging and coherent change detection via variational Bayesian inference,” *Signal Process.*, vol. 151, pp. 144–154, 2018.
- [103] O. Dikmen and A. T. Cemgil, “Gamma Markov random fields for audio source modeling,” *IEEE Trans. Audio, Speech, Language Process.*, vol. 18, no. 3, pp. 589–601, 2010.

- [104] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, “Variational inference: A review for statisticians,” *J. Am. Stat. Assoc.*, vol. 112, no. 518, pp. 859–877, 2017.
- [105] D. G. Tzikas, A. C. Likas, and N. P. Galatsanos, “The variational approximation for Bayesian inference,” *IEEE Signal Process. Mag.*, vol. 25, no. 6, pp. 131–146, 2008.
- [106] A. T. Cemgil and O. Dikmen, “Conjugate gamma Markov random fields for modelling nonstationary sources,” in *Proc. Int. Conf. Independent Compon. Anal. Signal Separation*, 2007, pp. 697–705.
- [107] Y. Altmann, M. Pereyra, and S. McLaughlin, “Bayesian nonlinear hyperspectral unmixing with spatial residual component analysis,” *IEEE Trans. Comput. Imag.*, vol. 1, no. 3, pp. 174–185, 2015.
- [108] Y. Altmann, A. Maccarone, A. McCarthy, G. Newstadt, G. S. Buller, S. McLaughlin, and A. Hero, “Robust spectral unmixing of sparse multispectral lidar waveforms using gamma Markov random fields,” *IEEE Trans. Comput. Imag.*, vol. 3, no. 4, pp. 658–670, 2017.
- [109] C. Oliver and S. Quegan, *Understanding Synthetic Aperture Radar Images*. Raleigh, NC, USA: SciTech, 2004.
- [110] V. Cevher, M. Duarte, C. Hegde, and R. Baraniuk, “Sparse signal recovery using Markov random fields,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2009, pp. 257–264.
- [111] G. E. Hinton, “Training products of experts by minimizing contrastive divergence,” *Neural Comput.*, vol. 14, no. 8, pp. 1771–1800, 2002.
- [112] O. Dikmen and A. T. Cemgil, “Inference and parameter estimation in gamma chains,” Univ. of Cambridge, Tech. Rep. CUED/F-IN-FENG/TR.596, Feb. 2008.
- [113] Z.-P. Liang, “A model-based method for phase unwrapping,” *IEEE Trans. Med. Imag.*, vol. 15, no. 6, pp. 893–897, 1996.
- [114] J. Zhang, “The mean field theory in EM procedures for Markov random fields,” *IEEE Trans. Signal Process.*, vol. 40, no. 10, pp. 2570–2583, 1992.
- [115] J. Besag, “On the statistical analysis of dirty pictures,” *J. R. Stat. Soc. Series B Stat. Methodol.*, vol. 48, no. 3, pp. 259–302, 1986.
- [116] R. C. Daudt, B. Le Saux, A. Boulch, and Y. Gousseau, “Urban change detection for multispectral earth observation using convolutional neural networks,” in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, 2018, pp. 2115–2118.

- [117] S. R. Cloude, *Polarisation: Applications in Remote Sensing*. New York, NY, USA: Oxford Univ. Press, 2009.
- [118] E. Ertin, C. D. Austin, S. Sharma, R. L. Moses, and L. C. Potter, “GOTCHA experience report: three-dimensional SAR imaging with complete circular apertures,” in *Proc. SPIE*, vol. 6568, 2007, pp. 9 – 20.
- [119] D. Tucker and L. C. Potter, “Polarimetric SAR despeckling with convolutional neural networks,” *IEEE Trans. Geosci. Remote Sens.*, vol. 60, pp. 1–12, 2022.
- [120] S. R. Cloude and E. Pottier, “An entropy based classification scheme for land applications of polarimetric SAR,” *IEEE Trans. Geosci. Remote Sens.*, vol. 35, no. 1, pp. 68–78, 1997.
- [121] I. Hajnsek, E. Pottier, and S. R. Cloude, “Inversion of surface parameters from polarimetric SAR,” *IEEE Trans. Geosci. Remote Sens.*, vol. 41, no. 4, pp. 727–744, 2003.
- [122] X. Huang, J. Wang, and J. Shang, “An integrated surface parameter inversion scheme over agricultural fields at early growing stages by means of C-band polarimetric RADARSAT-2 imagery,” *IEEE Trans. Geosci. Remote Sens.*, vol. 54, no. 5, pp. 2510–2528, 2016.
- [123] J.-S. Lee, T. L. Ainsworth, Y. Wang, and K.-S. Chen, “Polarimetric SAR speckle filtering and the extended sigma filter,” *IEEE Trans. Geosci. Remote Sens.*, vol. 53, no. 3, pp. 1150–1160, 2015.
- [124] J. Chen, Y. Chen, W. An, Y. Cui, and J. Yang, “Nonlocal filtering for polarimetric SAR data: A pretest approach,” *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 5, pp. 1744–1754, May 2011.
- [125] G. Vasile, E. Trounev, J.-S. Lee, and V. Buzuloiu, “Intensity-driven adaptive-neighborhood technique for polarimetric and interferometric SAR parameters estimation,” *IEEE Trans. Geosci. Remote Sens.*, vol. 44, no. 6, pp. 1609–1621, 2006.
- [126] C. Dong, C. C. Loy, K. He, and X. Tang, “Image super-resolution using deep convolutional networks,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 2, pp. 295–307, 2016.
- [127] B. Lim, S. Son, H. Kim, S. Nah, and K. M. Lee, “Enhanced deep residual networks for single image super-resolution,” in *Proc. IEEE Conf. Comput. Vision Pattern Recognit. Workshop*, 2017, pp. 1132–1140.
- [128] Z. Zhao, P. Zheng, S. Xu, and X. Wu, “Object detection with deep learning: A review,” *IEEE Trans. Neural Netw.*, vol. 30, no. 11, pp. 3212–3232, 2019.

- [129] M. Dahmane, S. Foucher, M. Beaulieu, F. Riendeau, Y. Bouroubi, and M. Benoit, "Object detection in Pleiades images using deep features," in *2016 IEEE Int. Geosci. Remote Sens. Symp.*, 2016, pp. 1552–1555.
- [130] C. Tian, L. Fei, W. Zheng, Y. Xu, W. Zuo, and C.-W. Lin, "Deep learning on image denoising: An overview," *Neural Netw.*, vol. 131, pp. 251–275, 2020.
- [131] K. Zhang, W. Zuo, and L. Zhang, "FFDNet: Toward a fast and flexible solution for CNN based image denoising," 2017, *arXiv:1710.04026*.
- [132] G. Chierchia, D. Cozzolino, G. Poggi, and L. Verdoliva, "SAR image despeckling through convolutional neural networks," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, 2017, pp. 5438–5441.
- [133] P. Wang, H. Zhang, and V. M. Patel, "SAR image despeckling using a convolutional neural network," *IEEE Signal Process. Lett.*, vol. 24, no. 12, pp. 1763–1767, 2017.
- [134] Q. Zhang, Q. Yuan, J. Li, Z. Yang, X. Ma, H. Shen, and L. Zhang, "Learning a dilated residual network for SAR image despeckling," *Remote Sens.*, vol. 10, pp. 1–18, Feb. 2018.
- [135] F. Lattari, B. Gonzalez Leon, F. Asaro, A. Rucci, C. Prati, and M. Matteucci, "Deep Learning for SAR Image Despeckling," *Remote Sens.*, vol. 11, no. 13, p. 1532, Jun. 2019.
- [136] X. X. Zhu, S. Montazeri, M. Ali, Y. Hua, Y. Wang, L. Mou, Y. Shi, F. Xu, and R. Bamler, "Deep learning meets SAR," Jun. 2020, *arXiv:2006.10027*.
- [137] L. Lin, J. Li, Q. Yuan, and H. Shen, "Polarimetric SAR image super-resolution via deep convolutional neural network," in *2019 IEEE Int. Geosci. Remote Sens. Symp.*, 2019, pp. 3205–3208.
- [138] H. Shen, L. Lin, J. Li, Q. Yuan, and L. Zhao, "A residual convolutional neural network for polarimetric SAR image super-resolution," *J. Photogrammetry Remote Sens.*, vol. 161, pp. 90–108, 2020.
- [139] Z. Zhang, H. Wang, F. Xu, and Y. Jin, "Complex-valued convolutional neural network and its application in polarimetric SAR image classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 55, no. 12, pp. 7177–7188, 2017.
- [140] M. Castelluccio, G. Poggi, C. Sansone, and L. Verdoliva, "Land Use Classification in Remote Sensing Images by Convolutional Neural Networks," 2015, *arXiv:1508.00092*.

- [141] F. Xu, H. Wang, and Y. Jin, “Deep learning as applied in SAR target recognition and terrain classification,” *J. Radars*, vol. 6, pp. 136–148, Apr. 2017.
- [142] R. Haensch and O. Hellwich, “Complex-valued convolutional neural networks for object detection in PolSAR data,” in *8th Eur. Conf. Synthetic Aperture Radar*, 2010, pp. 1–4.
- [143] S. Foucher, M. Beaulieu, M. Dahmane, and F. Cavayas, “Deep speckle noise filtering,” in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, 2017, pp. 5311–5314.
- [144] S. Foucher, M. Beaulieu, F. Cavayas, and M. Dahmane, “Revisiting an iterative speckle filtering technique,” in *2019 IEEE Int. Geosci. Remote Sens. Symp.*, 2019, pp. 5213–5216.
- [145] X. Yang, T. Pan, W. Yang, and H. Li, “PolSAR image despeckling using trained models on single channel SAR images,” in *Proc. 6th Asia-Pacific Conf. Synthetic Aperture Radar*, 2019, pp. 1–4.
- [146] T. Pan, D. Peng, W. Yang, and H.-C. Li, “A filter for SAR image despeckling using pre-trained convolutional neural network model,” *Remote Sens.*, vol. 11, Oct. 2019.
- [147] E. Dalsasso, X. Yang, L. Denis, F. Tupin, and W. Yang, “SAR image despeckling by deep neural networks: from a pre-trained model to an end-to-end training strategy,” *Remote Sens.*, vol. 12, no. 16, 2020.
- [148] D. Cozzolino, L. Verdoliva, G. Scarpa, and G. Poggi, “Nonlocal CNN SAR image despeckling,” *Remote Sens.*, vol. 12, no. 6, p. 1006, Mar. 2020.
- [149] A. B. Molini, D. Valsesia, G. Fracastoro, and E. Magli, “Speckle2Void: Deep self-supervised SAR despeckling with blind-spot convolutional neural networks,” *IEEE Trans. Geosci. Remote Sens.*, pp. 1–17, 2021.
- [150] E. Dalsasso, L. Denis, and F. Tupin, “SAR2SAR: a self-supervised despeckling algorithm for SAR images,” 2020, *arXiv:2006.15037*.
- [151] G. Fracastoro, E. Magli, G. Poggi, G. Scarpa, D. Valsesia, and L. Verdoliva, “Deep learning methods for synthetic aperture radar image despeckling: An overview of trends and perspectives,” *IEEE Geosci. Remote Sens. Magazine*, vol. 9, no. 2, pp. 29–51, 2021.
- [152] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. IEEE Conf. Comput. Vision Pattern Recognit.*, 2016, pp. 770–778.

- [153] D. Nagar and A. Gupta, “Expectations of functions of complex Wishart matrix,” *Acta Applicandae Math.*, vol. 113, pp. 265–288, Mar. 2011.
- [154] J.-S. Lee, J.-H. Wen, T. L. Ainsworth, K.-S. Chen, and A. J. Chen, “Improved sigma filter for speckle filtering of SAR imagery,” *IEEE Trans. Geosci. Remote Sens.*, vol. 47, no. 1, pp. 202–213, 2009.
- [155] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *Int. Conf. Mach. Learn.*, ser. ICML’15, 2015, p. 448–456.
- [156] X. Yang, L. Denis, F. Tupin, and W. Yang, “SAR image despeckling using pre-trained convolutional neural network models,” in *Proc. Joint Urban Remote Sens. Event*, May 2019, pp. 1–4.
- [157] P. A. Rosen, S. Hensley, K. Wheeler, G. Sadowy, T. Miller, S. Shaffer, R. Muellerschoen, C. Jones, S. Madsen, and H. Zebker, “UAVSAR: New NASA airborne SAR system for research,” *IEEE Aerosp. Electron. Syst. Mag.*, vol. 22, no. 11, pp. 21–28, 2007.
- [158] A. Lapini, T. Bianchi, F. Argenti, and L. Alparone, “Blind speckle decorrelation for SAR image despeckling,” *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 2, pp. 1044–1058, 2014.
- [159] A. Arienzo, F. Argenti, L. Alparone, and M. Gherardelli, “Accurate despeckling and estimation of polarimetric features by means of a spatial decorrelation of the noise in complex PolSAR data,” *Remote Sens.*, vol. 12, no. 2, p. 331, 2020.
- [160] K. He, X. Zhang, S. Ren, and J. Sun, “Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification,” in *2015 IEEE Int. Conf. Comput. Vision (ICCV)*, 2015, pp. 1026–1034.
- [161] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, “Image quality assessment: from error visibility to structural similarity,” *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [162] G. E. P. Box and D. A. Pierce, “Distribution of residual autocorrelations in autoregressive-integrated moving average time series models,” *J. Am. Stat. Assoc.*, vol. 65, no. 332, pp. 1509–1526, 1970.
- [163] S. R. Cloude and E. Pottier, “A review of target decomposition theorems in radar polarimetry,” *IEEE Trans. Geosci. Remote Sens.*, vol. 34, no. 2, pp. 498–518, 1996.
- [164] J. J. van Zyl, H. A. Zebker, and C. Elachi, “Imaging radar polarization signatures: Theory and observation,” *Radio Sci.*, vol. 22, no. 4, pp. 529–543, 1987.

- [165] S. N. Anfinsen, A. P. Doulgeris, and T. Eltoft, “Estimation of the equivalent number of looks in polarimetric synthetic aperture radar imagery,” *IEEE Trans. Geosci. Remote Sens.*, vol. 47, no. 11, pp. 3795–3809, 2009.
- [166] A. Vedaldi and K. Lenc, “MatConvNet: Convolutional neural networks for MATLAB,” in *Proc. 23rd ACM Int. Conf. Multimedia*, ser. MM ’15, 2015, pp. 689–692.
- [167] Y. Yamaguchi, T. Moriyama, M. Ishido, and H. Yamada, “Four-component scattering model for polarimetric SAR image decomposition,” *IEEE Trans. Geosci. Remote Sens.*, vol. 43, no. 8, pp. 1699–1706, 2005.
- [168] M. Markl, A. Frydrychowicz, S. Kozerke, M. Hope, and O. Wieben, “4D flow MRI,” *J. Magn. Reson. Imaging*, vol. 36, no. 5, pp. 1015–1036, 2012.
- [169] A. Ghasemmehdi and E. Agrell, “Faster recursions in sphere decoding,” *IEEE Trans. Inf. Theory*, vol. 57, no. 6, pp. 3530–3536, 2011.
- [170] X.-W. Chang and T. Zhou, “MILES: MATLAB package for solving mixed integer least squares problems,” *GPS Solutions*, vol. 11, pp. 289–294, 11 2007.
- [171] D. Micciancio and P. Voulgaris, “A deterministic single exponential time algorithm for most lattice problems based on voronoi cell computations,” *SIAM Journal on Computing*, vol. 42, no. 3, pp. 1364–1391, 2013.
- [172] M. Pohst, “On the computation of lattice vectors of minimal length, successive minima and reduced bases with applications,” *SIGSAM Bull.*, vol. 15, no. 1, p. 37–44, Feb. 1981.
- [173] U. Fincke and M. Pohst, “Improved methods for calculating vectors of short length in a lattice, including a complexity analysis,” *Math. Comput.*, vol. 44, no. 170, pp. 463–471, 1985.
- [174] B. Hassibi and H. Vikalo, “On the sphere-decoding algorithm I. Expected complexity,” *IEEE Trans. Signal Process.*, vol. 53, no. 8, pp. 2806–2818, 2005.
- [175] J. Conway and N. Sloane, “Fast quantizing and decoding and algorithms for lattice quantizers and codes,” *IEEE Trans. Inf. Theory*, vol. 28, no. 2, pp. 227–232, 1982.
- [176] R. G. McWilliam, W. D. Smith, and I. V. L. Clarkson, “Linear-time nearest point algorithms for Coxeter lattices,” *IEEE Trans. Inf. Theory*, vol. 56, no. 3, pp. 1015–1022, 2010.

- [177] S. Sahraei and M. Gastpar, “Polynomially solvable instances of the shortest and closest vector problems with applications to compute-and-forward,” *IEEE Transactions on Information Theory*, vol. 63, no. 12, pp. 7780–7792, 2017.
- [178] A. Akhlaq, R. McKilliam, R. Subramanian, and A. Pollok, “Selecting wavelengths for least squares range estimation,” *IEEE Trans. Signal Process.*, vol. 64, no. 20, pp. 5205–5216, 2016.
- [179] W. Li, X. Wang, and B. Moran, “Wireless signal travel distance estimation using non-coprime wavelengths,” *IEEE Signal Process. Lett.*, vol. 24, no. 1, pp. 27–31, 2017.
- [180] D. Tucker, S. Zhao, and L. C. Potter, “Multivariate maximum likelihood estimation with the Chinese remainder theorem,” *IEEE Trans. Signal Process.*, 2024, to be submitted.
- [181] O. Knill, “A multivariable Chinese remainder theorem,” *arXiv preprint arXiv:1206.5114*, 2012.
- [182] B. Sury, “Multivariable Chinese remainder theorem,” *Resonance*, vol. 20, pp. 206–216, Mar. 2015.
- [183] O. Goldreich, D. Ron, and M. Sudan, “Chinese remaindering with errors,” *IEEE Trans. Inf. Theory*, vol. 46, no. 4, pp. 1330–1338, 2000.
- [184] X.-G. Xia and G. Wang, “Phase unwrapping and a robust Chinese remainder theorem,” *IEEE Signal Process. Lett.*, vol. 14, no. 4, pp. 247–250, 2007.
- [185] L. Xiao, X.-G. Xia, and W. Wang, “Multi-stage robust Chinese remainder theorem,” *IEEE Trans. Signal Process.*, vol. 62, no. 18, pp. 4772–4785, 2014.
- [186] G. Maze, J. Rosenthal, and U. Wagner, “Natural density of rectangular unimodular integer matrices,” *Linear Algebra Appl.*, vol. 434, no. 5, pp. 1319–1324, 2011.
- [187] W. Li, X. Wang, X. Wang, and B. Moran, “Resolving phase measurement ambiguity in presence of coloured noise,” *Electron. Lett.*, vol. 49, pp. 1188–1190, 2013.
- [188] D. Prasanna, C. Sriram, and C. R. Murthy, “On the identifiability of sparse vectors from modulo compressed sensing measurements,” *IEEE Signal Process. Lett.*, vol. 28, pp. 131–134, 2021.
- [189] D. Prasanna, C. R. Murthy, and C. Sriram, “On the application of modulo-ADCs for compressed sensing,” in *Proc. Asilomar Conf. Signals, Syst., Comput.*, 2021, pp. 852–856.

- [190] A. Benedetti, “Phase unwrapping with multiple wavelengths on the flat torus,” in *Proc. 54th Asilomar Conf. Signals, Syst., Comput.*, 2020, pp. 1301–1304.
- [191] D. Micciancio and S. Goldwasser, *Complexity of Lattice Problems: A Cryptographic Perspective*. Boston, MA, USA: Kluwer, 2002.
- [192] M. R. Bremner, *Lattice Basis Reduction: An Introduction to the LLL Algorithm and Its Applications*. Boca Raton, USA: CRC Press, 2011.
- [193] J. Martinet, *Perfect Lattices in Euclidean Spaces*. Berlin, Germany: Springer-Verlag, 2003.
- [194] A. Schrijver, *Theory of Linear and Integer Programming*. Chichester, U.K.: Wiley, 1998.
- [195] C. Ding, D. Pei, and A. Salomaa, *Chinese Remainder Theorem*. Singapore: World Scientific, 1996.
- [196] N. Koblitz, *A Course in Number Theory and Cryptography*, 2nd ed. New York, NY, USA: Springer-Verlag, 1994.
- [197] L. Babai, “On Lovász’ lattice reduction and the nearest lattice point problem,” *Combinatorica*, vol. 6, pp. 1–13, 1986.
- [198] J. H. Silverman, J. Pipher, and J. Hoffstein, *An Introduction to Mathematical Cryptography*. New York, NY, USA: Springer, 2008.
- [199] H. Edelsbrunner, J. O’Rourke, and R. Seidel, “Constructing arrangements of lines and hyperplanes with applications,” *SIAM Journal on Computing*, vol. 15, no. 2, pp. 341–363, 1986.
- [200] N. Sleumer, “Output-sensitive cell enumeration in hyperplane arrangements,” in *Proc. 6th Scand. Workshop Algorithm Theory*, vol. 1432, 1998, pp. 300–309.
- [201] D. Micciancio, “The hardness of the closest vector problem with preprocessing,” *IEEE Trans. Inf. Theory*, vol. 47, no. 3, pp. 1212–1215, 2001.
- [202] M. Pohst, “A modification of the LLL reduction algorithm,” *J. Symbolic Comput.*, vol. 4, pp. 123–127, 1987.
- [203] C. C. Sims, *Computation with Finitely Presented Groups*. Cambridge, U.K.: Cambridge Univ. Press, 1994.
- [204] A. K. Lenstra, H. W. Lenstra, and L. Lovász, “Factoring polynomials with rational coefficients,” *Math. Ann.*, vol. 261, p. 515–534, 1982.

- [205] H. Cohen, *A Course in Computational Algebraic Number Theory*. Berlin, Germany: Springer-Verlag, 1993.
- [206] X.-W. Chang and G. H. Golub, “Solving ellipsoid-constrained integer least squares problems,” *SIAM J. Matrix Anal. Appl.*, vol. 31, no. 3, pp. 1071–1089, 2010.
- [207] X.-W. Chang, “MILES: MATLAB package for solving mixed integer least squares problems,” <http://www.cs.mcgill.ca/~chang/software/MILES.php>, School of Computer Science, McGill University, 2006, last updated: Dec., 2022. [Online]. Available: <http://www.cs.mcgill.ca/~chang/software/MILES.php>
- [208] C. Schnorr and M. Euchner, “Lattice basis reduction: Improved practical algorithms and solving subset sum problems,” *Math. Prog.*, vol. 66, p. 181–199, 1994.
- [209] A. Pruitt, A. Rich, Y. Liu, N. Jin, L. Potter, M. Tong, S. Rajpal, O. Simonetti, and R. Ahmad, “Fully self-gated whole-heart 4D flow imaging from a 5-minute scan,” *Magn. Reson. Med.*, vol. 85, no. 3, pp. 1222–1236, 2021. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.28491>
- [210] S. Sahraei and M. Gastpar, “Polynomially solvable instances of the shortest and closest vector problems with applications to compute-and-forward,” *IEEE Trans. Inf. Theory*, vol. 63, no. 12, pp. 7780–7792, 2017.
- [211] F. Athley, “Optimization of element positions for direction finding with sparse arrays,” in *Proc. 11th IEEE Signal Process. Workshop Statistical Signal Process.*, 2001, pp. 516–519.
- [212] D. Malioutov, M. Cetin, and A. S. Willsky, “A sparse signal reconstruction perspective for source localization with sensor arrays,” *IEEE Trans. Signal Process.*, vol. 53, no. 8, pp. 3010–3022, 2005.
- [213] T. Ballal and C. J. Bleakley, “Phase-difference ambiguity resolution for a single-frequency signal,” *IEEE Signal Process. Lett.*, vol. 15, pp. 853–856, 2008.
- [214] H. Chen, T. Ballal, and T. Y. Al-Naffouri, “DOA estimation with non-uniform linear arrays: A phase-difference projection approach,” *IEEE Wireless Commun. Lett.*, vol. 10, no. 11, pp. 2435–2439, 2021.
- [215] Y. Hu and S. Gannot, “Closed-form single source direction-of-arrival estimator using first-order relative harmonic coefficients,” in *IEEE Int. Conf. Acoust., Speech Signal Process.*, 2022, pp. 726–730.

- [216] D. Tucker, S. Zhao, R. Ahmad, and L. C. Potter, "Alias-free arrays," *IEEE Signal Process. Lett.*, vol. 29, pp. 2457–2461, 2022.
- [217] J. Jaldén and B. Ottersten, "On the complexity of sphere decoding in digital communications," *IEEE Trans. Signal Process.*, vol. 53, no. 4, pp. 1474–1484, 2005.
- [218] Y. Nasser, S. Aubert, and F. Nouvel, "A simplified hard output sphere decoder for large MIMO systems with the use of efficient search center and reduced domain neighborhood study," *EURASIP Journal on Wireless Communications and Networking*, 10 2015.
- [219] K. M. Johnson and M. Markl, "Improved SNR in phase contrast velocimetry with five-point balanced flow encoding," *Magn. Reson. Med.*, vol. 63, no. 2, pp. 349–355, 2010.
- [220] C. Binter, V. Knobloch, R. Manka, A. Sigfridsson, and S. Kozerke, "Bayesian multipoint velocity encoding for concurrent flow and turbulence mapping," *Magn. Reson. Medicine*, vol. 69, no. 5, pp. 1337–1345, 2013.
- [221] D. A. Holdsworth, "Angle of arrival estimation for all-sky interferometric meteor radar systems," *Radio Sci.*, vol. 40, no. 6, pp. 1–8, 2005.
- [222] P. Stoica and A. Nehorai, "Performance study of conditional and unconditional direction-of-arrival estimation," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 10, pp. 1783–1795, 1990.
- [223] F. Athley, "Threshold region performance of maximum likelihood direction of arrival estimators," *IEEE Trans. Signal Process.*, vol. 53, no. 4, pp. 1359–1373, 2005.
- [224] H. Carrillo, A. Osses, S. Uribe, and C. Bertoglio, "Optimal dual-VENC unwrapping in phase-contrast MRI," *IEEE Trans. Med. Imag.*, vol. 38, no. 5, pp. 1263–1270, 2019.
- [225] M. Loecher and D. B. Ennis, "Velocity reconstruction with nonconvex optimization for low-velocity-encoding phase-contrast MRI," *Magn. Reson. Med.*, vol. 80, no. 1, pp. 42–52, 2018.
- [226] H. Gazzah and S. Marcos, "Cramer-Rao bounds for antenna array design," *IEEE Trans. Signal Process.*, vol. 54, no. 1, pp. 336–345, 2006.
- [227] H. Gazzah and K. Abed-Meraim, "Optimum ambiguity-free directional and omnidirectional planar antenna arrays for DOA estimation," *IEEE Trans. Signal Process.*, vol. 57, no. 10, pp. 3942–3953, 2009.

- [228] A. Ghani, F. Keyvani, and S. H. Sedighy, “Antenna array placement on limited bound for isotropic and optimal direction-of-arrival estimation,” *IET Signal Process.*, vol. 12, no. 3, pp. 277–283, 2018.
- [229] K.-C. Tan, S. S. Goh, and E.-C. Tan, “A study of the rank-ambiguity issues in direction-of-arrival estimation,” *IEEE Trans. Signal Process.*, vol. 44, no. 4, pp. 880–887, 1996.
- [230] H. L. Van Trees, *Optimum Array Processing: Part IV of Detection, Estimation, and Modulation Theory*. New York, NY, USA: Wiley, 2002.
- [231] B. Friedlander, “On the Cramér-Rao bound for sparse linear arrays,” in *Proc. Asilomar Conf. Signals, Syst., Comput.*, 2020, pp. 1255–1259.
- [232] K.-C. Tan, G.-L. Oh, and M. H. Er, “A study of the uniqueness of steering vectors in array processing,” *Signal Process.*, vol. 34, no. 3, pp. 245–256, 1993.
- [233] A. Moffet, “Minimum-redundancy linear arrays,” *IEEE Trans. Antennas Propag.*, vol. 16, no. 2, pp. 172–175, 1968.
- [234] P. Pal and P. P. Vaidyanathan, “Nested arrays: A novel approach to array processing with enhanced degrees of freedom,” *IEEE Trans. Signal Process.*, vol. 58, no. 8, pp. 4167–4181, 2010.
- [235] P. P. Vaidyanathan and P. Pal, “Sparse sensing with coprime arrays,” in *44th Asilomar Conf. Signals, Systems and Computers*, 2010, pp. 1405–1409.
- [236] H. Chen, Y. Wang, and S. Wan, “Performance improvement of estimation direction-of-arrival via array geometry arrangement,” in *Proc. IEEE Int. Symp. Antennas and Propag. Soc.*, vol. 3, 1999, pp. 1600–1603.
- [237] M. D. Zoltowski and C. P. Mathews, “Real-time frequency and 2-D angle estimation with sub-Nyquist spatio-temporal sampling,” *IEEE Trans. Signal Process.*, vol. 42, no. 10, pp. 2781–2794, 1994.
- [238] T. Ballal and C. J. Bleakley, “DOA estimation of multiple sparse sources using three widely-spaced sensors,” in *17th Eur. Signal Process. Conf.*, 2009, pp. 1978–1982.
- [239] N. J. Pelc, R. J. Herfkens, A. Shimakawa, and D. R. Enzmann, “Phase contrast cine magnetic resonance imaging,” *Magn. Reson. Q.*, vol. 7, no. 4, pp. 229–254, 1991.
- [240] S. Zhao, R. Ahmad, and L. C. Potter, “Maximizing unambiguous velocity range in phase-contrast MRI with multipoint encoding,” in *IEEE Int. Symp. Biomed. Imag.*, 2022, pp. 1–5.

- [241] O. Lange and B. Yang, “Optimization of array geometry for direction-of-arrival estimation using a priori information,” *Adv. Radio Sci.*, vol. 8, pp. 87–94, 2010.
- [242] C. Gauss, *Disquisitiones Arithmeticae*, ser. Translated by A.A. Clarke. New Haven, CT: Yale University Press, 1966.
- [243] A. Barvinok, “Math 669: Combinatorics, Geometry, and Complexity of Integer Points,” math.lsa.umich.edu/~barvinok/c20.html, accessed 2021-10-11.
- [244] J. Chen, Y. Feng, Y. Liu, and W. Wu, “On the probability of generating a primitive matrix,” *J. Syst. Sci. Complex.*, 2024.
- [245] I. Reiner, “Unimodular complements,” *Amer. Math. Monthly*, vol. 63, no. 4, pp. 246–247, 1956.
- [246] A. Kintz and I. J. Gupta, “Airborne antenna array calibration with signals of opportunity,” in *IEEE Antennas Propag. Soc. Int. Symp. (APSURSI)*, 2013, pp. 1264–1265.
- [247] J. L. Volakis, A. J. O’Brien, and C.-C. Chen, “Small and adaptive antennas and arrays for GNSS applications,” *Proc. IEEE*, vol. 104, no. 6, pp. 1221–1232, 2016.
- [248] H. Liao and A. Shamim, “A low-profile and high-aperture-efficiency hexagonal circularly polarized microstrip antenna array,” *IEEE Antennas Wireless Propag. Lett.*, vol. 21, no. 3, pp. 615–619, 2022.
- [249] B. Rafaely, B. Weiss, and E. Bachmat, “Spatial aliasing in spherical microphone arrays,” *IEEE Trans. Signal Process.*, vol. 55, no. 3, pp. 1003–1010, 2007.
- [250] J. Gebbie, M. Siderius, P. L. Nielsen, and J. Miller, “Passive localization of noise-producing targets using a compact volumetric array,” *J. Acoust. Soc. Am.*, vol. 136, no. 1, pp. 80–89, 2014.
- [251] D. Kastinen and J. Kero, “Probabilistic analysis of ambiguities in radar echo direction of arrival from meteors,” *Atmos. Meas. Tech.*, vol. 13, no. 12, pp. 6813–6835, 2020.
- [252] D. Avis, H. Miyata, and S. Moriyama, “Families of polytopal digraphs that do not satisfy the shelling property,” *Computational Geometry*, vol. 46, no. 3, pp. 382–393, 2013.
- [253] J. E. Goodman, J. O’Rourke, and C. D. Tóth, *Handbook of Discrete and Computational Geometry*, 3rd ed. Boca Raton, USA: CRC Press, 2017.
- [254] K. Kreutz-Delgado, “The Complex Gradient Operator and the CR-Calculus,” Jun. 2009, *arXiv:0906.4835*.