

Optical System Identification for Passive Electro-Optical Imaging

by

Joel W. LeBlanc

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Electrical Engineering: Systems)
in The University of Michigan
2019

Doctoral Committee:

Professor Alfred Hero, Chair
Professor Jeffrey Fessler
Professor Stephen Rand
Professor John Schotland

Joel W. LeBlanc
jwleblan@umich.edu

ORCID iD: 0000-0001-8708-0194

© Joel W. LeBlanc 2019
All Rights Reserved

To all my teachers

TABLE OF CONTENTS

DEDICATION	ii
LIST OF FIGURES	vi
LIST OF APPENDICES	xiv
ABSTRACT	xv

CHAPTER

I. Introduction	1
1.1 Motivation	1
1.2 Overview of Incoherent Imaging	2
1.3 Main Challenges of System Identification	3
1.4 Contributions	7
1.5 Limitations	8
1.6 Document Overview	9

II. Problem Setup	11
2.1 Introduction	11
2.2 Physical Processes	11
2.2.1 Object Model	12
2.2.2 Imaging Geometry	12
2.2.3 Optical System	15
2.2.4 Detection	18
2.3 Previous Approaches	22
2.4 Forward Model Components	23

2.4.1	Object Model	24
2.4.2	Application to Image Recovery and Guard Banding . . .	25
2.4.3	Warp Operator (\mathbf{W})	29
2.4.4	Blur Operator (\mathbf{B})	32
2.4.5	Sampling Operator (\mathbf{S})	39
2.5	Statistical Inverse-Problem Approach	47
2.6	Non-Linear Optimization	49
2.6.1	Quasi-Newton Direction Updates	50
2.6.2	Preconditioning and Transformation Invariance	52
2.6.3	Initial Step	55
III. System Identification with Known Objects		58
3.1	Introduction	58
3.2	Forward Imaging Model	61
3.3	Maximum Likelihood Estimation	62
3.4	Cramér-Rao Bound	64
3.5	Bias Sensitivity	70
3.6	Identifiability	74
3.7	Calibration Target Construction	75
3.8	Simulation Experiments	79
3.9	Conclusions & Ongoing Work	80
IV. Point Spread Function Identification Under Large Aberrations		83
4.1	Introduction	83
4.2	Testing for a Local Optimum of the Likelihood	85
4.2.1	A Simple Motivating Example	87
4.2.2	1-sided Biernacki Test	88
4.2.3	Reparameterized Embedding	94
4.3	Identifying Alternate Candidate Solutions	102
4.4	Simulation Experiments	105
4.4.1	Computational Comparison	110
4.5	Hardware Experiments	118
4.5.1	Literature Review: Depth from Defocus	118
4.5.2	Estimating Depth from Defocus	121

4.5.3	Restoration	126
V.	Conclusions and Ongoing Work	128
5.1	Summary	128
5.2	Ongoing Work	130
5.2.1	Continuous Representation	130
5.2.2	Sufficient Representations For Imaging	132
5.2.3	Evaluating Line-Spread Functions	135
5.2.4	Joint Model	135
5.2.5	Experiment	140
5.3	Future Directions	143
APPENDICES	144
BIBLIOGRAPHY	192

LIST OF FIGURES

Figure

2.1	An overview of the basic sensing process involved in optical digital imaging.	12
2.2	Illustration of the key elements defining the projective relationship between the object and detection planes.	13
2.3	Diagram and relevant parameters needed to relate a physical camera to the generalized imaging model.	16
2.4	(a) 4 Blocks of a 2x2 Bayer pattern with each of the 4 data channels numbered and the corresponding object channels are indicated through the use of color, (b) Elements of channel 1, (c) Elements of channel 2 . . .	19
2.5	Key dimensions for determining the detector sample rate s_r , element sampling rate relative to Nyquist R_N , pixel pitch p_p , and element fill-factor f_f ($f_f \approx 50\%$ as shown)	20
2.6	The continuous object is represented though a regular grid of points defined over a toroid, with observations only occurring a smaller circumscribed subset.	24
2.7	Partitioning of \mathbf{x}	27
2.8	Example of padding using the proposed technique.	28
2.9	General in-painting example using the proposed approach.	28
2.10	Example reconstruction from 50% of the data using the proposed approach.	29
2.11	Example reconstruction from 25% of the data using the proposed approach.	29
2.12	Example reconstruction from 10% of the data using the proposed approach.	30
2.13	Optimal bicubic interpolating kernel (a) and its corresponding frequency response (b).	31
2.14	Action of a 45° rotation operator and the associated derivative	32
2.15	First 10 Zernike basis functions as defined by (2.39) along with their common names	34
2.16	Apertures for a narrow-band RGB system sampled at λ_{ref} in the exit pupil, and the associated 2X oversampled narrow-band PSFs.	36

2.17	Comparison between the PSF of a 3-channel monochromatic system (a), and the PSF for the same system when collected 3 ideal single-channel detectors (b), or an ideal 4-channel Bayer detector (c).	38
2.18	Continuous periodic extension of a signal sampled over the toroid.	39
2.19	Illustration of the splitting of the energy at the line spectra corresponding to $w^{(n/2 + 1)}$ for a length n representation of $f_c(t)$ (n even) when upsampling to $m > n$	43
2.20	A real, discrete, signal sampled critically with an even number of samples along with both of its consistent implied continuous functions.	45
2.21	Cartoon illustrating the key blocks in the matrix representation of the upsampling operator \mathbf{U} in the Fourier domain. The block of zeros corresponding to the familiar zero-padding in the Fourier domain is explicitly called out.	46
2.22	Action of the sampling (aliasing) operator taking an image form (a) 256×256 to (b) 100×100	46
2.23	Inverse problem flow diagram showing the relationship between modeling components (green) and physical processes (gray)	47
2.24	Basic structure of quasi-Newton optimization methods	50
2.25	An example of the BFGS technique applied to a low-dimensional non-linear problem.	51
3.1	Failed attempt to measure a complex imaging system using an interferometer by retro-reflecting off the detector.	59
3.2	Diagram illustrating key elements of the physical imaging process. The corresponding block diagram relates physical processes to the mathematical operators, and corresponding unknowns, used to model them.	62
3.3	Connectivity of the bipartite graph relating FIM blocks to is underlying sufficient statistics	68
3.4	Graphical description of the computational stages used to compute the FIM in a memory efficient fashion. Each computational stage is represented as a bipartite graph between the FIM terms (blue), and the terms sufficient for its calculation (black). The transition of an items color to green indicates it is no longer needed for future calculations.	69
3.5	A plot of constant contours of the quadratic objective function and associated maximum bias-gradient constraint involved in computing the uniform CRB.	72

3.6	Uniform CRB for the defocus parameter of a system with an overall aberration strength of 0.25 WRMS based 6.25 times undersampled data.	73
3.7	Sierpinski pre-carpet of order 4	76
3.8	Standard tri-bar target	78
3.9	RMS wavefront-error bound vs noise	78
3.10	The true object (a) is shown alongside a series of data frames generated with all parameters fixed other than the SNR. The read-noise standard deviation corresponding to (b)-(d) were 0.68, 1.41, and 7.15 counts respectively.	80
3.11	Total wavefront estimation error (a) and bias (b) shown as a function of SNR for initial estimates within 0.01 waves RMS of the true 0.25 wave RMS perturbation.	81
3.12	Example data-frame collected with iPad 2	82
3.13	iPad2 wavefront estimates	82
4.1	A local minima of the log-likelihood that corresponds to a wavefront solution 0.19 WRMS from the true solution despite its associated PSF having a similar dispersion structure. The PSFs are shown on a grid upsampled 10X relative to the Nyquist-rate of the detector.	84
4.2	Diagram illustrating how a test for suboptimal solutions can be used in conjunction with an approach for identifying alternative candidates to perform global optimization.	85
4.3	(a) The true signal achieving global minimum (blue solid) and signal achieving a sub-optimal local minimum (red dotted) of the likelihood function (4.8) when noise variance σ^2 equals zero. (b) The measured data when the noise variance is $\sigma^2 = 1$	89
4.4	For a one-sided test to improve upon Biernacki's more general form, the set of stationary points comprising a problem's local minima would produce bootstrapped log-likelihood values like those shown in (a). The somewhat unintuitive situation of an incorrect model parameterization resulting in smaller log-likelihood values than the true parameterization is shown in (b). In both figures, the semi-transparent box emphasizes the fact that the curves are expectations of random functions, and that individual bootstrap tests will involve variation.	91

4.5 Detection performance of Biernacki’s two-sided test of a local minimum compared with the one-sided variant for the example problem described by (4.5)-(4.7). For this example problem the data distribution is in the location family and Thm. 1 applies leading to significantly improved performance for the one-sided test. 93

4.6 The reparameterized embedding leads to a relaxation that permits the ML estimator $\hat{\theta}$ to improve within the set $S(\hat{\theta})$. The test for global convergence exploits the distribution differences in the log-likelihood gap between $H_0 : \hat{\theta} = \theta_0$ and $H_1 : \hat{\theta} \neq \theta_0$ 95

4.7 Diagram illustrating how a locally, but not globally, optimal solution $\hat{\theta}$ can be identified by relaxing the parameter space from Θ to $\tilde{\Theta}$. Under H_1 , minimizing the negative log-likelihood under the relaxation often leads to a relaxed solution $\hat{\tilde{\theta}}$ with a substantially larger gap between its log-likelihood value and the bootstrap estimate (shown in red). 96

4.8 ROC curve illustrating the potential of reparameterized embeddings when applied to our simple example. The performance of Biernacki’s test is compared with the proposed approach when the naively chosen relaxation has 1 and 3 additional degrees of freedom respectively. 97

4.9 Algorithm for identifying relaxation. 99

4.10 Additional relaxation dimension suggested by the procedure given in Figure 4.9 100

4.11 Expected negative log-likelihood associated with the simple example given by (4.23)-(4.24) under the 1-dimensional relaxation provided by the algorithm described in Algorithm 4.9. This figure is analogous to the conceptual plot shown in Algorithm 4.7, and illustrates how a well-chosen relaxed embedding can be used to detect convergence to local minima. . . 100

4.12 ROC curve comparing Biernacki’s test to the reparameterized embedding approach when using a single relaxation dimension suggested by the procedure described in Algorithm 4.9 101

4.13 A 2D embedding of the points within the space of the first 12 Zernike modes that maximize the Strehl ratio for shells of a fixed RMS wave-front deviation. Shells up to 0.09 waves RMS are shown, with the points numbered to indicate correspondence between phase-screens in neighboring shells. Note that the asymmetries are not unusual when embedding high-dimensional surfaces into 2D. 104

4.14	A PSF corresponding to a random phase-aberration of 0.25 waves RMS is shown in (a). (b)–(f) show the first few perturbations of 0.2 waves RMS drawn from the set given by (4.39). Each of the PSF’s are displayed at 2X the Nyquist rate for the optical system.	104
4.15	The true object (a) is shown alongside a series of data frames generated with all parameters fixed other than the SNR. The read-noise standard deviation corresponding to (b)-(d) were 0.68, 1.41, and 7.15 counts respectively.	106
4.16	Wavefront estimation error is shown as a function of SNR for 0.25 wave perturbations where the starting point is a clear aperture, and the proposed global optimization procedure described in Section 4.2 is used.	107
4.17	(a) Shows the true PSF on a Nyquist sampled grid, and phase-screen associated with data simulated with and SNR of 18. (b) shows the local minima which failed to reject the null hypothesis at a power of 1E-4 due to the elevated noise levels. The actual p-value associated with the local solution was 0.019, and (c) shows the resulting PSF and phase-screen errors.	108
4.18	(a)-(c) shows a typical succession of PSF’s and phase-screen estimates associated with local minima found using the proposed global optimization procedure under high SNR conditions. (d) and (e) show the true solution and the associated errors respectively	109
4.19	Monte-Carlo study of optimization runtimes as a function of the number of aberration modes in the model.	111
4.20	A point is said to improve upon a search line if it improves upon the line’s best log-likelihood value using a wavefront solution within 0.04 WRMS of the associated wavefront description.	112
4.21	Final objective function value plotted against the creation time for each of the 54 unique lines identified by the SA algorithm. Wavefront solutions corresponding to the global optimum as well as the nearest two local minima are compared with the true solution.	113
4.22	Wavefront error plotted versus runtime over the first hour of the simulated annealing run.	114
4.23	Comparison between the 3 most likely solutions identified using simulated annealing and the true solution.	115
4.24	Wavefront error and objective function value plotted as a function of time. This plot is analogous to Figure 4.22, where the addition of the objective function value helps clarify the locations of the local minima.	116

4.25	Comparison between all unique solutions identified using the proposed optimization approach.	117
4.26	Experimental setup involving a Nikon D7000 imaging a calibration target a 3 ranges. The devices was auto-focused at the nearest range and then the lens configuration was held fixed for the remaining collections.	118
4.27	The data collected at 3.0, 3.2, and 3.3 meters is shown in (a)-(c) with the corresponding PSF estimates shown on the data grid in (d)-(f).	123
4.28	(a)-(c) The PSF estimates corresponding to the 3 target locations and 3 color channels are plotted on a common grid that is Nyquist sampled for the shortest wavelength of light (450 nm). (d)-(f) provides the corresponding 12 Zernike coefficients under the Noll ordering starting from defocus.	124
4.29	A zoom in of the red channel data at 3.2 m (a), the predicted data (b) shown on the same color axis, the residual (c), and a Q-Q plot for the residual after accounting for the intensity dependency.	125
4.30	Predicted RMS defocus vs target depth based on the proposed camera model for each of the 3 RGB color channels assuming best-focus positions of 3.024, 2.979, and 3.058 meters respectively and a pupil magnification of 0.985. The 9 defocus terms from the PSF estimates are plotted as squares.	125
4.31	Single frame, red-channel restoration results based on the defocused data (a) collected at 3.2 meters. The Richardson-Lucy technique is shown operating both blindly (b) and with the aid of the PSF information from the proposed blur estimation algorithm (c). The proposed ML solution of (4.41) for $\lambda = 300$ is shown in (f), and similar results using the strict ML estimate as well as regularized estimation with only the defocus term are shown in (d) and (e) respectively. RMSE's are shown for all restorations in the lower-left corner.	127
5.1	Geometry of an edge	131
5.2	Illustration of the projections used to solve the integer sampling sub-problem	134
5.3	Algorithm for finding data sizes for downsampling	135
5.4	The data used to estimate edge-target ranges based on relative fiducial defocus amounts is shown in (a). Subfigure (b) shows the observed experimental error along with the theoretical error had the wavefront estimation accuracy been 0.01 waves RMS.	141

5.5	Comparison of maximum theoretical ranging uncertainty for common consumer cameras relative to the Nikon D7000 assuming defocus estimation accuracy of 0.01 waves RMS.	142
A.1	The original test image (a), is shown alongside the mask used to best inpainting (b). To avoid unnecessarily penalizing methods that implicitly assume inpainting regions are surrounded by known values, errors were computed over both the whole image and the interior set indicated by the white dashed line.	147
A.2	Comparison between minimum bandwidth continuous function extension and various inpainting techniques. Errors are reported in terms of PSNR, measured in units of dB, over the interior portion of the image. The error for the entire image is also reported in parentheses. All of the images are shown on the same scale as the truth.	148
A.3	This figure illustrates the same experiment shown in Figure A.2 with the exception that the exterior pixels were assumed known. For this reason, the interior, and whole image, PSNR calculations match in every instance.	149
B.1	Overall structure showing group homomorphisms (solid lines), and set maps (dashed lines). The dotted line indicates the diagrams don't, in general, commute along compositions of solid lines (e.g. $(F_n \chi_{n,r}) \psi_n \neq F \chi_j$).	151
B.2	Linear transforms relating $\mathbb{C}(\mathbb{Z}_n)$, $\mathbb{C}(\mathbb{T})$, and $\mathbb{C}(\mathbb{Z})$	153
B.3	Relationship between homomorphisms and set maps for functions defined over the relevant groups. Solid lines indicate group homomorphisms, and dashed lines set maps. As before, the dotted line indicates that the diagrams don't commute along compositions of solid lines except for under the restrictions that $r = \psi_n(j)$ and $s = \psi_m(j)$	157
E.1	Diagram and relevant parameters needed to relate a physical camera to the generalized model	165
E.2	Rays tracing diagram corresponding to the near- and far-focus positions that yield a fixed Circle-of-Confusion (CoC) of diameter c	166
E.3	The relationship between Peak-to-Valley (PF) defocus and the Circle-of-Confusion (CoC) is given by the difference in sagitta between the in- and out- of focus converging waves.	167
F.1	PSF variation induced by 1×10^{-3} waves RMS in the exit pupil over 7 Zernike modes	170
F.2	Blur operator error as a function of exit-pupil diameter	171

F.3 Blur operator error as a function of exit-pupil diameter for shaded and binary aperture approximations 172

G.1 Area of intersection between (shown in blue) between a disc of radius r and a pixel centered about the point P . The problem is naturally parameterized by the boundary points A and B 173

G.2 Pupil diameter of 9 with boundary pixels (with respect to a circle of radius 4.5) shown in light red 174

LIST OF APPENDICES

Appendix

A.	Minimum Bandwidth Function Extension and Inpainting Comparison	145
B.	Discrete to Continuous Fourier Transform	150
C.	Subgroups of Perspective Transformations	158
D.	Warp Partials	161
E.	Geometric Optics for the Asymmetric Generalized Imaging Model	164
F.	Numerical PSF Evaluation	169
G.	Shaded Aperture Derivation	173
H.	OTF Partials	176
I.	Derivations for Known Objects	178
J.	Known Object CRB Terms	181
K.	Smoothing Regularizer	188

ABSTRACT

A statistical inverse-problem approach is presented for jointly estimating camera blur from aliased data of a known calibration target. Specifically, a parametric Maximum Likelihood (ML) PSF estimate is derived for characterizing a camera's optical imperfections through the use of a calibration target in an otherwise loosely controlled environment. The unknown parameters are jointly estimated from data described by a physical forward-imaging model, and this inverse-problem approach allows one to accommodate all of the available sources of information jointly. These sources include knowledge of the forward imaging process, the types and sources of statistical uncertainty, available prior information, and the data itself. The forward model describes a broad class of imaging systems based on a parameterization with a direct mapping between its parameters and physical imaging phenomena. The imaging perspective, ambient light-levels, target-reflectance, detector gain and offset, quantum-efficiency, and read-noise levels are all treated as nuisance parameters. The Cramér-Rao Bound (CRB) is derived under this joint model, and simulations demonstrate that the proposed estimator achieves near-optimal MSE performance. Finally, the proposed method is applied to experimental data to validate both the fidelity of the forward-models, as well as to establish the utility of the resulting ML estimates for both system identification and subsequent image restoration.

CHAPTER I

Introduction

1.1 Motivation

The prevalence of passive optical imaging sensors is higher today than at any other time in human history. The UN estimated that the world population at the end of 2015 was slightly greater than 7.2 billion people [191], and distributed amongst this population was an even larger number of active mobile phones [85]; most of which included a camera. This statistic does not include dedicated consumer imaging devices, commercial cameras, and so forth. The data collected by these systems is not the best possible representation of the scene. Inherent physical limitations, as well as manufacturing imperfections, both limit the ultimate image fidelity. By understanding the physical relationships between the environment, the imaging system, and the collected data, one can make improved inferences about the sensor, the scene, or both simultaneously. This work is about a systematic approach for making these types of inferences, as well as understanding the limitations imposed on them by uncertainties from the environment and imaging system.

The physical relationship between the scene and the collected data is represented through a forward imaging model. This model describes the imaging process in terms of physical optics; starting from light leaving the scene and terminating in data collected from the sensor array. Image reconstruction from the collected data is said to result from an inverse-problem because one is attempting to invert this forward imaging process. All inference problems from data collected on physical systems share this structure, and if one is explicit about the forward mode, assumptions exploited by the inverse-problem approach are clarified. Furthermore, expressions for the stochastic data allow the application of standard statistical estimation techniques for solving the resulting inverse-problems.

Another motivation for using a statistical inverse-problem approach is the use of information-theoretic tools to fundamentally addressing questions of solubility. By expressing the quantities of interest as an estimation problem over a stochastic data model, one

may use information-theoretic techniques to establish lower bounds on estimation accuracy. Such bounds apply to the data model itself, and thus are independent of the methods used to construct estimators. The utility of these bounds is twofold: First, a problem can be declared insoluble if the level of error ensured by the bound is unacceptable from an engineering perspective. Second, the accuracy of specific estimators can be compared to lower-bounds to quantify the potential for future improvements.

Finally, and perhaps most importantly, this work lays a groundwork for improved image restoration approaches. Model-based image reconstruction has been a topic of interest since at least the 1950's, and interest was renewed with the widespread availability of digital cameras in the 1990's [49, 50, 55–57, 89, 90, 108, 124, 139, 140, 148, 168, 185, 193, 194, 210]. A review of the history, as well as current state-of-the-art, in image restoration can be found in [127, 133]. The fidelity of the underlying imaging model fundamentally limits the effectiveness of image restoration techniques, yet the literature has disproportionately studied scene reconstruction under assumed imaging models as opposed to sensor model-fidelity under assumed scene models. This work emphasizes the latter through the construction of a flexible, and physically motivated, forward imaging model. One's ability to infer the unknown parameters within this model is studied through the joint Maximum Likelihood (ML) estimator of the parameters assuming a known calibration target. The Cramér-Rao Bound (CRB) is derived, and simulations demonstrate that the proposed estimator can achieve near-optimal MSE performance. The proposed method is applied to experimental data to validate both the fidelity of the forward-models, as well as the applicability of the resulting ML estimates for both system identification and subsequent image restoration applications.

1.2 Overview of Incoherent Imaging

In passive image formation, a continuous scene is typically illuminated by an incoherent light source, and the reflected electromagnetic radiation is sensed through a coherent interference process first described by Ernst Abbe in 1873 [1]. When the object being imaged is a sufficient distance from the physical collection aperture, the impinging radiation may be treated as a planar wavefront undergoing Fraunhofer diffraction at the sensor's aperture [77]. Optical elements are used to coherently interfere the electromagnetic field to form an image, and that modern sensors measure this image using a semi-conducting medium. Imperfections in these optics result in wavefront distortions that manifest themselves as a blur at the detection plane. The spatial sampling rate of the wavefront depends on the physical configuration of the detector, but this rate is often beneath the Nyquist sampling rate

given the optics, and when this is the case, aliasing occurs. Furthermore, the photo-electric effect and dark current induced by biasing of the semi-conducting material causes the signal to contain a filtered Poisson component [64, 172], and thus the measurement apparatus is inherently uncertain. Finally, thermal noise from the read-out electronics adds a filtered Gaussian component to the signal prior to it being recorded.

In addition to being intrinsically uncertain, incoherent imaging is wavelength-dependent. Diffraction at the sensing aperture, optical lenses, and the detector itself all exhibit chromatic effects. Given a common optical path, aliasing at the detector is a wavelength-dependent phenomenon. It is also common practice to use chromatically selective filters at the detection plane to simultaneously collect multi-channel color images. These filters are usually spatially distributed, and thus the wavelength-dependent channels do not coincidentally sample the scene. A consequence of this chromatic dependency, high-fidelity forward models must accommodate these effects.

While the majority of modern digital imaging systems undergo these basic steps, the relationship between the sensor's physical configuration and the mathematical description of the operators that describe the imaging process is typically unknown. An accurate description of a system's incoherent Point-Spread Function (PSF) is fundamental to image reconstruction, especially reconstruction from aliased data [49, 127]. A PSF model based on a description of the optical system's wavefront aberrations is *useful* in image formation because it parsimoniously describes the space of PSF's likely to be encountered [136, 164], but rapidly becomes *essential* when datasets are fundamentally a function of changes in these aberrations. Intentionally introducing such changes is the basis of phase-diversity imaging [75, 97, 142, 186], and unintentional changes serve as the basis for depth-from-defocus techniques, certain 3D shape reconstruction algorithms, and even sophisticated image restoration approaches [58, 103, 120]. These application areas all motivate the characterization of imaging systems through a physical wavefront description of their blur.

1.3 Main Challenges of System Identification

The main challenges of system identification all revolve around the inherent complexity of incoherent imaging. One must construct models that accurately reflect the physical phenomena of interest while simultaneously maintaining tractability. A physically accurate forward model that is too complex to be part of an inverse-problem formulation is not useful in making inferences from image data. Conversely, overly simplistic models that do not capture key physical effects may not adequately describe the physical devices they aim to explain. These models must also include the degrees-of-freedom inherent in their

physical counterparts. Consider, for example, a camera model that does not fundamentally accommodate changes in focus. A camera description under such a model would have limited utility because changes in the device that altered this parameter would render the description invalid. Modern cameras permit changes in their aperture, focus, exposure time, and gain, therefore accommodating these degrees of freedom is of particular import. Finally, the implicitly defined estimators that provide the quantities of interest must be solvable. Even if solutions are known to exist, one must also identify a practical means at finding them. This thesis makes contributions toward each of these key challenges.

The most direct means of characterizing an imaging system is by explicitly measuring its transfer function, but direct wavefront sensing approaches require modification, or partial disassembly, of the sensor [73]. In many situations, this is not possible, and even when disassembly can be performed, it is difficult to know if the reassembly process has altered the state of the sensor. This desire for nondestructive testing procedures has led to an interest in inverse-problem formulations using canonical sensor outputs.

Methods that estimate blur directly from a sensor’s measurements can be broadly categorized by how much information is known about the scene being imaged, and what type of data is used in the subsequent inverse estimation problem. When an inverse-problem approach makes very weak assumptions about the scene, it is referred to as a “blind” approach (cf. [25, 62, 87, 99, 101, 102, 167]). By contrast, “non-blind” or “known-object” methods [8, 21, 36, 44–46, 83, 87, 101, 115, 141, 144, 155, 169, 178, 213] leverage more stringent assumptions on prior scene knowledge. Most approaches are neither entirely blind nor wholly non-blind, but this categorization is both useful and relatively easy to apply in practice. Furthermore, these two categories of blur estimation procedures typically target different use cases. When one only anticipates intrinsic camera blur, blur fundamentally due to properties or imperfections of the imaging device itself, non-blind estimation techniques provide the benefit of improved accuracy brought about by the additional knowledge of the calibration object. On the other hand, when one anticipates extrinsic blur, blur caused by relative motion between the camera and the scene, this additional blur is usually better addressed by blind approaches because prior knowledge of the scene is often unavailable when the extrinsic blurring occurs. While the two classes of techniques are often studied independently, they are not mutually exclusive, and one can in principle use them in conjunction with each other.

Blind methods inherently require either more data than non-blind techniques or depend on strong a priori assumptions. Most single-frame methods attempt to exploit the existence of sharp edges in the latent unblurred scene, as these often exist due to the projection of multiple 3D objects into the 2D imaging plane. A common implicit approach is to leverage

assumptions on the object gradients. For example, Fergus [62] uses a variational Bayes approach that assumes a Gaussian mixture distribution, Chan exploits a Laplacian prior [10, 30], and Kotera uses heavy-tail priors [106]. The sparsity of the occurrence of such edges can also be exploited in the gradient space [167]. One can identify edges using convolutional edge detectors [101], Hough transform techniques [199], or patch priors [183]. An alternative approach is to forgo explicit latent scene assumptions altogether for those learned indirectly from a training set (cf. [116]). Finally, one should note the significant overlap between this body of work and the blind deconvolution literature [26, 32]. A key feature that distinguishes the references above from the blind deconvolution literature is the author's preeminent treatment of the blur as a quantity of interest. Comparing the relative merits of blind blur identification techniques is significantly hampered by the difficulty of verifying the validity of their underlying assumptions a posteriori, and without strong assumptions, blind deconvolution is fundamentally ill-posed. For these reasons, this work emphasizes intrinsic camera blur estimation from known objects.

Known-object approaches exploit the presence of a calibration target to aid in blur identification, yet despite this additional information, accurate system characterization remains a challenging problem. One way of classifying previous approaches is through the restrictions they place on the calibration object. Imaging a point source results in a direct measurement of the PSF, and Grisan [83] has achieved very accurate results by replacing the traditional wavefront sensor inside an adaptive optics control loop with a direct signal processing approach. The inherent challenge in using a point source, however, is that the informative region within the collection array becomes small as the PSF approaches the diffraction limit. This results in a lower signal-to-noise ratio (SNR) than comparable extended object approaches. Imaging an ideal edge, commonly referred to as a knife-edge object because of the use of razors in lab experiments, provides an extended object whose measured response asymptotically becomes an integration of the PSF along the edge direction as both the edge and field-of-view become large. Reichenbach [155] and Claxton [36] both use a knife-edge object within a continuous-to-discrete imaging model to characterize digital imaging devices. Each of the authors notes the additional complications brought about by sensor aliasing effects, and attempt to mitigate this issue through registering and averaging multiple measurements. Exploiting larger patches within the collected data helps to overcome SNR challenges, and numerous authors have considered pseudo-random regions with a known power spectral. Gaussian [144] and uniform distributed targets [21] are natural choices; however, once again, aliasing complicates the analysis. Backman [8] generalized these ideas to allow targets with zeros in the power spectrum. Discretely generated random targets naturally exhibit these gaps, and Backman suggests that they can

be carefully designed to mitigate difficulties associated with aliasing. Delbracio [46] helped justify the use of pseudo-random targets by showing that Bernoulli distributed objects permit near-optimal PSF recover, and Mosleh [131] was later able to demonstrate a reliable PSF recovery framework based on these ideas. Mosleh’s approach, like other contemporary work (cf. [46, 87, 101]) is also able to accommodate arbitrary calibration objects. The early pioneering work by Pavlovic [141] is also notable not only because it incorporates arbitrary calibration targets, but also because it was one of the earliest approaches to use a continuous-to-discrete forward imaging model. Previous work has identified many of the key challenges in system identification, yet satisfactory solutions applicable to modern hardware remains a significant challenge.

The trade-off between model fidelity and computational tractability for the fundamentally continuous-to-discrete forward imaging process remains a substantial challenge in system identification. Early attempts to quantify an optical system’s ability to form images emphasized metrics derived relative to the aforementioned known objects under a continuous imaging model [27, 163]. While the author’s were undoubtedly aware of the impact finite-sized photocells had on final resolving capability of the sensor, these issues were treated as secondary effects. Subsequent signal-processing approaches emphasized image restoration over optical physics [109], leading to discrete-to-discrete models between the scene and collected data. Reichenbach and Pavlovic [141, 155] focused on the importance of continuous-to-discrete models, which naturally leads to questions surrounding the impact of aliasing at the detector. Nevertheless, tractable blur modeling and estimation remained the primary concern, and aliasing was either omitted [141, 144] or addressed indirectly [8, 155] within forward models. The push to achieve practical system characterization techniques lead researchers toward computationally simpler discrete-to-discrete imaging models [45, 101, 131]. Under these models performance bounds were derived demonstrating the fundamental tractability of the inverse-problem [46, 87, 190], but wavefront blur descriptions, diffraction effects, and detector effects remain largely unaddressed.

The use of simplified forward models and image priors has led to reliable blur estimation strategies (c.f. [45, 101, 131]), however, these models do not maintain a direct correspondence between the physical system configuration and forward-modeling parameters. For example, most imaging devices have an adjustable focus setting, yet the aforementioned blur estimation strategies do not use imaging models that succinctly capture this effect. As a consequence, a change in focus leads to an altogether different system description despite the tight correspondence between the sensor before and after the focus change. A corresponding argument can be made for imaging aperture (pupil diameter). Diffraction at the aperture induces a blur that is distinct from defocus, yet without modeling this effect

directly, it is difficult to distinguish the source of imaging system non-idealities. Such modeling limitations restrict the utility of simplified models when such perturbations are expected, and these concerns motivate maintaining a direct correspondence between the imaging physics and inference model when possible.

1.4 Contributions

A statistical inverse-problem approach is presented for jointly estimating camera blur from aliased data of a known calibration target. Specifically, a parametric Maximum Likelihood (ML) PSF estimate is derived for characterizing a camera’s optical imperfections through the use of a calibration target in an otherwise loosely controlled environment. The unknown parameters are jointly estimated from data described by a physical forward-imaging model, and this inverse-problem approach allows one to accommodate all of the available sources of information jointly. These sources include knowledge of the forward imaging process, the types and sources of statistical uncertainty, available prior information, and the data itself. The forward model describes a broad class of imaging systems based on a parameterization with a direct mapping between its parameters and physical imaging phenomena. The imaging perspective, ambient light-levels, target-reflectance, detector gain and offset, quantum-efficiency, and read-noise levels are all treated as nuisance parameters. The Cramér-Rao Bound (CRB) is derived under this joint model, and simulations demonstrate that the proposed estimator achieves near-optimal MSE performance. Finally, the proposed method is applied to experimental data to validate both the fidelity of the forward-models, as well as to establish the utility of the resulting ML estimates for both system identification and subsequent image restoration.

This work extends existing known-object system identification techniques through its explicit incorporation of physical modeling within the continuous-to-discrete forward model. By contrast, Delbracio [45] also obtains accurate sub-pixel PSF estimates with performance bounds, but takes a model-free estimation approach. First, this work uses a phase-screen description of the optics to parameterize the PSF under a Fraunhofer diffraction assumption. Many previous approaches have not explicitly modeled diffraction effects [8, 21, 36, 45, 87, 101, 115, 131, 141, 144, 213], or modeled it only in conjunction with defocus [178]. Second, the proposed approach explicitly models aliasing. Aliasing is common amongst commercial devices, and previous approaches have either ignored this effect [8, 21, 83, 87, 115, 141, 144, 178, 213], or address it in an incomplete way [36, 45, 101]. Notably, Joshi [101] and Delbracio [45] both explicitly model aliasing, but only accommodate integer downsampling factors. Third, the proposed technique explicitly models charge-

coupled detector effects. Very few previous works (cf. [178]) have accommodated the intensity-variant uncertainty that is inherent when measuring electromagnetic fields using semi-conductors. This work accounts for this source of uncertainty. Taken together, the resulting high-fidelity forward model accurately reflects the physics of imaging with modern digital devices and accommodates the main sources of uncertainty in the resulting data while maintaining computational tractability. Furthermore, the physically motivated forward model is inherently modular, and thus can be adapted to study numerous questions that are not directly addressed by this document. To summarize, the main contributions of this thesis are:

- A flexible, physically-motivated, high-fidelity, forward-imaging model is developed for digital imaging.
- An approach for joint PSF and camera-pose estimation in the presence of aliasing is presented.
- CRB bounds for the joint-estimation approach are derived that are applicable to any calibration object.
- A bootstrap test for global convergence is presented, enabling joint PSF estimation in the presence of large wavefront aberrations.
- The forward model and inverse-problem approach are validated through laboratory experiments, and their utility is demonstrated in the context of image restoration using commercial devices.

1.5 Limitations

This dissertation aims to develop and explore a simple model capable of addressing the aforementioned physical effects. Potential limitations of the proposed model are as follows. First, the imaging system’s blur is assumed to be spatially invariant, and like many previous works [45, 104, 131] spatially variant blurring is only addressed by applying the proposed approach locally throughout the field of view. Second, lens distortions such as “barrel” or “pincushion” effects are not modeled. Third, multi-channel (color) systems are addressed, but, these channels are assumed to be well approximated by a narrow-band approximation at their center wavelengths. Fourth, each sensor element is assumed to be identical. In principle, one could jointly estimate non-uniformity effects due to inter-element variations, but such a modification would undoubtedly require multi-frame data. This work applies to multi-frame data, but an emphasis is placed on single-channel results to ease the exposition

of an already complex topic. These limitations can all be addressed within the proposed framework, but have been left for future work.

Similarly, broad-band or panchromatic channels are conceptually a straightforward modification to the proposed model, but the details of such an extension would undoubtedly uncover new challenges. Finally, many systems also do not provide raw-data access, and when this is the case, any compression algorithms should be directly modeled in the forward operator. The impact of compression can be large relative to many of the effects considered in this work, but this challenge remains unaddressed. The proposed inverse-problem approach fundamentally depends on a sequence of operators that are non-linear in their parameters, and in principle, each of these limitations can be addressed using the same approach taken throughout the dissertation. In practice, these would all be substantial extensions to this work. These issues highlight the vast amount of effort needed to bridge the gap between the contributions made by this dissertation and the tools currently available.

1.6 Document Overview

The emphasis of this dissertation is system identification from a known calibration target. Given a calibration target, described through its relative reflectance, the problem of joint camera-pose and blur is explored through the development of the necessary components of the aforementioned inverse-problem framework. Chapter II develops a high-fidelity, physically-motivated, forward imaging model by describing image acquisition through a series of constituent physical processes. This physical imaging model is compared and contrasted with previous approaches from the literature. The constituent mathematical operators necessary to implement the forward imaging model are then developed, and the problem of system identification is formulated in terms of these models. Finally, the general philosophical approach leading to this mathematical formulation is discussed, and an overview of the optimization tools used in its solution described.

In Chapter III the Maximum Likelihood (ML) estimator and Cramér-Rao Bound (CRB) are derived for the resulting joint inverse-problem. The implicit estimator is stated in terms of the solution to a non-linear, non-convex optimization problem. The identifiability of the resulting estimator is addressed for small deviations of the unknowns, and bias-variance trade-offs are explored using the uniform CRB. Despite finite sample sizes, the solutions are shown to be essentially unbiased given small deviations of the parameters, and these conclusions are verified through simulation.

Chapter IV revisits the question of identifiability in the context of large parameter deviations and the impact of low SNR. The physics of the imaging process ensure that large

parameter deviations can occur under reasonable imaging conditions. This implies that these questions must be answered if one is to develop a practical inverse-problem approach. The issue of spurious estimates converging to inconsistent roots of the log-likelihood is explicitly addressed through the development of a new bootstrap test for global convergence. This test is used in conjunction with a heuristic for identifying candidate wavefront solutions based on minimizing the L1-norm of the induced deviation of the Point-Spread Function (PSF). Taken together, this results in a practical approach for high-fidelity system identification in the presence of aliasing, and this technique is verified through simulation. Finally, the efficacy of the proposed approach is demonstrated through experiments performed using a commercial camera, and estimation performance is compared with the CRB.

Chapter V discusses ongoing and future work related to the general inverse-problem framework. At its core, this dissertation is based on exploiting knowledge of the physical imaging process within an inverse-problem framework to develop substantially improved techniques for imaging. Chapter V illustrates how one can extend these ideas to build estimation methods that exploit fiducials rather than wholly known objects. A relaxation of the known-object model is used to derive a new inverse-problem approach applicable to edges within a scene, and the resulting estimator is demonstrated through a hardware experiment. In doing so, this chapter takes one additional step toward blind methods while demonstrating the flexibility of the proposed inverse-problem approach.

Publications associated with this work are:

- Joel W. LeBlanc, Brian J. Thelen, and Alfred O. Hero. Testing for local minima of the likelihood using reparameterized embeddings. *Mathematical Imaging and Vision (In Review)*, 2019
- Joel W. LeBlanc, Brian J. Thelen, and Alfred O. Hero. Joint camera blur and pose estimation from aliased data. *J. Opt. Soc. Am. A*, 35(4):639–651, Apr 2018
- Brian J. Thelen, John R. Valenzuela, and Joel W. LeBlanc. Theoretical performance assessment and empirical analysis of super-resolution under unknown affine sensor motion. *J. Opt. Soc. Am. A*, 33(4):519–526, Apr 2016

CHAPTER II

Problem Setup

2.1 Introduction

This chapter develops a physics-based forward model that is applicable to commercial incoherent imaging devices, and describes the methodology that will be used to invert the resulting equations. The desire to use the models to make inferences about the world as measured through commercially available sensors precludes many of the simplifying assumptions that have historically been made for mathematical convenience. The models proposed in this chapter are more complex than traditional approaches, but also directly tied to the underlying physical processes. This proximity leads to better traceability of modeling assumptions, a clearer description of the impact of physical unknowns, and ultimately results in a more flexible framework. For clarity, the model development is broken into 3 broad sections. Section 2.2 discusses the model components in relationship to the physical imaging processes; the forward model. Section 2.3 compares previous techniques to this physical motivated approach, and Section 2.4 discusses the individual modeling components used to implement the forward model. Given such a framework, inference problems regarding modeled physical parameters can be explicitly stated in terms of the solution to a statistical estimation problem.

2.2 Physical Processes

All incoherent imaging devices may be modeled in terms of the physical processes that constitute image acquisition. A scene reflects and radiates light toward an imaging device that records a function of the impinging electric field as observed through a finite physical aperture. Behind the aperture of a typical imaging device, a series of optical elements attempt to form an image of the apparent emitted intensity onto a detector. Throughout this work, the measured electric fields will be assumed to have been emitted from an object plane

at some unknown distance from the sensor, and the intervening medium between the object plane and the sensor may attenuate the electric field, but otherwise has no effect. After passing through the optical device, the object plane is re-imaged onto a semi-conducting detector that measures the time-averaged intensity of the field over a small series of small, spatially-extended, sensors. Figure 2.1 illustrates this basic configuration.




Figure 2.1: An overview of the basic sensing process involved in optical digital imaging.

2.2.1 Object Model

The object, or scene, is modeled as a 2D plane relative to the camera's optical axis, with each point within the plane possibly emitting radiation. For the purposes of this work, we will not distinguish between reflected and emitted light. The sensor is assumed to only measure the time- and spatially-averaged intensity of the field and thus the objects description needn't include polarization information. Many sensors, however, do collect light from multiple wavelengths, so the object will be described through the spatial- and wavelength-dependent distribution of energy over the object plane $f(x, y, \lambda)$. Object modeling assumptions of this nature are ubiquitous within the physical imaging community (c.f. [58, 69, 77, 145, 184]).

2.2.2 Imaging Geometry

The imaging geometry is described through the relationship between the object and detection planes, assuming an ideal imaging device (pin-hole camera). An ideal imaging device is defined as a device that measures a projection of the object plane onto the detector plane. This relationship is mathematically described as the relative orientation between the

object plane and its ideal projection onto the detection plane. Figure 2.2 illustrates this imaging geometry, where the origin of the system is the camera center, and the optical axis is defined as the line orthogonal to the detector plane that passes through the origin. This idealized projective camera model will serve as the starting point for describing less ideal systems. The excellent texts by Hartley and Zisserman [91], Prince [145], Szeliski [184], and Shapiro [165], as well as the references therein, provide a good overview of both the mathematics behind this model as well as its application within common computer vision problems.

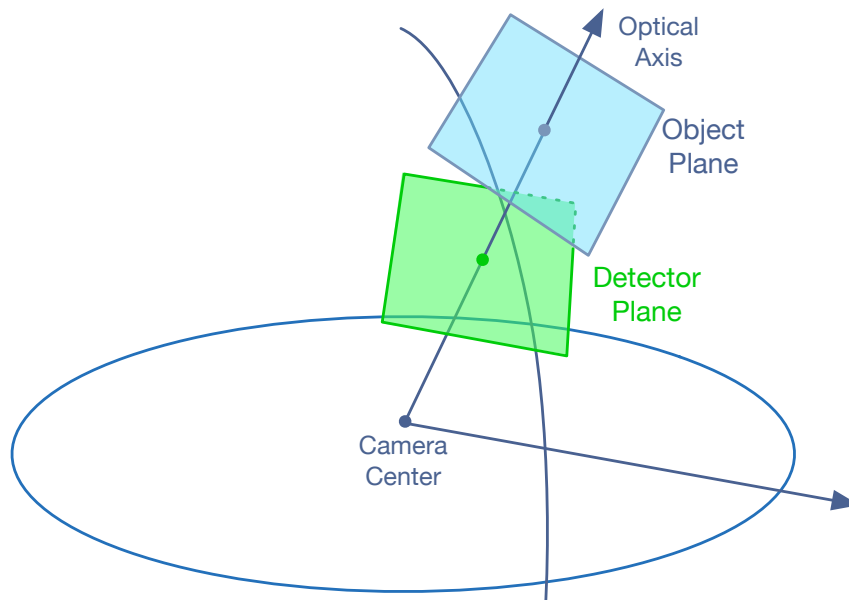


Figure 2.2: Illustration of the key elements defining the projective relationship between the object and detection planes.

The imaging geometry is characterized by a description of the isomorphism formed between the object and detector planes, where points between the planes coincide if they are collinear with the origin. The description of the isomorphism is characterized by a 3×3 matrix acting on a homogeneous representation of the points within the reference planes. That this representation is unique, follows from the fundamental theorem of projective geometry. In this geometry, points in homogeneous space represent lines in \mathbb{R}^3 passing through the origin via the equivalence relation

$$\mathbf{x} \stackrel{\text{def}}{=} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} : \mathbf{x} \equiv k\mathbf{x} \quad \forall k \in \mathcal{K} \quad (2.1)$$

where \mathcal{K} is a division ring ($\mathcal{R} \setminus \{0\}$ for the rest of this discussion). Thus we have an imbed-

ding $\mathbb{R}^2 \mapsto \{\mathbb{R}^3 \setminus \mathbf{0}\} \stackrel{\text{def}}{=} \mathbb{P}^2$; the real projective plane. This embedding naturally separates points into two classes through the incidence relation. Sets of points satisfying the relation are collinear, and thus any of these points is sufficient to define a line. As one would expect, two distinct points define a line, and any two distinct lines intersect at exactly one point. Furthermore, parallel lines intersect at ideal points (points at infinity). In this way, the perspective plane induces stronger symmetries than the Euclidean plane by mapping points onto lines, and lines onto planes, through the center.

$$\left\{ \begin{bmatrix} x_1 & x_2 \end{bmatrix}^T \right\} \mapsto \left\{ \begin{bmatrix} x_1 & x_2 & 1 \end{bmatrix}^T \right\} \quad \text{“Ordinary points”} \quad (2.2)$$

$$\left\{ \begin{bmatrix} x_1 & x_2 & 0 \end{bmatrix}^T \right\} \quad \text{“Ideal point”} \quad (2.3)$$

The ideal points are the complement of the image of the embedding, and represent the set of intersections of parallel lines. These points form a line themselves that is commonly referred to as the horizon.

The group element \mathbf{x} describes the line $x_1x + x_2y + x_3 = 0$ in \mathbb{R}^3 , with the point \mathbf{p} lying on the line iff $\mathbf{x}^T \mathbf{p} = 0$. Two such lines \mathbf{x} and \mathbf{y} intersect at the point $\mathbf{x} \times \mathbf{y}$, and similarly the intersection of \mathbf{p} and \mathbf{q} is given by $\mathbf{p} \times \mathbf{q}$. This duality between points and lines is a fundamental property of the perspective plane (a special case of a cyclic group of order 2), and will greatly simplify our calculations.

Perspective transforms are defined as the set of transforms mapping lines onto lines in \mathbb{P}^2 , thus the set of all linear transforms in the space.

$$\mathbf{y} = \mathbf{V}\mathbf{x}: \quad \mathbf{V} \in \mathbb{R}^3, |\mathbf{V}| \neq 0 \quad (2.4)$$

The equivalence relation implies that $\{\mathbf{V}\}$ is an 8-dimensional space, and w.l.o.g. it will be assumed that $\mathbf{V}_{33} = 1$. Thus, given a vector describing the 8-degrees of freedom in \mathbf{V} , an ideal imaging system is the unique projection of the function $f(x, y, \lambda)$ into the detector plane. A set of subgroups to this class of transformation that are commonly used in imaging are discussed in the appendices.

Next consider two important transformations; transformations of the global coordinate system, and transformations of a projective mapping with respect to a change in the detector

plane ordinate system. A transformation of the global ordinate system is given by

$$\mathbf{y} = \underbrace{\mathbf{V}\mathbf{L}^{-1}}_{\text{Transform}} \underbrace{\mathbf{L}\mathbf{x}}_{\text{Ordinates}} \quad (2.5)$$

where \mathbf{L} is any invertible transform. More commonly, however, one wishes to relabel the ordinates within the image (or object) plane. This is useful because it is often numerically inconvenient to treat the origin as lying in the center of an image despite the fact that image samples are typically indexed by row and column. Such transformations are also necessary for finding separable (“chipped”) representations of global transformations. This second situation arises when one wish to process a sub-region of an image under a common global description of the imaging geometry.

A simple way to think about coordinate transforms is as a through a similarity transformation. Suppose one has a relationship between two planes in euclidean space described by $\mathbf{y} = \mathbf{V}\mathbf{x}$, and it is desirable to relabel one (or both) spaces via $\tilde{\mathbf{y}} = \mathbf{L}_y\mathbf{y}$, and $\tilde{\mathbf{x}} = \mathbf{L}_x\mathbf{x}$, then the new transformation is given by

$$\mathbf{L}_y\mathbf{y} = \tilde{\mathbf{V}}\mathbf{L}_x\mathbf{x} \quad \text{“Transformed space”} \quad (2.6)$$

$$\therefore \mathbf{y} = \mathbf{L}_y^{-1}\tilde{\mathbf{V}}\mathbf{L}_x\mathbf{x} \implies \mathbf{V} = \mathbf{L}_y^{-1}\tilde{\mathbf{V}}\mathbf{L}_x \quad (2.7)$$

In the context of imaging with an ideal camera, one has the additional constraint that $\mathbf{L}_y = \mathbf{L}_x$ because the data and the object are coupled via physical space. This implies that a relabeling of either the data or the object via a transformation \mathbf{L} results in

$$\mathbf{y} = \underbrace{\mathbf{L}^{-1}\tilde{\mathbf{V}}\mathbf{L}}_{\text{Transform}} \underbrace{\mathbf{x}}_{\text{Ordinates}}. \quad (2.8)$$

2.2.3 Optical System

The optical system is composed of a series of lenses and/or mirrors that attempt to re-image the object plane onto a physical detector. This system is constructed to approximate the aforementioned ideal camera while simultaneously using finite apertures to ensure sufficient light collection. The modeling fidelity needed to compliment the type of wavefront-estimation techniques proposed in this work cannot be achieved though typical blur models

applied to the geometric description from a pin-hole camera (c.f. [113]). At the other extreme is physical ray-tracing approaches that lead to computationally challenging expressions for the optical wavefront in terms of a very precise material description of the imaging system [19]. While exquisitely precise, such approaches require a physical description of the imaging system that is rarely available. This work will apply a blur description resulting from the generalized imaging model proposed by Goodman [77], coupled with a corresponding geometric optics model. This approach is based on a Fraunhofer diffraction approximation, explicitly accounts for the finite imaging aperture, but only models the aggregate effects of the optical elements through a description of their action on a reference sphere passing through an entrance and exit-pupil. This 2-port model results in a space-invariant description of the Point-Spread Function (PSF) in terms of the deviations of the optical wavefront from a reference sphere measured in the exit-pupil. This blur model is then applied to a geometric description of the camera that is also based on an entrance- and exit-pupil description, and which allows for these two apertures to be of differing sizes. This is necessary if one is interested in modeling non-telecentric devices.

Figure E.1 illustrates the key elements of this generalized imaging model, as well as the relationships to the relevant camera and lens parameters. The aggregate impact of the optical

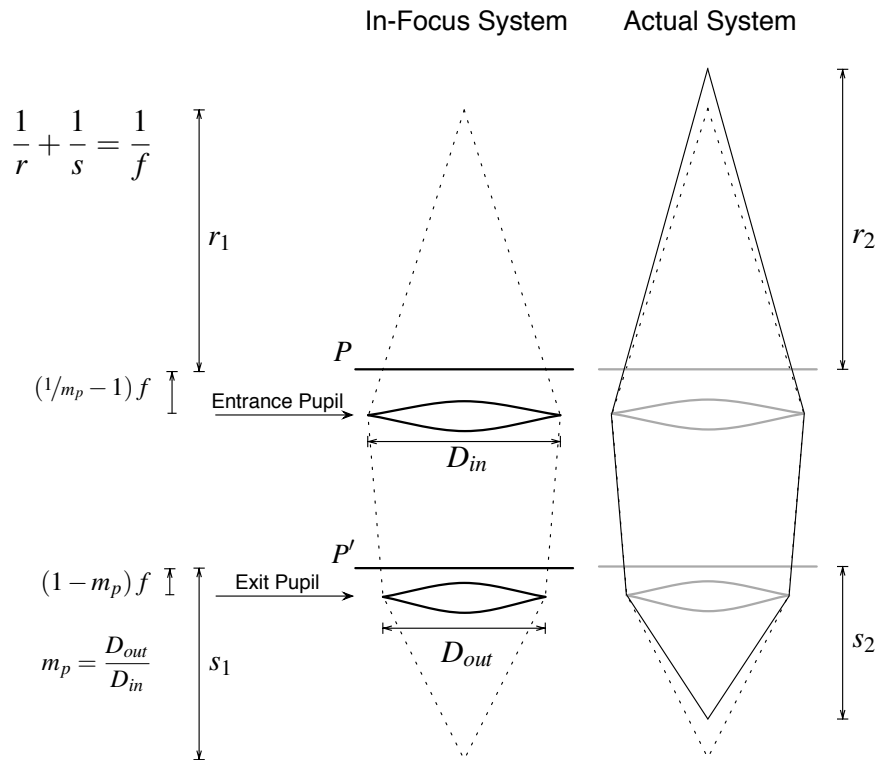


Figure 2.3: Diagram and relevant parameters needed to relate a physical camera to the generalized imaging model.

system is described in terms of its effect on a spherical reference wavefront entering through the entrance-pupil and exiting through the exit-pupil. The diameters of these pupils, assumed to be circular unless specified otherwise, are denoted by D_{in} and D_{out} . The entrance- and exit-pupils are the images of the aperture stop as seen from the object- and image- side of the optical system respectively. The entrance- and exit- principal planes, denoted P and P' , in conjunction with the effective focal length f , determines the geometric properties of the camera. These planes are positioned orthogonal to the optical axis where the Gaussian lens formula (E.1) is valid [93].

$$\frac{1}{r} + \frac{1}{s} = \frac{1}{f}, \quad (2.9)$$

This implies that for the purposes of ray-tracing, the two principal planes behave as though they were collocated.

The size of entrance and exit-pupils, along with their positions relative to the principal planes, can be determined by considering how light passes through the system when focused at infinity. When focused at infinity, taking the limit as $r \rightarrow \infty$ in (E.1), the beam of light passing through the entrance-pupil intersects P with a diameter of D_{in} and converges to a point on the optical axis at a distance $s = f$ behind P' . In practice, the entrance-pupil diameter is often specified in terms of the system's "F-Number", defined to as $f\# = f/D_{in}$. Applying a similar calculation using a plane-wave emitted from the detector gives the diameter of the exit-pupil. The pupil magnification of the system is determined by the ratio of exit- to entrance-pupil diameters $m_p = D_{out}/D_{in}$. Applying the law of similar triangles to a plane-wave entering the system reveals that the exit principal plane is located a distance $(1 - m_p)f$ in front of the exit-pupil. By tracing a plane-wave through the system in the opposite direction one finds that the entrance principal plane P lies a distance $(1/m_p - 1)f$ in front of the entrance-pupil. These relative relationships are summarized in Figure E.1, and entirely specify the imaging geometry in terms of common camera parameters.

The aggregate PSF resulting from diffraction at the aperture stop and imperfections in the optics may be described in terms of the deviations of the optical wavefront from a reference sphere measured in the exit-pupil [77]. A key feature of this model is that it accounts for these effects without requiring explicit knowledge of the layout or materials of the optical elements. In fact, the lens system needn't ever even form an image of the exit pupil for this technique to be applicable. These deviations are represented by a single complex function that describes changes in the reference sphere's intensity through its magnitude, and relative temporal advances or delays through its phase. Before proceeding it is helpful to introduce

the following standard notation.

$$\begin{aligned}
h(x,y) &\stackrel{\text{def}}{=} \text{Point Spread Function (PSF)} \\
H(w_x, w_y) &\stackrel{\text{def}}{=} \text{Optical Transfer Function (OTF)} \\
C(w_x, w_y) &\stackrel{\text{def}}{=} \text{Coherent Transfer Function (CTF)} \\
M(w_x, w_y) &\stackrel{\text{def}}{=} \text{Modulation Transfer Function (MTF)} \\
A(w_x, w_y) &\stackrel{\text{def}}{=} \text{Shaded aperture function} \\
\Psi(w_x, w_y) &\stackrel{\text{def}}{=} \text{Aperture phase function}
\end{aligned}$$

All of these quantities are sometimes used to describe this single complex function, and their relationships to one another are given as follows

$$h(x,y) = \mathcal{F}^{-1}\{H(\tilde{w}_x, \tilde{w}_y)\} \quad (2.10)$$

$$= c_0 |\mathcal{F}^{-1}\{C(\tilde{w}_x, \tilde{w}_y)\}|^2 \quad (2.11)$$

$$M(\tilde{w}_x, \tilde{w}_y) = |H(\tilde{w}_x, \tilde{w}_y)| \quad (2.12)$$

$$C(\tilde{w}_x, \tilde{w}_y) = A(\tilde{w}_x, \tilde{w}_y) \exp\left[j2\pi \frac{\lambda}{\lambda_{ref}} \Psi(\tilde{w}_x, \tilde{w}_y)\right] \quad (2.13)$$

$$\tilde{w} \stackrel{\text{def}}{=} \frac{\lambda}{\lambda_{ref}} w \quad (2.14)$$

where c_0 is a normalizing constant that ensures the PSF integrates to 1, the 2π in the exponent of the CTF implies that phase errors are being represented in units of waves, and the ratio λ/λ_{ref} accounts for the scaling between the wavelength being used to represent the optical system and the wavelength where the linear forward model is being evaluated. Our preferred description of the blur will be through the magnitude and phase decomposition given by A and Ψ .

2.2.4 Detection

For digital imaging devices, the intensity of the continuous electromagnetic wavefront impinging upon the detector plane is typically measured through an electrically biased semi-conducting substrate [172]. A series of detectors arranged on a focal plane array each measure the total energy received over their spacial support, the integration time, as well as their spectral sensitivity. Many modern multi-channel devices utilize a Bayer-filter, named after American scientist Bryce Bayer [11], to simultaneously collect multi-wavelength data. This is achieved by applying differing color (wavelength selective) filters throughout the

sensor. The filters are usually collocated in blocks that are repeated throughout the larger aggregate array. A camera using a Bayer filter is able to simultaneously collect multiple spectral bands at the cost of spatial resolution. Figure 2.4 illustrates part of a commonly used Bayer pattern, along with the corresponding elements for two of its four channels. Currently 2×2 patterns are most commonly used, however, some systems employ larger

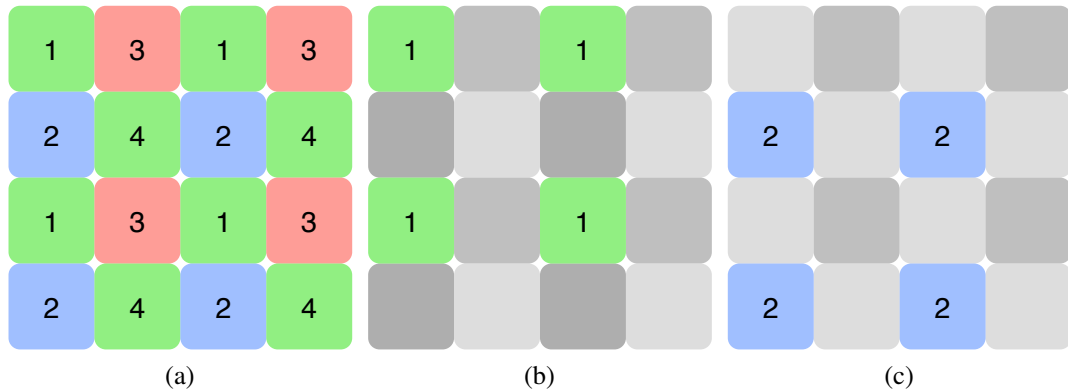


Figure 2.4: (a) 4 Blocks of a 2×2 Bayer pattern with each of the 4 data channels numbered and the corresponding object channels are indicated through the use of color, (b) Elements of channel 1, (c) Elements of channel 2

patterns with possibly overlapping bands in wavelength. The detector model will assume a regular grid of rectangular pixels within a larger Bayer block, but will permit arbitrary orderings of an arbitrary number of channels. The term “pixel” is used to indicate a single contiguous active silicon region. Each of the Bayer pattern blocks is assumed to be identical, and a data channel is defined to be the collection of pixels in a common position within their associated Bayer blocks. Notice that single-channel sensors may be treated as a special case of this more general detector model.

Now consider the relationship between the detector geometry, the impinging wavefront, and the data. In practice, the active silicon regions may not entirely fill a pixel’s physical support. When this is the case, we say that the pixel has a fill-factor f_f that is less than one. Figure 2.5 illustrates this situation along with the other key measurements that parameterize how a detector’s geometry influences the discrete data resulting from a measurement of the continuous electromagnetic field. Using these definitions, a Bayer pattern can be described by an $m_1 \times m_2$ matrix \mathbf{M} , whose size corresponds to the size of the Bayer filter, and whose entries correspond to the *object channel* associated with the *data channel* referenced by the

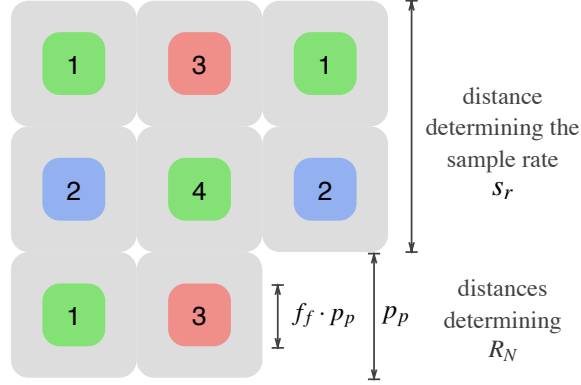


Figure 2.5: Key dimensions for determining the detector sample rate s_r , element sampling rate relative to Nyquist R_N , pixel pitch p_p , and element fill-factor f_f ($f_f \approx 50\%$ as shown)

row and column indices. For example, the detector illustrated in Figure 2.4 would result in

$$M \stackrel{\text{def}}{=} \begin{bmatrix} 2 & 1 \\ 3 & 2 \end{bmatrix}, \quad (2.15)$$

where we have chosen, without loss of generality, to label the object channels in RGB (Red, Green, Blue) order.

The spatial sampling rate of the electromagnetic field may be expressed in terms of the geometric parameters shown in Figure 2.5. The electromagnetic field striking the sensing aperture is band-limited by the aperture stop in the pupil, and the associated cutoff frequency \tilde{f}_c is given in terms of D_{in} and f for monochromatic light with a wavelength of λ as [77]

$$\tilde{f}_c = \frac{D_{in}}{2\lambda f}. \quad (2.16)$$

The maximum bandwidth of the spatial variation of the *intensity* of the electric field at the detector is thus given by

$$f_c = \frac{D_{in}}{\lambda f}. \quad (2.17)$$

For a polychromatic system, it's convenient to describe the common physical aperture-stop in units of a reference wavelength λ_{ref} that is equal to the smallest detectable wavelength. By convention, the sampling rate of the polychromatic detector is described by R_N ; the rate that individual detector sensing elements sample the intensity variation of the optical wavefront relative to its Nyquist rate at λ_{ref} (see Figure 2.5). For example, a system with $R_N = 1/3$ has individual sensing elements that will measure the intensity variation of the component of the

electric field at λ_{ref} at $1/3$ the Nyquist rate. Continuing with our aforementioned example, the reference wavelength will be associated with the blue channel ($\lambda_{ref} \approx 450$ nm), and due to the Bayer pattern and complete channel isolation, the collected data ultimately be at $1/6$ the Nyquist rate. The green and red channels measure longer wavelengths of light through a common physical aperture, and thus will sample their respective components of the electric field at slightly faster rates relative to the Nyquist rate for the corresponding wavelengths. The sampling rate along the i^{th} direction of the detector plane may be written in terms of R_N and λ_{ref} as

$$s_{r,i}(\lambda) = \frac{R_N \lambda}{m_i \lambda_{ref}}. \quad (2.18)$$

This relative sampling *rate* is independent of the corresponding sampling fill-factor $s_{f,i}$, which is a function of both the per-element fill factor f_f as well as the size (in pixels) of the Bayer filter. This aggregate sensing fill factor along the i^{th} direction is given by

$$s_{f,i} = \frac{f_f}{m_i}. \quad (2.19)$$

Letting p_p represent the pixel pitch, R_N can be expressed in terms of the physical camera parameters as

$$R_N = \frac{\lambda_{ref} f}{2 p_p D_{in}} \quad (2.20)$$

Taken together, these parameters provide a complete geometric description of how the detector samples the intensity of the electromagnetic field.

Finally, consider the statistical aspects of the sensing process. The absorption of the electromagnetic field by the photo-sensitive region of the detector results in the scattering of electrons from the valance band into the conduction band [172], and this electron flow is detected by an electric circuit. The electron generation process is Poisson distributed with a rate parameter that is proportional to the intensity of the impinging electric field [64]. The proportionality constant is known as the quantum efficiency of the device, and is effected by unwanted generation of electron-hole pairs (dark current) in the detector. The electronics that measure the induced current flow also produce a gaussian distributed signal (read-noise). From a signal processing perspective, dark current is indistinguishable from an addition bulk illumination (stray-light), or even self emission. Furthermore, read-noise levels are typically a function of the bias of semiconducting material (the gain) as well as the ambient temperature. Because these external factors are rarely reported with collected data, it will

be necessary to model them to ensure a correct absolute scaling of the noise distribution. In this work, the aggregate effect of these unknowns is accounted for through an unknown affine transformation of a normalized representation of the scene radiance along with an explicit estimate of the read-noise variance. This allows the lexicographically ordered data \mathbf{d} to be written in terms of a continuous-to-discrete linear mapping $\mathbf{A}(\boldsymbol{\theta})$ acting on the scene $f(x, y)$ as follows

$$\mathbf{d} = \mathcal{P}\{\mathbf{A}(\boldsymbol{\theta}) \circ [c_1 f(x, y) + c_2]\} + \mathcal{N}(\mathbf{0}, \sigma_r^2 \mathbf{I}) \quad (2.21)$$

where \mathcal{P} is the Poisson operator, and \mathbf{A} will depend on the unknowns $\boldsymbol{\theta}$ in a non-linear way. The expression (2.21) is written with c_1 , c_2 , and σ_r explicitly called out for clarity, but it should be understood that these values are members of $\boldsymbol{\theta}$. The particular form of \mathbf{A} is addressed in Section 2.4.

2.3 Previous Approaches

Section 2.2 describes the physical imaging process of a digital camera, and is sufficiently flexible to allow for joint estimation of the nuisance parameters likely to be unknown outside of a laboratory setting. Depending on which parameters are treated as unknowns, this forward model can be used to solve numerous problems in optical imaging. The current work, however, will focus on the estimation of blur and camera-pose using a known calibration target and possibly aliased data. Price [144] also considered the problem of blur estimation in limited-control environments, and is one of the only a few researches to also apply a continuous-to-discrete forward-model (cf. [141, 211]). Price, however, only considered the impact of blur on the data, and did so over a much more restrictive set of blurring functions. Our more general parameterization, through an unknown phase function in the exit-pupil (2.10)-(2.14), is made possible because of the explicit modeling of the band-limiting effects of the aperture-stop. Of the existing discrete-to-discrete approaches, Strong [178] uses a blur model based on Fraunhofer diffraction, but only considers defocus. Like previous works [21, 45, 46, 101], the model accommodates an unknown imaging perspective with respect to the calibration target. Of these works, however, only Delbracio [45] and Joshi [101] consider aliasing, and neither do so in conjunction with a diffraction-based optics model. None of these previous works account for the intensity-dependent variations caused by the physics of the detection process, and as a result, these models are unable to provide estimates of the read-noise variance. Finally, this work accommodates chromatic effects, including the impact of using a Bayer filter [11]. Jointly optimizing over

the resulting high-fidelity model leads to substantially better system characterization, and direct modeling the physical imaging processes ensures the technique is flexible enough to accommodate a broad class of imaging systems.

2.4 Forward Model Components

This section further develops the forward modeling components necessary to *implement* the physical effects outlined in Section 2.2. Some of the constituent processes that compose $\mathbf{A}(\boldsymbol{\theta})$ from (2.21) commute, and this fact will be exploited to produce a series of operators that are mathematically equivalent to a more naive implementation, but that are faster and more accurate from a numerical perspective. The discrete representation of continuous objects, especially as they pertain to edge-effects, is also addressed. The resulting model describes the lexicographically ordered observations \mathbf{d} in terms of \mathbf{f} ; a discrete lexicographically ordered description of the continuous object over a finite support. The action of $\mathbf{A}(\boldsymbol{\theta})$ is then described in terms of the constituent operator series $\mathbf{T} \circ \mathbf{S} \circ \mathbf{B} \circ \mathbf{W}$, where conceptually \mathbf{W} represents a warp, \mathbf{B} a blur, \mathbf{S} a resampling, and \mathbf{T} spatial truncation.

A direct description of the forward imaging model for the j^{th} frame and k^{th} channel can be written in matrix form as

$$\mathbf{A}_{j,k} = \mathbf{T}\mathbf{S}(\mathbf{s}_r) \mathbf{W}_{det}(\mathbf{M}, R_N) \mathbf{B}_{det}(\mathbf{s}_r, \mathbf{s}_f) \mathbf{B}_{opt}(\boldsymbol{\alpha}, \lambda_k) \mathbf{W}_{pos}(\mathbf{v}_j), \quad (2.22)$$

where \mathbf{W}_{det} describes the geometric transform associated with the sensor's Bayer pattern, \mathbf{B}_{det} describes the blur induced by sampling the wavefront using a detector with finite spatial support [17], \mathbf{B}_{opt} accounts for the optical blur induced by wavefront distortion from the optical system [77], and \mathbf{W}_{pos} describes the geometric transform that accounts for the camera pose relative to the object [91]. These operators are parameterized by \mathbf{s}_r , a vector describing the sampling rate, \mathbf{s}_f , a vector defining the aggregate fill-factor, $\boldsymbol{\alpha}$, a vector that parameterizes the blur, λ_k , the optical wavelength of the k^{th} channel, and \mathbf{v}_j , a vector parameterizing the camera pose during the j^{th} frame. In the remainder of this section it will be shown that \mathbf{W}_{det} represents a translation that commutes with any blur operator. This follows from the fact that both pure translation and blur are diagonalized by the Fourier transform. As a consequence, this series of physical effects can be implemented by the simpler aggregate series $\mathbf{T}\mathbf{S}(\mathbf{s}_r) \mathbf{B}(\mathbf{s}_r, \mathbf{s}_f, \boldsymbol{\alpha}) \mathbf{W}(\mathbf{v}_j, \mathbf{M}, R_N)$.

2.4.1 Object Model

The continuous object model $f(x, y, \lambda)$ described in Section 2.2.1 is represented by a discrete lexicographically ordered vector \mathbf{f} supported over a toroid defined through a regular grid of points. The elements of this vector represent independent sources of radiation in both space and wavelength. The object is assumed to be circularly symmetric (represented over a toroid), and is only observed over a smaller circumscribed sub-region \mathcal{Y} that represents the projection of the observation support \mathcal{X} into the object plane. Figure 2.6 illustrates this situation where dots are used to show elements within \mathbf{f} . A light dashed line represents the circularly symmetric boundary, and the dark region shows the support of observations when projected into the object plane. Recall from Section 2.2.2 that the transformation




Figure 2.6: The continuous object is represented through a regular grid of points defined over a toroid, with observations only occurring a smaller circumscribed subset.

between $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} \in \mathcal{Y}$ is described by a perspective transformation matrix \mathbf{V} acting on the points described through a homogeneous coordinate representation; $\mathbf{y} = \mathbf{V}\mathbf{x}$. For every point within the larger object support, the continuous object is defined through a Fourier series whose coefficients are given by its finite representation \mathbf{f} (c.f. [76] Chapter 3). The use of a toroidal support in conjunction with a Fourier basis allows both the continuous object and its continuous Fourier transform to be computed efficiently and exactly through the Discrete Fourier Transform (DFT). That the object support extends beyond \mathcal{Y} , the projection of the observations into the object plane, models the influence of unobserved object content, and allows these unknowns to be directly accommodated.

Extending the object model beyond the support of the adjoint of the idealized pinhole camera model introduces a guard region that resolves the incongruity between the toroidal support assumption needed to ensure a finite continuous representation of both the intensity

and its Fourier transform, as well as the Euclidean plane assumption used in Sections 2.2.1 and 2.2.2. Referring to Figure 2.6, one can see that this object model indeed describes a region beyond the field-of-view of the sensor. For the forward model to be valid, this must always be the case. In practice, however, two questions naturally arise. Given sampled data, how should one form its continuous extension, and what is the associated Fourier transform?

2.4.2 Application to Image Recovery and Guard Banding

We illustrate the object modeling approach through an image recovery application. Consider the common practice of tapering as a bias reduction technique in spectral estimation. Tapering measured data unequivocally reduces the information available for subsequent processing tasks. This result follows immediately from the data processing inequality [37, p. 32] for processing data generated from a random process (e.g. images obeying the forward model 2.21). Following the proof given by Cover, suppose one is interested in inferring some underlying statistical quantity X given a measurement D . If a non-trivial taper is applied resulting in $G = g(D)$, then a Markov chain is formed $X \leftrightarrow D \leftrightarrow G$ where

$$I(X;G) = I(X;D,G) - I(X;D|G) \quad (2.23)$$

This implies information is lost in any taper where $I(X;D|G) > 0$. It also implies that a necessary and sufficient condition for preventing information loss is to ensure that X and D given G are conditionally independent. Letting g be a function that maps D into G which also preserves D through the canonical mapping is clearly a sufficient condition for preventing information loss. This result immediately suggests the construction of a model whose support exceeds the observed data because bias reduction can be achieved through extrapolation, which has the potential to preserve all the information of the measured data.

The harmonic nature of the Fourier basis, in conjunction with the desire for a discrete representation, suggests the use of a toroidal object support (c.f. [195]). In light of the data processing inequality, the only remaining question is how the available data ought to be extrapolated. Considering the uncertainty principle [16, 173], and the fact that one does not expect a natural scene to be “temporally” concentrated, a natural goal is to minimize the bandwidth of the overall signal subject to the observed data. This objective function is given in 1D by

$$\Lambda(f(x)) = \int (w - w_0)^2 |\hat{f}(w)|^2 dw, \quad (2.24)$$

where w_0 is a free variable related to where one labels the origin of the frequency axis. The

signals of interest are always real, implying $|\hat{f}|^2$ is a symmetric positive function about zero. It follows from symmetry, that the minimizer with respect to w_0 is given by $w_0 = 0$. By absorbing w into the modulus, and applying Parseval's theorem [82, 138], this objective function can be seen to be equivalent to minimizing the integrated squared derivative over the circle.

$$\Lambda(f) = \int |w\hat{f}(w)|^2 dw \quad (2.25)$$

$$= \int \left| \frac{\partial f(x)}{\partial x} \right|^2 dx \quad (2.26)$$

This is a classic measure of roughness [88, 152], and a member of the more general class of functions

$$R_m(f) = \int \left(\frac{\partial^m f}{\partial x^m} \right)^2 dx. \quad (2.27)$$

Thus, one can interpret the class of functions that minimize this objective function as the maximally smooth periodic extensions of the data subject to the support limitations. Taking a viewpoint common in physics [128], one might note that the non-commutable induced differential operators associated with the tangent-space of the Heisenberg group are related by first-order differentiation, also suggesting the use of $m = 1$ in (2.27). From this viewpoint, maximizing the smoothness of the continuous extension is equivalent to minimizing the unobserved member of the two incompatible observables.

Given a set of observations $\tilde{x} \in \mathbb{R}^M$, whose support is a subset of $x \in \mathbb{R}^N$, edge effects are mitigated through the data-dependent extension given by the sub-problem

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{W}\mathbf{F}\mathbf{x}\|^2 \quad \text{s.t.} \quad (2.28)$$

$$\mathbf{T}\mathbf{x} - \tilde{\mathbf{x}} = \mathbf{0}, \quad (2.29)$$

where \mathbf{F} is the Fourier transform operator, \mathbf{W} is a diagonal matrix of weights, and $\mathbf{T} : \mathbb{R}^N \mapsto \mathbb{R}^M$ is a spatial truncation operator. “ \mathbf{W} ” is commonly used to represent weighting matrices, and its distinction from the aforementioned warping operator should be clear from context. Let $\mathbf{U} : \mathbb{R}^N \mapsto \mathbb{R}^{N-M}$ be the canonical mapping from \mathbb{R}^N to the null-space of $\mathbf{T}^T\mathbf{T}$. Then \mathbf{T}^T , and \mathbf{U}^T form a partition on \mathbb{R}^N as shown in Figure 2.7. Back-substituting results in an




Figure 2.7: Partitioning of \mathbf{x}

equivalent unconstrained optimization problem

$$\hat{\mathbf{u}} = \arg \min_{\mathbf{u}} \Lambda(\mathbf{u}) \quad (2.30)$$

$$= \arg \min_{\mathbf{u}} \frac{1}{2} \|\mathbf{W}\mathbf{F}(\mathbf{U}^T \mathbf{u} + \mathbf{T}^T \tilde{\mathbf{x}})\|^2 \quad (2.31)$$

The non-commutability of the space and frequency truncation operators preclude a closed form solution, but even for large systems the problem is efficiently solved using quasi-newton techniques. The necessary gradient is given by

$$\frac{\partial \Lambda}{\partial \mathbf{u}} = \mathbf{U}\mathbf{F}^H \mathbf{W}^2 \mathbf{F}(\mathbf{U}^T \mathbf{u} + \mathbf{T}^T \tilde{\mathbf{x}}). \quad (2.32)$$

Figure 2.8 illustrates this technique used to find the minimal bandwidth pad for a typical image. Noting that extrapolation and interpolation on a toroid are in fact the same, Figures 2.9-2.12 show the technique applied with a more general time-truncation operators. In all cases, the results provide the expected behavior, but do so while exactly preserving the observations, and therefore address the aforementioned edge-effect issue (bias) without unnecessarily destroying the information content in the observations. The examples over more general truncation operators illustrate that this method achieves our aim of bias mitigating, but does so based on a reasonable criterion.

Viewing the minimal bandwidth objective as a prior, we say that this approach is reasonable to the extent that it is predictive of the missing observations. It should also be noted that the proposed approach directly optimizes the objective over the continuous manifold, is unique as a consequence of the convexity of the resulting objective function, and may be efficiently computed (3X slower than bicubic interpolation [105] for examples shown).

Continuous, and more recently smooth, function extension is a classical problem that has been well studied. Well known results include Urysohn's Lemma, the Tietze Extension




Figure 2.8: Example of padding using the proposed technique.

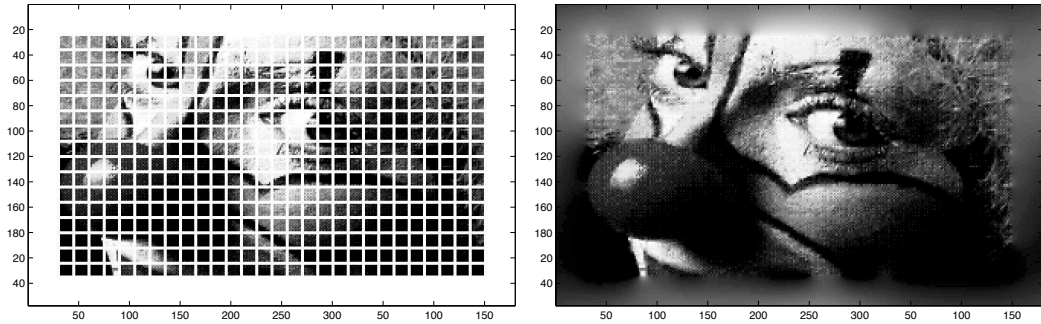


Figure 2.9: General in-painting example using the proposed approach.

Theorem, and the Whitney Extension Theorem. More recently, Fefferman addressed the problem of smooth function extension through a series of related works [59–61]. Chan [29] and Chui [33] study a number of smooth function extension techniques, and provide bounds on their performance. Chui [34] goes on to study multi-resolution function extensions in \mathbb{R}^N using a local harmonic basis. This work is clearly related to the proposed technique, however, using a local basis over \mathbb{R}^N rather than a global basis over \mathbb{T}^N results in substantially different algorithms. Hoang [96] discusses the implementation details of the technique suggested in [33, 34], and provides a performance comparison. The image inpainting literature (see, e.g. [22], and the references therein) is another related field distinguished by its more aggressive use of heuristics, with the aim of creating images that appear natural to the human observer. While this work has a subtly different aim, the ideas are close enough to warrant a comparison. A number of inpainting approaches [34, 43, 51, 54, 117] are compared to the proposed smooth function extension in Appendix A.

The bandlimited nature of imaging ensures that a bandlimited continuous object model is sufficient for solving inverse problems. For computational reasons, we require both a finite object representation and a finite representation of that object’s Fourier transform. These requirements immediately suggest that we are interested in the center of a discrete Heisenberg group. This group structure will, however, imply a continuous object defined over




Figure 2.10: Example reconstruction from 50% of the data using the proposed approach.




Figure 2.11: Example reconstruction from 25% of the data using the proposed approach.

a toroid, and this must be reconciled with our plane-to-plane imaging geometry. This section proposes a resolution by defining the object over a toroidal support that exceeds the support of the adjoint of the idealized pinhole camera model used to define the plane-on-plane imaging geometry. What results is an implicit guard-band in the continuous object model that accommodates these differences while maintaining the necessary finite representations. In the absence of measurements, a minimum bandwidth extension is proposed, and this extension is shown to be unique and computationally efficient to estimate. Section 2.4.5 expands upon these ideas in the context of upsampling, downsampling, and aliasing. This section also explicitly defines the mappings between discrete object representations and their implied continuous forms.

2.4.3 Warp Operator (\mathbf{W})

Given a continuous object $f(x, y, \lambda)$ represented though \mathbf{f} , the warping operator \mathbf{W} is an 8 degree-of-freedom projective transformation parameterized through the 3×3 matrix \mathbf{V} , and describes the mapping of \mathbf{f} from the object plane into an imaging plane. This situation is illustrated in Figure 2.2. The space of continuous objects representable by a finite Fourier series is not closed under perspective transformations, so the warping operator must be carefully constructed to simultaneously maintain object fidelity and computational efficiency.




Figure 2.12: Example reconstruction from 10% of the data using the proposed approach.

This is achieved through a carefully controlled bicubic interpolation operator acting on \mathbf{f} . Let \mathbf{Q} be the unique separable bicubic interpolation operator that achieves $O(n^3)$ convergence [105], and \mathbf{v} a vector that encodes the degrees of freedom within \mathbf{V} as given in Appendix C, then the approximation to the exact perspective transform of the continuous object may be written as

$$\mathbf{W}(\mathbf{v}) \circ f(x, y, \lambda) \stackrel{\text{def}}{=} f(x(\mathbf{v}), y(\mathbf{v}), \lambda) \quad (2.33)$$

$$\approx \mathbf{Q}(x(\mathbf{v}), y(\mathbf{v})) \mathbf{f}, \quad (2.34)$$

where the coordinates in the right-hand side of (D.1) are given in terms of their homogeneous representations by (2.8). This approximation can be made sufficient due to the band limiting nature of the camera's entrance pupil, and its fidelity can be precisely controlled by ensuring the object is sufficiently upsampled with respect to \mathbf{v} .

The band-limiting nature of a camera's aperture-stop, in conjunction with the Nyquist sampling theorem, allows the discrete circulant representation of the object to be adequate if the sampling grid is sufficiently fine. From the perspective of the camera, one can work backwards from (E.1) to establish that the Ground Sample Distance (GSD) in the object plane is given by

$$GSD = (R - f) \frac{p_p}{f}. \quad (2.35)$$

Applying the definition of R_N , one finds that the Nyquist rate of the object plane *with respect to the sensor* is given by GSD/R_N . This analysis, however, does not account for possible geometric transformations acting on a pre-defined object. One could imagine a sufficiently strong transformation (say a magnification), that causes a particular object representation to become inadequate with respect to the camera pose parameterization \mathbf{v} . Specifically, given an object represented at the Nyquist rate with respect to a sensor configuration we require

that

$$\min_{\mathbf{x} \in \mathcal{X}} \frac{\sigma^{(1)}\left(\left[\mathbf{V}^{-1}(\mathbf{v})\right]_{1:2,1:2}\right)}{\left|\left[\mathbf{V}^{-1}(\mathbf{v})\right]_{3,1:2} \mathbf{x} + \mathbf{1}\right|} \leq \frac{1}{\gamma} \leq 1, \quad (2.36)$$

where $\sigma^{(1)}(\cdot)$ is the maximum singular value, and γ is an engineering safety factor that ensures the bicubic interpolation sufficiently approximates sinc interpolation. The sampling rate requirement imposed by (2.36) requires that the local change in topology induced by the warping operator compresses the object representation by a factor of at least γ for every point and every direction vector in \mathcal{Y} . In practice, letting $\gamma = 2$ results in a reasonable compromise between computational efficiency and accuracy. Figure 2.13 shows the optimal bicubic interpolation kernel [105], as well as the magnitude of its transfer function in the frequency domain. The magnitude of the transfer function shown in (b) motivates the use of $\gamma = 1/2$, and illustrates the impact of alternative choices of γ .

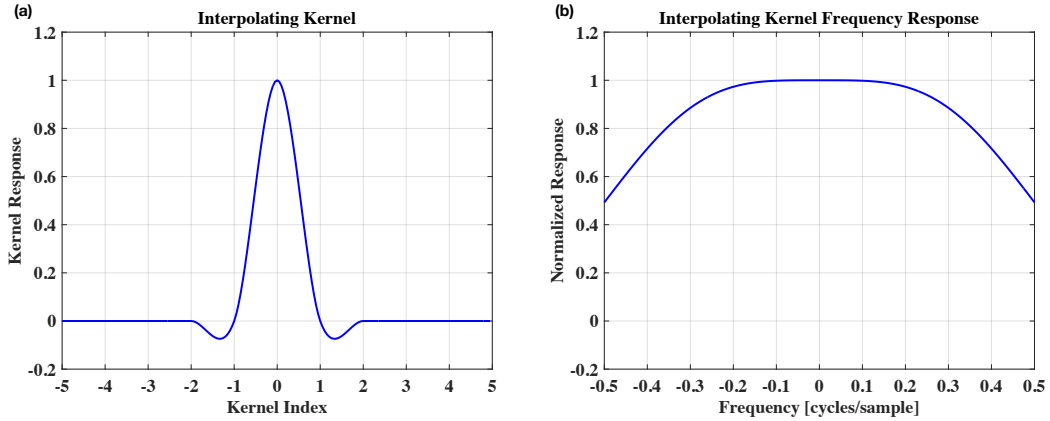


Figure 2.13: Optimal bicubic interpolating kernel (a) and its corresponding frequency response (b).

The use of homogeneous coordinates in (D.2), in conjunction with the fact that the warping operator acting on \mathbf{f} is a matrix function of two vector arguments that are themselves functions of a vector argument \mathbf{v} , make the derivative of this operator w.r.t. \mathbf{v} somewhat unusual. For completeness, details of this derivative are given in Appendix D.

For clarity, consider the application of (D.2) where the camera perspective parameterization \mathbf{v} contains only a single degree of freedom describing rotation about the center of the object description. Defining the origin of the object plane to be at the center of the object representation is very natural, but requires the use of the similarity transform given in (2.8).

The associated 3×3 perspective transformation is given by

$$\mathbf{V}_{45^\circ} = \begin{bmatrix} 0.70711 & 0.70711 & -\frac{N_r+1}{2} \\ -0.70711 & 0.70711 & -\frac{N_c+1}{2} \\ 0 & 0 & 1 \end{bmatrix}, \quad (2.37)$$

where the discrete object representation is $N_r \times N_c$ and is indexed by a grid of natural numbers. The code for implementing this situation is shown below, and Figure 2.14 shows the corresponding image and its gradient with respect

```

1 % Simple warping demonstration
2 load('clown','X');
3 theta = pi/4;
4 [imgWarp,imgWarpDeriv] = warp2D(X,theta);

```




Figure 2.14: Action of a 45° rotation operator and the associated derivative

2.4.4 Blur Operator (\mathbf{B})

Conventional imaging systems contain at least two sources of blur. The first is related to the diffraction effects induced by a finite imaging aperture, and the second is from recording the continuous electromagnetic field using individual detector elements of finite physical dimension. Additionally, the proposed camera model accommodates imperfections in the individual optical elements through an aggregate description of the wavefront-aberrations they induce. Unless otherwise noted, each of these three effects will be represented by the aggregate blur operator $\mathbf{B}(\boldsymbol{\alpha})$, where it is understood that this corresponds to the more

precise set of operators

$$\mathbf{B}(\boldsymbol{\alpha}) \stackrel{\text{def}}{=} \mathbf{B}_{det}(\mathbf{s}_r, \mathbf{s}_f) \mathbf{B}_{opt}(\boldsymbol{\alpha}, \lambda) \quad (2.38)$$

with \mathbf{B}_{det} representing the detector blur as a function of the sampling-rate relative to Nyquist \mathbf{s}_r , and the detector fill-factor \mathbf{s}_f . \mathbf{B}_{opt} represents the blur induced by optical aberrations modeled in the exit pupil. Recall that the relationship between the Point Spread Function (PSF) and wavefront aberration description in the exit pupil $\Psi(\boldsymbol{\alpha})$ is given at a particular wavelength λ by (2.10). The resulting PSF is convolved with a band-limited representation of the image of the object at the detector plane to determine the field intensity at each sensing element. Recall also that this PSF description is given in terms of normalized spatial frequency variable (2.14), or equivalently, a normalized spacing in the detector plane. A sequence simple of examples is provided at the end of this section.

The ratio λ/λ_{ref} accounts for the scaling between the wavelength being used to represent the optical system, and the wavelength where the model is being evaluated. The reference wavelength must be smaller than the smallest wavelength of light collected by the system to ensure the exit-pupil representation is not aliased.

The forward camera model is flexible enough to support an arbitrary finite wavefront aberration description, however, a Zernike basis [214] is commonly used because it has been shown to parsimoniously represent common manufacturing errors [164]. This basis is orthonormal over the unit circle, and given in polar coordinates in terms of the integral indices m, n as

$$\mathbf{Z}_n^m(\rho, \theta) = \begin{cases} N_n^m R_n^{|m|}(\rho) \cos(m\theta) & m \geq 0 \\ -N_n^m R_n^{|m|}(\rho) \sin(m\theta) & m < 0 \end{cases} \quad \text{s.t.} \quad (2.39)$$

$$R_n^{|m|}(\rho) = \sum_{k=0}^{(n-|m|)/2} \frac{(-1)^k (n-1)!}{k! [n+|m|/2-k]! [n-|m|/2-k]!} \rho^{n-2k} \quad (2.40)$$

$$N_n^m = \sqrt{\frac{2(n+1)}{1+\delta(m)}}, \quad (2.41)$$

where $\delta(\cdot)$ is the Kronecker delta function, $m \leq n$, and $n - |m|$ is even. In this decomposition, n describes the highest order of the radial polynomial and m describes the azimuthal frequency. By convention, the basis functions are ordered according to a single index over the natural numbers (the Noll ordering [136]), allowing the overall wavefront to be parameterized by the vector $\boldsymbol{\alpha}$. A consequence of the orthonormality of the Zernike basis is that when one approximates a continuous wavefront using these functions, each

subsequent term minimizes the RMS error of the wavefront description up to that radial order and azimuthal frequency. It is common practice to use all azimuthal frequencies up to a particular radial order, and continuous wavefronts are often expressed up to radial order 10 (55 terms) [66]. Navarro [134] discusses of the impact of the number of Zernike modes on modeling fidelity as a function of the aperture amplitude function. The first few Zernike basis functions are shown in Figure 2.15. An analysis of this optical modeling approach to a thin lens is given by Welford [201, Chap. 12].

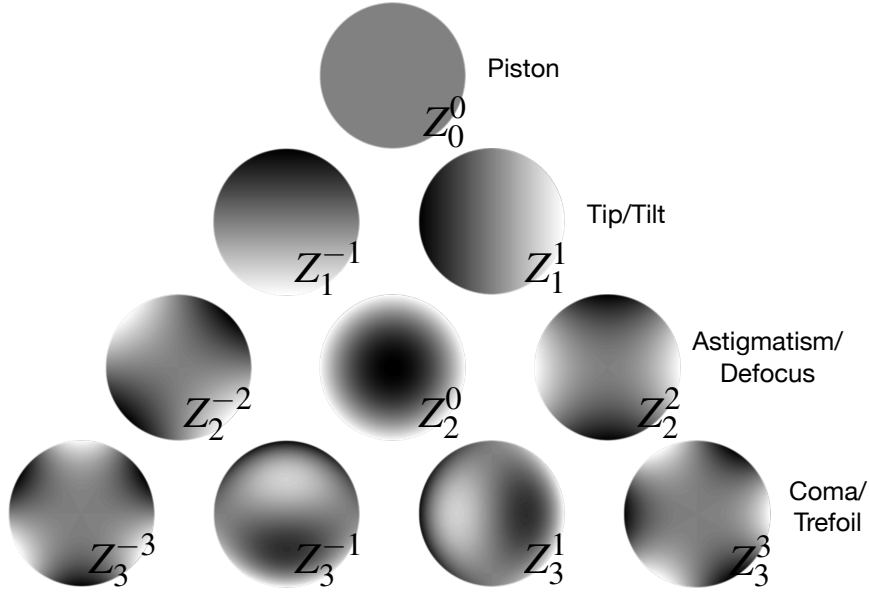


Figure 2.15: First 10 Zernike basis functions as defined by (2.39) along with their common names

The circulant object model allows the PSF to be efficiently applied in the Fourier domain, however, the autocorrelation function implicit in (2.10) must still be evaluated numerically. For the purposes of this work, it will be necessary to distinguish between PSF variations resulting from wavefront descriptions that differ by only 1×10^{-3} Waves RMS (WRMS). This accuracy requirement places lower limits on the size of the discrete representation of the coherent transfer function (2.13), and for computational reasons, also mandates the use of a shaded aperture approximation. These two issues are addressed in the Appendices F and G respectively.

The detector blur \mathbf{B}_{det} models the deviations from an ideal impulse response brought about by sampling a continuous signal with individual sensing elements of finite physical support. Before considering the transfer function in terms of the physical camera parameters,

it is helpful to recall the following property of the Fourier transform.

$$\frac{1}{2\pi} \int_{-\infty}^{\infty} \text{rect}(t) e^{-j\omega t} dt = \frac{1}{\sqrt{2\pi}} \text{sinc}\left(\frac{\omega}{2\pi}\right) \quad (2.42)$$

$$\text{sinc}(x) \stackrel{\text{def}}{=} \frac{\sin(\pi x)}{\pi x} \quad (2.43)$$

$$\text{rect}(x) \stackrel{\text{def}}{=} \left\{ \begin{array}{ll} 1 & , \quad |x| \leq 1/2 \\ 0 & , \quad \text{Otherwise} \end{array} \right\} \quad (2.44)$$

The final form of the detector transfer function is separable, so it is sufficient to study the 1D transfer function. Referring to Figure 2.5, one can see that the final form of the detector transfer function for a single data channel can be computed by applying the Fourier scaling property to account for the sampling rate, and an energy scaling to accommodate the inactive regions. Applying these two requisite factors, given by (2.18) and (2.19), one finds that the transfer function in the i^{th} direction is given by

$$B_{det,i}(\omega, s_{r,i}, s_{f,i}) = \frac{s_{f,i}}{s_{r,i}} \text{sinc}\left(\frac{\omega}{2\pi s_{r,i}}\right). \quad (2.45)$$

By construction, this transfer function is centered on the data channel to which it refers. Jointly accommodating multiple, simultaneous, data channels requires one to account for their positions relative to one another. This is precisely \mathbf{W}_{det} in (2.22), and its description in homogeneous coordinates relative to λ_{ref} is given by

$$\mathbf{V}_{det}(k_d) = \begin{bmatrix} 1 & 0 & \frac{-\Delta_{k_d,1}}{R_N} \\ 0 & 1 & \frac{-\Delta_{k_d,1}}{R_N} \\ 0 & 0 & 1 \end{bmatrix}, \quad (2.46)$$

where k_d is the k^{th} data channel, and $\Delta_{k_d,i}$ is its offset from the upper-leftmost channel in the Bayer filtering block along the i^{th} direction. It is sometimes useful to express $H_{det,i}$ entirely in terms of physical sensor quantities. Substituting from R_N from (2.20) one finds

$$B_{det,i}(\omega, s_{r,i}, s_{f,i}) = \frac{2f_f p_p D_{in}}{f\lambda} \text{sinc}\left(\frac{2\omega m_i p_p D_{in}}{2\pi f\lambda}\right) \quad (2.47)$$

To make the previous section more concrete, consider the PSF associated with a narrow-band RGB system that is represented at half its minimum wavelength.

```

1 % Generate a 3-channel narrow-band RGB system
2 psfOpt = getPSFOptions;
3 psfOpt.lambdaVec = [650 550 450]; % in nm
4 psfOpt.lambdaRef = 225; % 2X oversampled grid
5 psfOpt.aberVec = 0;
6 psfOpt.broadband = false ;
7
8 psfSize = [1024 1024];
9 [psf , ctf , config ] = getPSF(psfSize , psfOpt);

```

The resulting aperture functions in the exit pupil are shown in Figure 2.16a. Due to the use

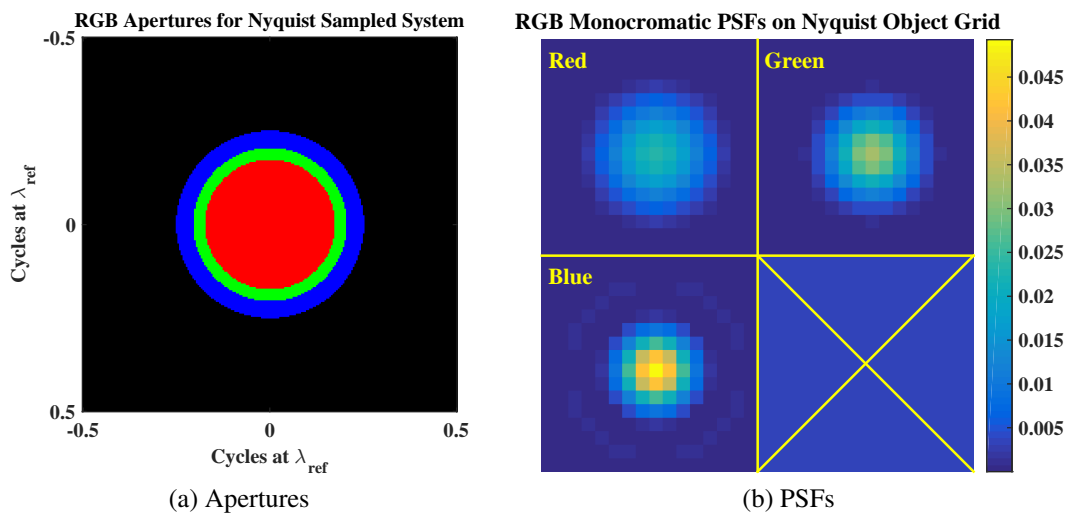


Figure 2.16: Apertures for a narrow-band RGB system sampled at λ_{ref} in the exit pupil, and the associated 2X oversampled narrow-band PSFs.

of a common reference wavelength, one can easily see the relative scaling of the pupil with respect to the color of light passing through the system. The reference wavelength is half that of the smallest wavelength of light, and we say that such a system is oversampled by a factor of 2. It is understood that the actual oversampling amount for any particular wavelength will be *at least* this large. Similarly, the PSF representation is said to be oversampled by a factor of 2 (relative to the Nyquist rate for the reference wavelength) where the PSF at any individual wavelength is once again oversample by *at least* this amount.

Now consider the addition blur imposed by collecting this PSF on an ideal ($f_f = 1$, $s_{r,i} = 1$, $s_{f,i} = 1$), 3-channel detector. In practice this could be achieved with beam-splitters and 3 geometrically identical detectors that have each been doped to respond equally within their respective bands.

```

1 % Account for an ideal 3-detector system
2 f_f = 1;
3 s_f = f_f;
4 for chanInd = 1:3
5     s_r = psfOpt.lambdaVec(chanInd)/psfOpt.lambdaRef;
6     HDet(:, :, chanInd) = sampleTransfer(psfSize, s_r, s_f);
7 end
8 psf2 = real( ifft2 ( fft2 ( psf ).* HDet)); % real for round-off

```

The difference between the PSF collected on this system and an ideal delta-function detector is imperceptible to the human eye and shown in Figure 2.17b. Similarly, consider an ideal Bayer detector using a common RRGB pattern ($f_f = 1$, $s_{r,i} = .5$, $s_{f,i} = .5$).

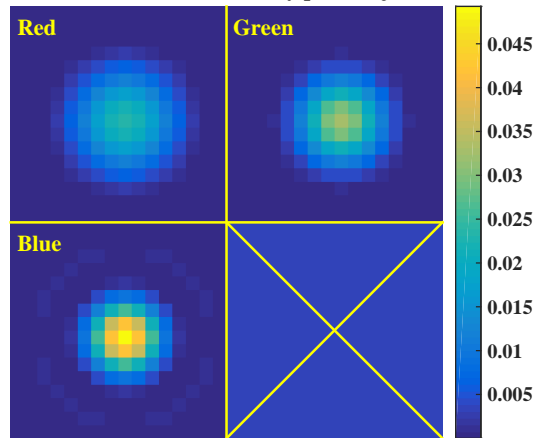
```

1 % Account for an ideal Bayer system
2 BayerPatternSize = [2 2];
3 s_f = f_f ./ BayerPatternSize;
4 for chanInd = 1:3
5     s_r = psfOpt.lambdaVec(chanInd)./(psfOpt.lambdaRef*BayerPatternSize);
6     HDet2(:, :, chanInd) = sampleTransfer(psfSize, s_r, s_f);
7 end
8 psf3 = real( ifft2 ( fft2 ( psf ).* HDet2));

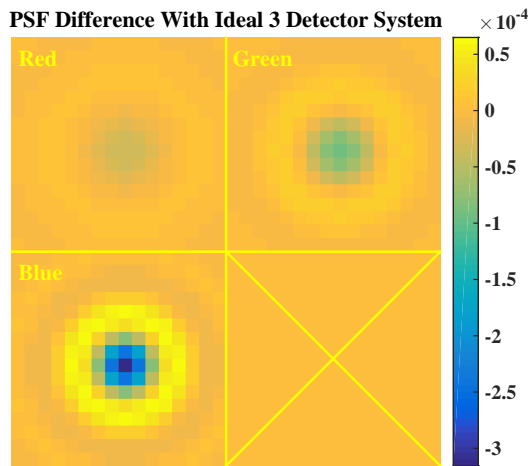
```

The individual pixel geometries of the Bayer detector are the same as the 3 sensor system, but one can see in Figure 2.17c that the PSF differs substantially due to the compromise between spatial and frequency resolution.

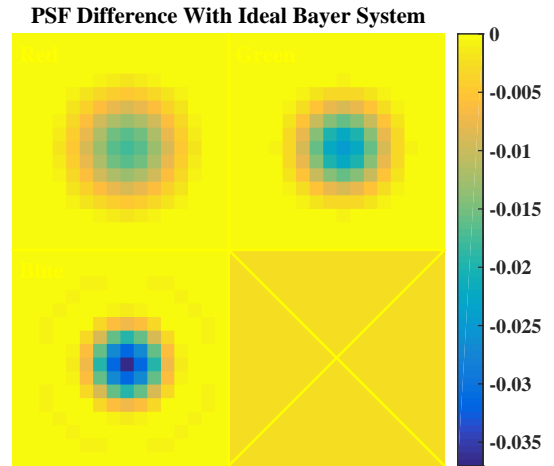
RGB Monochromatic PSFs on Nyquist Object Grid



(a) Ideal 3-channel monochromatic system



(b) PSF difference using 3 ideal detector-arrays



(c) PSF difference using ideal Bayer filter

Figure 2.17: Comparison between the PSF of a 3-channel monochromatic system (a), and the PSF for the same system when collected 3 ideal single-channel detectors (b), or an ideal 4-channel Bayer detector (c).

2.4.5 Sampling Operator (\mathcal{S})

Many imaging systems, including most commercially available cameras, are under-sampled with respect to the optical wavefront they measure. Section 2.4.1 developed a discrete representation for continuous bandlimited signals, which given the band-limiting nature of the sensors aperture, is sufficient for describing the object being imaged. The sampling operator \mathcal{S} allows one to model the effects of sampling this class of continuous, bandlimited, signal at rates different from its canonical representation. This section builds upon the signal description by providing a general mathematical formulation for sampling, interpolation, differentiation, and integration of this class of signals.

The purpose of this mathematical formulation is for modeling the types of bandlimited signals that result from collecting data using finite sensing apertures. Such collections always result in band-limited measurements; a concept that fundamentally involves a Fourier description of the signal in question. Furthermore, any system only collects a finite number of measurements, and this mandates the use of discrete signal representations. This mandate pushes us away from the traditional, analysis-based, perspective of harmonic analysis [20],[48] and toward a group-theoretic viewpoint [159],[188]. By basing our model on these much more fundamental mathematical structures, an enormous body of work can be utilized to establish properties of the signal model.

Let $\{(t_i, f_i)\}_{i=1}^n$ be a uniformly spaced series of N measurements on an interval $[a, b]$ that is contained within a periodic domain whose fundamental period is given by $T = \frac{N}{N-1}(b-a)$. It follows immediately that the sample spacing is given by $\Delta = T/n$, and without loss of generality, let $t_1 = 0$. The relationship between these samples some periodic continuous extension $f_c(t)$ is shown in Figure 2.18. If $\forall i$ one associates the coordinates t_i with a member




Figure 2.18: Continuous periodic extension of a signal sampled over the toroid.

of $\mathbb{Z}_n \stackrel{\text{def}}{=} \mathbb{Z}/n\mathbb{Z}$, the ordinates f_i with a member of the field \mathbb{C} , then the signal can be thought

of as a point in the usual complex vector space. \mathbb{Z}_n is an additive finite abelian group, and its character is defined as the homomorphism to the multiplicative group of non-zero complex numbers $\chi : \mathbb{Z}_n \mapsto \mathbb{C}^\times$.

$$\chi(a+b) = \chi(a)\chi(b) \quad \forall a, b \in \mathbb{Z}_n \quad (2.48)$$

Notice that $\chi(na) = \chi(a)^n = \chi(0) = 1$, implying that the values of χ are the n -roots of unity. It follows immediately that this map is given by

$$\chi(a) = e^{2\pi ja/n}, \quad (2.49)$$

and the associated character group is simply $\hat{\mathbb{Z}}_n \stackrel{\text{def}}{=} \{\chi_0, \chi_1, \dots, \chi_{n-1}\}$. By constructing the usual complex vector-space over this character group as well, the Discrete Fourier Transform (DFT) can be seen to be the mapping $F_n : L^2(\mathbb{Z}_n) \mapsto L^2(\overline{\mathbb{Z}}_n)$

$$(F_n f)(\bar{a}) = \sum_{a \in \mathbb{Z}_n} f(a) \langle -a, \bar{a} \rangle \quad \text{s.t.} \quad \bar{a} \in \overline{\mathbb{Z}}_n \quad (2.50)$$

$$= \sum_{a \in \mathbb{Z}_n} f(a) \langle \chi(-a), \bar{a} \rangle \quad (2.51)$$

$$= \sum_{a \in \mathbb{Z}_n} f(a) e^{-2\pi j\bar{a}a/n} \quad (2.52)$$

The scaling in this definition of the DFT has intentionally been chosen to be consistent with the DFT implementation within FFTW¹ to simplify the translation to software. This abstract view of the Fourier transform is useful because it motivates replacing the underlying groups to establish new mathematical objects with similar properties, and this idea will be exploited to find a natural extension of the observed samples to a continuous function. The completely general construction is provided in Rudin [159] theorem 1.2.2, and details associated with the discussion that follows can be found in B.

From a group-theoretic perspective, the original aim of identifying discrete representations to continuous bandlimited signals with a set of convenient properties becomes much simpler. Let \mathbf{f}_n be the vector of measurements, and a discrete representation of a continuous bandlimited signal $f_c(t)$ through the mapping $M_n : \mathbb{C}(\mathbb{Z}_n) \mapsto \mathbb{C}(\mathbb{T})$. Let $\mathbb{R}(\mathbb{Z}_n)$ be the real-valued members of $\mathbb{C}(\mathbb{Z}_n)$, and similarly, $\mathbb{R}(\mathbb{T})$ the real-valued members of $\mathbb{C}(\mathbb{T})$. Then the following desirable properties will drive our construction of the continuous signal extension.

1. $M_n(\mathbb{R}^n) \subset \mathbb{R}(\mathbb{T})$

¹<http://www.fftw.org>

2. $\forall \mathbf{f}_n \in \mathbb{C}^n, \exists \mathbf{g} \in \mathbb{C}^n : M_n(\mathbf{g}) = \int_0^t [(M_n(\mathbf{f}))](s) ds$
3. $\forall \mathbf{f}_n \in \mathbb{C}^n, \exists \mathbf{g} \in \mathbb{C}^n : M_n(\mathbf{g}) = \frac{\partial}{\partial t} [(M_n(\mathbf{f}))](t)$
4. $\forall m > n, M_n(\mathbb{C}^n) \subset M_m(\mathbb{C}^m)$

Property (1) says that real representations always map to real continuous extensions, while the (2) and (3) ensure that the space of representable signals is closed under integration and differentiation. Property (4) requires the finite signal space to form a flag. Specifically, given two representations \mathbf{f}_m and \mathbf{g}_n satisfying (4), the upsampling operator is defined to be $U_{m,n} : U_{m,n} \circ \mathbf{f}_m = \mathbf{g}_n$. An application of the Pontryagin duality theorem to the previous abstract description of the Fourier transform immediately suggests the natural form of M_n , and with one caveat, the desired properties will all follow from basic results in Fourier analysis. Specifically, for all representable continuous functions $f_c(t)$, let n^* be the smallest integer such that $f_c(t) \in M_{n^*}(\mathbb{C}^{n^*})$, then properties (1)-(3) will not hold when n^* is even, and $f_c(t)$ is represented by $\mathbf{f} \in \mathbb{C}^{n^*}$. This situation, however, can be easily remedied by applying property (4) and the upsampling operator to express $f_c(t)$ through a slightly larger finite representation.

Transform	Domain	Dual
Discrete Fourier Transform	\mathbb{Z}_n	\mathbb{Z}_n
Fourier series	\mathbb{T}	\mathbb{Z}
Discrete-time Fourier transform	\mathbb{Z}	\mathbb{T}
Fourier transform	\mathbb{R}	\mathbb{R}

Table 2.1: Relationship between common instantiations of the Fourier transform and the group structure of their corresponding domains.

The Pontryagin duality theorem states that every locally compact abelian group is the dual of its dual group, and as a consequence, a discrete group implies a compact dual and visa versa. Table 2.1 indicates commonly used versions of the Fourier transform as well as the group structures associated with their domain and dual. To treat \mathbf{f}_n as samples from a continuous signal, and maintain basic properties of the DFT in the continuous extension, one must then support it on \mathbb{Z} in the dual domain. The discrete Fourier series has precisely this property, implying that the association between the finite observations and their continuous extension can be viewed in the dual domain as the mapping $\gamma : \mathbb{Z}_n \mapsto \mathbb{Z}$. Thus the structure of the natural continuous extension is then given by

$$f_c(t) = \frac{1}{n} \sum_{\bar{a}=0}^{n-1} (F_n f_n)(\bar{a}) e^{2\pi j \gamma(\bar{a}) \frac{t}{T}}. \quad (2.53)$$

To help identify γ , first consider $n = 1$. For property (1) to hold given (2.53) it immediately follows that $\gamma(0) = 0$. More generally, the conjugate symmetry of the DFT of real-valued signals in conjunction with property (1) implies that the set $\gamma(\mathbb{Z}_n)$ must be closed under the modular additive inverse. The necessary use of a modular additive inverse is made clearer by considering $f_c(t)$ evaluated at the observed samples

$$f_c\left(\frac{aT}{n}\right) = \frac{1}{n} \sum_{\bar{a}=0}^{n-1} (F_n f_n)(\bar{a}) e^{2\pi j \frac{\gamma(\bar{a})}{n} a} \quad (2.54)$$

By choosing a minimum-bandwidth description of the observations, closure under the modular additive inverse can be replaced with the usual additive inverse and the constraint that $|\gamma(\bar{a})| \leq n/2$. Notice, however, that the Nyquist frequency component is its own additive inverse, implying that property (1) cannot be met under the construction given by (2.53) because this frequency component lacks a conjugate pair. Addressing this issue by allocating the measured at the Nyquist frequency between the corresponding conjugate pair in \mathbb{Z} will be addressed in the next section. This fully specifies γ up to the sign of the Nyquist frequency for n even. This choice is arbitrary because the exponent in (2.54) will be one of $\pm\pi ja$, however, this work will follow the engineering convention of treating even length signals as having more negative than positive frequency components. Thus γ is given by

$$\{\gamma(\bar{a})\}_{\bar{a}=0}^{n-1} = \left\{ 0, 1, \dots, \left\lfloor \frac{n-1}{2} \right\rfloor, -\left\lfloor \frac{n}{2} \right\rfloor, -\left\lfloor \frac{n}{2} \right\rfloor + 1, \dots, -1 \right\}, \quad (2.55)$$

which is simply the identity mapping where those values greater than $n/2$ have been replaced with their value less n . Substituting this definition of γ and (2.52) into (2.54), one can easily verify that f_c interpolates $\{(t_i, f_i)\}_{i=1}^n$

$$f_c\left(\frac{aT}{n}\right) = \frac{1}{n} \sum_{\bar{a}=0}^{n-1} \sum_{a' \in \mathbb{Z}_n} f(a') e^{-2\pi j \frac{\bar{a}}{n} a'} e^{2\pi j \frac{\gamma(\bar{a})}{n} a} \quad (2.56)$$

$$= \sum_{a' \in \mathbb{Z}_n} f(a') \frac{1}{n} \sum_{\bar{a}=0}^{n-1} e^{-2\pi j \frac{\bar{a}}{n} a'} e^{2\pi j \frac{\gamma(\bar{a})}{n} a} \quad (2.57)$$

$$= \sum_{a' \in \mathbb{Z}_n} f(a') \frac{1}{n} \sum_{\bar{a}=0}^{n-1} \left\langle e^{2\pi j \frac{\gamma(\bar{a})}{n} a}, e^{2\pi j \frac{\bar{a}}{n} a'} \right\rangle \quad (2.58)$$

$$= \sum_{a' \in \mathbb{Z}_n} f(a') \frac{1}{n} \sum_{\bar{a}=0}^{n-1} \left\langle e^{2\pi j \frac{\bar{a}}{n} a}, e^{2\pi j \frac{\bar{a}}{n} a'} \right\rangle \quad (2.59)$$

$$= f(a), \quad (2.60)$$

where (2.59) follows from (2.58) by the fact that $\gamma(\bar{a}) \equiv \bar{a} \pmod{n}$, and the desired result is a consequence of the complex exponentials forming an orthonormal basis.

For n odd, one can easily verify that properties (1)-(4) hold under the continuous extension defined through (2.53) and (2.55). When n is even, one can see from the magnitudes of (2.55) that introducing an additional term will not alter the bandwidth of the continuous extension. Therefore, the problem of identifying the desired continuous extension reduces to the problem of determining the upsampling operator $U_{n,n+1}$ for n even.

Let $w(\bar{a}) = 2\pi\gamma(\bar{a})t/T$ be the frequencies associated with the Fourier series of a real function f_c with two discrete representations of length n and m , with $m > n$ and n even. This situation is illustrated in Figure 2.19. Equating the continuous basis functions from (2.53),

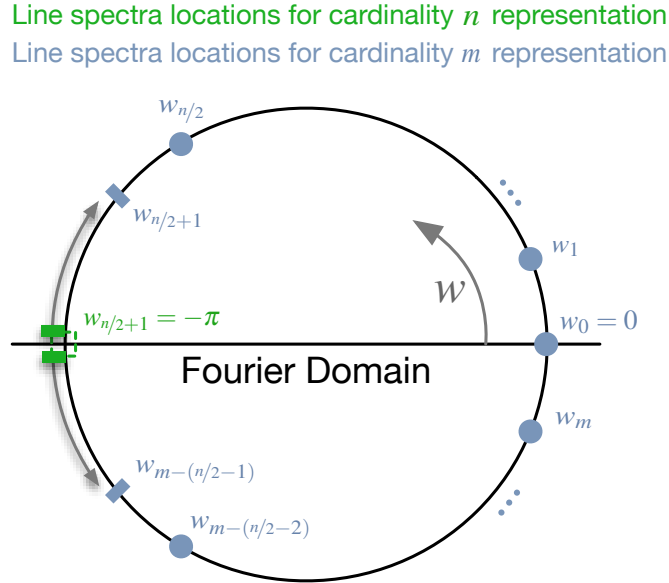


Figure 2.19: Illustration of the splitting of the energy at the line spectra corresponding to $w^{(n/2 + 1)}$ for a length n representation of $f_c(t)$ (n even) when upsampling to $m > n$

one finds that the two discrete representations correspond to a common $f_c(t)$ if and only if they each sample the continuous function uniformly over T ($T = n\Delta_n = m\Delta_m$), and the corresponding DFT coefficients have the following relationships

$$(F_m f)(\bar{a}) = (F_n f)(\bar{a}) \quad \forall |\gamma(\bar{a})| < n/2 \quad (2.61)$$

$$\sum_{|\gamma(\bar{a})|=n/2} (F_m f)(\bar{a}) = (F_n f)(n/2 + 1) \quad (2.62)$$

$$(F_m f)(\bar{a}) = 0 \quad \forall |\gamma(\bar{a})| > n/2 \quad (2.63)$$

This implies that for the two finite representations to correspond to the same continuous

function, the upsampling operator must appropriately zero-pad the shorter signal in the Fourier domain with an exception at the Nyquist frequency. Let the two terms on the LHS of (2.62) be denoted by $c_1 = r_1 e^{j\theta_1}$ and $c_2 = r_2 e^{j\theta_2}$. Then from (2.62), $c_1 + c_2 = (F_n f)(n/2 + 1)$. Hermitian symmetry implies $r_1 = r_2$ and $\theta_1 \equiv -\theta_2 \pmod{2\pi}$, and combining these two constraints results in a single remaining degree of freedom.

$$c_1 = r_1 e^{j\theta_1} \text{ and } c_2 = r_1 e^{-j\theta_1} \quad \text{s.t.} \quad r_1 = \frac{(F_n f)(n/2 + 1)}{2\cos(\theta_1)} \quad (2.64)$$

A natural additional constraint is to require that both the continuous extension and the observations contain the same amount of energy at the Nyquist frequency. Notice that the orthogonality of the characters (2.49) implies that this same constraint will hold for every other frequency. This additional constraint implies $|c_1 + c_2|^2 = |c_1|^2 + |c_2|^2$, from which one finds that $2\Re\{c_1 c_2^*\} = 0$, or equivalently, that $\theta_1 - \theta_2 \equiv \frac{\pi}{2} \pmod{\pi}$. Incorporating this new constraint determines the continuous representation up to a single bit.

$$r_1 = r_2 = \frac{(F_n f)(n/2 + 1)}{\sqrt{2}} \quad (2.65)$$

$$(\theta_1, \theta_2) \in \left\{ \left(\frac{\pi}{4}, -\frac{\pi}{4} \right), \left(-\frac{\pi}{4}, \frac{\pi}{4} \right) \right\} \quad \text{when } \theta_1, \theta_2 \in [-\pi, \pi) \quad (2.66)$$

$$\therefore c_1 = (F_n f)(n/2 + 1) \frac{1 \pm j}{2}, \quad c_2 = (F_n f)(n/2 + 1) \frac{1 \mp j}{2} \quad (2.67)$$

These two solutions correspond to interpreting the frequency components that aliased to the Nyquist frequency as corresponding to the basis function $\cos(\pi a) - \sin(\pi a)$ (positive sign in c_1), or $\cos(\pi a) + \sin(\pi a)$. Figure 2.20 shows 10 randomly drawn real values, and the two real continuous functions that are consistent with these 10 samples assuming they corresponded to a real, critically sampled, continuous function with period 11. Obviously the generation of this figure did not involve forming the entire continuous function, however, the upsampling operator allows one to compute a set of $m > n$ points that exactly coincide with samples from the continuous function as given by the Fourier synthesis equation acting on the DFT coefficients corresponding to the length n representation. For the remainder of this work, the positive sign will be taken in c_1 .

Under this signal model, the upsampling operator $U_{m,n}$ is linear and precisely equivalent to the sampling operator \mathbf{S} when upsampling. Let \mathbf{S}_U be the matrix representation of the upsampling operator, and \mathbf{F}_n the DFT matrix associated with a signal of length n , then

$$\mathbf{S}_U = \frac{M}{N} \mathbf{F}_M^H \mathbf{U} \mathbf{F}_N, \quad (2.68)$$




Figure 2.20: A real, discrete, signal sampled critically with an even number of samples along with both of its consistent implied continuous functions.

The structure of \mathbf{U} is indicated in Figure 2.21, where the identity matrices are replicating the existing spectral lines, and the familiar zero-padding is explicitly called out as a band of zeros. Notice that the “ $n \bmod 2$ ” is simply 1 when n is odd and 0 otherwise, and the general form of (c_1, c_2) from (2.67) is given for both odd and even length signals by

$$(c_1, c_2) = \begin{cases} \left(\frac{1+j}{2}, \frac{1-j}{2} \right) & N \text{ even} \\ (0, 0) & N \text{ odd} \end{cases}. \quad (2.69)$$

The adjoint follows immediately from the transpose of this matrix form. Downsampling, and the multi-dimensional extensions, all follow in a straightforward manner from classical signal processing concepts [48, 146, 147].

Figure 2.22 illustrates the use of these concepts to alias a two dimensional image. The sharp edges most clearly exhibit the aliasing phenomena, however, one can also see that impact of aliasing is predominately a local effect. The function `fftResampleN` implements periodic, band-limited resampling with and without aliasing (the sampling operator) as well as the associated derivatives and adjoints for N-Dimensional signals.

```

1 % Simple aliasing demonstration
2 load('cameraMan','X');
3 imgAlias = fftResampleN(X,[100 100],'-alias');
```




Figure 2.21: Cartoon illustrating the key blocks in the matrix representation of the upsampling operator \mathbf{U} in the Fourier domain. The block of zeros corresponding to the familiar zero-padding in the Fourier domain is explicitly called out.

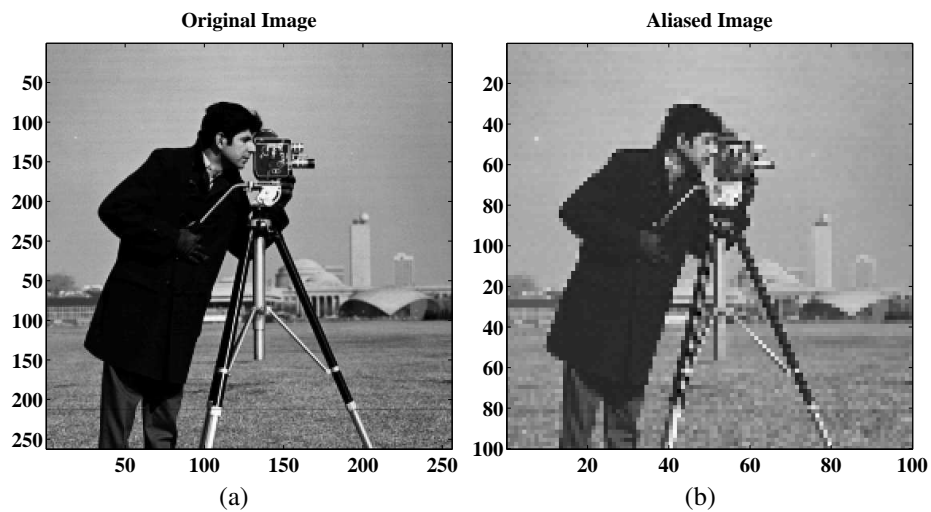


Figure 2.22: Action of the sampling (aliasing) operator taking an image from (a) 256×256 to (b) 100×100

2.5 Statistical Inverse-Problem Approach

This work takes a statistical inverse-problem approach to estimating sensor and environmental parameters from measured data. The estimation questions are said to be part of an inverse-problem because information regarding the unknown parameters is only observed through its influence on the output of a forward process; the data. A statistical inverse-problem approach to parameter estimation is one which properly accounts for all of the available sources of information in the generation of the estimates. This includes knowledge of the forward processes, the types and sources of statistical uncertainty, available prior information, and the data itself.

Strictly speaking, a model of the physical process is inverted, rather than the process itself. This model is known as the forward model, and in most modern inverse-problem approaches, is actually composed of two models. Using the terminology of Rosenblueth and Wiener [157], the forward model always includes a formal or intellectual model. This “symbolic assertion in logical terms” is an abstract description of the forward process such as that described in Sections 2.2 and 2.4. When such a model does not permit a closed form solution, a material model is also needed. Today, these are generally computer simulations. The associated material model for this work has been implemented in the computer language MATLAB[®], and key components corresponding to the formal model have been demonstrated throughout this chapter. When closed form solutions are unavailable, the inverse-problem solution is expressed as an optimization problem in terms of the material model. This situation is illustrated in Figure 2.23.

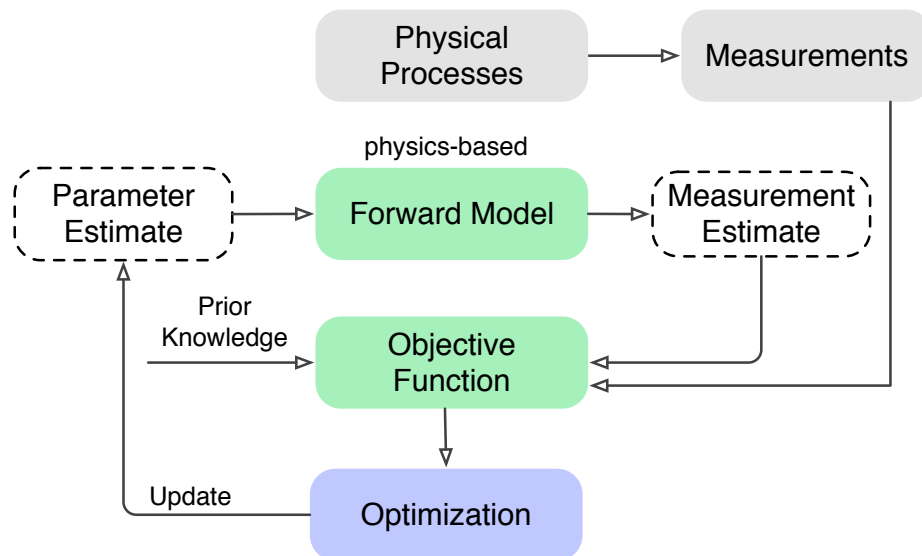


Figure 2.23: Inverse problem flow diagram showing the relationship between modeling components (green) and physical processes (gray)

All inverse-problems share this common structure, even when they are not explicitly stated in such formal terms. The benefit of expressing inverse-problems within a formal framework is that it more clearly delineates the underlying assumptions and problem structure. Outside of the pedagogical benefits, this modular approach simplifies the application of standard statistical and optimization techniques for studying inverse-problem solutions. To clarify this point, consider the distinction between prior information and regularization. Prior information, or a priori information about the distribution of an unknown parameter, typically results in an additional penalty term in the objective function that expresses information gained through an assumption that is believed to be true. The proposed framework readily accommodates the incorporation of such information. Regularization, on the other hand, is any alteration of the structure shown in Figure 2.23 for the purpose of changing the posterior distribution of the unknowns. Regularization is often used to generate point estimates in ill-posed problems; those which are not fundamentally well described by point estimates (c.f. [151]).

Unlike prior information, which makes a clear and scientifically verifiable assertion about the unknowns, regularization is often subtle, and indicative of an inverse problem with open questions. As an example, consider the common practice of using stopping criteria in the optimization stage of inverse-problems. Sometimes these criteria are simply an acknowledgment of the diminishing returns associated with increased estimate precision, and yet other times they are used to fundamentally alter the distribution of the estimates. In the latter case, application of standard statistical tools becomes difficult, and regularization obfuscates underlying questions regarding the forward model and/or objective function. Because of these difficulties, regularization should only be used when its impact on the posterior estimate distribution is well understood.

For the specific problem of camera system-identification from a known calibration target, one seeks to estimate the camera's aggregate wavefront imperfections (blur) as described by their Zernike coefficients in the pupil $\boldsymbol{\alpha}$. This estimation is performed subject to a set of unknown nuisance parameters $\{c_1, c_2, c_3, c_4, \sigma_r, \mathbf{v}\}$. The variables c_1 and c_2 accommodate an unknown radiance level emanating from the object, c_3 , c_4 , and σ_r account for the unknown detector gain, offset, and read-noise level respectively, and the vector parameter \mathbf{v} accounts for the unknown imaging geometry between the camera and the calibration target. Assuming the data is a single image of the calibration target, one can expand (2.21) in terms of the operators described in Section 2.4 to express the statistical distribution of the lexicographically ordered data under the formal model as

$$d = c_3 \mathcal{P}\{TSB(\boldsymbol{\alpha}) \mathbf{W}(\mathbf{v})(c_1 \mathbf{f} + c_2)\} + c_4 + \mathcal{N}(0, \sigma_r^2) \quad (2.70)$$

Grouping the unknowns into a single vector $\boldsymbol{\theta}$, the formal description of the maximum likelihood estimate can be written in terms of the log-likelihood function $\ell(\cdot)$ as

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} -\ell(\mathbf{d}|\boldsymbol{\theta}) \quad \text{s.t.} \quad (2.71)$$

$$c_1 \geq 0, c_2 \geq 0, c_3 \in [0, 1], \sigma_r > 0. \quad (2.72)$$

This formal statement of the solution is itself an optimization problem in terms of a material model. The next two chapters address the practical challenges of finding such solution(s), and their study their properties.

2.6 Non-Linear Optimization

This section provides a brief introduction to large-scale unconstrained optimization, with special attention paid to preconditioning and solver initialization. Such techniques are invaluable in engineering, and the estimators presented in this work will be generated using this class of techniques. More details regarding this approach to large-scale, non-convex optimization may be found in [135, Chapters 2-3,6,8-9] and the references therein.

Given an objective function mapping a set of variables onto \mathbb{R} , one is often interested in finding

$$\mathbf{x}_* = \arg \min_{\mathbf{x}} f(\mathbf{x}) \quad \text{s.t.} \quad \mathbf{x} \in \mathbb{R}^N. \quad (2.73)$$

Under mild conditions on f it can be shown that \mathbf{x}_* is locally unique. This section is concerned with efficiently determining a sequence $\{\mathbf{x}_k\}$ such that $\lim_{k \rightarrow \infty} \{\mathbf{x}_k\} = \mathbf{x}_*$ given \mathbf{x}_0 sufficiently close to \mathbf{x}_* . Assuming $f \in \mathbb{C}^2$, a necessary condition on \mathbf{x}_* is

$$\mathbf{g}(\mathbf{x}_*) \stackrel{\text{def}}{=} \nabla f(\mathbf{x}_*) = 0 \quad (2.74)$$

$$\mathbf{G}(\mathbf{x}_*) \stackrel{\text{def}}{=} \nabla^2 f(\mathbf{x}_*) : \mathbf{s}^T \mathbf{G}(\mathbf{x}_*) \mathbf{x} \geq 0 \quad \forall \mathbf{s} \in \mathbb{R}^N \quad (2.75)$$

Thus the local derivative is zero (\mathbf{x}_* is a stationary point) and the curvature is positive in any direction ($\mathbf{G}(\mathbf{x}_*)$ is positive definite).

Given the requirement of continuous 2nd derivative, it is natural to locally model the unknown function through its Taylor-series expansion because this simple expansion involves

the objective function value as well as both of the terms involved in the optimality criteria.

$$f(\mathbf{x}^{(k)} + \boldsymbol{\delta}) = f^{(k)} + \boldsymbol{\delta}^T \mathbf{g}^{(k)} + \frac{1}{2} \boldsymbol{\delta}^T \mathbf{B}^{(k)} \boldsymbol{\delta} \quad \text{s.t.} \quad (2.76)$$

$$\mathbf{g}^{(k)} \stackrel{\text{def}}{=} \nabla f(\mathbf{x}^{(k)}) \quad (2.77)$$

$$\mathbf{B}^{(k)} \stackrel{\text{def}}{=} \nabla^2 f(\mathbf{x}^{(k)}) \quad (2.78)$$

$$\approx \mathbf{G}^{(k)}(\mathbf{x}^{(k)}) \quad (2.79)$$

where the superscript k indicates the k^{th} element in a sequence of evaluations. This has historically been the approach to such problems (citations here), and remains so to this day (citations here). Though many variations exist, nearly all approaches share the same basic structure shown in Figure 2.24

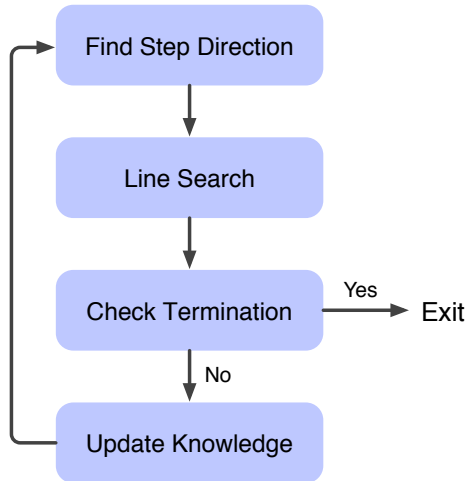


Figure 2.24: Basic structure of quasi-Newton optimization methods

2.6.1 Quasi-Newton Direction Updates

For high-dimensional problems, (2.76) becomes prohibitively expensive to evaluate due to the large amount of memory needed to represent the Hessian, however, this same equation clearly couples changes in $\boldsymbol{\delta}$ (a change in \mathbf{x}) to changes in $\boldsymbol{\gamma}$ (a change in \mathbf{g}). This is known as the secant condition, and is given by

$$\boldsymbol{\gamma}^{(k)} = \bar{\mathbf{B}}^{(k)} \boldsymbol{\delta}^{(k)} \quad (2.80)$$

$$\bar{\mathbf{B}}^{(k)} = \int_0^1 \mathbf{B}(\mathbf{x}^{(k)} + \theta \boldsymbol{\delta}^{(k)}) d\theta.$$

In other words, one should be able to infer from the derivatives how to move in the high-dimensional space without explicitly forming the entire Hessian. In fact, when we implicitly update the Hessian approximation through a set of rank-two updates using previous observations of the function value and gradient, what results is the Broyden update family, given by

$$\mathbf{B}^{(k+1)} = \mathbf{B}^{(k)} - \frac{\mathbf{B}^{(k)} \boldsymbol{\delta}^{(k)} \left(\boldsymbol{\delta}^{(k)} \right)^T \mathbf{B}^{(k)}}{\left(\boldsymbol{\delta}^{(k)} \right)^T \mathbf{B}^{(k)} \boldsymbol{\delta}^{(k)}} + \frac{\boldsymbol{\gamma}^{(k)} \left(\boldsymbol{\gamma}^{(k)} \right)^T}{\left(\boldsymbol{\delta}^{(k)} \right)^T \boldsymbol{\gamma}^{(k)}} + \theta^{(k)} \mathbf{w}^{(k)} \left(\mathbf{w}^{(k)} \right)^T \quad (2.81)$$

$$\mathbf{w}^{(k)} = \sqrt{\left(\boldsymbol{\delta}^{(k)} \right)^T \mathbf{B}^{(k)} \boldsymbol{\delta}^{(k)}} \left(\frac{\boldsymbol{\gamma}^{(k)}}{\left(\boldsymbol{\delta}^{(k)} \right)^T \boldsymbol{\gamma}^{(k)}} - \frac{\mathbf{B}^{(k)} \boldsymbol{\delta}^{(k)}}{\left(\boldsymbol{\delta}^{(k)} \right)^T \mathbf{B}^{(k)} \boldsymbol{\delta}^{(k)}} \right). \quad (2.82)$$

Using these expressions, one can define the change in \mathbf{x} implied by the local structure which has been itself implied by previous observations. Thus all we need to do is be able to differentiate the objective function with respect to the unknowns, and apply the action of this transformation to determine the next search direction. This is not a trivial task by any means, but it can be done, and what results is an efficient approach for solving non-linear optimization problems over high-dimensional spaces. The popular BFGS technique is a member of the Broyden family of updates with $\theta^{(k)} = 0$ in (2.81). The method is named after Charles Broyden, Roger Fletcher, Donald Goldfarb and David Shanno who independently discovered the approach. A graphical example of convergence using the Broyden update is shown in Figure 2.25. In this example each dot represent a point evaluated by the solver,




Figure 2.25: An example of the BFGS technique applied to a low-dimensional non-linear problem.

where unconnected points indicate those rejected by the line-search. The total number of function evaluations can typically be dramatically reduced through preconditioning; applying problem specific knowledge to the search procedure.

2.6.2 Preconditioning and Transformation Invariance

It is desirable to find optimization methods that are invariant under affine transformations of the search space. This property not only makes intuitive sense (i.e. redefining a search variable from meters to centimeters should not fundamentally alter how the search is performed), but is also a prerequisite for “preconditioning”. Preconditioning is the process of altering the objective function such that the surrogate function can be more easily minimized while still relating in a known way to the original objective function. To begin, consider the set of affine transformations

$$T : \mathbb{R}^N \rightarrow \mathbb{R}^N \quad (2.83)$$

$$\tilde{\mathbf{x}} = T(\mathbf{x}) \quad (2.84)$$

$$= \mathbf{A}(\mathbf{x} - \mathbf{b}) \quad (2.85)$$

where \mathbf{A} is invertible. Then one gets

$$\mathbf{x} = T^{-1}(\tilde{\mathbf{x}}) \quad (2.86)$$

$$= \mathbf{A}^{-1}\tilde{\mathbf{x}} + \mathbf{b} \quad (2.87)$$

$$\tilde{f}(\tilde{\mathbf{x}}) = f(T^{-1}(\tilde{\mathbf{x}})) \quad (2.88)$$

$$= f(\mathbf{x}) \quad (2.89)$$

where $\tilde{\mathbf{x}}$ and \tilde{f} are the transformed (preconditioned) vector and objective function respectively. The relationship of the derivatives of this new function to the original can be found by differentiating.

$$\nabla f(\mathbf{x}) = \mathbf{A}^T \nabla \tilde{f}(\tilde{\mathbf{x}}) \quad (2.90)$$

$$\nabla^2 f(\mathbf{x}) = \mathbf{A}^T \nabla^2 \tilde{f}(\tilde{\mathbf{x}}) \mathbf{A} \quad (2.91)$$

A search method is defined to be invariant w.r.t. T iff

$$\tilde{\mathbf{x}}_k = T(\mathbf{x}_k) \implies \tilde{\mathbf{x}}_{k+1} = T(\mathbf{x}_{k+1}) \quad \forall k \quad (2.92)$$

Preconditioning of such a method then becomes a matter of finding transformations that minimized the differences between the presumed initial hessian (typically I), and that of the surrogate objective function. Preconditioning is also useful because it can be used to remove a subspace of a problem. This happens when the natural parameterization leads to an invariance which can be removed prior to solving. Doing so lowers the effective search dimension, simplifying the overall optimization problem. The previous discussion can be extended to handle this case by redefining the class of transformations as follows

$$T : \mathbb{R}^N \rightarrow \mathbb{R}^M \quad M \leq N \quad (2.93)$$

$$\tilde{\mathbf{x}} = T(\mathbf{x}) = \mathbf{A}(\mathbf{x} - \mathbf{b}) \quad (2.94)$$

and the inverse operator is defined as the minimum L_2 norm expansion about \mathbf{b} . Therefore,

$$T^{-1} : \mathbb{R}^M \rightarrow \mathbb{R}^N \quad (2.95)$$

$$T^{-1}(\tilde{\mathbf{x}}) \stackrel{\text{def}}{=} \mathbf{A}_L^\dagger \tilde{\mathbf{x}} + \mathbf{b} \quad (2.96)$$

where \mathbf{A}_L^\dagger and \mathbf{A}_R^\dagger are the minimum norm left and right inverses respectively

$$\underbrace{\mathbf{A}\mathbf{A}^T}_{\mathbf{A}_R^\dagger} (\mathbf{A}\mathbf{A}^T)^\dagger = \mathbf{I}_M \quad (2.97)$$

$$\mathbf{A}_R^\dagger \stackrel{\text{def}}{=} \text{Right Inverse}$$

$$\underbrace{(\mathbf{A}^T \mathbf{A})^\dagger}_{\mathbf{A}_L^\dagger} \mathbf{A}^T \mathbf{A} = \mathbf{I}_N. \quad (2.98)$$

$$\mathbf{A}_L^\dagger \stackrel{\text{def}}{=} \text{Left Inverse}$$

The fact that this is a minimum norm solution about \mathbf{b} follows from the fact that $\mathbf{A}_L^\dagger \mathbf{A}$ is a projection operator onto the complement of the null-space of \mathbf{A} . Applying the inverse to the forward operator one gets

$$T^{-1}(T(\mathbf{x})) = \mathbf{A}_L^\dagger \mathbf{A}\mathbf{x} - \mathbf{A}_L^\dagger \mathbf{A}\mathbf{b} + \mathbf{b} \quad (2.99)$$

$$= \mathbf{A}_L^\dagger \mathbf{A}\mathbf{x} + \left(\mathbf{I} - \mathbf{A}_L^\dagger \mathbf{A} \right) \mathbf{b} \quad (2.100)$$

which is the projection of \mathbf{x} into the null-space complement of \mathbf{A} , with the projection of \mathbf{b} into the null-space of \mathbf{A} added back. This obviously minimizes the norm of the representation of \mathbf{x} through $\tilde{\mathbf{x}}$ about \mathbf{b} when \mathbf{A}_L^\dagger is defined to be the minimum L_2 projector. When \mathbf{A}_L^\dagger is not uniquely defined, one zeros the left singular values of any matrix that satisfies the left-inverse definition over the null-space complement.

The aforementioned relationships between the derivatives of the original objective function and its surrogate remain unchanged, however, it is more natural to think of the surrogate in terms of the original function because it is of a lower dimension. For this reason T is referred to as the “preconditioner,” because it moves the problem into a space with a dimensionality less than or equal to the original problem while simultaneously normalizing the hessian. Thus one may write

$$\nabla \tilde{f}(\tilde{\mathbf{x}}) = \mathbf{A}_R^{-T} \nabla f(\mathbf{x}) \quad (2.101)$$

$$\nabla^2 \tilde{f}(\tilde{\mathbf{x}}) = \mathbf{A}_R^{-T} \nabla^2 f(\mathbf{x}) \mathbf{A}_R^\dagger \quad (2.102)$$

The hessian is symmetric and non-negative definite, implying that its SVD may be written

$$\nabla^2 f(\mathbf{x}) = \mathbf{V} \mathbf{\Sigma} \mathbf{V}^T \quad (2.103)$$

where $\mathbf{V} \in \mathbb{R}^{N \times M}$ and $\mathbf{\Sigma} \in \mathbb{R}^{M \times M}$. Equating the result to 2.91 and setting $\nabla^2 \tilde{f}(\tilde{\mathbf{x}}) = I$ results in

$$\mathbf{A} = \mathbf{\Sigma}^{-\frac{1}{2}} \mathbf{V}^T \quad (2.104)$$

$$\therefore \mathbf{A}_R^{-T} = \mathbf{\Sigma}^{\frac{1}{2}} \mathbf{V}^T \quad (2.105)$$

In practice we always want to work in the space with the lowest possible dimension, and this implies that the most natural way to specify a preconditioner is as the transformation from the unconditioned space to the preconditioned one. Equation (2.105) suggests the rotational invariance of the preconditioned space and the utility of the SVD representation. For example, when only \mathbf{A} is provided one can write

$$\mathbf{A} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T \implies \mathbf{A}^\dagger = \mathbf{V} \mathbf{\Sigma}^{-1} \mathbf{U}^T \quad (2.106)$$

For the purposes of preconditioning, \mathbf{U} is unimportant because an arbitrary rotation will not effect the solver. In practice, however, one must be careful to properly rotate the starting value \mathbf{x}_0 within the preconditioned space.

2.6.3 Initial Step

The problem of initial step length selection arises in a number of contexts during a search. The first is truly an initial step, and the problem of interest may be formulated as

$$\text{Given: } \{\mathbf{x}_0, f_0, \mathbf{g}_0\} \quad (2.107)$$

$$\text{Find: } \hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \|\mathbf{x} - \mathbf{x}_*\|_2^2 \quad (2.108)$$

The standard solution to this problem has been motivated by an assumption of “perfect preconditioning”, that is, to assume that the hessian is the twice identity matrix. Making this assumption results in

$$\hat{\mathbf{x}} = -\frac{1}{2}\mathbf{g}_0 \quad (2.109)$$

To see why this is the case, consider the minimizer of the quadratic approximation of f about \mathbf{x}_0

$$f(\mathbf{x}_k + \delta) \approx q_k(\delta) = f_k + \delta^T \mathbf{g}_k + \frac{1}{2} \delta^T \mathbf{B}_k \delta \quad (2.110)$$

Assuming \mathbf{G}_k is Positive Definite (P.D.), one can find the minimizer by differentiating with respect to δ and setting the result equal to zero. Doing so results in

$$\mathbf{0} = \mathbf{g}_0 + \frac{1}{2} (\mathbf{G}_0 + \mathbf{G}_0^T) \delta \quad (2.111)$$

$$\therefore \delta = -2 (\mathbf{G}_0 + \mathbf{G}_0^T)^{-1} \mathbf{g}_0 \quad (2.112)$$

This implies that “perfect conditioning” is being defined as a problem where the preconditioner causes the hessian to be $2\mathbf{I}_{n \times n}$, or perhaps more initiatively, causes the leading coefficient of the quadratic term to be 1. This implies that the gradient both points in the correct direction, and does so with a magnitude that is related to the distance to \mathbf{x}_* in a known way.

Let’s now consider two relaxations on this rather stringent assumption based on the use of “non-expert” advice. Here “non-expert” advise is defined as advise that will likely be given by someone not knowledgeable in optimization theory or the underlying structure of their particular objection function. To do this, we will view finding the initial step as an

estimation problem under a stochastic framework. That is

$$\text{Given } \{\mathbf{x}_0, \gamma\}, \{f_0, \mathbf{g}_0\} \quad (2.113)$$

$$\text{Find } \hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \mathbb{E} \left[\|\hat{\mathbf{x}} - \mathbf{x}_*\|_2^2 \right] \quad (2.114)$$

where γ is some non-expert supporting information. The initial function value and gradient are separated to emphasize the fact that they are inferred by the algorithm *after* $\{\mathbf{x}_0, \gamma\}$ are given. To begin, assume the initial step should be made along the steepest-descent direction.

$$\hat{\mathbf{x}} = \mathbf{x}_0 + \alpha(-\mathbf{g}_0) \quad \text{s.t.} \quad \alpha > 0 \quad (2.115)$$

The risk associated with this step is given by

$$\mathbb{R}[\alpha, \mathbf{x}_*] = \alpha^2 \mathbf{g}_0^T \mathbf{g}_0 + 2\alpha \mathbf{g}_0^T \mathbb{E}[(\mathbf{x}_* - \mathbf{x}_0)] + \mathbb{E}[\|\mathbf{x}_* - \mathbf{x}_0\|_2^2] \quad (2.116)$$

The assumption that the $-\mathbf{g}_0$ is a “good” direction implies that it points into the quadrant of \mathbf{x}^* and thus

$$\text{sgn}(-\mathbf{g}_0)_{[i]} = \text{sgn}(\mathbb{E}[\mathbf{x}_* - \mathbf{x}_0])_{[i]} \quad \forall i \quad (2.117)$$

$$\therefore (-\mathbf{g}_0)^T \mathbb{E}[\mathbf{x}_* - \mathbf{x}_0] = |\mathbf{g}_0|^T |\mathbb{E}[\mathbf{x}_* - \mathbf{x}_0]| \quad (2.118)$$

Using this fact, and differentiating w.r.t α we get

$$\alpha_{min} = -\frac{\mathbb{E}[\mathbf{g}_0^T (\mathbf{x}_* - \mathbf{x}_0)]}{\mathbf{g}_0^T \mathbf{g}_0} \quad (2.119)$$

$$= \frac{|\mathbf{g}_0|^T |\mathbb{E}[\mathbf{x}_* - \mathbf{x}_0]|}{\mathbf{g}_0^T \mathbf{g}_0} \quad (2.120)$$

$$\therefore \hat{\mathbf{x}} = \mathbf{x}_0 - \frac{|\mathbf{g}_0|^T |\mathbb{E}[\mathbf{x}_* - \mathbf{x}_0]|}{\mathbf{g}_0^T \mathbf{g}_0} \mathbf{g}_0 \quad (2.121)$$

This implies that a minimum risk step can be made if the user provides $\gamma = |\mathbb{E}[\mathbf{x}_* - \mathbf{x}_0]|$. The advise is “non-expert” in the sense that the quantity is very intuitive, “The expected deviation of each of the individual parameters.”² This initial step also provides a number of good properties; it never perturbs any of the parameters beyond their expected deviation, and the length of the step is automatically adjusted according to how well the relative scaling of the derivative corresponds to the relative scaling of the expected deviations.

²Users have great difficulty when the $|\cdot|$ is inside the expectation

Now let's revisit the same problem where we no longer trust the magnitude of the derivative, but rather trust only its sign.

$$\hat{\mathbf{x}} = \mathbf{x}_0 + \delta \quad \text{s.t.} \quad \text{sgn}(\delta_{[i]}) = -\text{sgn}(\mathbf{g}_{0[i]}) \quad \forall i \quad (2.122)$$

As before, we are concerned with minimizing the risk of the initial step

$$\mathbf{R}[\delta, \mathbf{x}_*] = \delta^T \delta - 2\delta^T \mathbf{E}[(\mathbf{x}_* - \mathbf{x}_0)] + \mathbf{E}[\|(\mathbf{x}_* - \mathbf{x}_0)\|^2] \quad (2.123)$$

Here we are defining “trust” in $\text{sgn}(\mathbf{g}_0)$ through the assumption

$$\text{sgn}(\mathbf{x}_* - \mathbf{x}_0)_{[i]} = \text{sgn}(-\mathbf{g}_0)_{[i]} \quad \forall i \quad (2.124)$$

Using this fact, and differentiating the risk w.r.t. δ gives

$$|\delta_{min}| = |\mathbf{E}[\mathbf{x}_* - \mathbf{x}_0]| \quad (2.125)$$

$$\therefore \hat{\mathbf{x}} = \mathbf{x}_0 - \text{sgn}(\mathbf{g}_0) \odot |\mathbf{E}[\mathbf{x}_* - \mathbf{x}_0]| \quad (2.126)$$

Once again we find that the same information is needed, and this identifies the largest step that could have been taken when we *did* trust the direction of the gradient.

CHAPTER III

System Identification with Known Objects

3.1 Introduction

Optical imaging systems are designed to produce data that is highly correlated with the scene being imaged, but the collected data is not the best possible representation of the physical scene. Diffraction caused by the finite sensing aperture, lens imperfections, and physical limitations of the detector, all result in deviations from an ideal camera-model. By understanding the physical relationships between the environment, the imaging system, and the collected data, one can more effectively use available measurements to make inferences about the sensor, the scene being imaged, or both. This chapter presents a systematic approach for making these types of inferences, as well as understanding the limitations imposed on them by uncertainties from the environment and the imaging system.

A statistical inverse-problem approach is presented for jointly estimating camera blur from aliased data of a known calibration target. The estimation is said to be part of an inverse-problem, because information regarding the unknown parameters is only observed through their influence on the output of a forward process; the data. A statistical inverse-problem approach to parameter estimation is one that properly accounts for all of the available sources of information in the generation of the estimates. This includes knowledge of the forward processes, the types and sources of statistical uncertainty, available prior information, and the data itself.

An accurate description of a system's incoherent Point-Spread Function (PSF) is fundamental to image reconstruction, especially reconstruction from aliased data [49, 127]. A PSF model based on a description of the optical system's wavefront aberrations is *useful* in image exploitation because it parsimoniously describes the space of PSF's likely to be encountered [136, 164], but rapidly becomes *essential* when one is interested in processing datasets that contain changes in these aberrations. Intentionally introducing such changes is the basis of phase-diversity imaging [75, 97, 142, 186], and unintentional changes serve as

the basis for depth-from-defocus techniques, certain 3D shape reconstruction algorithms, and even sophisticated image restoration approaches [58, 103, 120]. These application areas all motivate the characterization of imaging systems through a physical wavefront description of their blur.

Direct wavefront sensing approaches require modification, or partial disassembly, of the sensor [73], and this has led to interest in less invasive approaches. Figure 3.1 illustrates a recent attempt to measure a Nikon D7000 and Nikkor lens using a wavefront interferometer. The large number of surfaces and various optical coatings prevented the interferometer from




Figure 3.1: Failed attempt to measure a complex imaging system using an interferometer by retro-reflecting off the detector.

making a reliable measurement. Removing the camera body to measure only the lens was also unsuccessful because the lens automatically stops down its aperture when disconnected from the body. While disappointing, these complications serve to reinforce the value of techniques that can be applied using only measured data.

Blur estimation approaches that depend only on collected data can be broadly separated into blind and non-blind procedures. Blind techniques attempt to estimate the PSF directly from an unknown image, and often focus on per-collection variations such as camera shake or object motion [62, 167, 207, 211]. If the object is treated as the principle quantity of interest, these techniques are also referred to as blind deconvolution [25, 62, 67, 87, 99, 101, 106, 119, 156]. Blind approaches are naturally less effective at estimating system imperfections, but are necessary when one expects collection-dependent variations in the PSF. By contrast, non-blind methods have the potential to more accurately characterize the system through the exploitation of additional scene knowledge. This class of techniques has traditionally sought to exploit either sharp edges [36, 87, 101], or use random targets with

a known power-spectra [8, 21, 45, 115, 141, 144]. In this chapter, a physically motivated, statistical inverse-problem approach is taken toward system identification using known calibration objects.

The proposed approach differs from existing known-object calibration techniques in a number of ways. First, the camera’s PSF is parameterized through a series of physical blur sources under a Fraunhofer diffraction assumption. By contrast, most previous work has not modeled diffraction effects [8, 21, 36, 45, 87, 101, 115, 141, 144, 213], or modeled only defocus [178]. Second, the proposed approach explicitly models aliasing. Modern commercial devices are often aliased, and previous approaches have either ignored this effect [8, 21, 87, 115, 141, 144, 178, 213], or address it in an incomplete way [36, 45, 101]. Joshi [101] and Delbracio [45] both explicitly model aliasing, but only accommodate integer downsampling factors. Third, the proposed technique explicitly models key detector effects. To the author’s knowledge, no previous work has explicitly addressed the impact of Bayer filters, and very few [178] have accommodated the intensity variant uncertainty that is inherent when measuring electromagnetic fields using semi-conductors [64]. Taken together, the resulting high-fidelity forward-model accurately reflects the physics of imaging with modern digital devices, and accommodates the main sources of uncertainty in the resulting data.

The remainder of this chapter is organized as follows. In Section 3.2 a rigorous, physics-based, forward-imaging model is proposed for estimating camera blur from a known calibration target. This model accommodates unknown ambient light levels, target reflectance, imaging perspective, detector gain and offset, quantum-efficiency, and read-noise variance. The forward model also accounts for diffraction due to the sensing aperture, detector fill-factors, aliasing, and the use of chromatic filters (Bayer patterns [11]). Section 3.3 utilizes the forward imaging model within a statistical inverse-problem framework to simultaneously estimate the unknown camera parameters from only the collected data. The ML estimator is shown to be essentially unbiased through the calculation of the bias-sensitivity index [95], and under moderate lighting conditions, achieves wavefront-estimation performance within 2dB RMS of the Cramér-Rao Bound (CRB). This corresponds to PSF estimates with less than 1% error under the L1 norm, and wavefront estimates under 0.01 Wavelengths RMS (WRMS) subject to the identifiability conditions discussed in Section 3.6. Section 3.7 discusses practical considerations related to calibration target construction, and simulation experiments are described in Section 3.8 that demonstrate the efficacy of the proposed approach.

3.2 Forward Imaging Model

Using the physical processes described in Section 2.2, and the modeling components developed in Section 2.4, one can accurately describe forward imaging process of a modern digital camera. This model treats the imaging geometry, optical blur, sensor gain and offset, scene illumination, and read-noise variance all as unknowns. The *known* calibration target at a particular wavelength λ is represented by the vector \mathbf{f} . This model is described at a single wavelength to ease exposition, however, it is easily extended to polychromatic case. Extending the previous ideas to accommodate the collection of J frames, one can model the unknown relationship between the object plane and the imaging plane through the vector of perspective parameters $\mathbf{v}^T = [\mathbf{v}_1^T, \dots, \mathbf{v}_J^T]^T$. The unknown aberrations in the optical wavefront are modeled by $\boldsymbol{\alpha}$, a vector of Zernike wavefront coefficients in the exit pupil. The constants c_1 and c_2 account for the scaling between calibration-target units (typically reflectance) and photons, while c_3 and c_4 account for the sensor gain and offset respectively. These constants will sometimes be represented together as \mathbf{c} . The process of measuring light intensity with a semiconductor results in a Poisson process [172] that is further subject to normally distributed read-noise parameterized by its variance σ_r^2 . Letting n_f represent the number of elements in \mathbf{f} , the distribution of the J lexicographically ordered data frames may be expressed in terms of these two independent statistical processes as

$$\mathbf{d} = c_3 \mathcal{P} \left\{ \text{diag} \left(\left\{ \mathbf{T} \mathbf{S} \mathbf{B}(\boldsymbol{\alpha}) \mathbf{W}(\mathbf{v}_j) \right\}_{j=1}^J \right) (\mathbf{1}_J \otimes \mathbf{I}_{n_f}) (c_1 \mathbf{f} + c_2) \right\} + c_4 + \mathcal{N}(\mathbf{0}_{Jn_f}, \sigma_r^2 \mathbf{I}_{Jn_f}) \quad (3.1)$$

where \mathbf{W} is the warping operator, \mathbf{B} is the blur operator, \mathbf{S} is the sampling operator, \mathbf{T} is the spatial truncation operator. The relationship between the physical processes, and these mathematical operators, is shown in Figure 3.2. Because the number of photons received is large under typical conditions [35], one may model the data as normally distributed where both the mean and variance are functions of the unknowns.

$$\mathbf{d} \sim \mathcal{N}(\boldsymbol{\mu}(\mathbf{v}, \boldsymbol{\alpha}, \mathbf{c}), \boldsymbol{\Sigma}(\mathbf{v}, \boldsymbol{\alpha}, \mathbf{c}, \sigma_r)) \quad (3.2)$$

$$\boldsymbol{\mu}(\mathbf{v}, \boldsymbol{\alpha}, \mathbf{c}) = c_3 \text{diag} \left(\left\{ \mathbf{T} \mathbf{S} \mathbf{B}(\boldsymbol{\alpha}) \mathbf{W}(\mathbf{v}_j) \right\}_{j=1}^J \right) (\mathbf{1}_J \otimes \mathbf{I}_{n_f}) (c_1 \mathbf{f} + c_2) + c_4 \quad (3.3)$$

$$= c_3 \mathbf{H}(\mathbf{v}, \boldsymbol{\alpha}) (c_1 \mathbf{f} + c_2) + c_4$$

$$\boldsymbol{\Sigma}(\mathbf{v}, \boldsymbol{\alpha}, \mathbf{c}, \sigma_r) = \text{diag}(c_3^2 \mathbf{H}(\mathbf{v}, \boldsymbol{\alpha}) (c_1 \mathbf{f} + c_2) + \sigma_r^2) \quad (3.4)$$

The data model described by (3.2) is a consequence of the key physical processes involved in imaging with a digital camera, and an identification of those parameters likely to

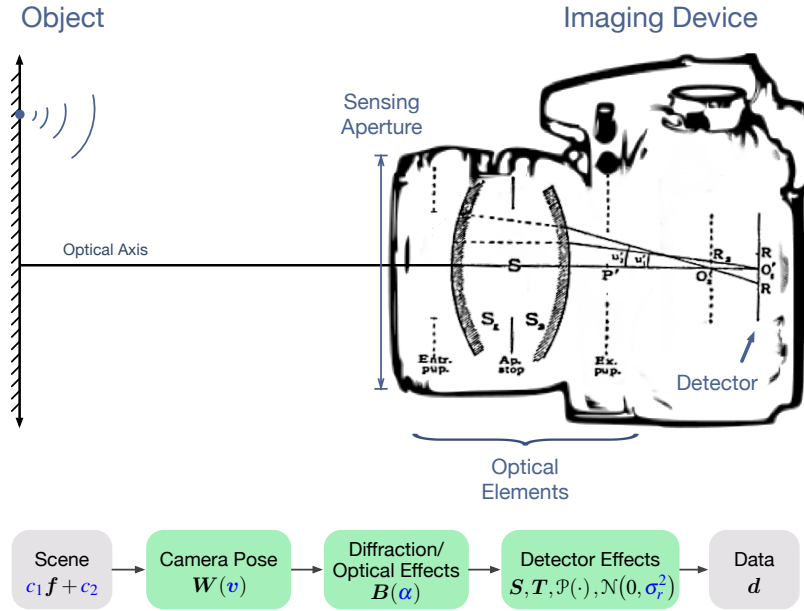


Figure 3.2: Diagram illustrating key elements of the physical imaging process. The corresponding block diagram relates physical processes to the mathematical operators, and corresponding unknowns, used to model them.

be unknown at the time of data acquisition. Collectively, we will refer to these unknowns as $\theta = [v^T, \alpha^T, c^T, \sigma_r]^T$, where the aberration parameters α are of primary interest, and the remaining unknowns are considered nuisance parameters; parameters not of immediate interest, but nevertheless must be estimated because of their influence the data distribution. Jointly estimating over the resulting high-fidelity model leads to substantially improved system characterization, and the direct modeling approach ensures the technique is applicable to a broad class of imaging systems. The difficulty in using the model proposed by (3.2) is that both the mean and covariance of the data are highly non-linear functions of the unknowns. This unfortunate property of the data model accurately reflects the nature of the forward imaging process, and therefore cannot be avoided.

3.3 Maximum Likelihood Estimation

Using the forward model (3.2)-(3.4), the goal is to characterized by the unknown wavefront aberrations α from the aliased data d by jointly estimating with respect to the nuisance parameters $\{v, c, \sigma_r\}$. The log-likelihood function for the collected data is given in

terms of the unknowns $\boldsymbol{\theta} = [\mathbf{v}^T, \boldsymbol{\alpha}^T, \mathbf{c}^T, \sigma_r]^T$ by

$$\ell(\mathbf{d}|\boldsymbol{\theta}) = -\frac{1}{2}\mathbf{a}^T \text{diag}(\mathbf{b})^{-1} \mathbf{a} - \frac{1}{2}\mathbf{1}_{Jn_f}^T \ln(\mathbf{b}) - \frac{Jn_f}{2} \ln(2\pi) \quad (3.5)$$

$$\mathbf{a}(\boldsymbol{\theta}) = \mathbf{d} - c_3 \mathbf{H}(\mathbf{v}, \boldsymbol{\alpha}) (c_1 \mathbf{f} + c_2) - c_4 \quad (3.6)$$

$$\mathbf{b}(\boldsymbol{\theta}) = c_3^2 \mathbf{H}(\mathbf{v}, \boldsymbol{\alpha}) (c_1 \mathbf{f} + c_2) + \sigma_r^2. \quad (3.7)$$

The log-likelihood expression (3.5) follows immediately from the fact that the data is jointly normal. Notice, however, that the mean and covariance are each affine functions of the calibration object \mathbf{f} , and non-linear in the unknowns $\boldsymbol{\theta}$. The dependency of \mathbf{a} and \mathbf{b} on parameter vector $\boldsymbol{\theta}$ has been suppressed in (3.5) to improve readability.

The joint ML estimator associated with (3.5) may be written as a constrained non-linear optimization problem. Assuming one chooses \mathbf{f} to be a linear function of the targets reflectance, the non-negativity of the intensity of the electromagnetic field emanating from the object ensures that c_1 and c_2 are non-negative. The model treats the signal impinging the detector as being in units of photons, and this implies $0 \leq c_3 \leq 1$. Finally, the read-noise variance must be positive, and under these constraints the joint ML estimator becomes

$$\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta}} \ell(\mathbf{d}|\boldsymbol{\theta}) \quad \text{s.t.} \quad c_1 \geq 0, c_2 \geq 0, 0 \leq c_3 \leq 1 \text{ and } \sigma_r \geq 0 \quad (3.8)$$

The foundation of our approach for solving for this estimator will be the class of quasi-Newton optimization techniques described in Section 2.6. When paired with modern line-search approaches [4, 79, 129, 130, 135, 205], quasi-Newton methods provide an efficient way of identifying locally optimal solutions. This chapter will address identifiability and fundamental information bounds for this estimation problem *assuming one is able to identify the neighborhood of the true solution*. The next chapter will build upon these ideas to explicitly address the problem of global convergence. As both the optimization techniques, and evaluation of the Cramér Roa bound, depend on the gradient of the log-likelihood

function, it is provided below

$$\nabla_{\boldsymbol{\theta}} \ell(\mathbf{d}; \boldsymbol{\theta}) = \sum_{j=1}^J \left(\frac{\mathbf{a}_j}{\mathbf{b}_j} \right)^T \begin{bmatrix} c_1 c_3 \nabla_{\mathbf{v}} (\mathbf{H}_j \mathbf{f})^T \\ c_1 c_3 \nabla_{\boldsymbol{\alpha}} (\mathbf{H}_j \mathbf{f})^T \\ c_3 (\mathbf{H}_j \mathbf{f})^T \\ c_3 \mathbf{1}_{n_f}^T \\ c_1 (\mathbf{H}_j \mathbf{f})^T + c_2 \\ \mathbf{1}_{n_f}^T \\ \mathbf{0}_{n_f}^T \end{bmatrix} - \frac{1}{2} \sum_{j=1}^J \left(\frac{1}{\mathbf{b}_j} - \left(\frac{\mathbf{a}_j}{\mathbf{b}_j} \right)^2 \right)^T \begin{bmatrix} c_1 c_3^2 \nabla_{\mathbf{v}} (\mathbf{H}_j \mathbf{f})^T \\ c_1 c_3^2 \nabla_{\boldsymbol{\alpha}} (\mathbf{H}_j \mathbf{f})^T \\ c_3^2 (\mathbf{H}_j \mathbf{f})^T \\ c_3^2 \mathbf{1}_{n_f}^T \\ 2c_3 (c_1 \mathbf{H}_j \mathbf{f} + c_2)^T \\ \mathbf{0}_{n_f}^T \\ 2\sigma_r \mathbf{1}_{n_f}^T \end{bmatrix}, \quad (3.9)$$

where $\mathbf{H}_j = \mathbf{T} \mathbf{S} \mathbf{B}(\boldsymbol{\alpha}) \mathbf{W}(\mathbf{v}_j)$ is the component of \mathbf{H} associated with the j^{th} frame such that $\mathbf{H} = \text{diag} \left(\{ \mathbf{H}_j \}_{j=1}^J \right) (\mathbf{1}_J \otimes \mathbf{I}_{n_f})$, and $\mathbf{a}_j, \mathbf{b}_j$ are similarly defined. Details of the derivation of these gradients is provided in Appendix I.

3.4 Cramér-Rao Bound

The Fisher information matrix describes one's fundamental ability to estimate a set of unknowns, and characterizes the influence of the nuisance parameters on the resulting estimation problem. Given the complexity of the forward imaging model (3.2), and the large number of unknowns, it is not reasonable to assume that the MAP estimator will converge to a useful solution. To help answer this question we compute Cramér-Rao Bound (CRB). This bound establishes lower limits on the performance of the set of minimum variance unbiased estimators, and given the asymptotic unbiasedness of the maximum likelihood estimator, this bound also often establishes limits on its mean squared error. In Section 3.8 it will be shown that when the specific optimization problem given by (3.8) is started in the neighborhood of the true solution, the proposed estimator is essentially unbiased. In the context of this

problem, essentially unbiased implies that a bias-variance trade-off cannot be exploited to yield an estimator resulting in a substantial improvement in wavefront estimation accuracy. Specifically, in Section 3.8 it will be demonstrated that use of the bias-gradient [95], in conjunction with the Fisher information matrix, leads to a biased estimator with whose wavefront estimation performance is increased by less than 1×10^{-4} WRMS of phase error.

The Fisher information matrix \mathbf{J} is given by the negative expected curvature of the log-likelihood about the true unknown parameters

$$\mathbf{J} \stackrel{\text{def}}{=} -\mathbb{E}_{\boldsymbol{\theta}, \mathbf{d}} [\nabla^2 L(\mathbf{d}; \boldsymbol{\theta})] \quad (3.10)$$

The subsequent bound on the MSE matrix is then given by

$$\mathbb{E}_{\boldsymbol{\theta}, \mathbf{d}} \left[\left(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta} \right) \left(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta} \right)^T \right] \geq \mathbf{J}^{-1} \quad (3.11)$$

where \geq implies that the difference of the MSE and inverse information matrices is positive definite. When prior information is available, the information matrix can be decomposed into two matrices; one term containing the contribution of information from the data, and a second term containing the contribution from the prior. Specifically,

$$\mathbf{J} = \mathbf{J}_D + \mathbf{J}_P \quad (3.12)$$

where

$$\mathbf{J}_D = -\mathbb{E}_{\mathbf{d}, \boldsymbol{\theta}} [\nabla^2 L(\mathbf{d} | \boldsymbol{\theta})] \quad (3.13)$$

$$\mathbf{J}_P = -\mathbb{E}_{\boldsymbol{\theta}} [\nabla^2 L(\boldsymbol{\theta})] \quad (3.14)$$

In the current blur estimation setting, prior information will not be available, so an emphasis will be placed on estimation bounds brought about entirely by the data-dependent term.

The data is jointly normal with mean and covariance each nonlinear functions of the unknowns. In this situation the Fisher Information is given by [9]

$$[\mathbf{J}_D]_{ij} = \frac{\partial \boldsymbol{\mu}^T}{\partial \boldsymbol{\theta}_i} \boldsymbol{\Sigma}^{-1} \frac{\partial \boldsymbol{\mu}}{\partial \boldsymbol{\theta}_j} + \frac{1}{2} \text{tr} \left(\boldsymbol{\Sigma}^{-1} \frac{\partial \boldsymbol{\Sigma}}{\partial \boldsymbol{\theta}_i} \boldsymbol{\Sigma}^{-1} \frac{\partial \boldsymbol{\Sigma}}{\partial \boldsymbol{\theta}_j} \right) \quad (3.15)$$

Expanding the second term of Equation 3.15 in terms \mathbf{a} and \mathbf{b} one gets

$$\text{tr}\left(\boldsymbol{\Sigma}^{-1} \frac{\partial \boldsymbol{\Sigma}}{\partial \boldsymbol{\theta}_i} \boldsymbol{\Sigma}^{-1} \frac{\partial \boldsymbol{\Sigma}}{\partial \boldsymbol{\theta}_j}\right) \quad (3.16)$$

$$= \text{tr}\left[\text{diag}\left(\frac{1}{\mathbf{b}}\right) \text{diag}(\nabla_{\boldsymbol{\theta}_i} \mathbf{b}) \text{diag}\left(\frac{1}{\mathbf{b}}\right) \text{diag}(\nabla_{\boldsymbol{\theta}_j} \mathbf{b})\right] \quad (3.17)$$

$$= \left(\frac{\nabla_{\boldsymbol{\theta}_i} \mathbf{b}}{\mathbf{b}}\right)^T \left(\frac{\nabla_{\boldsymbol{\theta}_j} \mathbf{b}}{\mathbf{b}}\right). \quad (3.18)$$

By similarly substituting into the first term, and collecting the results in matrix form one may write

$$\mathbf{J}_D = (\nabla_{\boldsymbol{\theta}} \mathbf{a})^T \text{diag}(\mathbf{b})^{-1} \nabla_{\boldsymbol{\theta}} \mathbf{a} + \frac{1}{2} \left(\frac{\nabla_{\boldsymbol{\theta}} \mathbf{b}}{\mathbf{1}_{n_{\theta}}^T \otimes \mathbf{b}}\right)^T \left(\frac{\nabla_{\boldsymbol{\theta}} \mathbf{b}}{\mathbf{1}_{n_{\theta}}^T \otimes \mathbf{b}}\right), \quad (3.19)$$

where n_{θ} is the total number of elements in $\boldsymbol{\theta}$. The necessary constituent derivatives are given through the expressions

$$\nabla_{\boldsymbol{\theta}} \mathbf{a} = -c_1 c_3 \nabla_{\boldsymbol{\theta}} \mathbf{H} \mathbf{f} - \begin{bmatrix} \mathbf{0}_{(n_v+n_a) \times J n_f} \\ c_3 (\mathbf{H} \mathbf{f})^T \\ c_3 \mathbf{1}_{J n_f}^T \\ c_1 (\mathbf{H} \mathbf{f})^T + c_2 \\ \mathbf{1}_{J n_f}^T \\ \mathbf{0}_{J n_f}^T \end{bmatrix}^T \quad (3.20)$$

$$\nabla_{\boldsymbol{\theta}} \mathbf{b} = c_1 c_3^2 \nabla_{\boldsymbol{\theta}} \mathbf{H} \mathbf{f} + \begin{bmatrix} \mathbf{0}_{(n_v+n_a) \times J n_f} \\ c_3^2 (\mathbf{H} \mathbf{f})^T \\ c_3^2 \mathbf{1}_{J n_f}^T \\ 2c_3 (c_1 \mathbf{H} \mathbf{f} + c_2)^T \\ \mathbf{0}_{J n_f}^T \\ 2\sigma_r \mathbf{1}_{J n_f}^T \end{bmatrix}^T \quad (3.21)$$

$$[\nabla \mathbf{H} \mathbf{f}]_{(j,:)} = \mathbf{T} \mathbf{S} [\mathbf{B}(\boldsymbol{\alpha}) \nabla \mathbf{W}(\mathbf{v}_j) \mathbf{f} + \nabla \mathbf{B}(\boldsymbol{\alpha}) \mathbf{W}(\mathbf{v}_j) \mathbf{f}] \quad (3.22)$$

This form of the CRB is particularly useful because it highlights the relationship between the key terms \mathbf{a} and \mathbf{b} , and the data-component of the Fisher information matrix. Given the potentially large size of the object representation with aliased systems, in conjunction with the multi-frame formulation, memory efficiency is essential for practical implementations. By identifying the relationships between the aforementioned terms and the individual

data frames, a computationally efficient way to compute Equation 3.19 can be identified. Decomposing the constituent terms in a manner similar to Equation 3.5, one may write

$$\mathbf{a} = \text{vec} \left(\left\{ \mathbf{d}_j - c_3 \mathbf{H}_j (c_1 \mathbf{f} + c_2) - c_4 \right\}_{j=1}^J \right) \quad (3.23)$$

$$= \text{vec} \left(\left\{ \mathbf{a}_j \right\}_{j=1}^J \right) \quad (3.24)$$

$$\mathbf{b} = \text{vec} \left(\left\{ c_3^2 \mathbf{H}_j (c_1 \mathbf{f} + c_2) + \sigma_r^2 \right\}_{j=1}^J \right) \quad (3.25)$$

$$= \text{vec} \left(\left\{ \mathbf{b}_j \right\}_{j=1}^J \right) \quad (3.26)$$

$$\boldsymbol{\Sigma}^{-1} = \text{diag} \left(\left\{ \mathbf{b}_j^{\odot -1} \right\}_{j=1}^J \right) \quad (3.27)$$

which leads the very intuitive result regarding the statistical independence of the individual frames

$$\mathbf{J}_D = \sum_{j=1}^J \left[\nabla \mathbf{a}_j^T \boldsymbol{\Sigma}_j^{-1} \nabla \mathbf{a}_j + \frac{1}{2} \nabla \mathbf{b}_j^T \boldsymbol{\Sigma}_j^{-2} \nabla \mathbf{b}_j \right] \quad (3.28)$$

In this expression the differentials are partitioned in the natural way

$$\nabla \mathbf{a}_j \stackrel{\text{def}}{=} [\nabla \mathbf{a}]_{((j-1)n_f+1:(j-1)n_f,:)} \quad (3.29)$$

with a similar expression for $\nabla \mathbf{b}_j$. One can further simplify the computations by exploiting the natural boundaries in the parameters. To do this, begin by writing

$$\nabla \cdot = \left[\nabla_{\mathbf{v}}(\cdot) \quad \nabla_{\boldsymbol{\alpha}}(\cdot) \quad \nabla_{\mathbf{c}}(\cdot) \quad \nabla_{\sigma_r}(\cdot) \right] \quad (3.30)$$

from which it follows that

$$\nabla \mathbf{a} = - \begin{bmatrix} \mathbf{C}_{\mathbf{v}}^T \\ \mathbf{C}_{\boldsymbol{\alpha}}^T \\ \mathbf{C}_{c_1}^T \\ \mathbf{C}_{c_2}^T \\ \mathbf{C}_{c_3}^T \\ \mathbf{C}_{c_4}^T \\ \mathbf{0}_{Jn_f \times 1}^T \end{bmatrix}^T \quad \nabla \mathbf{b} = c_3 \begin{bmatrix} \mathbf{C}_{\mathbf{v}}^T \\ \mathbf{C}_{\boldsymbol{\alpha}}^T \\ \mathbf{C}_{c_1}^T \\ \mathbf{C}_{c_2}^T \\ 2\mathbf{C}_{c_3}^T \\ \mathbf{0}_{Jn_f \times 1}^T \\ 2\frac{\sigma_r}{c_3} \mathbf{C}_{c_4}^T \end{bmatrix}^T \quad (3.31)$$

Combining this result with the separability along the frame boundaries allows one to write

$$\mathbf{C}_v = \text{diag}\left(\{c_1 c_3 \nabla_{v_j}(\mathbf{H}_j) \mathbf{f}\}_{j=1}^J\right) \quad (3.32)$$

$(Jn_f \times n_v)$

$$\mathbf{C}_\alpha = \text{vec}\left(\{c_1 c_3 \nabla_\alpha(\mathbf{H}_j) \mathbf{f}\}_{j=1}^J\right) \quad (3.33)$$

$(Jn_f \times n_\alpha)$

$$\mathbf{C}_{c_1} = \text{vec}\left(\{c_3 \mathbf{H}_j \mathbf{f}\}_{j=1}^J\right) \quad (3.34)$$

$(Jn_f \times 1)$

$$\mathbf{C}_{c_2} = c_3 \mathbf{1}_{Jn_f} \quad (3.35)$$

$(Jn_f \times 1)$

$$\mathbf{C}_{c_3} = \text{vec}\left(\{c_1 \mathbf{H}_j \mathbf{f} + c_2\}_{j=1}^J\right) \quad (3.36)$$

$(Jn_f \times 1)$

$$\mathbf{C}_{c_4} = \mathbf{1}_{Jn_f} \quad (3.37)$$

$(Jn_f \times 1)$

Note that Equation 3.28 may be interpreted as weighted inner-products of these terms. By exploiting their commonality, the data component of the Fisher information matrix can be computed using scalar multiples of inner-products taken over only 9 unique terms, plus the ones vector. The relationship between these terms and the necessary blocks within the FIM is given in Figure 3.3. By constraining the number of times a constituent term of the CRB

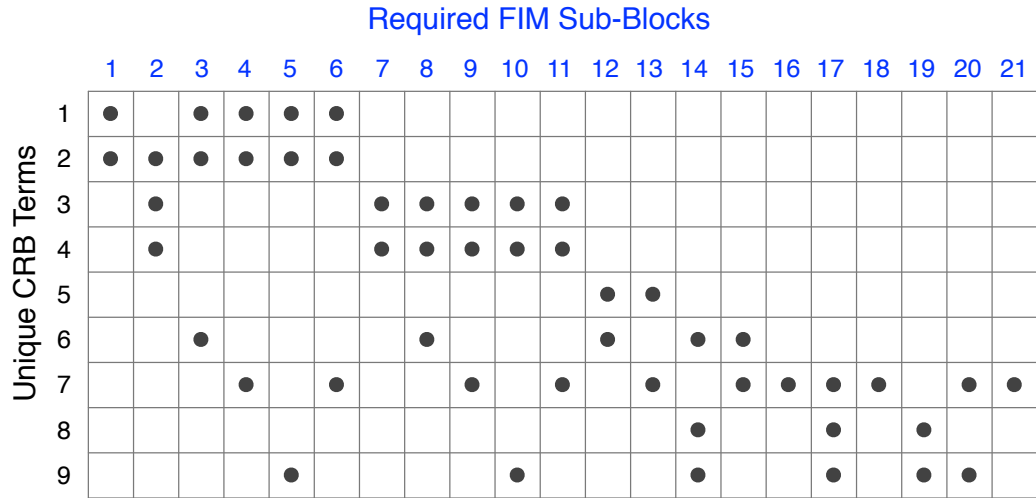


Figure 3.3: Connectivity of the bipartite graph relating FIM blocks to is underlying sufficient statistics

may enter/exit memory, one can reduce the overall memory footprint to around 10 times that of the Nyquist sampled object size. The result of this optimization is the prescription for computing the FIM in the 3 stages shown in Figure 3.4 and implemented in the file

sysIDFisherInfo.m.

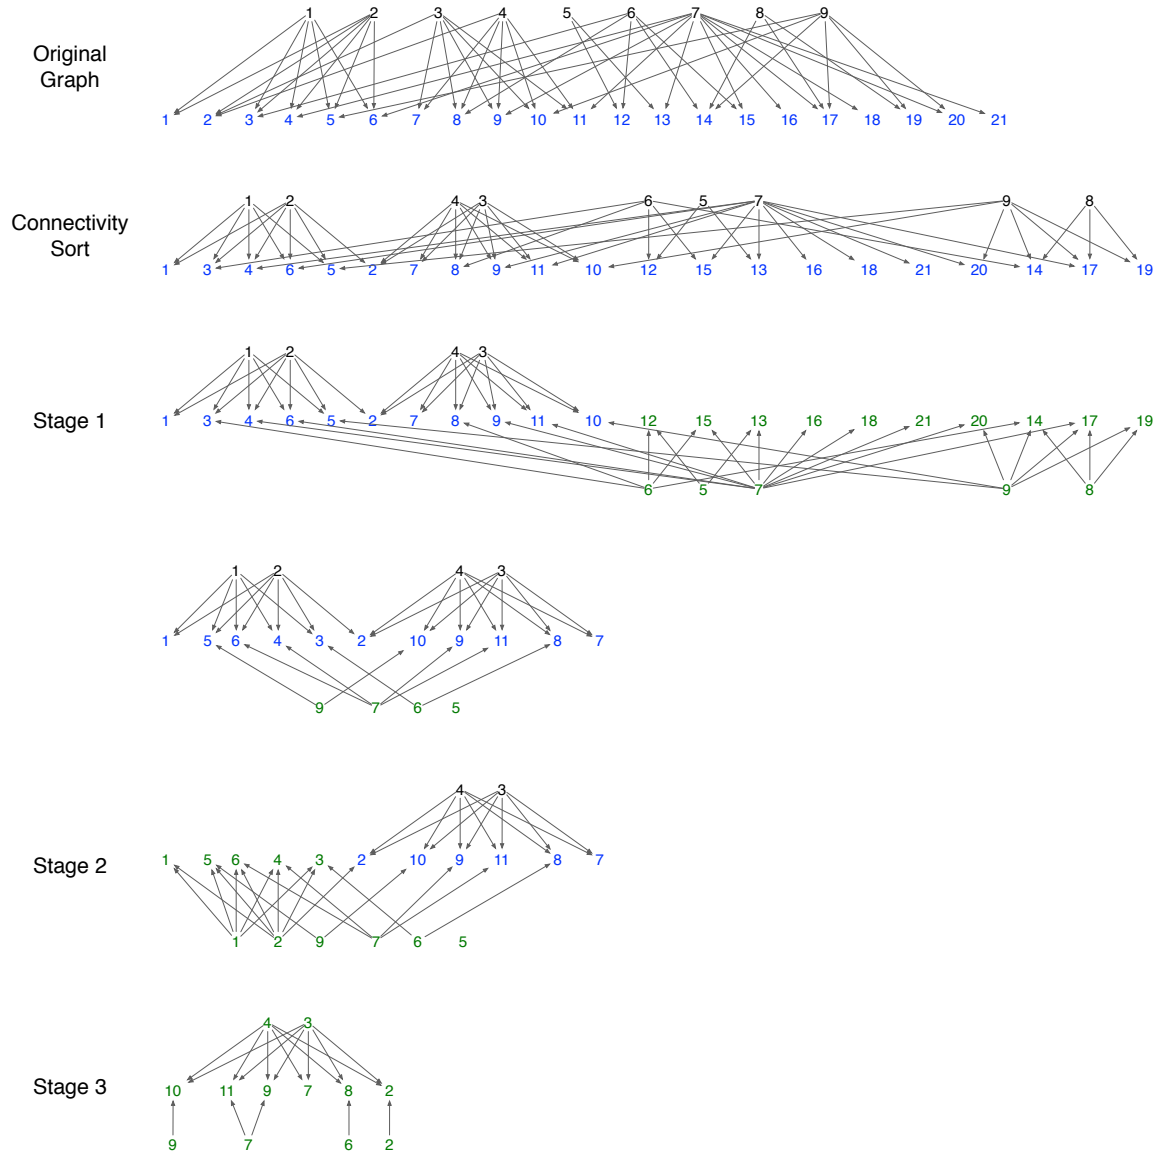


Figure 3.4: Graphical description of the computational stages used to compute the FIM in a memory efficient fashion. Each computational stage is represented as a bipartite graph between the FIM terms (blue), and the terms sufficient for its calculation (black). The transition of an items color to green indicates it is no longer needed for future calculations.

By numerically evaluating the FIM, one finds that the problem of wavefront estimation using the proposed model is well-posed under reasonable imaging conditions. Furthermore, the CRB suggests that the PSF can in theory be estimated with great accuracy.

3.5 Bias Sensitivity

Using the Uniform Cramér-Rao bound (CRB), this section establishes that the implicitly defined estimator given by (3.8) both reliable, and essentially unbiased, with a bias sensitivity index less than 1×10^{-3} . While asymptotically unbiased, the use of a finite-size observation ensures that any non-trivial estimator will incur some bias over any single experiment. Given such an estimator, small, otherwise insignificant biases, can result in the unbiased CRB being an unreliable measure of overall performance [177]. In this section we apply the uniform Cramér-Rao bound, and its associated bias-sensitivity index [94, 95], to establish that the unbiased CRB derived in Section 3.4 is a reliable predictor of estimator performance. The validity of these findings are further confirmed through a direct Monte-Carlo study of the proposed estimator's bias and variance described in Section 3.8.

The uniform CRB provides a way to study the reliability of the unbiased CRB in the presence of a small estimator bias. Given an observation \mathbf{d} , whose probability density function $f_D(\mathbf{d}; \boldsymbol{\theta})$ is a function of a vector-valued unknown $\boldsymbol{\theta}$, let $\hat{y} = \hat{y}(D)$ be an estimate of some function of interest $y_{\boldsymbol{\theta}} : \Theta \mapsto \mathbb{R}$. The squared error of this estimator can be described through its bias $b_{\boldsymbol{\theta}} = E_{\boldsymbol{\theta}}[\hat{y}] - y_{\boldsymbol{\theta}}$, and variance $\sigma_{\boldsymbol{\theta}}^2 = E_{\boldsymbol{\theta}}[(\hat{y} - m_{\boldsymbol{\theta}})^2]$, where $m_{\boldsymbol{\theta}} = E_{\boldsymbol{\theta}}[\hat{y}]$. The bias-variance plane, however, is not a particularly useful description of estimator performance because for a given $\boldsymbol{\theta}$ one can always define an estimator whose bias and variance are simultaneously zero. Furthermore, the ensemble average of an estimator's bias over Θ may be non-zero. This fixed component can be corrected without knowledge of $\boldsymbol{\theta}$, and thus does not fundamentally contribute to overall performance. For these reasons, Hero [94] suggests the use of the uniform CRB and its associated $\delta\sigma$ -plane description of estimator error.

The $\delta\sigma$ -plane describes estimator performance through a bound on the length of the bias-gradient $\delta_{\boldsymbol{\theta}} = \|\nabla b_{\boldsymbol{\theta}}\|_{\mathbf{C}} \leq 1$, and the square root variance $\sigma_{\boldsymbol{\theta}} = \sqrt{\sigma_{\boldsymbol{\theta}}^2}$. The bias-gradient length is defined with respect to a symmetric, positive-definite matrix \mathbf{C} that permits one to define ellipsoidal neighborhoods consistent with expected parameter variation. This length is unaffected by a fixed bias, and for an estimator whose Fisher information matrix is guaranteed to exist, uniquely specifies the unremovable component in the local neighborhood of the parameter of interest up to an additive constant. For this reason, it is a more direct measure of the error trade-off we wish to study. Furthermore, when the Fisher information exists, the uniqueness of the bias-gradient as a descriptor of the unremovable bias component depends only on the differentiability of bias gradient. This quantity is always differentiable so long as the raw second moment of the estimator $E_{\boldsymbol{\theta}}[\hat{y}^2]$ is finite ([100] Lemma 7.2), thus the approach is very broadly applicable.

The widely known unbiased version of the Cramér-Rao bound [41, 70, 153] extends naturally to accommodate the biased estimator \hat{y} of $y_{\boldsymbol{\theta}}$ [192]. Let \mathbf{J}_Y be the Fisher information matrix (FIM) associated with the parameters of interest, then

$$\sigma_{\boldsymbol{\theta}}^2 \geq [\nabla m_{\boldsymbol{\theta}}]^T \mathbf{J}_Y^{-1} [\nabla m_{\boldsymbol{\theta}}] \quad (3.38)$$

$$= [\nabla y_{\boldsymbol{\theta}} + \nabla b_{\boldsymbol{\theta}}]^T \mathbf{J}_Y^{-1} [\nabla y_{\boldsymbol{\theta}} + \nabla b_{\boldsymbol{\theta}}]. \quad (3.39)$$

For the remainder of this discussion the FIM is assumed to be positive definite; however, this result can be further generalized to accommodate singular information matrices [95]. The primary limitation of this result is that it depends on the entire bias gradient, and one rarely has prior knowledge of this quantity. The uniform Cramér-Rao bound [94] extends these ideas by allowing one to study estimator reliability in the $\delta\sigma$ plane. The bound applies uniformly over a neighborhood delineated by the norm of the bias gradient, and provides a direct approach for accessing bias sensitivity.

Consider the bias sensitivity of any individual aberration parameter. To ease the exposition, assume without loss of generality that this is the first element in the vector of unknowns $\boldsymbol{\theta}_1$. Let \mathbf{b}_1 be the bias vector associated with this parameter, and \mathbf{e}_i the indicator vector associated with the i^{th} element. Then plugging into (3.39) one gets the following variance bound in terms the bias gradient

$$\text{Var}_{\boldsymbol{\theta}} [\hat{\boldsymbol{\theta}}_1] \geq \underbrace{[\mathbf{e}_1 + \nabla \mathbf{b}_1]^T \mathbf{J}^{-1} [\mathbf{e}_1 + \nabla \mathbf{b}_1]}_{B_{b_1}(\boldsymbol{\theta})} \quad (3.40)$$

$$\geq \min_{\mathbf{b}_1: \|\nabla \mathbf{b}_1\|_C \leq \delta} B_{b_1}(\boldsymbol{\theta}) \quad (3.41)$$

$$\geq \underbrace{\min_{\mathbf{d}: \|\mathbf{d}\|_C \leq \delta} [\mathbf{e}_1 + \mathbf{d}]^T \mathbf{J}^{-1} [\mathbf{e}_1 + \mathbf{d}]}_{B_{b_1}(\boldsymbol{\theta}, \delta)}. \quad (3.42)$$

Notice that $B_{b_1}(\boldsymbol{\theta}, 0)$ is by definition the unbiased CRB, and thus the relative improvement attainable by a biased estimator subject to the bias-gradient constraint is given by

$$\Delta B_{b_1}(\boldsymbol{\theta}, \delta) = \frac{B_{b_1}(\boldsymbol{\theta}, 0) - B_{b_1}(\boldsymbol{\theta}, \delta)}{B_{b_1}(\boldsymbol{\theta}, 0)}. \quad (3.43)$$

As such, this quantity is a direct measure of the bias sensitivity of the unbiased CRB. For any particular $\boldsymbol{\theta}$ and $\delta \geq 0$, $\sqrt{B_{b_1}(\boldsymbol{\theta}, \delta)}$ bounds the achievable estimator error. An estimator that achieves this bound can only reduce its variance at the cost of an increased bias gradient and visa-versa. The slope of $B_{b_1}(\boldsymbol{\theta}, \delta)$ evaluated at $\delta = 0$ is known as the bias sensitivity

index [94]. Large values of this quantity indicate that the unbiased CRB is not a reliable predictor of estimator performance even when the estimator's bias is small.

Computation of the uniform CRB, and the associated bias sensitivity index, depends upon solutions to the constrained quadratic optimization problem described in (3.42). Figure 3.5 illustrates the minimizing vector \mathbf{d}_{min} geometrically in terms of the unique point of intersection between the contour where the bias-gradient norm achieves its limit δ and the associated contour of the quadratic such that the gradients of these contours are collinear with opposing signs. Solving for this point algebraically using the method of Lagrange




Figure 3.5: A plot of constant contours of the quadratic objective function and associated maximum bias-gradient constraint involved in computing the uniform CRB.

multipliers one finds

$$\mathbf{d}_{min}(\lambda) = -\mathbf{C}^{-1} \left[\mathbf{C}^{-1} + \lambda \mathbf{J} \right]^{-1} \mathbf{e}_1 \quad (3.44)$$

$$\lambda : \mathbf{d}_{min}^T(\lambda) \mathbf{d}_{min}(\lambda) = \delta^2 \quad (3.45)$$

Recall that \mathbf{d}_{min} is the bias-gradient that minimizes the biased CRB over all bias vectors satisfying the constraint $\|\nabla \mathbf{b}_1\|_{\mathbf{C}} \leq \delta$. It follows immediately that the associated bound is given by

$$B_{\mathbf{b}_1}(\boldsymbol{\theta}, \delta) = [\mathbf{e}_1 + \mathbf{d}_{min}]^T \mathbf{J}^{-1} [\mathbf{e}_1 + \mathbf{d}_{min}]. \quad (3.46)$$

A more computationally efficient form of the couple equations (3.44)-(3.45) is derived in terms of the eigendecomposition of the inverse Fisher information matrix relative to

C. Following [95], let $\mathbf{B} = \mathbf{C}^{-\frac{1}{2}} \mathbf{F}^{-1} \mathbf{C}^{-\frac{1}{2}}$ with $\boldsymbol{\xi}_i$ and β_i an associated eigenvector and eigenvalue pair. Then one can sweep out the uniform CRB curve (δ, B_{b_1}) in terms of λ by evaluating

$$\delta^2 = \sum_{i=1}^n \frac{\beta_i^2}{(\lambda + \beta_i)^2} \left| \mathbf{e}_i^T \mathbf{C}^{\frac{1}{2}} \boldsymbol{\xi}_i \right|^2 \quad (3.47)$$

$$B_{b_1}(\boldsymbol{\theta}, \delta) = \sum_{i=1}^n \frac{\lambda^2 \beta_i}{(\lambda + \beta_i)^2} \left| \mathbf{e}_i^T \mathbf{C}^{\frac{1}{2}} \boldsymbol{\xi}_i \right|^2. \quad (3.48)$$

Evaluating the uniform CRB for the defocus parameter of an optical system whose total aberration strength is 0.25 WRMS based on data collected on a detector that is 2.5 times undersampled in each linear dimension, one finds that the unbiased CRB estimate is a reliable predictor of overall estimator performance. Figure 3.6 shows the uniform CRB computed from the expressions (3.47)-(3.48). Note that the unbiased CRB is given by the intersection of the bound with the y-axis, and the associated bias sensitivity is the gradient evaluated at this point (less than 1×10^{-3}). This suggests that the overall estimator uncertainty will not be sensitive to small biases. This conclusion is confirmed through direct measurement of the estimator bias based on a larger ensemble of simulation experiments is described in Section 3.8.




Figure 3.6: Uniform CRB for the defocus parameter of a system with an overall aberration strength of 0.25 WRMS based 6.25 times undersampled data.

3.6 Identifiability

The estimation problem proposed in (3.8) in terms of a non-linear optimization problem is non-identifiable. This identifiability problem is related to the fact that *globally* the log-likelihood function does not provide a one-to-one map between possible system parameterizations and likelihood values given an observation despite the fact that *locally* the Fisher information matrix is not singular. The proposed formulation almost always results in a positive definite Fisher information matrix implying that locally the estimation problem is well behaved. Globally, however, there exist alternative parameterizations associated with any local minima of the negative log-likelihood function with equal likelihood. Given a sufficiently ambiguous calibration target (e.g. a checker board), ambiguities will clearly enter through the perspective parameters \mathbf{v} . This ambiguity is well understood, and generally found not to be an issue in practice [21, 46, 101]. The mapping between a wavefront description of a sensor's blur and its PSF is a non-injective surjective function. This implies that the sub-problem of estimating only blur through a wavefront-parameterization is itself non-identifiable. The remainder of this section addresses identifiability with respect to these parameters. It is shown that the specific problem of interest is *set-identifiable*, and that the cardinality of this set is at most 2. This set is also characterized such that given either of the two optimal solutions, the other can immediately be identified.

Phase retrieval, the recovery of a function from the modulus of its Fourier transform, has 3 types of ambiguities collectively known as the trivial associates [160]. Recovering a blur's wavefront description from intensity measurements is a type of phase retrieval problem known as wavefront sensing [65], and generally this body of literature is interested in solutions outside this set. Using the PSF description given by (2.10) and (2.13), if $C(\mathbf{w}) = A(\mathbf{w}) e^{j\Psi(\mathbf{w})}$ is a wavefront sensing solution then the trivial associates are given by

1. $c_1 C(\mathbf{w})$ s.t. $c_1 \in \mathbb{C}, |c_1| = 1$
2. $C^*(-\mathbf{w})$
3. $C(\mathbf{w} + \mathbf{\Delta})$ s.t. $\mathbf{\Delta} \in \mathbb{R}^2, \text{supp}[C(\mathbf{w} + \mathbf{\Delta})] = \text{supp}[C(\mathbf{w})]$.

For a Zernike basis description of the wavefront's phase (2.39)-(2.41), the first trivial associate corresponds to the span of Z_0^0 . Knowledge of this Zernike coefficient is equivalent to knowledge of the phase of the received light, and for this reason it is typically omitted when estimating blur. The second trivial associate represents a rotation of the phase-screen by 180° , and this introduces at most one additional wavefront solution when the even part of the phase-scene is non-zero. The third trivial associate does not introduce alternative

solutions due to the support constraints imposed by the physical aperture. Thus wavefront estimation from a single intensity measurement is set-identifiable only up to a set of at least two.

Aside from the trivial associates, 2D phase retrieval is known to be essentially unique [92, 160]. Hayes showed that outside of this set, discrete phase retrieval problems of dimensionality ≥ 2 are essentially unique, and gave a complete description of the measure zero set of alternate solutions. Sanz [160, 161] provides a similar theorem for the continuous problem, although in this case, a complete characterization of the alternative space of functions remains elusive.

The present application requires estimation through a computer simulation of the forward model, and thus the results of Hayes ensure identifiability up to a set of cardinality two with probability 1. This result, however, is asymptotic in the sense that the definition of identifiability does not account for the uncertainty inherent in estimation from a finite collection of noisy measurements. For the practitioner, a more relevant question is the detectability of convergence to a local rather than global solution. In chapter IV, this problem is directly addressed through the development of a global estimation procedure. Simulation experiments suggest that global convergence is achievable under the SNRs consistent with a well lit calibration object (see Figure 4.16), and global convergence is also demonstrated through experimental results.

3.7 Calibration Target Construction

The proposed inverse-problem approach for blur estimation is agnostic to the specific calibration target used, however, this does not imply that all targets result in data containing the same level of information regarding the unknowns. In this section a novel, self-similar, calibration target is proposed which yields highly-informative calibration data over a wide range of imaging geometries. In addition to the targets many convenient properties, the previously derived Cramér-Rao bound is used to demonstrate that this new target is fundamentally more informative than commonly used calibration objects. In doing so, this section also illustrates how the CRB can be applied to check the information of existing objects.

A Sierpinski pre-carpet [5] is proposed as a calibration target because of its scale invariance and large spectral support. These two properties ensure a large spatial-bandwidth over a broad range of imaging geometries. This strictly self-similar fractal target is composed of white squares on a black background, and is the natural analog to the cantor set over a 2D plane. Beginning with a black square (the support of the final target), a white square of $1/3$ its side-length is removed from the center. This process is then repeated for each remaining

black square, and the result of each such iterate is a Sierpinski pre-carpet of order n . An example of an order 4 pre-carpet is shown in Figure 3.7. Such a target maintains the sharp




Figure 3.7: Sierpinski pre-carpet of order 4

edges, corners, and large flat regions desired by many practitioners, while also providing a uniform and nearly scale invariant power-spectrum. The target enables superior estimation performance relative to traditionally used objects, and is relatively insensitive to imaging geometry. A further benefit of this calibration target is that it is constructed over a low-rank Haar basis. Modern computer displays use this same basis, namely square pixels, making the target ideally suited for projection from digital displays. An example of estimation using data collected this way is given in Section 3.8.

The physical size of the target best suited for calibration is determined by the optical resolution of the sensor along with the minimum available feature size. Assuming a circular aperture, the diffraction limited PSF is given by a Bessel function of the first kind. Mapping this into a discrete representation of the image plane, the distance between nulls in the PSF is given by $2z_1 R_N$, where z_1 is the first positive zero of $J_1(x)$, and R_N is the ratio of sensors sampling to the Nyquist rate. Defining the minimal “printer” resolution as P_{res} one can solve for the minimal supported imaging range

$$z_1 \frac{\lambda}{D} R_{min} = \frac{1}{P_{res}} \implies R_{min} = \frac{D}{z_1 \lambda P_{res}} \quad (3.49)$$

This minimal range is valid for an arbitrary target, however, the pre-carpets are only defined for integer orders. Assuming one desires to image the entire object, this places an additional restriction on the minimal range. The pre-carpet order is related to the number of elements

via

$$3^N = \# \text{ elements} \quad (3.50)$$

$$N_{max} = \log_3 \# \text{ elements} \quad (3.51)$$

$$= \log_3 \left(\frac{p_p D \min S_d}{z_1 \lambda f_L} \right) \quad (3.52)$$

where S_d represents the size of the detector (a 2 element tuple), and N_{max} is not yet restricted to \mathbb{Z}^+ . Adding this restriction results in the updated (restricted) expression for R_{min}

$$R_{min} = \frac{2D}{z_1 \lambda P_{res}} \exp [\ln (3) (N_{max} - \lfloor N_{max} \rfloor)], \quad (3.53)$$

where $\lfloor N_{max} \rfloor$ is then the maximum pre-carpet order supported by both the object medium and the optical system. The maximum experiment range is then only limited by ones ability to generate large targets, and desire to fill a substantial portion of the field of view.

Other benefits aside, Sierpinski pre-carpets are a fundamentally a better choice for system identification than traditional objects due to their increased information content relative to the unknown blur parameters. To demonstrate this, consider a monochromatic optical system is simulated with 0.15 waves of wavefront aberration. The incoming light is sampled by a 228×228 focal-plane with a 60% quantum efficiency that under-samples the impinging wavefront by a factor of 2. Overall, the sensor averaged 1 count for every 41 photons received, and had a sensor bias of 18 counts. Two such frames were used to exercise the multi-frame formulation, however, similar performance is observed with single-frame data given a commensurate increase in signal level. The read noise was adjusted to achieve the desired SNR, which was defined as the expectation over the image of the standard deviation of noise free signal divided by the aggregate standard deviation of the noise (read + Poisson). The two frames differed by 5% in magnification, 5° of rotation, and had small perspective and translational misalignments. The information content of an order 5 Sierpinski pre-carpet is compared to the more tradition tri-bar target shown in Figure 3.8. The read-noise was used to vary the overall quality of the data, and Figure 3.9 shows the wavefront-error lower bound vs noise level for each of the two targets. This plot illustrates that other things being equal, the Sierpinski target provides more information toward estimating the unknown parameters, however, the informational difference between the two calibration objects is small. This experiment also demonstrates how the forward imaging model and CRB formulation can be used to evaluate the informational content of any proposed calibration target to ensure that it is sufficient for a particular estimation task.




Figure 3.8: Standard tri-bar target




Figure 3.9: RMS wavefront-error bound vs noise

3.8 Simulation Experiments

Simulation experiments were performed to compare performance of the proposed estimator in the neighborhood of the true solution to the information-theoretic limits defined by the Cramér-Rao bound. Results of these experiments are also used to study the estimators bias. Given a matched model, and assuming convergence to the neighborhood of the true solution, the proposed estimator is shown to be within 2dB of the CRB. It is also demonstrated that under typical imaging conditions, even modest data sizes result in a negligible amount of estimator bias. Finally, a the bias-gradient and Fisher information matrix are used to construct an improved estimator. Such estimators, however, are shown not provide a substantial improvement in estimated wavefront accuracy.

Data was generated from the model given by (3.1) over a range of SNR's. For all experiments, the wavefront aberration strength was 0.25 WRMS, and the detector was 2.5 times undersampled in each linear dimension. The SNR, defined here to the standard deviation of the noise-free scene divided by the standard deviation of the total noise, was adjusted by increasing or decreasing read-noise levels. An order 5 Sierpinski carpet was used as the calibration target, and for this object the lighting environment and exposure times were such that approximately 2.05×10^5 photons were received per pixel. The detector gain and offset were chosen so that the final recorded data was centered, and filled 70% of a typical 8-bit dynamic range. These gain and offset parameters resulted in an expectation of approximately 2000 photons/count. Figure 4.15 shows the true object alongside a sequence of data frames at varying SNRs generated under these conditions. Figure 3.11 shows the total RMS estimator error and bias compared to the Cramèr-Rao bound when the estimator is started in the neighborhood of the true solution. The Marèchal criterion, the level of RMS wavefront aberration where a system can reasonably be considered diffraction limited [18, 123], is also shown as a point of reference. From these two plots one can conclude that when global convergence is achieved, the estimator is nearly unbiased, and achieves a total RMS error within approximately 2dB of the Cramèr-Rao bound. This level of error corresponds to wavefront estimation errors on the order of 1×10^{-3} WRMS under model matched conditions.

The near unbiasedness of the estimator was also quantified though calculations of the bias sensitivity index associated with the uniform Cramèr-Rao bound [95]. Under moderate lighting conditions, the measured bias sensitivity index was of the order of 1×10^{-4} , and while use of the bias-gradient direction in conjunction with the Fisher information matrix did lead to a biased estimator with lower overall MSE, these improvements correspond to less than 1×10^{-4} WRMS of phase error.




Figure 3.10: The true object (a) is shown alongside a series of data frames generated with all parameters fixed other than the SNR. The read-noise standard deviation corresponding to (b)-(d) were 0.68, 1.41, and 7.15 counts respectively.

3.9 Conclusions & Ongoing Work

A joint-estimation algorithm is presented that for robust system identification from under-sampled and misaligned data in loosely controlled environments. Despite the relatively large body of literature regarding the use of calibration targets for this purpose, to the author’s knowledge this is the first time information-theoretic lower bounds have been computed for this problem over such a complete model. These bounds suggest that the problem is “solvable”, to the accuracies needed by modern restoration approaches, if one is able identify the neighborhood of the true solution. Simulation results show that under a matched model, the proposed technique is able to recover the wavefront to within 2dB of the CRB under moderate SNR conditions.

Studies regarding the sensitivity of this method to model mismatches are still ongoing. Shortly after the failed attempt to compare the proposed technique to wavefront interferometric measurements (see Figure 3.1), a simpler *differential* experiment was performed. In this experiment an iPad 2’s rear-facing camera was used to collect two sets of data, each




Figure 3.11: Total wavefront estimation error (a) and bias (b) shown as a function of SNR for initial estimates within 0.01 waves RMS of the true 0.25 wave RMS perturbation.

containing an order 5 pre-carpet displayed on an LCD screen. The two datasets differed by a near 90° rotation of the sensor, and estimation was independently performed over each dataset without use of this prior information. 55 Zernike polynomials, excluding the piston, tip, and tilt were used to describe the optical wavefront. Based on the aperture diameter and detector size, the sensor was estimated to be sampled just under half the Nyquist rate at a wavelength of 550 nm. The band was assumed to be 100 nm wide with uniform sensitivity. One of the collected frames is shown in Figure 3.12, and the resulting wavefront estimates are shown in Figure 3.13. The overall wavefront errors were estimated to be 0.072 and 0.076 waves respectively, with defocus and spherical errors dominating the wavefront. This is consistent with a well calibrated device containing an anti-aliasing filter. The asymmetrical terms were consistent with the rotation of the device, and after registration, the residual difference between the two estimates was around 0.01 of a wavelength. While this experiment cannot speak to the absolute accuracy of the method, the results are consistent with the predicted outcomes. Like many small mobile devices, the iPad also JPEG compresses collected imagery before saving it to memory. This experiment was performed over the compressed data, suggesting at least some level of robustness to this mismatch. These very encouraging results motivated the improvements given in the next chapter, as well as the development of a more definitive hardware experiment.




Figure 3.12: Example data-frame collected with iPad 2




Figure 3.13: iPad2 wavefront estimates

CHAPTER IV

Point Spread Function Identification Under Large Aberrations

4.1 Introduction

The maximum likelihood estimator (3.8) is stated in terms of an optimization problem over the likelihood function, which in our case is a non-convex function. Gradient descent techniques can result in convergence to stationary points that are local, but not necessarily global, optima. This challenge arises because two relatively large wavefront aberration descriptions can result in a set of Point Spread Functions (PSFs) that are similar when viewed on the scale of the detector's sampling element. Given a sufficiently large set of wavefront aberrations, two PSFs with similar diffusion patterns can have disparate wavefront descriptions. As a result, one cannot expect the true solution to lie within the basin of attraction of an aberration-free system. This problem could be addressed using global optimization strategies, but such approaches are rarely applied to high-dimensional problems because of the severe computational demands. This chapter develops a new strategy for identifying globally optimal solutions that is compatible with standard gradient-descent techniques. The efficacy of this approach is demonstrated through both simulation and hardware experiments.

Section 3.6 established that the ML estimator contains at most 2 optimal solutions, but the non-convexity of (3.8) potentially admits additional local minima. Given a PSF h , described through its phase-screen Ψ , there exist alternative solutions of the form $\tilde{\Psi} = \Psi + \beta$ that result in a PSF \tilde{h} with dispersion characteristics similar to h despite the fact that β is not small. The simplest example of this phenomena is the linear shifting of the PSF when β is in the span of the Zernike basis functions Z_1^{-1} and Z_1^1 . Such translations are indistinguishable from those represented by v_1 and v_2 in (D.3), and for this reason, these basis elements are typically removed when performing wavefront estimation. For higher-order aberrations,

this concept was described by Maréchal [123], and its use in optical design is known as aberration balancing [121]. Unsurprisingly, when Ψ describes a wavefront corresponding to a global minimum of (3.8), this phenomenon often leads to multiple stationary points.

In Section 3.8 simulation experiments were used to compare estimation performance in the neighborhood of the true solution to the Cramér-Rao lower bound. By repeating this same experiment, but starting the solver from an aberration-free system description rather than the true solution, one is able to confirm the existence of the aforementioned local minima immediately. In fact, with an optical aberration strength of 0.25 WRMS, the gradient-descent solver converges to a stationary point other than one of the two global optima 96% of the time. Figure 4.1 illustrates an example of this phenomenon from the demonstration software available with this thesis. In this example, the true aberration strength is 0.25 WRMS, the identified local minima of the log-likelihood corresponds to an aberration strength of 0.18 WRMS, but the wavefront estimate at this local minima remains 0.19 WRMS from the true solution despite resulting in a PSF with a structure similar to the true solution. The existence of such stationary points poses a substantial challenge to the development of practical wavefront estimation techniques.

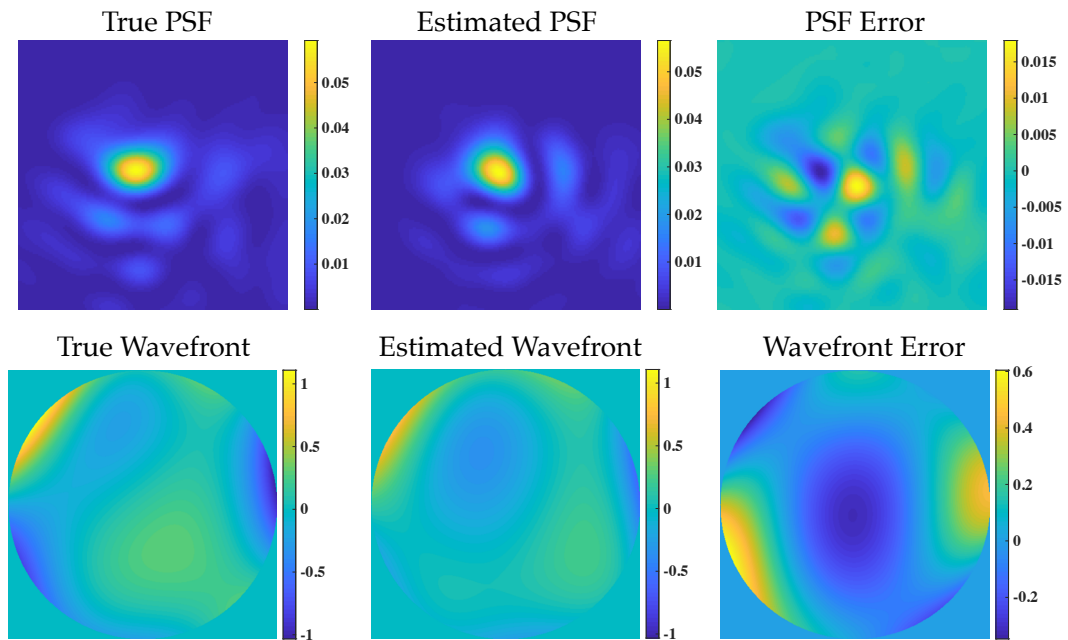


Figure 4.1: A local minima of the log-likelihood that corresponds to a wavefront solution 0.19 WRMS from the true solution despite its associated PSF having a similar dispersion structure. The PSFs are shown on a grid upsampled 10X relative to the Nyquist-rate of the detector.

This chapter addresses this challenge through an improved test for convergence of the

log-likelihood to suboptimal stationary points, as well as a technique for identifying alternate aberration descriptions that are dissimilar in the wavefront domain while resulting in PSFs with a similar dispersion structure. The improved test is based on an improvement to a global maximum validation function proposed by Biernacki [13], and is shown to uniformly improve upon this test when the statistical distribution of the data is in the generalized location family. In addition, a new reparameterization and embedding procedure is presented that exploits knowledge about the forward operator to further improve this global maximum validation function. Given a stationary point of the log-likelihood believed to be a local rather than global optimum, new candidate wavefronts are chosen from a precomputed set of perturbations that optimizes a point-wise minimax criterion over the resulting incoherent PSF. The result is an application-specific strategy for identifying global optimum, and the efficacy of this approach is numerically demonstrated in terms of increased detection accuracy and reduced computation. Figure 4.2 provides an overview of the resulting global optimization strategy.




Figure 4.2: Diagram illustrating how a test for suboptimal solutions can be used in conjunction with an approach for identifying alternative candidates to perform global optimization.

4.2 Testing for a Local Optimum of the Likelihood

The construction of improved tests for convergence to a local optimum of the likelihood is first studied in the general setting of maximum likelihood parameter estimation from multiple samples from a distribution that is in a known parametric family. The conceptual simplicity and tractability of the Maximum Likelihood (ML) principle have engendered its use for well over a century, yet questions surrounding its practical application remain open. Fisher [68] was an early advocate of the ML approach and is generally credited with its development despite earlier work based on similar concepts. Stigler [176] provides an excellent historical account of the theory’s maturation throughout the nineteenth and twentieth centuries. The local asymptotically normal theory for parametric models largely ended with the work of Le Cam (c.f. [110], ch. 6), yet the application of the approach to non-convex models can still result in computational difficulties. Standard asymptotic theory

ensures that under relatively mild conditions there exists a unique consistent root to the gradient of the log-likelihood [42, 197], and that this root will occur at the global optimum of the likelihood function [197]. The natural question then becomes, how one can identify consistent roots of the likelihood?

There exist general-purpose algorithms to address this question, e.g., simulated annealing and genetic algorithms. However, these algorithms are rarely applied to high-dimensional problems because of severe computational demands [6, 7, 166]. Stationary points of the gradient of the likelihood function can be readily found using iterative root-finding methods such as Quasi-Newton gradient descent [135]. Once a stationary point is located, it would be useful to have a simple test to determine whether or not it is global optimal without knowing the maximum value of the likelihood function. Several such tests have been proposed for this purpose [13, 14]. In this chapter, the focus is on testing local optima of the likelihood function for high dimensional inverse-problems occurring in signal processing and imaging applications.

Assume that the data \mathbf{d} is a realization of the random vector \mathbf{D} having a parametric density $f(\mathbf{d}; \boldsymbol{\theta}_0)$, where $\boldsymbol{\theta}$ is an unknown parameter vector taking values in a subset Θ of \mathbb{R}^p and \mathbf{d} takes values in the sample space \mathcal{D} . We define $\boldsymbol{\theta}_0$ as a fixed value, called the true value, of the parameter vector and for a function X of \mathbf{D} define the statistical expectation $E_{\boldsymbol{\theta}_0}[X] = \int X(\mathbf{d})f(\mathbf{d}; \boldsymbol{\theta}_0)d\mu(\mathbf{d})$. Assume that f is differentiable in $\boldsymbol{\theta}$ and that the Fisher information matrix $E_{\boldsymbol{\theta}_0}[(\nabla \ln f)(\nabla \ln f)^T]$ exists. Given independent and identically distributed observations $\{\mathbf{d}_1, \dots, \mathbf{d}_n\}$ the log-likelihood function is

$$\ell(\mathbf{d}; \boldsymbol{\theta}) = \frac{1}{n} \sum_{k=1}^n \ln f(\mathbf{d}_k; \boldsymbol{\theta}), \quad (4.1)$$

and the maximum likelihood estimator (MLE) is the global minimum

$$\hat{\boldsymbol{\theta}}_{\text{Global}} = \arg \min_{\boldsymbol{\theta} \in \Theta} -\ell(\mathbf{d}; \boldsymbol{\theta}). \quad (4.2)$$

An estimator $\hat{\boldsymbol{\theta}}$ is said to be consistent (statistically consistent in norm) when

$$\lim_{n \rightarrow \infty} E_{\boldsymbol{\theta}_0} \left[\left\| \hat{\boldsymbol{\theta}}(\mathbf{d}) - \boldsymbol{\theta}_0 \right\|^2 \right] \rightarrow 0 \quad (4.3)$$

In the problems of interest, the global minimum is unknown and only a local minimum $\hat{\boldsymbol{\theta}}$ is available, which is not necessarily equal to $\hat{\boldsymbol{\theta}}_{\text{Global}}$. For example, the local minimum could be the convergent limit of a gradient descent algorithm. Then, given $\hat{\boldsymbol{\theta}}$, the local

minimum testing problem is to decide between the two hypotheses

$$H_0 : \hat{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}}_{\text{Global}} \quad \text{vs.} \quad H_1 : \hat{\boldsymbol{\theta}} \neq \hat{\boldsymbol{\theta}}_{\text{Global}}. \quad (4.4)$$

A test between H_0 and H_1 is defined as a binary valued function $\phi : \mathcal{D} \rightarrow \{0, 1\}$ that maps the data sample \mathbf{d} to 0 or 1, indicating the decision H_0 or H_1 , respectively. For a given test function ϕ , a test statistic is any scalar function of the data $\phi(\mathbf{d})$. The accuracy of a test is measured by its probability of false alarm $\text{PFA} = E_{\boldsymbol{\theta}_0}[\phi|H_0]$ and its probability of detection $\text{PD} = E_{\boldsymbol{\theta}_0}[\phi|H_1]$. If for two tests ϕ_1 and ϕ_2 having identical PFA, the PD of ϕ_1 is greater than the PD of ϕ_2 , then ϕ_1 is said to be more powerful than ϕ_2 .

Many approaches to the general hypothesis testing problem (4.4) have been studied over the years. Blatt and Hero [14] present the historical context which is summarized here. The likelihood ratio test [203], Wald test [196], and Rao score test [154] are asymptotically equivalent tests as the number n of samples approaches infinity. The likelihood ratio and Wald tests require the distribution under H_0 to be known, which for (4.4) requires knowledge of the true parameter. On the other hand, the Rao score test, later independently discovered and popularized under the name Lagrange multiplier test [170], can be implemented when the true parameter is unknown. Rao's test measures the Euclidean norm of the score function $\mathbf{s}(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} \ell(\mathbf{d}, \boldsymbol{\theta})$ weighted by the inverse Fisher information evaluated at a local maximum $\boldsymbol{\xi}_R = \frac{1}{p} \mathbf{s}(\mathbf{d}, \hat{\boldsymbol{\theta}})^T \mathbf{F}^{-1}(\hat{\boldsymbol{\theta}}) \mathbf{s}(\mathbf{d}, \hat{\boldsymbol{\theta}})$. Gan and Jiang [71] propose a similar test for consistency of a stationary point of the likelihood based on White's information test [202]. White's original work was concerned with testing for model misspecification under the assumption that the global maximum of the likelihood function had been located, and Gan uses the same test statistic but in the converse situation.

The Rao test may be used to test for consistency of a local maximum of the log-likelihood function. Unfortunately, Monte Carlo experiments indicate that this test is not very powerful even in the univariate setting. Biernacki [13] suggested an improved test for the consistency of a stationary point following ideas presented by Cox [38, 39]. Biernacki's test uses a bootstrap estimate to directly compare the observed value of the locally maximized log-likelihood to its statistical expectation. Both the Rao score and the Biernacki tests fall under the more general M-testing framework described by Blatt and Hero [14] where additional types of tests of local maxima are proposed.

4.2.1 A Simple Motivating Example

It is useful to illustrate the testing of local optima of the log-likelihood through a simpler inverse problem. Let \mathbf{d} be data measured from a forward model with mean response $\boldsymbol{\mu}(\boldsymbol{\theta})$

and i.i.d. additive Gaussian noise ε of variance σ^2 :

$$\mathbf{d} = \boldsymbol{\mu}(\boldsymbol{\theta}) + \varepsilon \quad \text{s.t.} \quad \varepsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_{100 \times 100}). \quad (4.5)$$

The unknown true value of $\boldsymbol{\mu}(\boldsymbol{\theta}_0)$ is a vector of time samples of a sinusoidal signal

$$\boldsymbol{\mu}(\boldsymbol{\theta}_0) = \sin(\boldsymbol{\theta}_0 \mathbf{x}), \quad \boldsymbol{\theta}_0 = 3\pi \quad (4.6)$$

$$\mathbf{x} = \frac{1}{99} [0, 1, \dots, 99]^T. \quad (4.7)$$

We suppose that it is known *a priori* that $\boldsymbol{\mu}$ is in the signal class $\mathcal{C} = \{\boldsymbol{\mu} : \boldsymbol{\mu} = \sin(\boldsymbol{\theta} \mathbf{x}), \boldsymbol{\theta} \in [0, 4\pi]\}$. The maximum likelihood estimator is then the solution of the constrained optimization problem

$$\hat{\boldsymbol{\theta}}_{\text{Global}} \stackrel{\text{def}}{=} \arg \min_{\boldsymbol{\theta} \in [0, 4\pi]} \|\boldsymbol{\mu}(\boldsymbol{\theta}) - \mathbf{d}\|^2 \quad (4.8)$$

The constraint set \mathcal{C} is a 1-dimensional manifold parameterized by $\boldsymbol{\theta}$, and thus the objective function $\|\boldsymbol{\mu}(\boldsymbol{\theta}) - \mathbf{d}\|^2$ has sub-optimal local minima in addition to a global minimum. Figure 4.3(a) shows two of these local minima for the case that the noise variance σ^2 is zero. The blue curve in Figure 4.3(a) is the global minimum, which is the true signal, and the dotted red curve is another local minimum. Figure 4.3(b) shows the corresponding data observation \mathbf{d} for realizations of these two signals when the noise variance is $\sigma^2 = 1$. The perceptual similarity between these two realizations illustrates the potential difficulty of distinguishing a sub-optimal local minimum from the global minimum. This situation is a straightforward analog of the more complex optical inverse problem of interest. It was chosen because it simply illustrates the issue at hand while providing a challenge comparable to that encountered in the aforementioned inverse-imaging problem (c.f. 3.8).

4.2.2 1-sided Biernacki Test

The key to testing a local minimum of the likelihood function is to define a suitable *global maximum validation function* whose statistical distribution changes depending on whether the local minimum $\hat{\boldsymbol{\theta}}$ is global or not [14]. Define the validation function $\varphi(\mathbf{d}, \hat{\boldsymbol{\theta}}) \stackrel{\text{def}}{=} \ell(\mathbf{d}, \hat{\boldsymbol{\theta}}) - \mathbb{E}_{\hat{\boldsymbol{\theta}}}[\ell(\mathbf{D}, \hat{\boldsymbol{\theta}})]$. Under the null hypothesis H_0 , where $\hat{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}}_{\text{Global}}$, the distribution of $\varphi(\mathbf{D}, \hat{\boldsymbol{\theta}})$ will have approximately zero mean. For an i.i.d. data sample

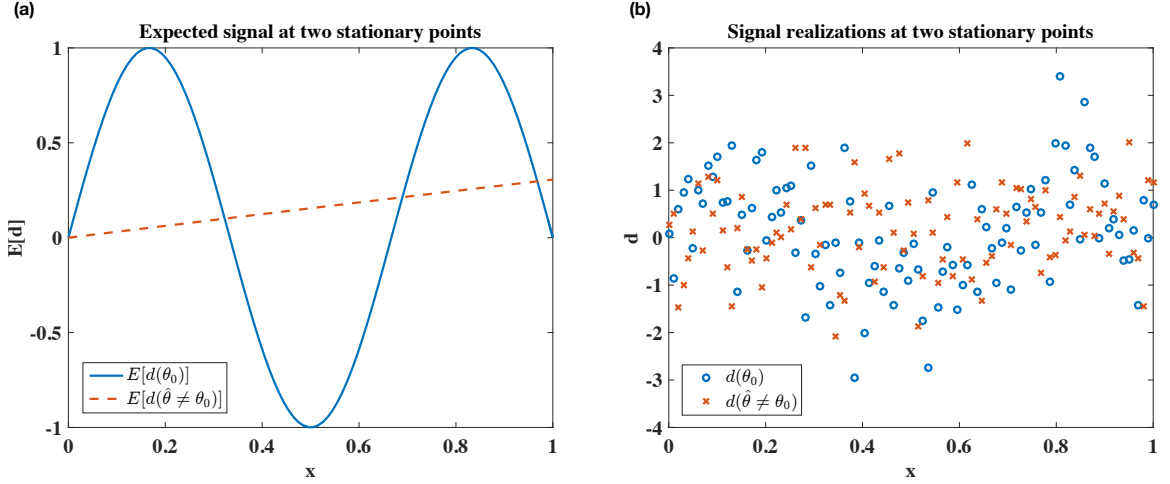


Figure 4.3: (a) The true signal achieving global minimum (blue solid) and signal achieving a sub-optimal local minimum (red dotted) of the likelihood function (4.8) when noise variance σ^2 equals zero. (b) The measured data when the noise variance is $\sigma^2 = 1$.

$\mathbf{D}_1, \dots, \mathbf{D}_n$ Biernacki showed the asymptotic result ([13], Theorem 2) under H_0 :

$$\frac{1}{\sqrt{n}} \sum_{t=1}^n \varphi(\mathbf{D}_t, \hat{\boldsymbol{\theta}}) \xrightarrow{D} \mathcal{N}(0, \text{Var}[\ell(\mathbf{D}_1, \boldsymbol{\theta}_0)]), \quad (4.9)$$

where $\text{Var}[\ell(\mathbf{D}_1, \boldsymbol{\theta}_0)]$ is the variance of the log-likelihood function for a single data sample ($n = 1$). Overloading the notation by designating the entire n data samples by $\mathbf{D} = \{\mathbf{D}_1, \dots, \mathbf{D}_n\}$, this Gaussian limit motivates us to define test

$$\frac{(\ell(\mathbf{D}, \hat{\boldsymbol{\theta}}) - \mathbb{E}_{\hat{\boldsymbol{\theta}}}[\ell(\mathbf{D}, \hat{\boldsymbol{\theta}})])^2}{\text{Var}_{\hat{\boldsymbol{\theta}}}[\ell(\mathbf{D}, \hat{\boldsymbol{\theta}})]} \begin{matrix} > & H_1 \\ < & H_0 \end{matrix} \eta, \quad (4.10)$$

where η is a threshold selected to fix the false alarm probability equal to a suitably small number $\alpha \in [0, 1]$. Under local asymptotically normal (LAN) conditions on the likelihood function [114] $\hat{\boldsymbol{\theta}} \xrightarrow{P} \boldsymbol{\theta}_0$ (a.s.) and the test statistic on the left hand side of (4.10) has an approximately Chi-Square distribution under H_0 . Hence η can be selected as the $1 - \alpha$ quantile of the Chi-square distribution. Biernacki implemented this test by approximating the mean $\mathbb{E}_{\hat{\boldsymbol{\theta}}}[\ell(\mathbf{D}, \hat{\boldsymbol{\theta}})]$ and the variance $\text{Var}[\ell(\mathbf{D}, \boldsymbol{\theta}_0)]$ using a parametric bootstrap estimator.

The test (4.10) is called a two-sided test as it can be equivalently expressed as

$$-\sqrt{\text{Var}_{\hat{\boldsymbol{\theta}}}[\ell(\mathbf{D}, \hat{\boldsymbol{\theta}})]} \eta \leq \varphi(\mathbf{D}, \hat{\boldsymbol{\theta}}) \leq \sqrt{\text{Var}_{\hat{\boldsymbol{\theta}}}[\ell(\mathbf{D}, \hat{\boldsymbol{\theta}})]} \eta.$$

This is thus a test for which, as compared to the global minimum $\hat{\boldsymbol{\theta}}_{\text{Global}}$, a sub-optimal local minimum $\hat{\boldsymbol{\theta}}$ will cause the test function to undergo a shift in mean, where the shift could either be in a positive or a negative direction.

If it were known *a priori* that a sub-optimal local minimum causes a negative shift in the mean of the global maximum validation function $\varphi(\mathbf{D}, \hat{\boldsymbol{\theta}})$ a one-sided test would be advantageous over a two-sided test. More specifically, a one-sided test would be expected to have higher power than the two-sided test (4.10) when, for all $\hat{\boldsymbol{\theta}} \neq \hat{\boldsymbol{\theta}}_0$,

$$\mathbb{E}_{\hat{\boldsymbol{\theta}}}[\ell(\mathbf{D}, \hat{\boldsymbol{\theta}})] \geq \mathbb{E}_{\boldsymbol{\theta}_0}[\ell(\mathbf{D}, \hat{\boldsymbol{\theta}})]. \quad (4.11)$$

When this condition is satisfied the two-sided test (4.10) can be replaced by the one-sided test

$$\frac{\ell(\mathbf{D}, \hat{\boldsymbol{\theta}}) - \mathbb{E}_{\hat{\boldsymbol{\theta}}}[\ell(\mathbf{D}, \hat{\boldsymbol{\theta}})]}{\sqrt{\text{Var}_{\hat{\boldsymbol{\theta}}}[\ell(\mathbf{D}, \hat{\boldsymbol{\theta}})]}} \underset{H_1}{>} \underset{H_0}{<} \eta_1. \quad (4.12)$$

The intuition behind a one-sided variant leading to an improvement is shown in Figure 4.4. Under the alternative hypothesis ($\hat{\boldsymbol{\theta}} \neq \hat{\boldsymbol{\theta}}_0$), one might intuitively expect that log-likelihood values computed as part of a correctly specified bootstrap procedure to be more likely than under an incorrect parameterization. In general, however, this is not the case. As a counter-example, consider a scalar problem whose incorrect parameterization ($\hat{\boldsymbol{\theta}} \neq \hat{\boldsymbol{\theta}}_0$) causes the presumed likelihood function to concentrate in the neighborhood of the observation. Such a problem would clearly permit the situation illustrated in Figure 4.4b. Nevertheless, there are many situations where this intuition is, in fact, correct.

The condition (4.11) is satisfied for many imaging and inverse problems. For example, consider the case where $\boldsymbol{\theta}$ is a clean image that one wishes to recover from samples \mathbf{D} of the output of an imaging sensor with known point spread function (forward operator) in additive correlated noise. When the point-spread function (PSF) and the covariance are known, this model will always satisfy the inequality (4.11), and the one-sided test might be expected to lead to a better test for local minima. Define $\boldsymbol{\theta}_0 \in \mathbb{R}^p$ the vectorized true image to be recovered, and $\mathbf{D} \in \mathbb{R}^q$ the vectorized the image acquired from the camera, which obeys the model:

$$\mathbf{D} = \mathbf{H}\boldsymbol{\theta}_0 + \boldsymbol{\varepsilon} \quad \text{s.t.} \quad \boldsymbol{\varepsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}), \quad (4.13)$$

where \mathbf{H} is a $q \times p$ matrix representing the forward operator and $\boldsymbol{\Sigma}$ is the $q \times q$ camera



(a) Ideal Situation Under H_1 (b) Unusual Situation Under H_1 

Figure 4.4: For a one-sided test to improve upon Biernacki's more general form, the set of stationary points comprising a problem's local minima would produce bootstrapped log-likelihood values like those shown in (a). The somewhat unintuitive situation of an incorrect model parameterization resulting in smaller log-likelihood values than the true parameterization is shown in (b). In both figures, the semi-transparent box emphasizes the fact that the curves are expectations of random functions, and that individual bootstrap tests will involve variation.

covariance matrix.

To show that (4.11) holds in this case, start with the log-likelihood function for the above model

$$\ell(\mathbf{D}, \boldsymbol{\theta}) = -\frac{1}{2} (\boldsymbol{\theta}_0 - \boldsymbol{\theta}) + \boldsymbol{\varepsilon})^T \mathbf{H}^T \boldsymbol{\Sigma}^{-1} \mathbf{H} (\boldsymbol{\theta}_0 - \boldsymbol{\theta}) + \boldsymbol{\varepsilon}) - \frac{1}{2} \ln(\det \boldsymbol{\Sigma}) - \frac{q}{2} \ln(2\pi) \quad (4.14)$$

$$= -\frac{1}{2} [(\boldsymbol{\theta}_0 - \boldsymbol{\theta}) + \boldsymbol{\varepsilon}]^T \boldsymbol{\Sigma}^{-1} [(\boldsymbol{\theta}_0 - \boldsymbol{\theta}) + \boldsymbol{\varepsilon}] - \frac{1}{2} \ln(|\boldsymbol{\Sigma}|) - \frac{p}{2} \ln(2\pi) \quad (4.15)$$

$$= -\frac{1}{2} [\boldsymbol{\Sigma}^{-1/2} (\boldsymbol{\theta}_0 - \boldsymbol{\theta}) + \boldsymbol{\xi}]^T [\boldsymbol{\Sigma}^{-1/2} (\boldsymbol{\theta}_0 - \boldsymbol{\theta}) + \boldsymbol{\xi}] - \frac{1}{2} \ln(|\boldsymbol{\Sigma}|) - \frac{p}{2} \ln(2\pi), \quad (4.16)$$

where $\boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. For any value of $\boldsymbol{\theta}$, the quadratic form in (4.14) has a non-central chi-squared distribution with non-centrality parameter λ :

$$\lambda = (\boldsymbol{\theta}_0 - \boldsymbol{\theta})^T \mathbf{H}^T \boldsymbol{\Sigma}^{-1} \mathbf{H} (\boldsymbol{\theta}_0 - \boldsymbol{\theta}). \quad (4.17)$$

The moment properties of the non-central chi-square distribution [107] thus specify the statistical expectation of the log-likelihood function (4.14) :

$$\mathbf{E}_{\boldsymbol{\theta}_0}[\ell(\mathbf{D}, \boldsymbol{\theta})] = -\frac{1}{2} (q + \lambda) - \frac{1}{2} \ln(\det \boldsymbol{\Sigma}) - \frac{q}{2} \ln(2\pi). \quad (4.18)$$

The difference $\mathbf{E}_{\boldsymbol{\theta}}[\ell(\mathbf{D}, \boldsymbol{\theta})] - \mathbf{E}_{\boldsymbol{\theta}_0}[\ell(\mathbf{D}, \boldsymbol{\theta})] = \lambda/2$, which is non-negative, establishing that

(4.11) as claimed. For this example, the unconstrained maximum likelihood estimator of $\boldsymbol{\theta}$ is a solution to a convex optimization problem, which is strictly convex when \mathbf{H} is full column rank, and thus there will be no sub-optimal isolated local minima of (4.2). As our simple example in Figure 4.3 illustrated, addition of constraints on $\boldsymbol{\theta}$ can give rise to additional local minima.

Below a stronger result is stated that implies the condition (4.11) is satisfied for any camera model of the form $\mathbf{D} = \boldsymbol{\mu}(\boldsymbol{\theta}_0) + \boldsymbol{\varepsilon}$, where $\boldsymbol{\mu}(\cdot)$ is a possibly non-linear function and $\boldsymbol{\varepsilon}$ is a possibly non-Gaussian noise as long as its distribution is independent of $\boldsymbol{\theta}_0$. This condition is equivalent to the condition that the distribution $f(\mathbf{d}, \boldsymbol{\theta})$ belong to the generalized location family of distributions, i.e., for all $\boldsymbol{\theta} \in \Theta$

$$f(\mathbf{d}, \boldsymbol{\theta}) = f(\mathbf{d} - \boldsymbol{\mu}(\boldsymbol{\theta}); 0) \quad (4.19)$$

for some function $\boldsymbol{\mu}(\cdot)$.

Theorem IV.1. *Assume that the distribution $f(\mathbf{d}, \boldsymbol{\theta})$ belongs to a generalized location family. Then the inequality (4.11) holds.*

Proof. By definition we have

$$E_{\boldsymbol{\theta}} \left[\ell(\mathbf{D}, \hat{\boldsymbol{\theta}}) \right] = \int \ln \left(f(\mathbf{d}; \hat{\boldsymbol{\theta}}) \right) \ell(\mathbf{d}; \boldsymbol{\theta}) d\mu(\mathbf{d}).$$

Since the distribution is in the location family, it is easily shown that

$$E_{\boldsymbol{\theta}_0} \left[\ell(\mathbf{D}, \hat{\boldsymbol{\theta}}) \right] = E_0 \left[\ell(\mathbf{D}, \hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \right]$$

and therefore

$$E_{\hat{\boldsymbol{\theta}}} \left[\ell(\mathbf{D}, \hat{\boldsymbol{\theta}}) \right] = E_0 [\ell(\mathbf{D}, 0)],$$

where $E_0[\ell]$ denotes $E_{\boldsymbol{\theta}_0}[\ell]$. Consider the expectation of the test function φ under the true model

$$\begin{aligned} E_{\boldsymbol{\theta}_0} \left[\varphi(\mathbf{D}, \hat{\boldsymbol{\theta}}) \right] &= E_{\boldsymbol{\theta}_0} \left[\ell(\mathbf{D}, \hat{\boldsymbol{\theta}}) \right] - E_{\hat{\boldsymbol{\theta}}} \left[\ell(\mathbf{D}, \hat{\boldsymbol{\theta}}) \right] \\ &= \int \ln \left(\frac{f(\mathbf{d}; \hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)}{f(\mathbf{d}; 0)} \right) f(\mathbf{d}; 0) d\mu(\mathbf{d}) \end{aligned}$$

As $\ln(x) \leq x - 1$, the integral on the right can't exceed zero so that

$$E_{\boldsymbol{\theta}_0} \left[\varphi(\mathbf{D}, \hat{\boldsymbol{\theta}}) \right] \leq 0.$$

□

We now return to the simple example presented in Section 4.2.1 to illustrate that the one-sided test (4.12) gives significant improvement in performance relative to the two-sided test (4.10) when the distribution $f(\mathbf{d}, \boldsymbol{\theta})$ is in the location family. Figure 4.5 shows the receiver operating characteristic (ROC) curves for both tests. The ROC of the one-sided test is uniformly better than the two-sided test since it achieves higher power (PD) for any level of false alarm (PFA). This example illustrates how one can exploit knowledge about the nature of the data distribution to implement a better local minimum test. In the next section, we show how additional improvements in performance can be achieved by exploiting application-specific information, for example, knowledge of the forward operator \mathbf{H} of an inverse problem.

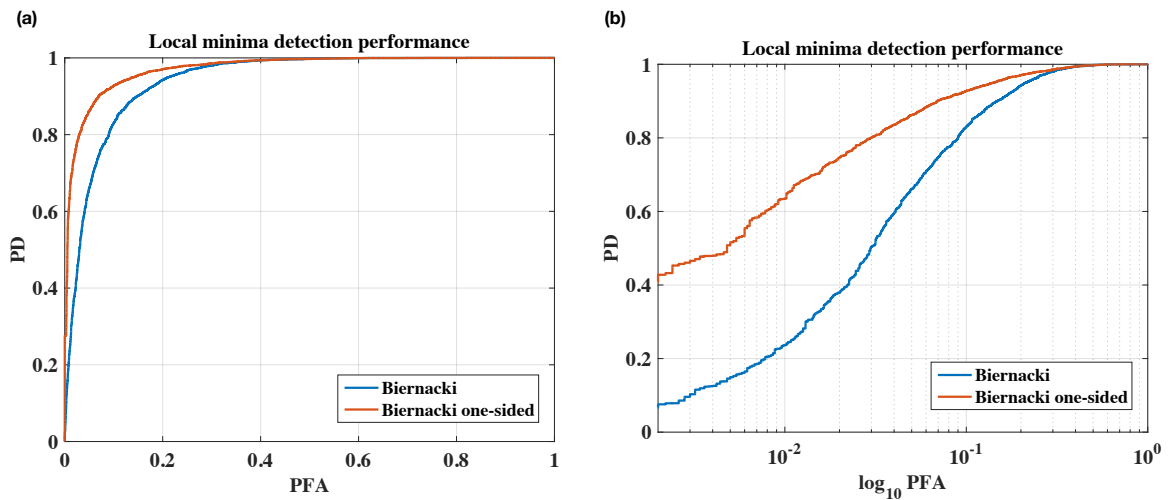


Figure 4.5: Detection performance of Biernacki's two-sided test of a local minimum compared with the one-sided variant for the example problem described by (4.5)-(4.7). For this example problem the data distribution is in the location family and Thm. 1 applies leading to significantly improved performance for the one-sided test.

4.2.3 Reparameterized Embedding

Consider a generalized version of the inverse problem (4.13) with forward model

$$\mathbf{D} = \boldsymbol{\mu}(\boldsymbol{\theta}) + \boldsymbol{\varepsilon},$$

where $\boldsymbol{\mu}(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{\theta}}[\mathbf{D}] \in \mathcal{U}$ is a possibly non-linear forward operator on the parameter vector $\boldsymbol{\theta}$, and $\boldsymbol{\varepsilon}$ is a zero mean noise whose distribution may depend on $\boldsymbol{\theta}$. The log-likelihood is written as $\ell(\mathbf{d}, \boldsymbol{\mu}(\boldsymbol{\theta}))$ to emphasize its dependence on the mean of the data \mathbf{D} parameterized by $\boldsymbol{\theta}$. We introduce the higher dimensional embedding $\tilde{\boldsymbol{\theta}} \in \tilde{\Theta}$ and the corresponding embedded data $\tilde{\mathbf{D}}$ with mean $\tilde{\boldsymbol{\mu}}(\tilde{\boldsymbol{\theta}}) = \mathbb{E}_{\tilde{\boldsymbol{\theta}}}[\tilde{\mathbf{D}}] \in \tilde{\mathcal{U}}$. This reparameterization is defined to be an embedding in the sense that $\mathcal{U} \subseteq \tilde{\mathcal{U}}$, and is a relaxation because the embedding directly implies that

$$\min_{\tilde{\boldsymbol{\theta}}} -\ell(\mathbf{d}, \tilde{\boldsymbol{\mu}}(\tilde{\boldsymbol{\theta}})) \leq \min_{\boldsymbol{\theta}} -\ell(\mathbf{d}, \boldsymbol{\mu}(\boldsymbol{\theta})). \quad (4.20)$$

More generally, the embedding implies a relaxation in the neighborhood of *all* local minima. Let $\hat{\boldsymbol{\theta}}$ be a local minima of $-\ell(\mathbf{d}, \boldsymbol{\mu}(\boldsymbol{\theta}))$, and let $S(\hat{\boldsymbol{\theta}}) \subseteq \tilde{\mathcal{U}}$ be the connected set containing $\boldsymbol{\mu}(\hat{\boldsymbol{\theta}})$ such that $\forall \mathbf{s} \in S, -\ell(\mathbf{d}, \mathbf{s}) \leq -\ell(\mathbf{d}, \boldsymbol{\mu}(\hat{\boldsymbol{\theta}}))$. Then the set S contains all connected points in $\tilde{\mathcal{U}}$ that improve upon ℓ in the neighborhood of $\hat{\boldsymbol{\theta}}$, and we will refer to a minimizer within this set as $\hat{\tilde{\boldsymbol{\theta}}}$. $\hat{\tilde{\boldsymbol{\theta}}}$ is clearly a function of $\hat{\boldsymbol{\theta}}$, however, this dependency is suppressed in the notation to aid readability. Similarly, let $\tilde{\boldsymbol{\theta}}_0$ represent a point in the relaxed space such that $\tilde{\boldsymbol{\mu}}(\tilde{\boldsymbol{\theta}}_0) = \boldsymbol{\mu}(\boldsymbol{\theta}_0)$. It is helpful to think of $\hat{\tilde{\boldsymbol{\theta}}}$ and $\tilde{\boldsymbol{\theta}}_0$ as unique, however, the ideas presented here can easily be modified to accommodate a more general case. The key idea behind using a reparameterized embedding to test for consistency of a root of the log-likelihood is to monitor the gap

$$g(\mathbf{d}, \hat{\boldsymbol{\theta}}, \tilde{\Theta}) = \ell(\mathbf{d}, \tilde{\boldsymbol{\mu}}(\hat{\tilde{\boldsymbol{\theta}}})) - \ell(\mathbf{d}, \boldsymbol{\mu}(\hat{\boldsymbol{\theta}})), \quad (4.21)$$

and exploit distributional differences in this quantity to test for $H_0 : \hat{\boldsymbol{\theta}} = \hat{\boldsymbol{\theta}}_0$ vs. $H_1 : \hat{\boldsymbol{\theta}} \neq \hat{\boldsymbol{\theta}}_0$. Notice that (4.21) takes the form of a generalized likelihood-ratio test between the original parameterization and its relaxation.

Figure 4.6 illustrates this relaxation when viewed relative to the space of expected measurements. For a relaxation that introduces only a few additional statistical degrees of freedom, one expects the relaxation to only permit a small improvement in the neighborhood of the true solution. As before, assume that the Fisher information at $\tilde{\boldsymbol{\theta}}_0$ both exists and is invertible. Then the asymptotic unbiasedness of the MLE, in conjunction with the fact that

the Fisher information increases proportional to the number of independent observations, ensures that in the neighborhood of the true solution $E[\mu(\hat{\theta}_0)] = E[\tilde{\mu}(\hat{\theta}_0)] = \mu(\theta_0)$. By contrast, no such properties exist in the neighborhoods of the local minima. In fact, the entire goal of using a relaxation is to alter the local minima structure away from θ_0 .




Figure 4.6: The reparameterized embedding leads to a relaxation that permits the ML estimator $\hat{\theta}$ to improve within the set $S(\hat{\theta})$. The test for global convergence exploits the distribution differences in the log-likelihood gap between $H_0 : \hat{\theta} = \theta_0$ and $H_1 : \hat{\theta} \neq \theta_0$.

Figure 4.7 illustrates the relaxation when viewed relative to the log-likelihood function's value. The solid and dashed blue lines shown in the foremost plane are the same two functionals described in Section 4.2.2. The solid green line shows the expectation of the log-likelihood after relaxation, where this functional is shown on the same domain as the first two functionals by taking a minimum over the subset of \tilde{U} orthogonal to U . Given a relaxed parameterization space $\tilde{\Theta}$, parametric bootstrap will once again be used to assess the significance of the observed gap, which immediately suggests the test

$$\frac{g(\mathbf{d}, \hat{\theta}, \tilde{\Theta}) - E_{\hat{\theta}}[g(\mathbf{D}, \hat{\theta}, \tilde{\Theta})]}{\sqrt{\text{Var}_{\hat{\theta}}[g(\mathbf{D}, \hat{\theta}, \tilde{\Theta})]}} \underset{H_0}{\overset{H_1}{>}} \tau. \quad (4.22)$$

Central limit theory can be used to show that the statistic on the LHS of (4.22) converges in distribution to $\mathcal{N}(0, 1)$, and thus choosing τ to be the $1 - \alpha$ quantile of the normal distribution will result in an expected false-alarm rate of α .

The form of the test given by (4.22) is often helpful in approximating the threshold

corresponding to particular significance level. Wilks' theorem [203] states that under the null hypothesis with a large number of measurements, $2g(\mathbf{d}, \hat{\boldsymbol{\theta}}, \tilde{\boldsymbol{\Theta}})$ is distributed chi-squared with a number of degrees of freedom corresponding to the difference in dimensionality between the original and relaxed parameterizations. When the relaxations involve non-linear parameterizations, however, determining the number of degrees of freedom generated by the relaxation is often difficult (c.f. [171], example 7.2.3). The suggested form of the test exploits the fact that as $k \rightarrow \infty$, $(\chi_k^2 - k)/\sqrt{2k} \xrightarrow{d} \mathcal{N}(0, 1)$, leading to the property that as the dimensionality of the relaxation increases, the proposed form converges to χ_1^2 . A direct application of the Wilson-Hiferty transform [204] leads to a less intuitive form, but one which converges more rapidly to χ_1^2 . Under the null hypothesis consistency of the MLE ensures that in the neighborhood of the true solution $E_{\theta_0} [g(\mathbf{D}, \hat{\boldsymbol{\theta}}, \tilde{\boldsymbol{\Theta}})] = E_{\hat{\boldsymbol{\theta}}} [g(\mathbf{D}, \hat{\boldsymbol{\theta}}, \tilde{\boldsymbol{\Theta}})]$, and thus the proposed framework for testing based on parameterized embeddings is also a special case of the more general M-testing framework proposed by Blatt and Hero [14].




Figure 4.7: Diagram illustrating how a locally, but not globally, optimal solution $\hat{\boldsymbol{\theta}}$ can be identified by relaxing the parameter space from $\boldsymbol{\Theta}$ to $\tilde{\boldsymbol{\Theta}}$. Under H_1 , minimizing the negative log-likelihood under the relaxation often leads to a relaxed solution $\hat{\boldsymbol{\theta}}$ with a substantially larger gap between its log-likelihood value and the bootstrap estimate (shown in red).

Let's now return to the simple example from Section 4.2.2, with the notation modified slightly to better conform to the discussion on reparameterized embeddings.

$$\hat{\boldsymbol{\theta}}_0 \stackrel{\text{def}}{=} \arg \min_{\boldsymbol{\theta}} \|\boldsymbol{\mu}(\boldsymbol{\theta}) - \mathbf{d}\|^2 \quad \text{s.t.} \quad (4.23)$$

$$\boldsymbol{\mu}(\boldsymbol{\theta}) = \sin(\boldsymbol{\theta}\mathbf{x}), \quad w \in [0, 4\pi] \quad (4.24)$$

Before discussing how one might identify good reparameterizations, first consider the naive choice of $\tilde{\boldsymbol{\mu}}(\tilde{\boldsymbol{\theta}}) = \sin(\tilde{\boldsymbol{\theta}}_0 \mathbf{x} + \tilde{\boldsymbol{\theta}}_1 \mathbf{x}^2 + \dots + \tilde{\boldsymbol{\theta}}_k \mathbf{x}^{k+1})$. This reparameterization permits spatial variation in the instantaneous frequency while implicitly assuming the phase is known. Figure 4.8 compares detection performance between Biernacki’s test and that given by (4.22) when the embedding contains one and three additional degrees of freedom ($k = 1$ and $k = 3$ respectively). This illustrates the potential of the proposed approach, however, many problems do not present themselves with an obvious choice of reparameterization.




Figure 4.8: ROC curve illustrating the potential of reparameterized embeddings when applied to our simple example. The performance of Biernacki’s test is compared with the proposed approach when the naively chosen relaxation has 1 and 3 additional degrees of freedom respectively.

The use of reparameterized embeddings to test for global convergence is based on the idea that the quasi-likelihood function disproportionately benefits from the relaxation under the alternative hypothesis. One is then naturally interested in the marginal cost of the constraints imposed by the intrinsic parameterization because if these costs were known, the directions leading to the greatest marginal benefit conditioned upon the alternative hypothesis would provide a reasonable basis of relaxation. The remainder of this section describes a practical approach for using a forward model to identify such a basis, and the efficacy of the resulting relaxation is illustrated using our simple example.

In constrained optimization, a common interpretation of Lagrange multipliers is the

direction of maximum marginal benefit subject to a set of equality constraints ([135] Chap. 12). When inequality constraints are permitted, this same concept generalizes, and is known as the Karush-Kuhn-Tucker (KKT) conditions. For linear equality constraints, it is trivial to prove that the Lagrange multipliers identify the direction of optimal cost improvement subject to the constraints, and Bertsekas ([12] Chapter 3) establishes this result for non-linear constraints. Silvey [2, 3, 170] studied the hypothesis testing problem for a restricted ML estimator, and developed a Lagrange multiplier test for the hypothesis that the true parameters lie in the subset defined by the restriction. It was later discovered [40] that the Lagrange multiplier tests is identical to Rao's score test [154].

Rao was interested in the locally most powerful test for detecting $H_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$ against $H_1 : \boldsymbol{\theta} = \boldsymbol{\theta}_0 + \boldsymbol{\delta}$. Rao first considered the case where $\boldsymbol{\delta}$ was known. Under this condition, the proportional change in the log-likelihood function when moving from $\boldsymbol{\theta}_0$ to $\boldsymbol{\theta}_0 + \boldsymbol{\delta}$ is given by $\boldsymbol{\delta}^T \nabla_{\boldsymbol{\theta}} \ell(\boldsymbol{\theta}_0) = \boldsymbol{\delta}^T \mathbf{s}(\boldsymbol{\theta}_0)$. This results in the test statistic

$$\xi(\boldsymbol{\delta}) = \frac{[\boldsymbol{\delta}^T \mathbf{s}(\boldsymbol{\theta}_0)]^2}{\boldsymbol{\delta}^T \mathbf{I}(\boldsymbol{\theta}_0) \boldsymbol{\delta}}, \quad (4.25)$$

which is distributed χ_1^2 under H_0 . When $\boldsymbol{\delta}$ is unknown, Rao proposed choosing $\boldsymbol{\delta}$ to maximize (4.25). The resulting test is given by

$$\max_{\boldsymbol{\delta}} \xi(\boldsymbol{\delta}) = \mathbf{s}(\boldsymbol{\theta}_0)^T \mathbf{I}(\boldsymbol{\theta}_0)^{-1} \mathbf{s}(\boldsymbol{\theta}_0). \quad (4.26)$$

Notice that Rao's score is the 2-norm of the log-likelihood gradient after a transformation of variables that causes the iso-likelihood contours in the neighborhood of $\boldsymbol{\theta}_0$ to become spheres. Our approach to identifying candidate relaxations is based upon this test, where we are interested in directions that maximally discriminate between local and global minima.

Rao's score test for constrained and unconstrained ML estimators immediately suggests an approach to this problem. Given a set of nominal conditions $\Theta_0 \subseteq \Theta$, for each $\boldsymbol{\theta}_0 \in \Theta_0$ there possibly exist a non-empty set of local minima $L(\boldsymbol{\theta}_0)$ that are stationary points of the ambiguity function which are not equal to $\boldsymbol{\theta}_0$. Our goal is to identify relaxed parameterizations that maximally discriminate between $H_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$ and $H_1 : \boldsymbol{\theta} \in L(\boldsymbol{\theta}_0)$. Consider a greatly relaxed reparameterized embedding $\tilde{\Theta}$, say the measurement domain. As before, let $\tilde{\boldsymbol{\theta}}_0$ represent a point in the relaxed space such that $\tilde{\boldsymbol{\mu}}(\tilde{\boldsymbol{\theta}}_0) = \boldsymbol{\mu}(\boldsymbol{\theta}_0)$, and let $\tilde{\mathbf{I}}(\tilde{\boldsymbol{\theta}}_0)$, and $\tilde{\mathbf{s}}(\tilde{\boldsymbol{\theta}}_0)$, represent the Fisher information matrix and the score evaluated at the restricted ML estimate respectively. For each local minima encountered, the score function in the relaxed space identifies the direction of greatest improvement of the log-likelihood had all

of the additional degrees of freedom of the relaxed parameterization been available. Using a collection of these points one can identify the single additional degree of freedom that maximally distinguishes between the encountered members in $L(\boldsymbol{\theta}_0)$ and $\boldsymbol{\theta}_0$. The proposed procedure for identifying candidate relaxation dimensions is described by Algorithm 4.9. If more than one additional relaxation dimension is desired, this procedure can be repeated with previously identified dimensions removed from the relaxed space, and included in the restricted estimator.

Input: Θ_0 : Set of nominal parameterizations
Input: Θ_s : Set of starting points
Input: $\tilde{\Theta}$: Relaxed embedding
Output: \mathbf{r} : Relaxation dimension

- 1 **for** $(\boldsymbol{\theta}_0^{(i)}, \boldsymbol{\theta}_s^{(i)}) \in \Theta_0 \times \Theta_s$ **do**
- 2 Generate noise-free data \mathbf{d}_{nf} Solve for the restricted ML estimator $\hat{\boldsymbol{\theta}}(\mathbf{d}_{nf})$ from $\boldsymbol{\theta}_s^{(i)}$ **if** $\hat{\boldsymbol{\theta}} \neq \boldsymbol{\theta}_0^{(i)}$ **then**
- 3 [Record the whitened score function: $[\Delta]_j = \tilde{\mathbf{I}}(\tilde{\boldsymbol{\theta}}_0^{(i)})^{-1/2} \tilde{\mathbf{s}}(\hat{\boldsymbol{\theta}})$
- 4 **Compute** the direction \mathbf{r} that maximizes the inner product with Δ $\mathbf{r} = \mathbf{u}^{(1)} \stackrel{\text{def}}{=} \text{First left singular vector of } \Delta$

Algorithm 4.9: Algorithm for identifying relaxation.

To demonstrate the efficacy of the proposed approach, once again consider our simple example of detecting convergence to a local minima when estimating the frequency of a sinusoid in noise (4.23)-(4.24). The algorithm described in Figure 4.9 was run with the nominal conditions Θ_0 chosen to be 100 equally spaced frequencies over the interval $[0, 4\pi]$, and the relaxed embedding $\tilde{\Theta}$ chosen to be the entire measurement domain $\tilde{\boldsymbol{\mu}}(\tilde{\boldsymbol{\theta}}) \in \mathcal{R}^{100}$. Figure 4.10 shows the additional relaxation dimension \mathbf{r} suggested by the proposed procedure. Using this additional degree of freedom, the relaxed embedding used in conjunction with the test given by (4.22) becomes $\tilde{\boldsymbol{\mu}}(\tilde{\boldsymbol{\theta}}) = \sin([\tilde{\boldsymbol{\theta}}]_1 \mathbf{x}) + [\tilde{\boldsymbol{\theta}}]_2 \mathbf{r}$. Figure 4.11 shows how the relaxed embedding alters the structure of the expected minima of the negative log-likelihood. This figure is analogous to the conceptual diagram shown in Figure 4.7, and the ROC curve associated with the resulting test shown in Figure 4.12.




Figure 4.10: Additional relaxation dimension suggested by the procedure given in Figure 4.9




Figure 4.11: Expected negative log-likelihood associated with the simple example given by (4.23)-(4.24) under the 1-dimensional relaxation provided by the algorithm described in Algorithm 4.9. This figure is analogous to the conceptual plot shown in Algorithm 4.7, and illustrates how a well-chosen relaxed embedding can be used to detect convergence to local minima.




Figure 4.12: ROC curve comparing Biernacki's test to the reparameterized embedding approach when using a single relaxation dimension suggested by the procedure described in Algorithm 4.9

4.3 Identifying Alternate Candidate Solutions

Given a locally optimal solution that fails the proposed test (believed not to be a global minimizer), one needs to identify alternate candidate solutions. The aim of this section is to exploit knowledge of a locally optimal solution to identify a set of alternative candidates that are highly probable to also be stationary points. The proposed approach for generating these candidates is based on identifying a set of perturbations to the locally-optimal blur solution that minimize a point-wise bound on the deviation of the PSF.

Consider the PSF h corresponding to the perturbed phase-screen $\Psi + \beta$. From (2.10) and (2.13) one may write [78]

$$h = c_0 \left| \mathcal{F}^{-1} \left\{ A e^{j\Psi} A_B e^{j\beta} \right\} \right|^2 \quad (4.27)$$

$$= c_0 \left| \mathcal{F}^{-1} \left\{ A e^{j\Psi} \right\} * \mathcal{F}^{-1} \left\{ A_B e^{j\beta} \right\} \right|^2 \quad (4.28)$$

$$= c_0 |g * d|^2, \quad (4.29)$$

where c_0 is a normalizing constant, $A_B = 1_{\text{supp}(A)}$ is the binary aperture corresponding to the support of A , and g and d are the coherent PSF's corresponding to the unperturbed PSF and the perturbing phase-screen under a binary aperture respectively. Letting δ be the Kronecker delta function, and a an arbitrary complex constant such that $|a| = 1$, the $(m, n)^{\text{th}}$ element of the discrete representation of h is then given by

$$[h]_{m,n} = c_0 |\langle g(m-x, n-y), d \rangle|^2 \quad (4.30)$$

$$= c_0 |\langle g(m-x, n-y), a\delta + (d - a\delta) \rangle|^2 \quad (4.31)$$

$$= c_0 |\langle g(m-x, n-y), a\delta \rangle + \langle g(m-x, n-y), d - a\delta \rangle|^2 \quad (4.32)$$

$$= c_0 \left[|g]_{m,n}|^2 + |\langle g(m-x, n-y), d - a\delta \rangle|^2 \right] + \quad (4.33)$$

$$2\Re(a[g]_{m,n}^H \langle g(m-x, n-y), d - a\delta \rangle).$$

The magnitude of the PSF change induced by β at the $(m, n)^{\text{th}}$ element is then given by

$$|[\mathcal{E}]_{m,n}| = \left| [h]_{m,n} - c_0 |g]_{m,n}|^2 \right| \quad (4.34)$$

$$= c_0 \left| \langle g(m-x, n-y), d-a\delta \rangle|^2 + 2\Re(a[g]_{m,n}^H \langle g(m-x, n-y), d-a\delta \rangle) \right| \quad (4.35)$$

$$\leq c_0 \left| \langle g(m-x, n-y), d-a\delta \rangle|^2 + 2|[g]_{m,n}| |\langle g(m-x, n-y), d-a\delta \rangle| \right| \quad (4.36)$$

$$\leq c_0 \left[\|g\|^2 \|d-a\delta\|^2 + 2|[g]_{m,n}| \|g\| \|d-a\delta\| \right] \quad (4.37)$$

$$\leq \|d-a\delta\| \left[\|d-a\delta\| + 2 \frac{|[g]_{m,n}|}{\|g\|} \right]. \quad (4.38)$$

This point-wise bound on the PSF perturbation \mathcal{E} , associated with the wavefront perturbation β , is clearly minimized when $\angle a = \angle [d]_{0,0}$. Under this condition, the right-hand side of (4.38) is monotonic in the Strehl ratio [125] associated with β , which we will denote as $c_0 |[d]_{0,0}(\beta)|^2$. Thus, the set of wavefronts that maximize the Strehl ratio for a fixed RMS strength, also minimizes the worst-case point-wise error in the perturbed PSF. These wavefronts are given by

$$\left\{ \beta = \arg \max_{\tilde{\beta}} \tilde{\beta} c_0 |[d]_{0,0}(\beta)|^2 : \|\tilde{\beta}\|^2 = \tau \right\} \quad (4.39)$$

Determining this set is related to the problem of wavefront balancing, and it is well known that such sets are discontinuous in aberration space, and have no closed form solution [125]. Fortunately, this set is independent of locally optimal phase screen Ψ . We have identified this set of points numerically under a basis containing the first 12 Zernike modes for perturbations up to 0.2 waves RMS, where none of the solutions are trivial associates of each other.

This set (plus its corresponding trivial associates) can be used to identify aberrations that are likely to result in similar values of the log likelihood function given an aberration estimate believed to be a local minima. Figure 4.13 shows a 2D embedding of the points given by (4.39) for RMS perturbations up to 0.09 waves, and Figure 4.14 shows a PSF corresponding to a phase-screen of 0.25 waves RMS as well as the first few perturbations drawn from the set given for $\tau = 0.2$. From this latter plot, it is clear that the proposed approach indeed has the desired property.




Figure 4.13: A 2D embedding of the points within the space of the first 12 Zernike modes that maximize the Strehl ratio for shells of a fixed RMS wavefront deviation. Shells up to 0.09 waves RMS are shown, with the points numbered to indicate correspondence between phase-screens in neighboring shells. Note that the asymmetries are not unusual when embedding high-dimensional surfaces into 2D.

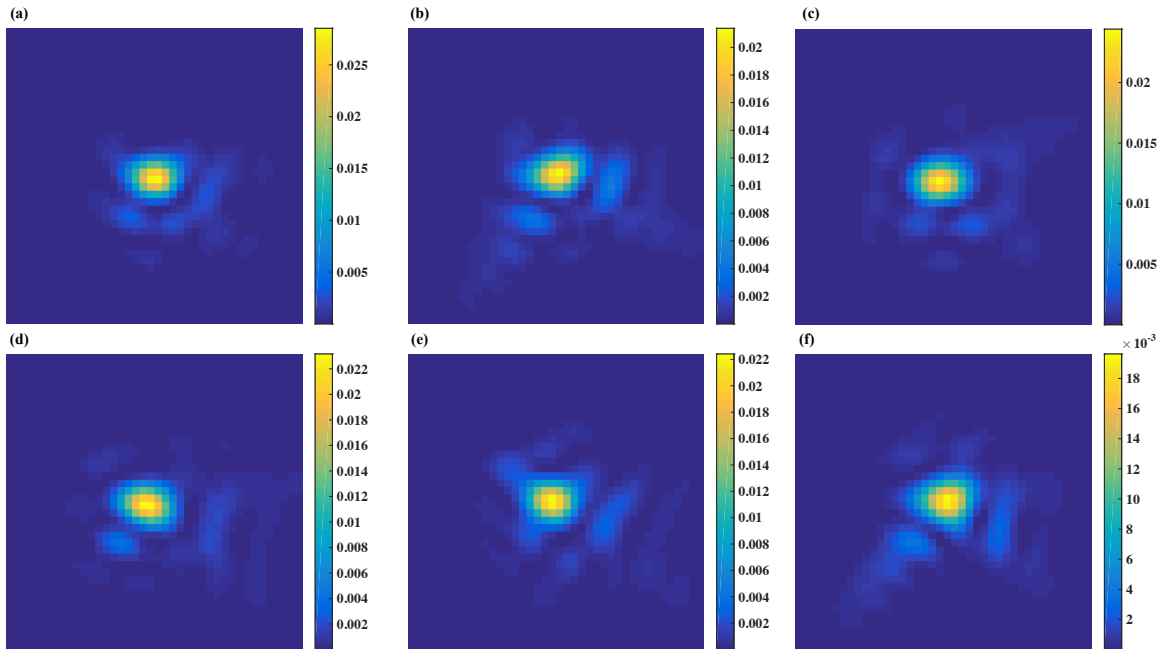


Figure 4.14: A PSF corresponding to a random phase-aberration of 0.25 waves RMS is shown in (a). (b)–(f) show the first few perturbations of 0.2 waves RMS drawn from the set given by (4.39). Each of the PSF's are displayed at 2X the Nyquist rate for the optical system.

4.4 Simulation Experiments

The simulation experiments described in Section 3.8 are repeated without assuming an initial blur estimate in the neighborhood of the true solution. By using initial blur estimates consistent with an ideal (unaberrated) optical system, the effectiveness of the proposed global optimization strategy is compared with both previous findings and the Cramèr-Rao bound. The proposed global optimization strategy is an application-specific implementation of simulated annealing. The computational benefits of incorporating problem-specific knowledge is assessed by comparing the proposed approach to a commercially available general-purpose simulated annealing implementation.

The aforementioned simulation experiments were repeated to study convergence behavior under large wavefront aberrations. Data was generated from the model given by (2.70) over a range of SNR's. For all experiments, the wavefront aberration strength was 0.25 WRMS, and the detector was 2.5 times undersampled in each linear dimension. The SNR, defined as the standard deviation of the noise-free scene divided by the standard deviation of the total noise, was adjusted by increasing or decreasing the sensor's read-noise level. A Sierpinski carpet was used as the calibration target, and for this object the lighting environment and exposure times were such that approximately 2.05×10^5 photons were received per pixel. The detector gain and offset were chosen so that the final recorded data was centered, and filled 70% of a typical 8-bit dynamic range. These gain and offset parameters resulted in an expectation of approximately 2000 photons/count. Figure 4.15 shows the true object alongside a sequence of data frames at varying SNRs generated under these conditions.

Figure 4.16 shows estimation error plotted against SNR for random blurs corresponding to 0.25 WRMS perturbations, where an ideal imaging system as the starting point of the inverse-problem. As expected, achieving global convergence is only possible given data of a sufficiently high SNR. Simulation experiments indicate, however, that this SNR requirement is met under typical imaging conditions. For data chips of only 106 pixels on a side, and SNRs exceeding approximately 50, the proposed estimator is reliably able to locate the globally optimal solution. This is a modest SNR requirement (cf. Figure 4.15), and could undoubtedly be relaxed if one were to use larger image chips. As such, the SNR requirement for achieving global convergence using the proposed application-specific simulated annealing method is not expected to be a practical limitation in most imaging problems.

It is important to note, however, that when the test for a local minima fails to reject the null hypothesis, one cannot conclude that a global optimum has been found. Instead,




Figure 4.15: The true object (a) is shown alongside a series of data frames generated with all parameters fixed other than the SNR. The read-noise standard deviation corresponding to (b)-(d) were 0.68, 1.41, and 7.15 counts respectively.

all that can be said is that the current solution is sufficiently close to the solution in the embedded space that it cannot be ruled out as a local minimum. Figure 4.17 shows a local solution that failed to meet a 1×10^{-4} significance level for detecting a local minimum on data with an SNR of 18. Similarly, Figure 4.18 shows the typical progression of local minima identified during the global search procedure in moderate to high SNR scenarios. The current implementation stops its search when a solution is found that fails to reject the null hypothesis of being a global minimum. If necessary, additional searching could be allowed, and the set of all local minima and their likelihood values returned.




Figure 4.16: Wavefront estimation error is shown as a function of SNR for 0.25 wave perturbations where the starting point is a clear aperture, and the proposed global optimization procedure described in Section 4.2 is used.




Figure 4.17: (a) Shows the true PSF on a Nyquist sampled grid, and phase-screen associated with data simulated with and SNR of 18. (b) shows the local minima which failed to reject the null hypothesis at a power of $1E-4$ due to the elevated noise levels. The actual p-value associated with the local solution was 0.019, and (c) shows the resulting PSF and phase-screen errors.

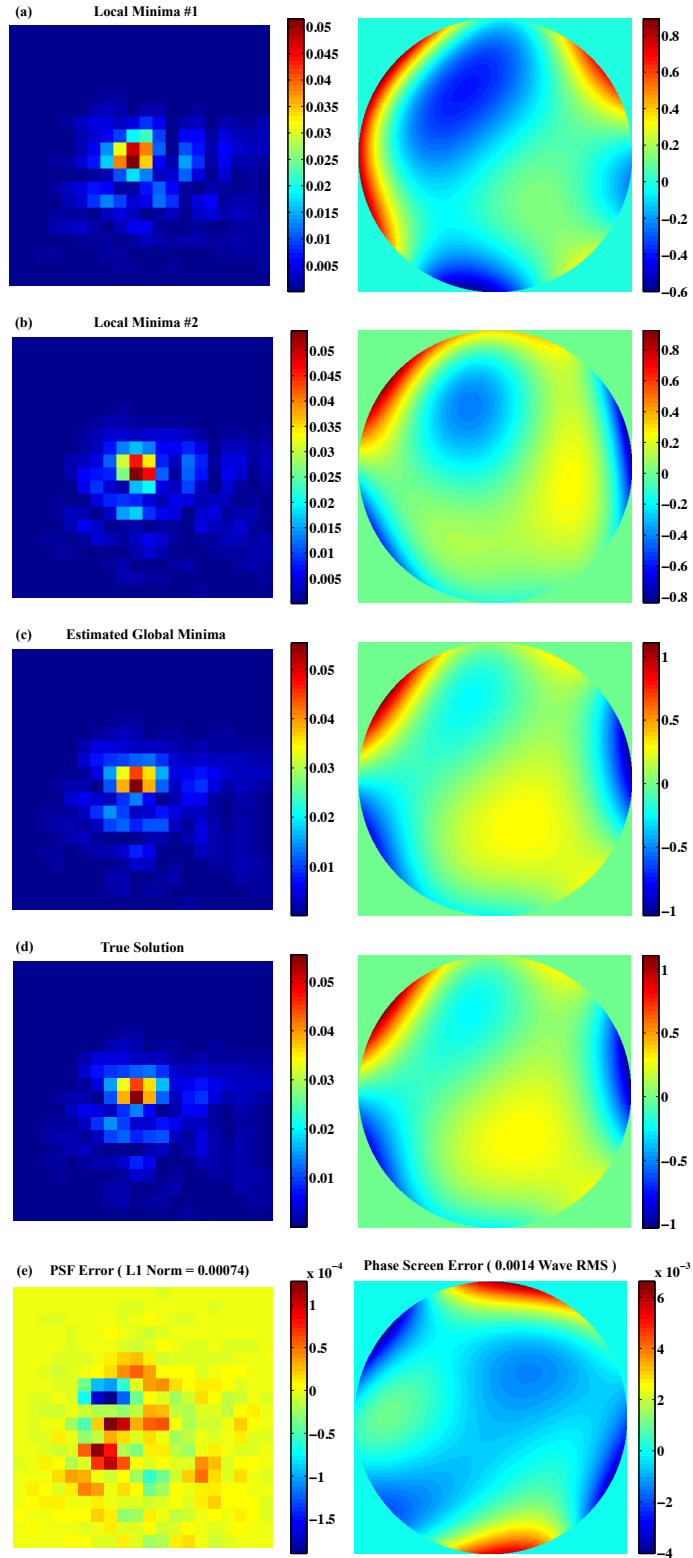


Figure 4.18: (a)-(c) shows a typical succession of PSF's and phase-screen estimates associated with local minima found using the proposed global optimization procedure under high SNR conditions. (d) and (e) show the true solution and the associated errors respectively

4.4.1 Computational Comparison

A natural point of comparison for the proposed problem-specific global optimization strategy (see Figure 4.2) is simulated annealing. Simulated annealing (SA) is a meta-heuristic for identifying heuristic global optimization procedures based on the Metropolis-Hastings algorithm [126]. As such, SA describes a large class of possible optimization approaches, and the performance of any individual SA derived method should not be used to make broader claims about the entire class of procedures. Be that as it may, comparison to widely used SA implementations are useful in providing a sense of scale. All of the simulation results presented in this section utilize the implementation provided by the MATLAB® Optimization Toolbox version 8.0.

Initial studies using simulated annealing indicated that additional problem-specific knowledge must be incorporated into the search strategy to ensure reasonable convergence times. Even after placing problem-specific bounds on the search-space, stock simulated annealing approaches are unlikely to be an efficient method of solving the optimization problem described by (3.8). Under the conditions described in Section 4.4, with an SNR of 100, and without the aid of gradient information, simulated annealing took more than 300 hours to place a single test point within 0.01 WRMS of the true global minima. During this experiment, the optimizer was stopped the first time a point was placed inside the 0.01 WRMS sphere surrounding the true solution, not when convergence was actually reached. Such long runtimes would preclude most practical applications of this work; however, practitioners sometimes use hybrid methods that exploit gradient information to improve convergence rates.

Hybrid simulated annealing approaches append an additional optimization procedure either during or at the end of each annealing stage. The Mathworks implementation applies the hybrid method to each randomly generated point before evaluating the objective function to determine that point's fitness. Consistent with our desire to both exploit gradient information and simultaneously incorporate parameter constraints, a bounded interior-point algorithm was used as part of the hybrid optimizer [198]. This method is designed for non-linear optimization problems and adaptively switches between a line-search method based on direct-factorization of the prime-dual equations and a trust-region method that uses conjugate gradient iterations to guarantee progress toward a stationary point. This is the same algorithm used in the well known KNITRO optimization suite [23, 24], and has been shown to perform well on the industry-standard CUTEst test suite [15, 80, 81]. This hybrid annealing strategy proved to be a reasonable point of comparison for the aforementioned application-specific approach.

A Monte Carlo simulation was performed under the conditions described in Section 4.4,

with an SNR of 100, and the number of aberration parameters varied to alter the problem difficulty. As the number of Zernike modes in the model increases, so too does the probability of encountering local minima, and thus overall search times increase. Figure 4.19 shows




Figure 4.19: Monte-Carlo study of optimization runtimes as a function of the number of aberration modes in the model.

the mean and standard errors of runtimes corresponding to 10 independent realizations of the same camera model with the number of consecutive Noll-ordered Zernike modes in the model ranging from 1 (defocus only) to a 12-mode wavefront description (all wavefront modes of radial orders 2 through 4).

The ability of the proposed application-specific optimization strategy to terminate as soon as a likely global optimum is detected leads to substantial runtime improvements. In the astronomical imaging community wavefront descriptions of optical systems typically include Zernike models up to at least radial-order 3 (7 Zernike modes). For models of this complexity, the proposed global optimization strategy results in 5X improvement in total runtimes. Furthermore, the proposed approach allows one to explicitly control the power of the test against convergence to local minima. In this study, the test was run at a power level of $\alpha = 0.01$. The advantage of a test for global convergence is the ability to terminate a SA procedure once one is confident that a global optimum has been reached. Generic simulated annealing algorithms, such as the one provided in MATLAB[®] Optimization Toolbox 8.0, do not exploit these types of tests leading to unnecessary computation.

To make this idea more concrete, consider a particular SA run under the aforementioned conditions with a 12 Zernike mode wavefront description. During this run 28,906 function evaluations were made over 11.4 hours before the default annealing schedule completed. The hybrid annealing procedure periodically invokes an interior-point solver, and this gradient-based optimization stage causes the test points to concentrate into basins of attraction. New test points are based on random perturbations of previously accepted solutions, and both the acceptance criteria and the hybrid interior-point solver lead to common “lines” of exploration. For the problem at hand, a test point was defined to improve upon an existing search line if it improved upon that line’s log-likelihood value while remaining within 0.04 WRMS of the associated wavefront solution. This situation is illustrated graphically in Figure 4.20.

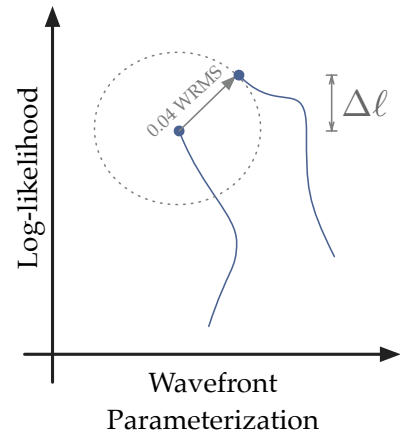


Figure 4.20: A point is said to improve upon a search line if it improves upon the line’s best log-likelihood value using a wavefront solution within 0.04 WRMS of the associated wavefront description.

If a point improves upon the greatest observed log-likelihood function but is not within the neighborhood of an existing line, a new line is formed. Similarly, two lines may merge into a single line if their best points converge to within this wavefront threshold. By defining the notion of a search line, all of the points tested by the SA algorithm may be associated with a much smaller number of lines. In this particular example, only 54 unique search lines were identified. In the limit, every line discovered will converge to a unique local optimum.

Figure 4.21 shows the final objective function value associated with each of the 54 uniquely identified lines plotted against their respective creation times. This figure illustrates that within the first hour of runtime, the simulated annealing algorithm identified all of the local basins of attraction that would eventually be explored. This particular run corresponds to the demonstration software associated with [111], where the random number generation seed was fixed to ensure reproducibility. In this example, the limited-memory quasi-Newton method used to initialize the algorithm happens to converge almost immediately. Figure 4.22 shows the wavefront error of the test points plotted against time. At around 18 minutes, the initialization routine is complete. After this point, the SA algorithm begins testing random points through simulated annealing. The hybrid method uses a gradient-based solver to explore a subset of the randomly generated points, and this repeats at roughly 20-minute intervals. In this example, however, the hybrid optimization technique identifies a point within 0.001 WRMS of the true solution after its first annealing cycle. Our test for global




Figure 4.21: Final objective function value plotted against the creation time for each of the 54 unique lines identified by the SA algorithm. Wavefront solutions corresponding to the global optimum as well as the nearest two local minima are compared with the true solution.

convergence indicates that a global minimum is achieved (p -value < 0.01) after the first annealing round (at 31 minutes), and had such a test been used the subsequent 11 hours of runtime could have been avoided. This example clearly illustrates the benefit of global convergence tests like those studied in [13, 14, 52, 71, 112, 174, 202].

One justification for permitting a SA technique to continue after a globally optimal solution has been identified is to be able to characterize the structure of the local minima. In Section 4.3 we demonstrated how disparate wavefront solutions could lead to PSF's with a similar structure. Figure 4.23 compares the true wavefront and its associated PSF to the 3 most likely solutions identified using simulated annealing. As expected, the most likely solution is indeed the global optimum, and the overall wavefront error of 0.0014 WRMS is consistent with the Cramér-Rao bound for this problem. The two nearest local minima, however, have PSFs that are structurally very similar despite wavefront descriptions that deviate from the true wavefront by 0.117 and 0.115 WRMS respectively. Recall that a random wavefront perturbation on the order of the Maréchal criteria (~ 0.07 WRMS) [123] is considered significant for a well-calibrated optical system. These perturbations, however, are far from random. Despite relatively large wavefront perturbations, the PSF's associated with the solutions at the first two local minima differ from the true solution by only $\sim 1\%$ in energy. The relative closeness of the locally optimal solutions measured through the PSF underscores the difficulty of blur estimation.

Figure 4.24 illustrates the convergence properties of the proposed approach outlined in Figure 4.2. A quasi-Newton method is used to identify local minima that are subsequently




Figure 4.22: Wavefront error plotted versus runtime over the first hour of the simulated annealing run.

tested for being a global minimum using the aforementioned reparameterized embedding technique. In this example, 6 local minima are reached before the global solution is located ($p < 0.01$). The 6 local minima correspond to only 3 unique wavefront solutions, of which two are trivial associates as discussed in Section 3.6. Figure 4.25 compares the wavefront solutions corresponding to the global minima and each of the unique local minima to the true solution. It should be noted that the proposed optimization approach *is* a simulated annealing technique, where both the annealing process and termination criteria are exploiting the specific problem structure.





Figure 4.23: Comparison between the 3 most likely solutions identified using simulated annealing and the true solution.



True	Global			
	Minimum	Local #1	Local #2	Local #3
-0.06241	-0.06203	0.02935	-0.02907	0.09032
0.11359	0.11351	0.05912	-0.05961	-0.01018
-0.07294	-0.07369	-0.01504	0.01517	0.10210
-0.10670	-0.10684	-0.11663	-0.11650	-0.10102
-0.08131	-0.08106	-0.09982	-0.09978	-0.09585
-0.05507	-0.05511	-0.03645	-0.03664	-0.01509
-0.05373	-0.05426	-0.03580	-0.03590	-0.06583
0.01715	0.01711	-0.03313	0.03364	0.04703
-0.01893	-0.01930	-0.06703	0.06680	0.01605
0.05640	0.05601	0.06712	-0.06738	0.02049
-0.08192	-0.08150	-0.03737	0.03746	-0.07944
0.07696	0.07676	0.05711	-0.05639	-0.05066

Figure 4.24: Wavefront error and objective function value plotted as a function of time. This plot is analogous to Figure 4.22, where the addition of the objective function value helps clarify the locations of the local minima.

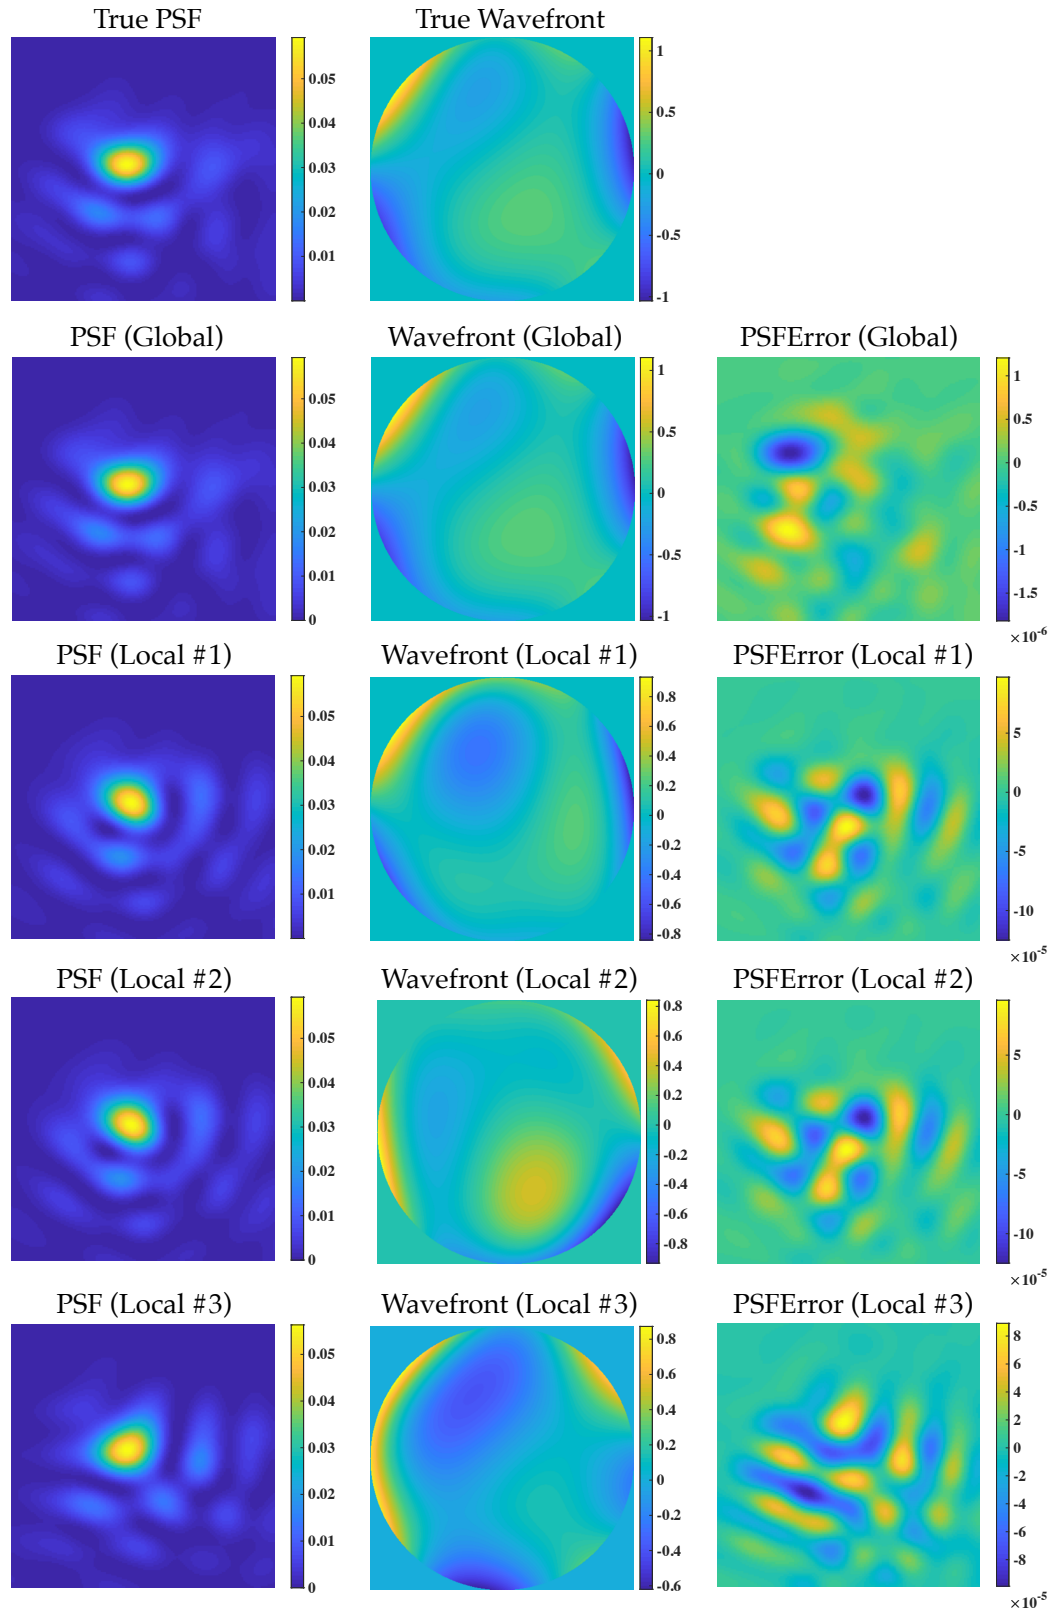


Figure 4.25: Comparison between all unique solutions identified using the proposed optimization approach.

4.5 Hardware Experiments

A hardware experiment was performed to validate the correctness of the forward-models, accuracy of the proposed inverse-problem approach under large wavefront perturbations, and to demonstrate the utility of the resulting system identification solutions for image restoration. A Nikon D7000 was modified to remove its anti-aliasing filter, and this camera was used in conjunction with a Nikkor 70-200 mm f/2.8 VR II lens to collect imagery of a Sierpinski carpet displayed on an iPad, model MC705LL/A, at ranges of 3, 3.2, and 3.3 meters. The camera autofocus was applied against the calibration target at 3 m, and subsequent images were collected with the lens configuration held fixed. A diagram of the experimental setup is shown in Figure 4.26. The focal length was 70 mm, and corresponded to an F-stop of 2.8. The key idea behind the collection geometry is to exploit changes in the calibration-target range to induce known levels of defocus at the focal plane. Estimation is then performed independently over the 3 color channels and 3 target depths to produce 9 independent wavefront estimates. The hypothesis is that the data will be consistent with the model described in (3.2), and that the resulting wavefront estimates will exhibit defocus terms consistent with that induced through the change in depth of field.




Figure 4.26: Experimental setup involving a Nikon D7000 imaging a calibration target a 3 ranges. The devices was auto-focused at the nearest range and then the lens configuration was held fixed for the remaining collections.

4.5.1 Literature Review: Depth from Defocus

Digital signal processing techniques for estimating range from defocus date back to the initial widespread availability of digital cameras in the mid- to late- 1980's. Our primary interest is in studying the value of defocus cues for depth estimation, and as such, we emphasize those works that treat this cue in isolation. Surprisingly, very little work has been done on single-frame depth estimation from data collected with traditional cameras. This is

clearly not due to a lack of interest, but rather is indicative of the very difficult nature of the problem.

Assuming a traditional (monocular, passive) imaging system operating under typical conditions the available depth perception mechanisms are [98]:

Absolute Scale Knowledge of the absolute scale of an object can be used to estimate its range from its perceived scale

Perspective The fact that parallel lines intersect at vanishing points (ideal points) can be used in conjunction with knowledge of the geometry of the scene to estimate depth

Perceived Intensity Lambertian surfaces will result in an image whose intensity is proportional to their solid angle less the atmospheric losses (scattering and absorption) along the intermediary path. The absorption and scattering is typically wavelength dependent, and with sufficient prior knowledge, this could be used to estimate range.

Occlusion Occlusions in the 3D to 2D projection operator provide ordering information related to depth

Shading Illumination effects in conjunction with prior knowledge about the object and its environment can be used to infer relative depth.

Distributional Prior For terrestrial imaging, the presence of the Earth generates a horizon in the image plane, and the distribution of object ranges is skewed about this contour resulting in statistical prior information about object range as a function of its relative position in the image.

Defocus The finite aperture of a traditional camera causes the sensor's point-spread-function to be dependent on range. Objects imaged in front of, or behind, the best focus plane will exhibit additional defocus blur.

Notice that of these range cues, only defocus provides absolute range without prior knowledge of the scene.

Nearly simultaneously, Pentland [143] and Grossmann [84] suggested using blur to estimate depth. Grossmann essentially demonstrates that a correlation exists between "edge width" and range, suggesting that the desired information content may exist in the collected

data. The proposed method for estimating the blur is “ill-defined and sensitive to noise,” leading the authors to ultimately advocate for the use of multiple images with varying sensor settings. Pentland provides a more rigorous theoretical development and illustrates that under good conditions, the human eye has a ranging accuracy based on defocus cues that are about half as accurate as stereopsis. Despite this, Pentland’s approach results in “considerable variability in the depth estimates,” and only relative depth estimation is actually presented. Ultimately, he too suggests a multi-frame approach based on a beam-splitter and varying imaging apertures. This type of multi-frame approach is known as phase-diverse phase-retrieval [75]. Subbarao subsequently pursued both the single-frame [180], and multi-frame ideas [179]. In Subbarao’s single-frame work he explicitly calls out the incongruous nature of using a Gaussian blur model in conjunction with simplistic geometric optics models, but ultimately uses a best-fit Gaussian with 1 unknown. Two more experimentally derived constants are introduced to fit the change of the defocus with range. Depth estimation errors on the order of a few percent were reported, but how much of this accuracy should be attributed to the fitting of the 3 unknowns is unclear. Subbarao also reported noise sensitivity but did not provide details on the “noise cleaning” procedure used to mitigate this problem. Perhaps because of these early difficulties, subsequent work in the literature is almost exclusively focused on multi-frame collections.

Ens [53] suggests a two-frame approach based on estimating the change in the camera’s transfer functions in each image using regularized circulant matrix approximation. Ens notes the relationship between his technique and the use of a windowed Fourier transform, and subsequent approaches will often be described as either “image domain” or “frequency domain”, despite the fact that they all will act on local regions within collected imagery. Subbarao and Surya [181] and Gokstorp [74] would nearly simultaneously release similar techniques referred to acting in the spatial- and frequency-domains respectively, and the bulk of the literature on depth from defocus would follow this same multi-frame approach (c.f. Rajagopalan [149, 150], Xiong [206], Watanabe [200], Chaudhuri [31], Ziou [216], Park [137], Joshi [101], Malik [122], Yang [209]).

Early work in single-frame depth from defocus convincingly suggested that object range can be recovered under some scene geometries when a sufficient amount of information is known about the camera. One could argue, however, that for typical imaging scenarios, this capability has not been demonstrated. Early work focused on producing “depth maps”; images correlated with actual scene ranges, but not sufficiently calibrated to be reported in units of length.

The more recent literature has approached this problem in two ways. One set of research findings has chosen to focus on *relative* depth, typically with increased segmentation

capabilities in mind (c.f. Namboodiri [132], Chan [28], Shuo [215], Sun [182]). The other set considers the impact of chromatic effects (c.f. Garcia [72], Guichard [86], and Trouvé [189]). Saxena [162] presents reasonably good results toward absolute range estimation based on an ad-hoc machine learning approach. This approach, however, needs to be trained for a particular camera. A severe limitation of Saxena’s technique is its inability to accommodate cameras with the ability to focus. Lin [118] is susceptible to this same criticism, and presents an approach that appears to depend critically on an ad-hoc regularization strategy. Exploiting chromatic dependence of the PSF to estimate depth from defocus could have considerable potential. Unfortunately, the aforementioned approaches either exploit or advocate for custom optics specifically designed to enhance chromatic differences between the channels. The degree to which this phenomenon can be exploited in conventional systems remains an open question.

4.5.2 Estimating Depth from Defocus

Before proceeding with the results of the experiment, a small amount of additional theory is needed to relate key geometric quantities of the camera system to the generalized aberration model describe in Section 2.2.3. Let D_{in} and D_{out} represent entrance and exit pupil diameters respectively, f the focal length, and r the range to best focus. Recall that the pupil magnification is defined by the ratio of exit- to entrance-pupil diameters $m_p = D_{out}/D_{in}$, and let s be defined with respect to r and f through the Gaussian lens formula $1/r + 1/s = 1/f$. Then the peak-to-valley (PV) defocus error induced by a change in range of the target is given by

$$\delta_{PV} = (s_d - s) + \sqrt{[s - (m_p - 1)f]^2 - \left(\frac{D_{out}}{2}\right)^2} - \sqrt{[s_d - (m_p - 1)f]^2 - \left(\frac{D_{out}}{2}\right)^2}, \quad (4.40)$$

where the relationship between peak-to-valley and RMS defocus error is further given by $\delta_{PV} = \delta_{RMS}\sqrt{12}$ for a circular aperture. A full derivation, and its relation to commonly used measures of defocus, is provided in Appendix E.

For the aforementioned experimental setup, a defocus at the Maréchal criterion corresponds to a beam converging $0.336 \mu\text{m}$ from best focus, which in turn corresponds to a translation of the calibration target of 13.2 mm along the optical axis. The individual detector elements within the Bayer filtered sensor will sample the impinging wavefront at 56.5X undersampled (7.518X undersampled in each linear dimension). Intensity variations induced by non-uniform transmission from the iPad’s display, and wavelength dependencies in both

the transmission of the optics, as well as the sensitivity of the CCD array, were calibrated by collecting flat-field images and then adjusting the known object in the forward-model. Figure 4.27 shows the raw sensor data as well as the corresponding PSF estimates on the aliased data grid. Figure 4.28 shows the independent PSF estimates on a common grid that is Nyquist sampled for the shortest wavelength (blue channel at 450 nm), and the corresponding Zernike coefficients are listed under the Noll ordering starting with defocus. Estimation was performed under a narrowband model with center wavelengths of 650, 550, and 450 nm corresponding to the red, green, and blue channels respectively. These estimates correspond to a predicted global optimum, and the resulting residuals indicate a good fit between the model and the measured data. Figure 4.29 shows a zoom in of the red channel at 3.2 meters, as well as the predicted data, its residual, and a Q-Q plot comparing the sample quantiles of the residuals to their associated theoretical distribution. Finally, the 9 independent defocus estimates were used to predict the range to best-focus for each of the 3 color channels, as well as the pupil magnification. The ranges to best focus for each of the 3 color channels most consistent with the wavefront estimates were 3.024, 2.979, 3.058 meters for the red, green, and blue channels respectively. The autofocus algorithm reported the best focus range to be 3.162 meters in the camera's meta-data. The corresponding pupil magnification estimate was 0.985, which is consistent with an essentially telecentric design. Nikon USA was contacted regarding the actual pupil magnification, and they responded by saying it is "not made available to the public."

Figure 4.30 shows the model predictions for defocus vs range based on the estimated best-focus distances for each color channel and pupil magnification. The 9 defocus terms from the aberration estimates are also plotted against this model. The deviations from the model, in terms of the mean absolute defocus perturbation across the 3 target positions, is 0.0067, 0.0145, 0.0007 WRMS at the reference wavelength for the red, green, and blue channels respectively. These errors are about 5 times the Cramèr-Rao bound for the given conditions, but correspond to relative depth from defocus errors of only 1.2, 3.0, and 0.1 mm along the optical axis.

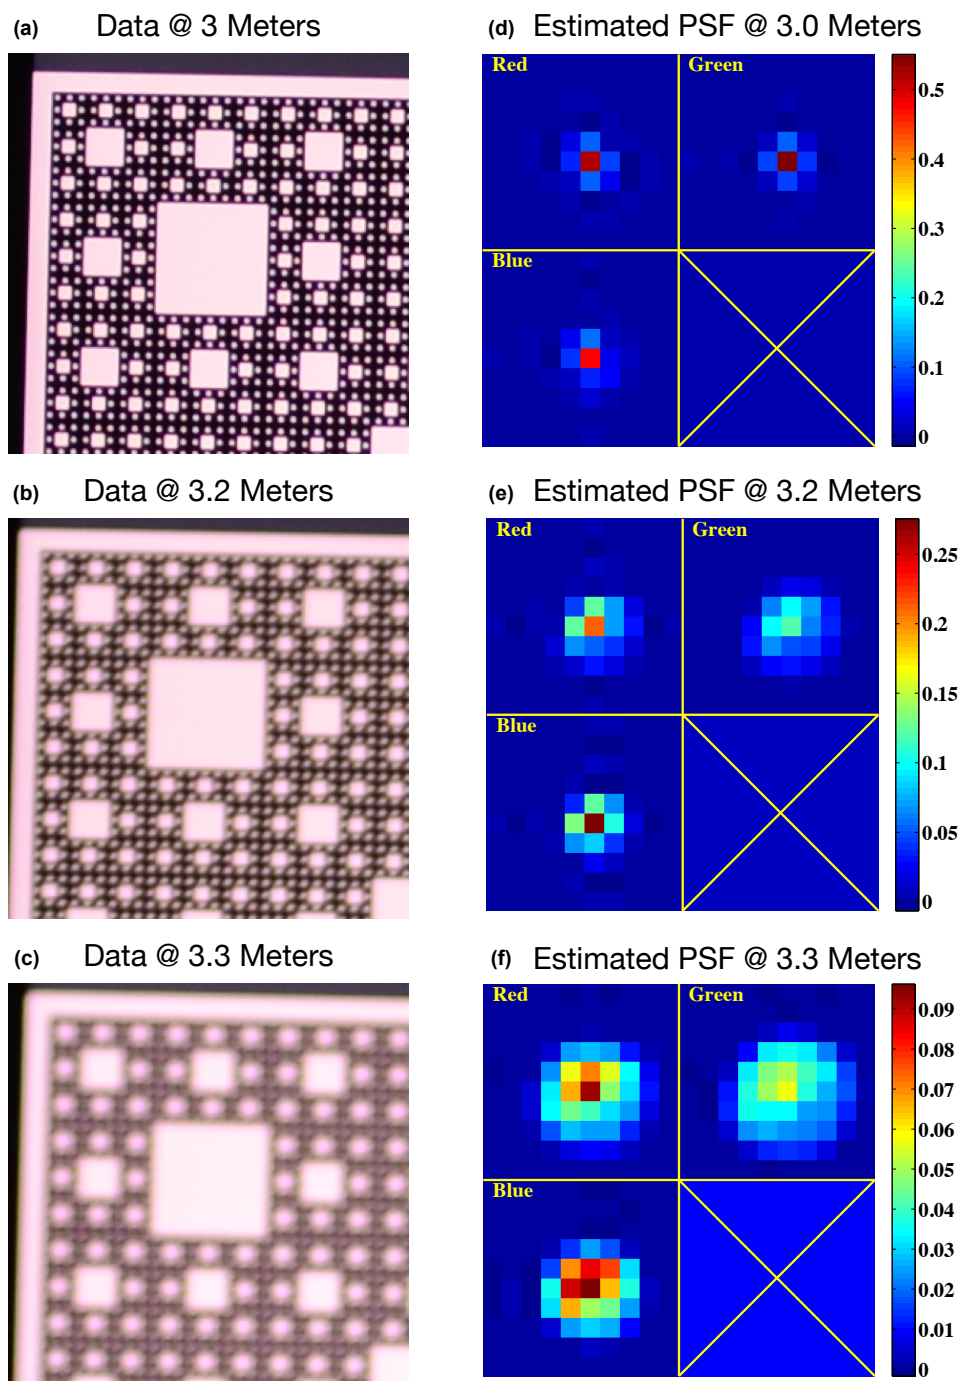



Figure 4.27: The data collected at 3.0, 3.2, and 3.3 meters is shown in (a)-(c) with the corresponding PSF estimates shown on the data grid in (d)-(f).


(a) Nyquist PSF @ 3 Meters



(d) Zernike Coefficients @ 3.0 m (RGB)

-0.13377	0.13304	-0.30905
-0.15575	-0.060195	-0.19207
0.00081206	0.067365	0.16844
-0.029302	-0.019711	-0.045141
-0.037424	-0.023149	-0.048895
-0.048346	-0.06558	-0.0062841
-0.0073797	0.12884	-0.024387
0.31027	-0.1366	-0.26275
-0.045268	-0.0064804	-0.033715
0.037648	-0.040146	0.070114
-0.0070822	0.038886	-0.01356
-0.014846	0.35761	-0.029844


(b) Nyquist PSF @ 3.2 Meters



(e) Zernike Coefficients @ 3.2 m (RGB)

0.89435	1.1186	0.7319
-0.08829	-0.10483	-0.068363
0.032858	0.041414	0.034394
-0.067338	-0.062135	-0.035213
-0.048391	-0.040263	-0.022473
-0.0042936	-0.01253	0.0030386
-0.0024232	-0.0059843	0.0029309
0.10342	0.078203	0.078216
-0.0067833	-0.0060826	-0.0055738
0.0064144	0.0085517	0.0079425
-0.0085396	-0.010338	0.0025812
-0.015188	-0.017207	-0.015635

(c) Nyquist PSF @ 3.3 Meters



(f) Zernike Coefficients @ 3.3 m (RGB)

1.354	1.6012	1.173
-0.080635	-0.098263	-0.05955
0.052402	0.070063	0.063945
-0.067565	-0.06915	-0.052975
-0.035451	-0.049503	-0.030681
0.0024025	-0.002559	0.0018034
-0.016745	-0.0050563	0.00068801
0.099293	0.074585	0.058477
-0.0043329	-0.0027635	-0.0087944
0.008114	0.0064652	0.013889
-0.01973	-0.015459	-0.017472
-0.015629	-0.015984	-0.011554

Figure 4.28: (a)-(c) The PSF estimates corresponding to the 3 target locations and 3 color channels are plotted on a common grid that is Nyquist sampled for the shortest wavelength of light (450 nm). (d)-(f) provides the corresponding 12 Zernike coefficients under the Noll ordering starting from defocus.

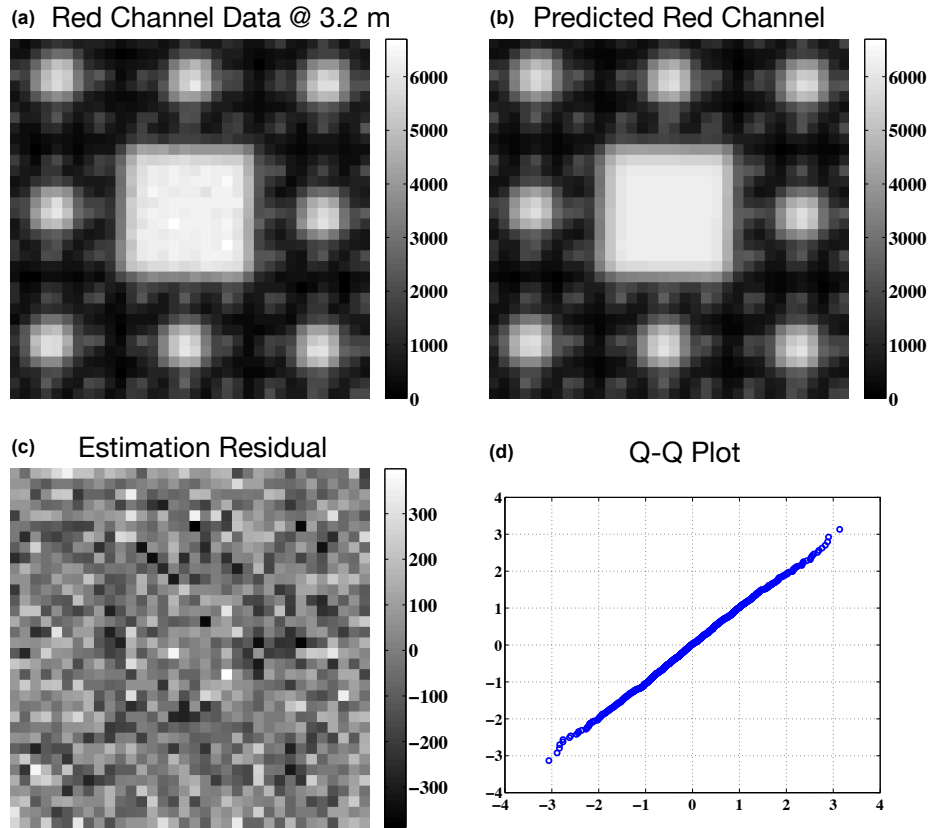


Figure 4.29: A zoom in of the red channel data at 3.2 m (a), the predicted data (b) shown on the same color axis, the residual (c), and a Q-Q plot for the residual after accounting for the intensity dependency.

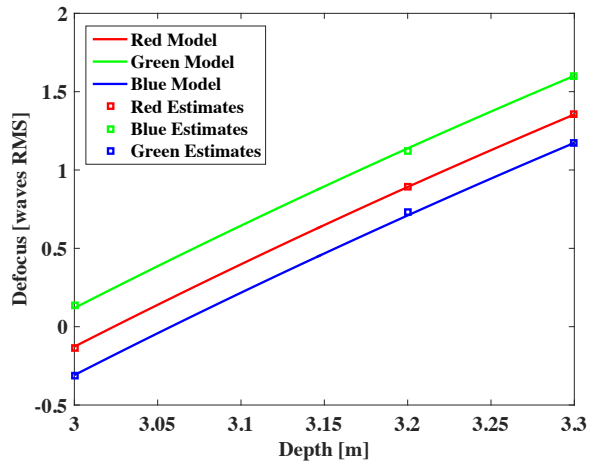


Figure 4.30: Predicted RMS defocus vs target depth based on the proposed camera model for each of the 3 RGB color channels assuming best-focus positions of 3.024, 2.979, and 3.058 meters respectively and a pupil magnification of 0.985. The 9 defocus terms from the PSF estimates are plotted as squares.

4.5.3 Restoration

This section uses data collected with the Nikon to demonstrate the utility of the recovered system information for image restoration. The solution to the system identification problem can be used in conjunction with the data model from (3.2) to perform image restoration by now treating the imaging system as *known*, and a Nyquist sampled representation of the object as *unknown*. Applying the same inverse-problem framework, but now treating \mathbf{f} as the only unknown, the regularized ML estimator for the object becomes

$$\hat{\mathbf{f}} = \arg \min_{\mathbf{f}} \ell(\mathbf{d}|\mathbf{f}) + \lambda R(\mathbf{f}) \quad \text{s.t.} \quad \mathbf{f} \geq 0, \quad (4.41)$$

where the log-likelihood function is once again given by (3.5). The regularizer $R(\mathbf{f})$ is a multi-dimensional variation of the classic roughness measure [88, 152]

$$R_m(f) = \int \left(\frac{\partial^m f}{\partial x^m} \right)^2 dx, \quad (4.42)$$

for $m = 2$. Details of this roughness penalty, and its invariance properties under the object description given in Section 2.4.1 are provided in Appendix K. Unlike system identification, regularization is required for object reconstruction because the number of degrees of freedom of the unknowns exceeds the number of measurements. For this experiment a single color channel (Bayer pattern element) from a single frame is being used to estimate the full Nyquist sampled object. The particular camera settings resulted in the individual sensing elements being 7.518X undersampled in each linear dimension, and the 2×2 , Bayer filter reduces any individual color channel by an additional factor of 2, resulting in data that is more than 15X undersampled in each linear dimension. This regularizer may be viewed as conservative compared to edge-preserving techniques (c.f. [212]), however, this approach ensures that recovered information is brought about by the inverse-problem framework rather than the regularization strategy.

The goal in performing restorations is to demonstrate the value of the information that results from the proposed joint inverse-problem approach for blur estimation when applied to image restoration, as well as the utility of the high-fidelity models for image restoration itself. To this end, Figure 4.31 shows single channel restoration results based on the recovered PSF information within both the Richardson-Lucy (RL) restoration algorithm [67, 119, 156], as well as the higher-fidelity models presented in this work. The red channel of a defocused image collected at 3.2 meters during the experiment described in Section 4.5.2. All restorations are shown on the Nyquist grid with RMS errors provided in the lower left

for reference. When the RL algorithm is provided the additional PSF information it clearly produces a better reconstruction, indicating that the recovered camera information can be useful within restoration techniques that use simpler camera models. The solution to full inverse problem from (4.41) is shown using strict ML ($\lambda = 0$), the defocus only PSF ($\lambda = 300$), and the full PSF description ($\lambda = 300$).

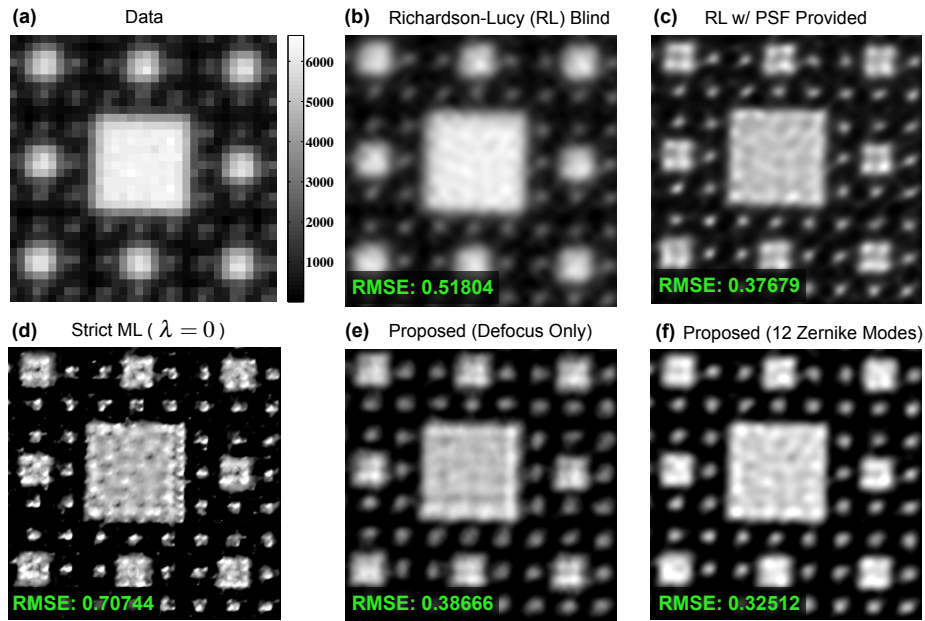


Figure 4.31: Single frame, red-channel restoration results based on the defocused data (a) collected at 3.2 meters. The Richardson-Lucy technique is shown operating both blindly (b) and with the aid of the PSF information from the proposed blur estimation algorithm (c). The proposed ML solution of (4.41) for $\lambda = 300$ is shown in (f), and similar results using the strict ML estimate as well as regularized estimation with only the defocus term are shown in (d) and (e) respectively. RMSE's are shown for all restorations in the lower-left corner.

CHAPTER V

Conclusions and Ongoing Work

5.1 Summary

The central tenet of this dissertation is that by understanding the physical relationships between the environment, the imaging system, and the collected data, one can more effectively use available measurements to make inferences about the sensor, the scene being imaged, or both. With this in mind, a natural question arises regarding existing image restoration approaches; “Given a model of the form $\mathbf{d} = \mathbf{A}\mathbf{f} + \boldsymbol{\epsilon}$, where the aim is to restore an unknown object \mathbf{f} from some collected data \mathbf{d} , to what extent can one reasonably assume \mathbf{A} is known?”. This question then leads to a series of subsequent questions that this work either answers, or provides a foundation for answering.

1. What is the correct structure of \mathbf{A} when working with modern cameras?
2. How should \mathbf{A} be parameterized?
3. Can the structure of \mathbf{A} be estimated by exploiting knowledge of the scene?
4. When is the parameterization of \mathbf{A} unique, and what is the structure of its ambiguities?
5. Can the non-linear relationship between the natural parameterization of \mathbf{A} and the data be overcome?
6. How does improved knowledge of \mathbf{A} improve ones ability to estimate ?
7. How much must one know about the scene \mathbf{f} to recover \mathbf{A} ?

The remainder of this chapter summarizes the progress to date, presents some ongoing work, and discusses future directions.

Chapter II lays the foundation for our approach to this problem. A continuous-to-discrete, physically motivated, forward imaging model is developed. Each component within this

model is tied directly to our understanding of the physical imaging process, and this close proximity leads to clear traceability of the modeling assumptions. Constructing the model in this way also ensures a one-to-one relationship between physical imaging devices and modeling parameters. For example, suppose one is interested in estimating defocus in a camera. Stopping the device down (reducing its collection aperture) is neatly accommodated, and does not interact with the wavefront description of this unknown. The proposed approach also accommodates physical effects that had previously been ignored, and which are critical for accurate system characterization. The most notable of these being diffraction, aliasing, and the use of Bayer filters. Finally, this chapter places the forward model within a statistical inverse-problem framework, and then formulates the system identification question as a joint inference problem.

Chapter III constructs a joint ML estimator for the unknown system parameters, and studies its properties. From the perspective of blur estimation, this results in a type of phase-retrieval problem known as wavefront-estimation. Surprisingly, wavefront estimation from single extended scene does not appear to have been previously studied. The estimator is stated in terms of a non-convex optimization problem, and is shown to be set identifiable up to a set of measure at most 2. The Cramér-Rao bound is computed for this estimator, and the bound is used to show that the data fundamentally includes enough information to permit use of wavefront estimation in the presence of the nuisance parameters. Simulation experiments demonstrate that the ML estimator is essentially unbiased, and is able to recover the unknown wavefront to within 2dB of the CRB. Finally, the estimation problem's bias-variance trade-off is explored through the uniform CRB. As expected, the estimator's bias-gradient was found to be very small. In conjunction with knowledge of the Fisher information matrix, an improved, biased estimator was constructed, but its performance improvements were shown to be minimal.

Chapter IV addresses the non-convexity of the ML estimator in the presence of large aberrations. For sufficiently small perturbations, or with sufficient prior information, one can ensure that the initial wavefront estimate starts in the neighborhood of the true solution. This condition, however, cannot be assured in many practical situations. Highly under-sampled systems can include large wavefront deviations while maintaining a PSF that is point-like with respect to its coarsely sampled detector. This challenge is addressed through two new contributions. An improved test for global convergence is constructed based on a relaxation of the problem's parameterization. For the problem at hand, this new test is shown to substantially outperform existing techniques, and is sufficiently accurate to enable a robust global-optimization strategy. This new test is then paired with a minimax L1 heuristic for identifying new candidate wavefront solutions given a local minima. The utility of this

approach is that one can pre-compute a series of perturbations which, with high-probability, will result in similar PSF structures, and hence similar log-likelihood values. Simulation experiments are used to confirm the efficacy of the resulting global optimization strategy, and it enables hardware simulations.

Using a commercially available camera, known defocus amounts are introduced by translating the calibration target along the sensor’s optical axis. The proposed wavefront estimation technique is then used to determine these relative defocus amounts by independently estimating (without exploiting knowledge of the type of perturbation) the blur from the collected data. This experiment not only lends credence to the proposed approach, but represents a depth-from-defocus capability well beyond what has been reported in the literature.

Finally, this chapter illustrates how the estimation framework can be used for restoration. The informational value of the recovered camera description is explored by comparing restorations made with standard deconvolution techniques (with and without the aid of the estimated blur) and by directly restoring using the full camera models proposed by this work. These results clearly indicate that (1) the recovered information is helpful in image restoration, and (2) image restoration is best performed using a full statistical inverse-problem approach.

5.2 Ongoing Work

The gestalt of the aforementioned work is the exploitation of low-dimensional object representations for system identification, and this presents ongoing work toward allowing not only ad-hoc environments, but also ad-hoc objects. So long as the object models have low-dimensional parameterizations, one can expect that for non-trivial scenes, similar types of performance can be achieved. In doing so, this section takes this work one additional step closer to blind system calibration when fiducials can be identified in the data.

5.2.1 Continuous Representation

Fiducial information can be used to blindly establish unknown system parameters when the object is otherwise unknown, but can be partially described by low-dimensional constituent objects. While a very broad class of potential fiducials could be used, the idea will first be explored through a set of restricted “edges”. For the purposes of this initial exploration, an edge will be defined to be the line which separates two regions of constant reflectance. Using this definition, an edge has a minimal representation given by $\boldsymbol{\theta}_f \stackrel{\text{def}}{=} (\theta_a, \|\mathbf{a}\|, \ell_1, \ell_2)$, where θ_a gives the angle to the vector orthogonal to the line, $\|\mathbf{a}\|$ is




Figure 5.1: Geometry of an edge

the distance to the line, and ℓ_1, ℓ_2 represent the radiance on the near and far side of the line respectively. It follows immediately, that the resulting object is given by

$$f(\boldsymbol{\theta}_f) = \left\{ \begin{array}{ll} \ell_1 & \forall \mathbf{p} : \langle \mathbf{p}, \hat{\mathbf{a}} \rangle - \|\mathbf{a}\| < 0 \\ \ell_2 & \text{otherwise} \end{array} \right\} \quad \text{s.t.} \quad \begin{array}{l} \mathbf{a} = \|\mathbf{a}\| (\cos(\theta_a), \sin(\theta_a)) \\ \mathbf{p} = \mathbf{a} + r_p \hat{\mathbf{a}} + t_p \hat{\mathbf{n}} \end{array} \quad (5.1)$$

$$= \ell_1 + (\ell_2 - \ell_1) u(\langle \mathbf{p}, \hat{\mathbf{a}} \rangle - \|\mathbf{a}\|) \quad (5.2)$$

$$= \ell_1 + (\ell_2 - \ell_1) u(r_p) \quad (5.3)$$

where $u(\cdot)$ is the unit-step function. Thus the forward model for a single fiducial becomes

$$\mathbf{d} = c_3 \mathcal{P} \{ TSB(\boldsymbol{\alpha}) f(\boldsymbol{\theta}_f) \} + c_4 + \boldsymbol{\varepsilon}_r, \quad (5.4)$$

where the warp operator may be omitted because its action on an edge remains in the class of edges.

For the moment, consider the differential of a single edge with respect to its unknowns.

$$\nabla_{\boldsymbol{\theta}} f = \mathbf{e}_3 + (\mathbf{e}_4 - \mathbf{e}_3) u(\langle \mathbf{p}, \hat{\mathbf{a}} \rangle - \|\mathbf{a}\|) + (\ell_2 - \ell_1) \delta(\langle \mathbf{p}, \hat{\mathbf{a}} \rangle - \|\mathbf{a}\|) \nabla_{\boldsymbol{\theta}} (\langle \mathbf{p}, \hat{\mathbf{a}} \rangle - \|\mathbf{a}\|) \quad (5.5)$$

$$\nabla_{\boldsymbol{\theta}} (\langle \mathbf{p}, \hat{\mathbf{a}} \rangle - \|\mathbf{a}\|) = \left[\langle \mathbf{p}, \hat{\mathbf{n}} \rangle \quad -1 \quad 0 \quad 0 \right]^T \quad (5.6)$$

back-substituting, and writing the result in vector form one gets

$$\nabla_{\boldsymbol{\theta}} \mathbf{f} = \begin{bmatrix} (\ell_2 - \ell_1) \langle \mathbf{p}, \hat{\mathbf{n}} \rangle \delta(\langle \mathbf{p}, \hat{\mathbf{a}} \rangle - \|\mathbf{a}\|) \\ -(\ell_2 - \ell_1) \delta(\langle \mathbf{p}, \hat{\mathbf{a}} \rangle - \|\mathbf{a}\|) \\ 1 - \mathbf{u}(\langle \mathbf{p}, \hat{\mathbf{a}} \rangle - \|\mathbf{a}\|) \\ \mathbf{u}(\langle \mathbf{p}, \hat{\mathbf{a}} \rangle - \|\mathbf{a}\|) \end{bmatrix}^T \quad (5.7)$$

5.2.2 Sufficient Representations For Imaging

The non-differentiability of the continuous edge-representation is not problematic in the context of imaging due to the smooth nature of the overall likelihood function associated the larger imaging model. This smoothness follows immediately from the band-limiting nature of the blur operator. This also implies that the proper place to numerically evaluate the continuous-to-discrete forward operator is at this blurred object. That is, for all intents and purposes our “object” will be defined through the discrete blurred representation of the edge. Following this idea through mathematically we may write

$$\mathbf{Bf}(\mathbf{p}) \stackrel{\text{def}}{=} \int \mathbf{f}(\mathbf{x}) \mathbf{h}(\mathbf{p} - \mathbf{x}) d\mathbf{x} \quad (5.8)$$

$$= \ell_1 + (\ell_2 - \ell_1) \int \int \mathbf{u}(r_x) \mathbf{h}((r_p - r_x)\hat{\mathbf{a}} + (t_p - t_x)\hat{\mathbf{n}}) dt_x dr_x \quad (5.9)$$

$$= \ell_1 + (\ell_2 - \ell_1) \int \mathbf{u}(r_x) \int \mathbf{h}(r_x, t_x) dt_x dr_x \quad (5.10)$$

Here it is understood that $\mathbf{u}(r_x)$ remains a 2D function despite the fact that it is being indexed by a single parameter. Furthermore, it is convenient to define $\mathbf{h}(\cdot)$ as the marginal of the PSF w.r.t. $\hat{\mathbf{n}}$. Using this new definition one may write

$$\mathbf{Bf}(\mathbf{p}) = \ell_1 + (\ell_2 - \ell_1) \int_0^{\infty} \mathbf{h}(r_p - r_x) dr_x \quad (5.11)$$

$$= \ell_1 + (\ell_2 - \ell_1) \int_{-\infty}^{r_p} \mathbf{h}(r_x) dr_x \quad (5.12)$$

Clearly this is a well-defined, differentiable function which, when combined with the imaging system’s parameters, will have a discrete representation. The derivative w.r.t. the object parameters also have well defined forms in terms of the line-spread function as given

below.

$$\mathbf{B}\nabla_{\theta_a}\mathbf{f} = (\ell_2 - \ell_1) \int \langle \mathbf{x}, \hat{\mathbf{n}} \rangle \delta(\langle \mathbf{x}, \hat{\mathbf{a}} \rangle - \|\mathbf{a}\|) \mathbf{h}(\mathbf{p} - \mathbf{x}) d\mathbf{x} \quad (5.13)$$

$$= (\ell_2 - \ell_1) \int \int t_x \delta(r_x) \mathbf{h}((r_p - r_x)\hat{\mathbf{a}} + (t_p - t_x)\hat{\mathbf{n}}) dt_x dr_x \quad (5.14)$$

$$= (\ell_2 - \ell_1) \int t_x \mathbf{h}(r_p \hat{\mathbf{a}} + (t_p - t_x)\hat{\mathbf{n}}) dt_x \quad (5.15)$$

$$= (\ell_2 - \ell_1) \left[t_p \mathbf{h}(r_p) - \int t_x \mathbf{h}(r_p \hat{\mathbf{a}} + t_x \hat{\mathbf{n}}) dt_x \right] \quad (5.16)$$

$$\mathbf{B}\nabla_{\|\mathbf{a}\|}\mathbf{f} = -(\ell_2 - \ell_1) \int \delta(\langle \mathbf{x}, \hat{\mathbf{a}} \rangle - \|\mathbf{a}\|) \mathbf{h}(\mathbf{p} - \mathbf{x}) d\mathbf{x} \quad (5.17)$$

$$= -(\ell_2 - \ell_1) \int \delta(r_x) \mathbf{h}((r_p - r_x)\hat{\mathbf{a}} + (t_p - t_x)\hat{\mathbf{n}}) dr_x dt_x \quad (5.18)$$

$$= -(\ell_2 - \ell_1) \int \mathbf{h}(r_p \hat{\mathbf{a}} + (t_p - t_x)\hat{\mathbf{n}}) dt_x \quad (5.19)$$

$$= -(\ell_2 - \ell_1) \mathbf{h}(r_p) \quad (5.20)$$

$$\mathbf{B}\nabla_{\ell_1}\mathbf{f} = \int (1 - \mathbf{u}(\langle \mathbf{x}, \hat{\mathbf{a}} \rangle - \|\mathbf{a}\|)) \mathbf{h}(\mathbf{p} - \mathbf{x}) d\mathbf{x} \quad (5.21)$$

$$= 1 - \int \int \mathbf{u}(r_x) \mathbf{h}((r_p - r_x)\hat{\mathbf{a}} + (t_p - t_x)\hat{\mathbf{n}}) dr_x dt_x \quad (5.22)$$

$$= 1 - \int \mathbf{u}(r_x) \mathbf{h}(r_p - r_x) dr_x \quad (5.23)$$

$$= 1 - \int_0^{\infty} \mathbf{h}(r_p - r_x) dr_x \quad (5.24)$$

$$= 1 - \int_{-\infty}^{r_p} \mathbf{h}(r_x) dr_x \quad (5.25)$$

$$\therefore \mathbf{B}\nabla_{\ell_2}\mathbf{f} = \int_{-\infty}^{r_p} \mathbf{h}(r_x) dr_x \quad (5.26)$$

Finally, differentiating w.r.t. the blur parameters one may write

$$\nabla_{\alpha}\mathbf{B}\mathbf{f} = (\ell_2 - \ell_1) \int_{-\infty}^{r_p} \nabla_{\alpha}\mathbf{h}(r_x) dr_x. \quad (5.27)$$

The blurred edges and their differentials can both be efficiently represented at the

Nyquist rate, however we are also interested in the action of a sampling operator on these representations. In practice one must represent both the Nyquist sampled object and its post-sampling representation on a computer using an integer number of samples. This poses a challenge given that some down-sampling rates will not be ratios of integers.

To find acceptably accurate representations one first notes that the underlying object representations can be formed at any integer sizes. Furthermore, we require that the post-sampled representation exceed the size of the measured data to ensure that the truncation operator is able to produce an estimate which “covers” the observations, and one also needs to ensure the Nyquist sampled object is sufficiently large to allow extrapolation to mitigate edge effects. Using this criteria, and imposing a constraint on the relative error of the sampling rate, results in the following subproblem. For reasonable downsampling factors, the provided algorithm converges to the optimal solution almost immediately. In essence, it works by projecting optimal but inadmissible solutions onto the admissible set until the projection yields a solution on the grid of integers.

$$\text{Given: } \tilde{r} \in \mathbb{R}^+ \text{ and } \tilde{z}_1, \tilde{z}_2 \in \mathbb{Z}^+ \quad (5.28)$$

$$\text{Find: } \arg \min_{z_1, z_2 \in \mathbb{Z}^+} z_2 \quad \text{s.t.} \quad (5.29)$$

$$\left| \frac{z_1}{z_2} - \tilde{r} \right| \leq \epsilon \tilde{r} \quad (5.30)$$

$$z_1 \geq \tilde{z}_1 \quad (5.31)$$

$$z_2 \geq \tilde{z}_2 \quad (5.32)$$

The specific algorithm for this projection is given in Figure 5.3.




Figure 5.2: Illustration of the projections used to solve the integer sampling sub-problem

Using this technique, one can apply the forward operator to continuous edges with arbitrary precision.

<p>Input: $\tilde{r}, \tilde{z}_1, \tilde{z}_2, \varepsilon$</p> <p>Output: z_1, z_2</p> <p>1 $z_1 = \tilde{z}_1$</p> <p>2 $z_2 = \max \left[\frac{z_1}{(1+\varepsilon)\tilde{r}}, \tilde{z}_2 \right]$</p> <p>3 while $z_2 > \frac{z_1}{(1-\varepsilon)\tilde{r}}$ do</p> <p>4 $z_1 = \max [z_2(1 + \varepsilon)\tilde{r}], z_1 + 1$</p> <p>5 $z_2 = \left\lceil \frac{z_1}{(1+\varepsilon)\tilde{r}} \right\rceil$</p>

Figure 5.3: Algorithm for finding data sizes for downsampling

5.2.3 Evaluating Line-Spread Functions

The previous section, while technically correct, leaves a lot to be desired in terms of implementation details. While the discontinuities in the underlying object have been addressed, and it should now be clear that only the object representing the action of the blur on the fiducial should ever be represented numerically, but how should one take the necessary integrals?

The naive approach is to simply upsample and integrate the PSF. Sinc-based upsampling of the PSF ensures that numerical precision is not lost during the subsequent integration, and the 1D integral can then be computed by first rotating the object such that integration is being computed along the first dimension, and then subsequently applying a 1D numerical integration scheme (I presently employ the trapezoidal rule). The rotation itself, also implies an interpolation, and here I presently use the bicubic technique referenced earlier. Taken together, this is an effective, albeit very numerically intensive, approach.

By contrast, the “correct” way of computing a line-spread function is to exploit the band-limited nature of the PSF, the volume preserving aspects of rotation (all eigenvalues are 1), and the Fourier slice theorem. This avoids the upsampling while simultaneously implementing the rotation using a perfect sinc-based interpolation. This rotation is performed on the Fourier transform of the PSF allowing one to simply read off the “slice” in the frequency domain, and hence the line integral in the spatial domain. While possible, exact, and much more numerically efficient than the previously described technique, this approach requires a number of signal processing steps that are not commonly used, and codes have not yet been developed and tested for this purpose.

5.2.4 Joint Model

A multi-frame formulation will be considered, where the gain and offset parameters (c_3, c_4) are allowed to change between frames, and the depth of field over the fiducial set will

be assumed unknown. We will not, however, allow for reconfigurations of the optics during the collection interval (non-defocus aberrations are assumed constant), nor will we attempt to relate the various fiducials between the frames. Combining these additional details, the full joint model is then given by

$$\mathbf{d} = \text{vec} \left(\left\{ \mathbf{c}_1^{(j)} \mathbf{1}_{n_f^{(j)}} \right\}_{j=1}^J \right) \odot \mathcal{P} \left\{ \text{vec} \left(\left\{ \mathbf{T} \mathbf{S} \mathbf{B}(\boldsymbol{\alpha}, \beta_k) \mathbf{f}_k(\boldsymbol{\theta}_f^{(k)}) \right\}_{k=1}^K \right) \right\} + \text{vec} \left(\left\{ \mathbf{c}_2^{(j)} \mathbf{1}_{n_f^{(j)}} \right\}_{j=1}^J \right) + \varepsilon_r(\boldsymbol{\sigma}_r) \quad (5.33)$$

where K is the total number of fiducials, J is the total number of frames, \mathbf{f}_k is the k^{th} fiducial with parameters $\boldsymbol{\theta}_f^{(k)}$, \mathbf{B} is the blur operator indexed by the common set of blur parameters $\boldsymbol{\alpha}$ and the individual defocus parameters $\boldsymbol{\beta} = [\beta_1, \dots, \beta_K]^T$, \mathbf{S} is the sampling operator, \mathbf{T} is the truncation operator, \mathcal{P} is the Poisson noise operator, \mathbf{c}_1 and \mathbf{c}_2 are the length J vectors of gains and offsets respectively, and ε_r is the read-noise as a function of the length J vector of read-noise standard deviations $\boldsymbol{\sigma}_r$. The total vector of unknowns then becomes $\boldsymbol{\theta} = [\boldsymbol{\alpha}, \mathbf{c}_1, \mathbf{c}_2, \boldsymbol{\sigma}_r, \boldsymbol{\beta}, \boldsymbol{\theta}_f]$. Notice that the warp operator does not appear because a warped edge falls within the class of warped edges, and for this discussion it is implicitly assumed (via the model) that the collection of edges each fall within a narrow depth range relative to the sensors depth of field. Finally, the image regions containing edges are allowed to take on arbitrary shapes, but are assumed to be simply connected. When tracking the sizes of the corresponding lexicographically ordered regions, and their relationship to frames, it is helpful to again abuse notation rather than introducing additional variables. By convention, $n_{(\cdot)}$ is used to denote the cardinality of (\cdot) , but depending on context we may wish to refer to the cardinality of an individual edge region, the cardinality of the edges in a frame, or the cardinality of all edge regions. To maximize the readability of the resulting expressions we introduce the notation

$$n_f = \sum_{k=1}^K n_f^{(k)} = \sum_{j=1}^J n_f^{(j)} \quad (5.34)$$

where n_f is the total number of samples across all edges, $n_f^{(j)}$ is the total number of edge samples across the j^{th} frame, and $n_f^{(k)}$ is the total number of edge samples in the k^{th} edge. The edge regions themselves are lexicographically ordered, and where the context is not clear additional annotation will be given. This allows one to consistently think of “ n_f ” as the number of object elements in the current context.

Applying the gaussian approximation to the Poisson process one may write

$$\mathbf{d} = \text{vec} \left(\left\{ \mathbf{c}_1^{(j)} \mathbf{1}_{n_f^{(j)}} \right\}_{j=1}^J \right) \odot \text{vec} \left(\left\{ \mathbf{TSB}(\boldsymbol{\alpha}, \beta_k) \mathbf{f}_k(\boldsymbol{\theta}_f^{(k)}) \right\}_{k=1}^K \right) + \quad (5.35)$$

$$\begin{aligned} & \text{vec} \left(\left\{ \mathbf{c}_2^{(j)} \mathbf{1}_{n_f^{(j)}} \right\}_{j=1}^J \right) + \boldsymbol{\varepsilon} \\ & = \boldsymbol{\mu}(\boldsymbol{\theta}) + \boldsymbol{\varepsilon} \end{aligned} \quad (5.36)$$

$$\boldsymbol{\varepsilon} \sim \mathcal{N} \left(\mathbf{0}, \text{diag} \left(\text{vec} \left(\left\{ \mathbf{c}_1^{(j)} \mathbf{1}_{n_f^{(j)}} \right\}_{j=1}^J \right) \odot^2 \text{vec} \left(\left\{ \mathbf{TSB}(\boldsymbol{\alpha}, \beta_k) \mathbf{f}_k(\boldsymbol{\theta}_f^{(k)}) \right\}_{k=1}^K \right) + \boldsymbol{\varepsilon}_r \right) \right) \quad (5.37)$$

$$\sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}(\boldsymbol{\theta})) \quad (5.38)$$

where $\boldsymbol{\varepsilon}_r$ relates to read noise for the individual frames via

$$\boldsymbol{\varepsilon}(\boldsymbol{\sigma}_r) = \text{vec} \left(\left\{ \left(\boldsymbol{\sigma}_r^{(j)} \right)^2 \mathbf{1}_{n_f^{(j)}} \right\}_{j=1}^J \right) \quad (5.39)$$

The aggregate log-likelihood function is proportional to¹

$$\ell(\boldsymbol{\theta}) = (\mathbf{d} - \boldsymbol{\mu}(\boldsymbol{\theta}))^T \boldsymbol{\Sigma}^{-1} (\mathbf{d} - \boldsymbol{\mu}(\boldsymbol{\theta})) + \mathbf{1}_{n_f}^T \ln(\text{diag}(\boldsymbol{\Sigma}(\boldsymbol{\theta}))) + n_f \ln(2\pi), \quad (5.40)$$

$$n_f \stackrel{\text{def}}{=} \sum_{j=1}^J n_f^{(j)} \quad (5.41)$$

The differential of the likelihood function is given by

$$\nabla \ell(\boldsymbol{\theta}) = 2\mathbf{a}^T \boldsymbol{\Sigma}^{-1} \nabla \mathbf{a} - \mathbf{a}^T \boldsymbol{\Sigma}^{-1} \nabla \boldsymbol{\Sigma} \boldsymbol{\Sigma}^{-1} \mathbf{a} + \mathbf{1}_{n_f}^T \boldsymbol{\Sigma}^{-1} \nabla \mathbf{b} \quad (5.42)$$

¹With proportionality constant $-1/2$

where for convenience the following variables are introduced

$$\mathbf{a} = \mathbf{d} - \boldsymbol{\mu}(\boldsymbol{\theta}) \quad (5.43)$$

$$\begin{aligned} &= \mathbf{d} - \text{vec} \left(\left\{ \mathbf{c}_1^{(j)} \mathbf{1}_{n_f^{(j)}} \right\}_{j=1}^J \right) \odot \text{vec} \left(\left\{ \mathbf{T} \mathbf{S} \mathbf{B}(\boldsymbol{\alpha}, \boldsymbol{\beta}_k) \mathbf{f}_k(\boldsymbol{\theta}_f^{(k)}) \right\}_{k=1}^K \right) - \\ &\text{vec} \left(\left\{ \mathbf{c}_2^{(j)} \mathbf{1}_{n_f^{(j)}} \right\}_{j=1}^J \right) \end{aligned} \quad (5.44)$$

$$\mathbf{b} = \text{diag}(\boldsymbol{\Sigma}(\boldsymbol{\theta})) \quad (5.45)$$

$$\begin{aligned} &= \text{vec} \left(\left\{ \mathbf{c}_1^{(j)} \mathbf{1}_{n_f^{(j)}} \right\}_{j=1}^J \right)^{\odot 2} \odot \text{vec} \left(\left\{ \mathbf{T} \mathbf{S} \mathbf{B}(\boldsymbol{\alpha}, \boldsymbol{\beta}_k) \mathbf{f}_k(\boldsymbol{\theta}_f^{(k)}) \right\}_{k=1}^K \right) + \\ &\text{vec} \left(\left\{ \left(\boldsymbol{\sigma}_r^{(j)} \right)^2 \mathbf{1}_{n_f^{(j)}} \right\}_{j=1}^J \right) \end{aligned} \quad (5.46)$$

Recall that the parameter ordering is given by $\boldsymbol{\theta} = [\boldsymbol{\alpha}, \mathbf{c}_1, \mathbf{c}_2, \boldsymbol{\sigma}_r, \boldsymbol{\beta}, \boldsymbol{\theta}_f]$, implying a total of $n_\alpha + J + J + J + K + 4K = n_\alpha + 3J + 5K$ unknowns, and let $F(k)$ map the edge index k onto its corresponding frame index. Utilizing $F(\cdot)$, \mathbf{a} and \mathbf{b} can be written in the more intuitive forms

$$\mathbf{a} = \mathbf{d} - \text{vec} \left(\left\{ \mathbf{c}_1^{(F(k))} \mathbf{T} \mathbf{S} \mathbf{B}(\boldsymbol{\alpha}, \boldsymbol{\beta}_k) \mathbf{f}_k(\boldsymbol{\theta}_f^{(k)}) + \mathbf{c}_2^{(F(k))} \right\}_{k=1}^K \right) \quad (5.47)$$

$$\mathbf{b} = \text{vec} \left(\left\{ \left(\mathbf{c}_1^{(F(k))} \right)^2 \mathbf{T} \mathbf{S} \mathbf{B}(\boldsymbol{\alpha}, \boldsymbol{\beta}_k) \mathbf{f}_k(\boldsymbol{\theta}_f^{(k)}) + \left(\boldsymbol{\sigma}_r^{(F(k))} \right)^2 \right\}_{k=1}^K \right) \quad (5.48)$$

It is important to note that both \mathbf{a} and \mathbf{b} , as well as their differentials, are all separable along the edge boundaries (the k index). Along these lines it is helpful to introduce the notation \mathbf{a}_k , \mathbf{b}_k , and so forth. Using this notation, the log-likelihood function can be written as

$$\ell(\boldsymbol{\theta}) = \sum_{k=1}^K \left[\begin{pmatrix} \mathbf{a}_k \\ \mathbf{b}_k \end{pmatrix}^T \mathbf{a}_k + \mathbf{1}_{n_f^{(k)}}^T \ln(\mathbf{b}_k) + n_f^{(k)} \ln(2\pi) \right] \quad (5.49)$$

Taking the differentials with respect to the parameter groups one finds

$$\nabla_{\boldsymbol{\alpha}} \mathbf{a} = -\text{vec} \left(\left\{ \mathbf{c}_1^{(F(k))} \mathbf{T} \mathbf{S} \nabla_{\boldsymbol{\alpha}} \mathbf{B} \mathbf{f}_k \right\}_{k=1}^K \right) \quad (5.50)$$

$$\nabla_{\mathbf{c}_1} \mathbf{a} = -\text{vec} \left(\left\{ [\mathbf{T} \mathbf{S} \mathbf{B} \mathbf{f}_k] \mathbf{e}_{F(k)}^T \right\}_{k=1}^K \right) \quad (5.51)$$

$$\nabla_{\mathbf{c}_2} \mathbf{a} = -\text{vec} \left(\left\{ \mathbf{1}_{n_f^{(j)}} \mathbf{e}_j^T \right\}_{j=1}^J \right) \quad (5.52)$$

$$\nabla_{\boldsymbol{\sigma}_r} \mathbf{a} = \mathbf{0} \quad (5.53)$$

$$\nabla_{\boldsymbol{\beta}} \mathbf{a} = -\text{vec} \left(\left\{ \mathbf{c}_1^{(F(k))} \mathbf{T} \mathbf{S} \nabla_{\boldsymbol{\beta}} \mathbf{B} \mathbf{f}_k \right\}_{k=1}^K \right) \quad (5.54)$$

$$\nabla_{\boldsymbol{\theta}_f} \mathbf{a} = -\text{vec} \left(\left\{ \mathbf{c}_1^{(F(k))} \mathbf{T} \mathbf{S} \mathbf{B} \nabla_{\boldsymbol{\theta}_f} \mathbf{f}_k \right\}_{k=1}^K \right) \quad (5.55)$$

$$\nabla_{\boldsymbol{\alpha}} \mathbf{b} = \text{vec} \left(\left\{ \left(\mathbf{c}_1^{(F(k))} \right)^2 \mathbf{T} \mathbf{S} \nabla_{\boldsymbol{\alpha}} \mathbf{B} \mathbf{f}_k \right\}_{k=1}^K \right) \quad (5.56)$$

$$\nabla_{\mathbf{c}_1} \mathbf{b} = \text{vec} \left(\left\{ \left[2\mathbf{c}_1^{(F(k))} \mathbf{T} \mathbf{S} \mathbf{B} \mathbf{f}_k \right] \mathbf{e}_{F(k)}^T \right\}_{k=1}^K \right) \quad (5.57)$$

$$\nabla_{\mathbf{c}_2} \mathbf{b} = \mathbf{0} \quad (5.58)$$

$$\nabla_{\boldsymbol{\sigma}_r} \mathbf{b} = \text{vec} \left(\left\{ 2\boldsymbol{\sigma}_r^{(j)} \mathbf{1}_{n_f^{(j)}} \mathbf{e}_j^T \right\}_{j=1}^J \right) \quad (5.59)$$

$$\nabla_{\boldsymbol{\beta}} \mathbf{b} = \text{vec} \left(\left\{ \left(\mathbf{c}_1^{(F(k))} \right)^2 \mathbf{T} \mathbf{S} \nabla_{\boldsymbol{\beta}} \mathbf{B} \mathbf{f}_k \right\}_{k=1}^K \right) \quad (5.60)$$

$$\nabla_{\boldsymbol{\theta}_f} \mathbf{b} = \text{vec} \left(\left\{ \left(\mathbf{c}_1^{(F(k))} \right)^2 \mathbf{T} \mathbf{S} \mathbf{B} \nabla_{\boldsymbol{\theta}_f} \mathbf{f}_k \right\}_{k=1}^K \right) \quad (5.61)$$

where based on our ordering of the unknowns the final differential is given by

$$\nabla(\cdot) = \left[\nabla_{\boldsymbol{\alpha}}(\cdot) \quad \nabla_{\mathbf{c}_1}(\cdot) \quad \nabla_{\mathbf{c}_2}(\cdot) \quad \nabla_{\boldsymbol{\sigma}_r}(\cdot) \quad \nabla_{\boldsymbol{\beta}}(\cdot) \quad \nabla_{\boldsymbol{\theta}_f}(\cdot) \right] \quad (5.62)$$

Using this notation, the sub-terms in Equation 5.42 can be more parsimoniously expressed as follows

$$\mathbf{a}^T \boldsymbol{\Sigma}^{-1} \nabla \mathbf{a} = \left(\frac{\mathbf{a}}{\mathbf{b}} \right)^T \nabla \mathbf{a} \quad (5.63)$$

$$= \sum_{k=1}^K \left(\frac{\mathbf{a}_k}{\mathbf{b}_k} \right)^T \nabla \mathbf{a}_k \quad (5.64)$$

$$\mathbf{a}^T \boldsymbol{\Sigma}^{-1} \nabla \boldsymbol{\Sigma} \boldsymbol{\Sigma}^{-1} \mathbf{a} = \left[\left(\frac{\mathbf{a}}{\mathbf{b}} \right)^{\odot 2} \right]^T \nabla \mathbf{b} \quad (5.65)$$

$$= \sum_{k=1}^K \left[\left(\frac{\mathbf{a}_k}{\mathbf{b}_k} \right)^{\odot 2} \right]^T \nabla \mathbf{b}_k \quad (5.66)$$

$$\mathbf{a}_{n_f}^T \boldsymbol{\Sigma}^{-1} \nabla \mathbf{b} = \left(\frac{1}{\mathbf{b}} \right)^T \nabla \mathbf{b} \quad (5.67)$$

$$= \sum_{k=1}^K \left(\frac{1}{\mathbf{b}_k} \right)^T \nabla \mathbf{b}_k \quad (5.68)$$

Combining terms results in following expression for the differential

$$\nabla \ell(\boldsymbol{\theta}) = 2 \sum_{k=1}^K \left(\frac{\mathbf{a}_k}{\mathbf{b}_k} \right)^T \nabla \mathbf{a}_k + \sum_{k=1}^K \left[\frac{1}{\mathbf{b}_k} - \left(\frac{\mathbf{a}_k}{\mathbf{b}_k} \right)^{\odot 2} \right]^T \nabla \mathbf{b}_k \quad (5.69)$$

5.2.5 Experiment

Edge fiducials were artificially inserted into an ad-hoc imaging scenario and used to estimate their range from the imaging device to illustrate the efficacy of these ideas. As with the hardware experiment described in Chapter IV, the wavefront description of the blur is used in conjunction with a physical imaging model to relate estimated defocus directly to the range. This section demonstrates that for the purposes of estimating defocus (range), the fiducial model leads to estimates that are only slightly less accurate than having a full object description.

A Nikon D7000 was modified to remove its anti-aliasing filter, and this camera was used in conjunction with a Nikkor 70-200 mm f/2.8 VR II lens to collect imagery of two people holding manilla envelopes with a white piece of printer paper used to induce an edge. The actual ranges to the targets were measured with a laser ranging device and found to be 5 m and 7.7m respectively. The camera was allowed to autofocus on the first edge, and the range to best focus was reported by the camera to be 5.012 m. Using (5.49), maximum likelihood estimates of the defocus blur at each edge were used in conjunction with (4.40) to estimate the ranges to the targets directly.

Estimation was performed over the blue channel, with a fully open aperture. The focal length was 70 mm, and corresponded to an F-stop of 2.8. With these settings, the detector undersampled the optical wavefront by 56.5X (7.518X in each linear dimension). The defocus at each of the two edges was jointly estimated alongside the edge geometry, intensity emanating from each side of the edge, and the sensor read noise variance. Using the focus range reported by the camera, and the estimated defocus amounts of 0.006 waves and 1.327 waves RMS at 450 nm. These correspond to target ranges of 5.021 m and 7.771 m respectively. Figure 5.4 compares the magnitude of the observed ranging errors

to those expected given wavefront estimation accuracy of 0.01 waves RMS at 450 nm. This experiment demonstrates the use of ad-hoc fiducials for wavefront estimation and suggests only a small loss in information compared to estimation performed using full object knowledge.

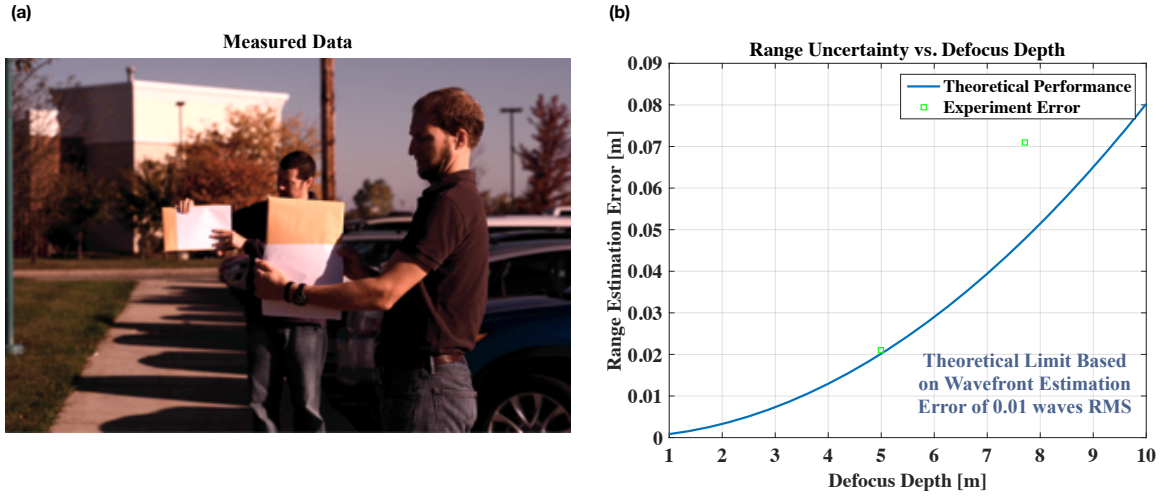


Figure 5.4: The data used to estimate edge-target ranges based on relative fiducial defocus amounts is shown in (a). Subfigure (b) shows the observed experimental error along with the theoretical error had the wavefront estimation accuracy been 0.01 waves RMS.

The ranging accuracy demonstrated during this experiment is predominantly a function of the large entrance aperture of the camera. Had the aperture been stopped down, or had a camera with a smaller aperture been used, the reduced depth of field would lead to substantially reduced ranging accuracy given the same blur estimation fidelity. Figure 5.5 compares the maximum possible ranging accuracy of various commercial cameras to the Nikon D7000 assuming a defocus estimation accuracy of 0.01 waves RMS. Overall, these experiments demonstrate the flexible nature of the proposed inverse-problem framework for solving inverse-imaging problems.




Figure 5.5: Comparison of maximum theoretical ranging uncertainty for common consumer cameras relative to the Nikon D7000 assuming defocus estimation accuracy of 0.01 waves RMS.

5.3 Future Directions

This thesis presents a modular and extensible inverse-problem approach for studying optical imaging systems that supports many future research directions. For example, given a prior probability on wavefront errors introduced by manufacturing imperfections, one could use the Cramér-Rao analysis described in Section 3.4 to design optimal calibration targets. Another interesting extension of this work would be the optimal sequential design of targets. One could design such sequences to optimize computational efficiency, solution stability, or other metrics designed to ensure fast, accurate, and reliable sensor characterizations.

APPENDICES

APPENDIX A

Minimum Bandwidth Function Extension and Inpainting Comparison

This section compares the minimum bandwidth criterion for smooth function extension (discussed in Section 2.4.1) to modern inpainting approaches. The aim of minimum bandwidth function extension is *not* the same as inpainting, but the two techniques are sufficiently related that a comparison is warranted. Minimum bandwidth function extension is based on constructing the smoothest continuous function that takes on certain values over a subset of its domain. By contrast, image inpainting approaches often combine similar energy functions with self-similarity or spatial coherence measures to produce results that appear more natural to a human observer. See [22] and the references therein for an overview this field. When successful, such heuristics can produce remarkable results. Unfortunately these techniques often dependent on tuning parameters that must be hand-picked and which must be changed on a per-image basis, require substantial computation, and occasionally fail to converge to reasonable solutions.

To illustrate the relative efficacy of smooth function extension it was compared to a number of modern inpainting approaches [34, 43, 51, 54, 117]. To eliminate the possibility of coding errors, or a misinterpretation of subtle algorithm details, all of the techniques are based on codes provided by the original authors. D’Errico [47] implements a finite-difference approximation to classical PDE inpainting. This technique is similar to that suggested by Chui [33, 34] and later studied by [96]. Of the inpainting approaches, this one is closest in spirit to the proposed function extension technique, but it does not immediately result in a continuous extension. Fadili et al. [54] uses a sparsity promoting prior within a Bayesian framework to perform inpainting using a redundant dictionary of wavelets, curvelets, and ridgelets [175]. In addition to the choice of dictionary, this algorithm includes

3 tuning parameters, and these were set to be consistent with the lena example shown in their publication. Elad et al. [51] also models the unknown image components through a sparse representation in an over-complete dictionary. The missing image components are assumed to be the composition of a smooth function with a texture component, where smoothness is enforced through a total variation (TV) penalty on the gradient, and both terms enter into their final objective function through mixed L_1 - and L_2 -norm objective. This approach also results in 3 tuning parameters, and these were again chosen to be consistent with the lena example given in their original work. Criminisi et. al. [43] uses block-based exemplars to iteratively inpaint missing regions based on a mixed objective that considers both texture and smoothness. This work introduces the most aggressive heuristics, and can produce remarkably consistent-looking results when large contiguous image regions are missing. Unfortunately, the approach does not lend itself to more spatially distributed masks, and for this reason the tight checkerboard pattern used in Figure 2.9 had to be relaxed slightly to allow this technique to properly function. Finally, Li [117] uses Bayesian model averaging in conjunction within a simulated annealing optimization strategy to enforce local smoothness with global structural consistency. Like the approach of Criminisi, the global similarity metric struggled with the complex masks used in this comparison. I was able to reproduce the results from the author's publication, but the technique failed to converge to reasonable solutions given complex masks. For this reason, the results are not shown. Additionally, this was the slowest of the approaches tested, running for well over an hour.

Figure A.1 illustrates the image used to exercise the inpainting algorithms and the associated inpainting mask. Errors in the reconstructions are reported in terms of PSNR measured in units of dB; a standard in the inpainting community. To avoid unnecessarily penalizing approaches that implicitly assume inpainting regions are surrounded by known values, errors were computed over the interior set indicated by the white dashed line in Figure A.1 (b). The result of applying the aforementioned techniques are shown in Figure A.2. Surprisingly, many of the more sophisticated inpainting approaches are outperformed by minimum bandwidth function extension when measured in terms of PSNR (or equivalently MSE). This experiment supports the use of the minimum bandwidth criterion, and the fact that the proposed technique has no tuning parameters and is extremely fast, suggests that it could also provide a valuable starting solution for many inpainting techniques. Figure A.3 repeats this experiment with exterior pixels known. The results are similar.

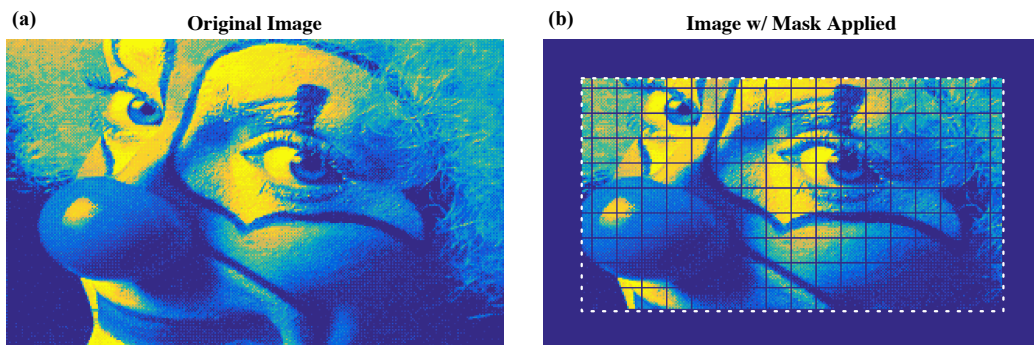


Figure A.1: The original test image (a), is shown alongside the mask used to best inpainting (b). To avoid unnecessarily penalizing methods that implicitly assume inpainting regions are surrounded by known values, errors were computed over both the whole image and the interior set indicated by the white dashed line.




Figure A.2: Comparison between minimum bandwidth continuous function extension and various inpainting techniques. Errors are reported in terms of PSNR, measured in units of dB, over the interior portion of the image. The error for the entire image is also reported in parentheses. All of the images are shown on the same scale as the truth.




Figure A.3: This figure illustrates the same experiment shown in Figure A.2 with the exception that the exterior pixels were assumed known. For this reason, the interior, and whole image, PSNR calculations match in every instance.

APPENDIX B

Discrete to Continuous Fourier Transform

This appendix revisits the ideas presented in Section 2.4.5 with a greater emphasis placed on the group-theoretic perspective. The aforementioned section is motivated by group-theoretic ideas, but uses language and notion aimed at an engineering audience. Some readers may find this appendix to be a helpful companion text.

Homomorphisms

A abelian group formed from a set G is a closure of G and its associated inverse elements under a commutative binary operation. We are interested in \mathbb{Z} with addition, $\mathbb{Z}_n = \{0, \dots, n-1\}$ with addition modulo n , and \mathbb{T} with multiplication, where \mathbb{T} is complex numbers with norm 1. A group homomorphism is a map between groups that commutes with the group operation. Let $\psi_n : \mathbb{Z} \rightarrow \mathbb{Z}_n$ be the map which takes an integer to its residue modulo n . For $\chi : \mathbb{Z}_n \rightarrow \mathbb{T}$ to be a group homomorphism, it must satisfy $\chi(1)^n = \chi(n) = f(0) = 1$. Therefore, we must have character $\chi_{n,k} \stackrel{\text{def}}{=} \chi(1) = e^{2\pi ik/n}$ for some $k \in \mathbb{Z}_n$.

Characters form a group with the operation of pointwise multiplication. The group of characters

$$\hat{\mathbb{Z}}_n = \{\chi_{n,k} \mid k \in \mathbb{Z}_n\}$$

is isomorphic to \mathbb{Z}_n as a group via the relationship $\chi_{n,k+l} = \chi_{n,k}\chi_{n,\ell}$. Similarly, the characters of \mathbb{T} are $\chi_j(z) = z^j$, with $j \in \mathbb{Z}$. The associated group of characters

$$\hat{\mathbb{T}} = \{\chi_j \mid j \in \mathbb{Z}\}$$

is isomorphic to \mathbb{Z} via $\chi_{k+l} = \chi_k\chi_\ell$.




Figure B.1: Overall structure showing group homomorphisms (solid lines), and set maps (dashed lines). The dotted line indicates the diagrams don't, in general, commute along compositions of solid lines (e.g. $(F_n \chi_{n,r}) \psi_n \neq F \chi_j$).

A group homomorphism $\phi : G \rightarrow \mathbb{Z}_n$ is a map such that for any g and $h \in G$, we have $\phi(gh) = \phi(g) + \phi(h)$ in \mathbb{Z}_n . It is injective if $\phi(g) = \phi(h)$ implies $g = h$, or equivalently, if the only element in $\phi^{-1}(0)$ is the identity element in G . Given an injective group homomorphism ϕ , it induces the homomorphism $\bar{\phi} : \mathbb{C}(\mathbb{Z}_n) \rightarrow \mathbb{C}(G)$, defined by $\bar{\phi}(f) = f \circ \phi$. The overarching structure being described is shown in Figure B.1, where the solid lines indicate group homomorphisms, the dashed lines set maps, and the dotted line indicates that the diagrams don't necessarily commute along compositions of solid lines. For example, $(F_n \chi_{n,r}) \psi_n \neq F \chi_j$, however, this relationship will hold under the restriction $r = \psi_n(j)$ and $s = \psi_m(j)$.

Consider the case when $n|m$ such that $\phi(r) = \frac{m}{n}r$ is an injective group homomorphism $\phi : \mathbb{Z}_n \rightarrow \mathbb{Z}_m$ inducing the $\frac{m}{n}$ -to-one surjection $\bar{\phi} : \hat{\mathbb{Z}}_m \rightarrow \hat{\mathbb{Z}}_n$. Let $m = dn$, and $\phi(r) = rd$ for all $r \in \mathbb{Z}_n$, then every element s in \mathbb{Z}_m can be written uniquely as $s = xn + r$ for some $r \in \mathbb{Z}_n$ such that $0 \leq r < n$. Then

$$\bar{\phi} \chi_{m,s}(k) = \chi_{m,s}(dk) = e^{2\pi i s dk/m} = e^{2\pi i s k/n} = e^{2\pi i r k/n} = \chi_{n,r}(k). \quad (\text{B.1})$$

Now let $\psi : \mathbb{Z}_m \rightarrow \mathbb{Z}_n$ be the quotient map $\psi(s) = r$ with n, m, s and r as above. This is a surjection, and together these two group maps provide a general description of integral upsampling and downsampling respectively. The difference between this description, and that typically taught in signal processing texts, is that the relationships are traditionally illustrated through functions on the groups.

Functions on \mathbb{Z}_n , \mathbb{T} , and \mathbb{Z}

The functions of interest are maps to \mathbb{C} that are not necessarily group homomorphisms. Let $\mathbb{C}(\mathbb{Z})$, $\mathbb{C}(\mathbb{Z}_n)$, and $\mathbb{C}(\mathbb{T})$ denote the complex vector spaces of functions on the corresponding groups. There is a linear isomorphism between $\mathbb{C}(\mathbb{T})$ and periodic functions on an interval $[0, T]$ defined through the span over basis elements of the form $\chi_j(t) = e^{2\pi i j t/T}$. The functions form a vector space with inner-products defined $\forall f, g \in \mathbb{C}(\mathbb{Z}_n)$,

$$\langle f, g \rangle = \frac{1}{n} \sum_{k=0}^{n-1} f(k)g(k)^*$$

The characters $\chi_{n,k} \in \mathbb{C}(\mathbb{Z}_n)$ form an orthonormal basis with respect to this inner product. Similarly, $\forall f, g \in \mathbb{C}(\mathbb{T})$,

$$\langle f, g \rangle = \frac{1}{T} \int_0^T f(t)g(t)^* dt$$

Fourier Transforms

The Fourier inversion formulas are simply an expression of functions decomposed with respect to the aforementioned orthonormal bases of characters:

$$\begin{aligned} \text{For } f \in \mathbb{C}(\mathbb{Z}_n) : \quad f(\ell) &= \sum_{k=0}^{n-1} \hat{f}(k)\chi_{n,k}(\ell) \\ \text{For } f \in \mathbb{C}(\mathbb{T}) : \quad f(t) &= \sum_{j \in \mathbb{Z}} \hat{f}(j)\chi_j(t) \end{aligned}$$

Exploiting the orthonormality of the characters, the Fourier transforms are given by:

$$\begin{aligned} \text{For } f \in \mathbb{C}(\mathbb{Z}_n) : \quad \hat{f}(k) &= \langle f, \chi_{n,k} \rangle = \frac{1}{n} \sum_{\ell=0}^{n-1} f(\ell)\chi_{n,k}^*(\ell) \\ &= \frac{1}{n} \sum_{\ell=0}^{n-1} f(\ell)\chi_{n,k}(-\ell) \\ \text{For } f \in \mathbb{C}(\mathbb{T}) : \quad \hat{f}(j) &= \langle f, \chi_j \rangle = \frac{1}{T} \int_0^T f(t)\chi_j^*(t) dt \\ &= \frac{1}{T} \int_0^T f(t)\chi_j(-t) dt \end{aligned}$$

$$\begin{array}{ccc}
\mathbb{C}(\mathbb{Z}_n) & \xleftarrow{\phi_n^*} & \mathbb{C}(\mathbb{T}) \\
F_n \downarrow & & \downarrow F \\
\mathbb{C}(\mathbb{Z}_n) & \xrightarrow{\bar{\psi}_n} & \mathbb{C}(\mathbb{Z})
\end{array}$$

Figure B.2: Linear transforms relating $\mathbb{C}(\mathbb{Z}_n)$, $\mathbb{C}(\mathbb{T})$, and $\mathbb{C}(\mathbb{Z})$.

The Fourier transform is a linear isomorphism of finite-dimensional complex vector spaces $F_n : \mathbb{C}(\mathbb{Z}_n) \rightarrow \mathbb{C}(\mathbb{Z}_n)$ by $F_n(f) = \hat{f}$. Due to the orthonormality of the characters one finds

$$\langle F_n(f), \hat{g} \rangle = \langle f, F_n^\dagger(\hat{g}) \rangle,$$

where $F_n^\dagger(g)$ is the adjoint transformation. For appropriately normalized definitions of the Fourier transform this relation becomes Parseval's equality. We may similarly define $F : \mathbb{C}(\mathbb{T}) \rightarrow \mathbb{C}(\mathbb{Z})$ by $F(f) = \hat{f}$, which is a linear isomorphism of infinite-dimensional complex vector spaces.

Discrete to Continuous Time Signals

Recall that $\psi_n : \mathbb{Z} \rightarrow \mathbb{Z}_n$ is the homomorphism sending an integer to its residue modulo n as shown in Figure B.1. This group homomorphism induces a linear transformation $\bar{\psi}_n : \mathbb{C}(\mathbb{Z}_n) \rightarrow \mathbb{C}(\mathbb{Z})$ defined by $\bar{\psi}_n f(k) = f(\psi_n(k))$. In particular, we have

$$\bar{\psi}_n \chi_j = \chi_{n,k}, \quad j \equiv k \pmod{n}.$$

Let $\phi_n : \mathbb{Z}_n \rightarrow \mathbb{T}$ be the homomorphism $\phi_n(k) = \chi_{n,1}(k) = e^{2\pi i k/n}$. This group homomorphism induces a linear transformation $\phi_n^* : \mathbb{C}(\mathbb{T}) \rightarrow \mathbb{C}(\mathbb{Z}_n)$ defined by $\phi_n^* f(k) = f(\phi_n(k)) = f(e^{2\pi i k/n}) = f(t_k)$ with $t_k = T k/n$ as a periodic function on $[0, T]$. In particular, $\forall j \in \mathbb{Z}$ we have

$$\phi_n^* \chi_j(k) = \chi_j(t_k) = \chi_{n,\psi_n(j)}(k) = \chi_{n,k}(\psi_n(j)) = \bar{\psi}_n \chi_{n,k}(j).$$

Thus we can relate the vector spaces $\mathbb{C}(\mathbb{Z}_n)$, $\mathbb{C}(\mathbb{T})$, and $\mathbb{C}(\mathbb{Z})$ with the linear transformations as shown in Figure B.2.

We want to define $M_n : \mathbb{C}(\mathbb{Z}_n) \rightarrow \mathbb{C}(\mathbb{T})$ with properties similar to the adjoint $(\phi_n^*)^\dagger$ while avoiding infinite sums. This will allow us to reconstruct a continuous signal from finite data.

First consider the formal adjoint $(\phi_n^*)^\dagger$. For $g \in \mathbb{C}(\mathbb{Z}_n)$ and $j \in \mathbb{Z}$,

$$\bar{\Psi}_n F_n(g)(j) = \langle g, \chi_{n,r} \rangle = \frac{1}{n} \sum_{k=0}^{n-1} g(k) \overline{\chi_{n,r}(k)} \quad \text{s.t.} \quad r = \psi_n(j).$$

Then,

$$(\phi_n^*)^\dagger g = \sum_{j \in \mathbb{Z}} \bar{\Psi}_n F_n(g)(j) \chi_j = \sum_{j \in \mathbb{Z}} \langle g, \chi_{n,r} \rangle \chi_j \quad \text{s.t.} \quad r = \psi_n(j).$$

Rewriting the adjoint to simplify comparing inner products, one has

$$(\phi_n^*)^\dagger g = \sum_{r=0}^{n-1} \sum_{\psi_n(j)=r} \langle g, \chi_{n,r} \rangle \chi_j.$$

For $f \in \mathbb{C}(\mathbb{T})$, we have

$$\langle f, (\phi_n^*)^\dagger g \rangle = \sum_{r=0}^{n-1} \sum_{\psi_n(j)=r} \langle g, \chi_{n,r} \rangle \langle f, \chi_j \rangle,$$

and

$$\begin{aligned} \langle \phi_n^* f, g \rangle &= \sum_j \langle f, \chi_j \rangle \langle \phi_n^* \chi_j, g \rangle \\ &= \sum_{r=0}^{n-1} \sum_{\psi_n(j)=r} \langle f, \chi_j \rangle \langle \chi_{n,r}, g \rangle = \langle f, (\phi_n^*)^\dagger g \rangle. \end{aligned}$$

The last equality uses the fact that $\forall k \in \mathbb{Z}_n$, we have $\phi_n^* \chi_j(k) = \chi_j(t_k) = e^{-2\pi i k j/n} = \chi_{n,r}(k)$, where $r = \psi_n(j)$. The adjunction of these characters is given by

$$(\phi_n^*)^\dagger \chi_{n,r} = \sum_{k=0}^{n-1} \sum_{\psi_n(j)=k} \langle \chi_{n,r}, \chi_{n,k} \rangle \chi_j = \sum_{\psi_n(j)=r} \chi_j.$$

Under the formal adjoint each, χ_j such that $\psi_n(j) = r$ is in the preimage of $\chi_{n,r}$ under ϕ_n^* . For reconstruction, however, we want $F(M_n \chi_{n,r}) = \bar{\Psi}_n F_n(\chi_{n,r})$. Since

$$\bar{\Psi}_n F_n(\chi_{n,r})(j) = \langle \chi_{n,r}, \chi_{n,\psi_n(j)} \rangle = \delta_{r,k},$$

M_n must be a truncation of $(\phi_n^*)^\dagger$. If we require $M_n \chi_{n,r}$ to be a single character, then there is

an injective map $\gamma_n : \mathbb{Z}_n \rightarrow \mathbb{Z}$ such that

$$M_n \chi_{n,r} = \chi_{\gamma_n(r)} \text{ and } \psi_n(\gamma_n(r)) = r.$$

Let Γ_n be the image of γ_n , and consider $\gamma_n : \mathbb{Z}_n \rightarrow \Gamma_n$ for all n . M_n must be a linear transformation, implying $M_n(\chi_0) = M_n(1) = 1$. This implies $\gamma_n(0) = 0$. We also know that $\psi_n \gamma_n$ must be the identity morphism on \mathbb{Z}_n , and equivalently, $\phi_n^* M_n$ must be the identity transformation on $\mathbb{C}(\mathbb{Z}_n)$. If $r + s \in \Gamma_n$, then since ψ_n is a homomorphism,

$$r + s = \gamma_n(\psi_n(r + s)) = \gamma_n(\psi_n(r) + \psi_n(s)).$$

When r and s are also individually in Γ_n , we must have $\gamma_n(\psi_n(r) + \psi_n(s)) = \gamma_n(\psi_n(r)) + \gamma_n(\psi_n(s))$. Therefore, γ_n evaluated at any other point (e.g. $\gamma_n(1)$) determines the entire map.

Reconstructibility

We say that a continuous signal $f_c \in \mathbb{C}(\mathbb{T})$ is “reconstructible” if there exists an n^* such that $f_c = M_{n^*}(\phi_{n^*}^* f_c)$. Given n measurements $f_n = \phi_n^* f_c$, it is not clear a priori whether n is sufficiently large for $M_n f_n = f_c$. We must define M_n collectively for all $n \in \mathbb{N}$. For every element of $\mathbb{C}(\mathbb{T})$ to be reconstructible we require

$$\mathbb{C}(\mathbb{T}) = \bigcup_{n=0}^{\infty} M_n(\mathbb{C}(\mathbb{Z}_n)) \text{ and } \bigcup_{n=0}^{\infty} \Gamma_n = \mathbb{Z}.$$

To ensure surjectivity, the maps γ_n must have the property that $\gamma_m \psi_m \gamma_n = \gamma_n$ for all $m > n$, and for every $\tilde{s} \in \bigcup_{n=0}^{\infty} \Gamma_n$, there is an n^* sufficiently large that $\tilde{s} = \gamma_{n^*}(s) = s \gamma_{n^*}(1) = s \gamma_3(1)$ for some $s \in \mathbb{Z}_{n^*}$.

$$\bigcup_{n=0}^{\infty} \Gamma_n = \{s \gamma_n(1) \mid s \in \mathbb{Z}\} = \gamma_3(1) \mathbb{Z}$$

is only equal to \mathbb{Z} when $\gamma_3(1) = \pm 1$. The multiplicative structure on \mathbb{Z} is such that $\psi_n(rs) = \psi_n(r)\psi_n(s)$. Therefore, for $k = 1, 2$, we have $\psi_3(\gamma_3(k)^2) = \psi_3(\gamma_3(k))^2 = k^2 \equiv 1 \pmod{3}$. This implies that if we set $\gamma_3(1) = 1$ by convention, we must have $\gamma_3(2) = -1$, and visa versa. It also follows that for n odd, we have $\gamma_n(n-k) = -\gamma_n(k)$ for all $k \in \mathbb{Z}_n$, and when $k \leq \frac{n-1}{2}$, we have $\gamma_n(k) = k$. That is, for $2n+1 \geq 3$, we have

$$\Gamma_{2n+1} = \{-n, -n+1, \dots, n-1, n\}.$$

Finally, for a vector of n measurements $f_n := \phi_n^* f_c$, if n is sufficiently large for Γ_n to contain the support of f_c , we have

$$f_c = \sum_{j \in \mathbb{Z}} F(f)(j) \chi_j = \sum_{r=0}^{n-1} (F_n f_n)(r) \chi_{\gamma_n(r)}.$$

Upsampling

The upsampling operator $U_{n,m} : \mathbb{C}(\mathbb{Z}_n) \rightarrow \mathbb{C}(\mathbb{Z}_m)$ is defined for every $m > n$ by the property that whenever $M_n(f_n) = M_m(f_m)$ we have $U_{n,m}(f_n) = f_m$. That is, $M_m \phi_m^* M_n f_n = M_n f_n$, and since $\phi_m^* M_m = id$ it follows immediately that

$$U_{n,m}(f_n) = \phi_m^* M_n(f_n).$$

If $m > n$, then $\Gamma_{2n+1} \subset \Gamma_{2m+1}$, and the associated characters are given by

$$\text{For } 0 \leq r \leq n-1 : \quad \psi_m(\gamma_n(r)) = \psi_m(r) = r \text{ and}$$

$$U_{n,m} \chi_{n,r} = \chi_{m,r}$$

$$\text{For } n \leq r \leq 2n : \quad \psi_m(\gamma_n(r)) = \psi_m(r-n) = m-n+r \text{ and}$$

$$U_{n,m} \chi_{n,r} = \chi_{m,m-n+r} = \chi_{m,m-n} \chi_{m,r}.$$

The upsampling operator is typically viewed with respect to the Fourier transform because

$$\text{For } r = \psi_n(\gamma_m(s)) : \quad F_m(U_{n,m} f_n)(s) = \langle U_{n,m} f_n, \chi_{m,s} \rangle = \langle f_n, \chi_{n,r} \rangle = (F_n f_n)(r),$$

$$\text{and otherwise : } \quad F_m(U_{n,m} f_n)(s) = 0.$$

This completes the connections relative to functions on groups, and Figure B.3 summarizes the key relationships. The adjunction of the $U_{n,m}$ defines the aliasing operator, and with the exception of the sign convention necessary in γ_n , the usual results follow. For example, for $f \in \mathbb{C}(\mathbb{Z}_m)$, if $s = nd + r$ and $m = nd$, we have

$$\bar{\phi} f = \sum_{s=0}^{m-1} (F_m f)(s) \bar{\phi} \chi_{m,s} = \sum_{r=0}^{n-1} \left(\sum_{x=0}^{d-1} (F_m f)(xn+r) \right) \chi_{n,r}$$

and therefore

$$(\psi^* F_n \bar{\phi} f)(s) = (F_n \bar{\phi} f)(\psi(s)) = \sum_{x=0}^{d-1} (F_m f)(xn+r).$$




Figure B.3: Relationship between homomorphisms and set maps for functions defined over the relevant groups. Solid lines indicate group homomorphisms, and dashed lines set maps. As before, the dotted line indicates that the diagrams don't commute along compositions of solid lines except for under the restrictions that $r = \psi_n(j)$ and $s = \psi_m(j)$.

APPENDIX C

Subgroups of Perspective Transformations

Within the class of projective transformations there are a number of important subgroups. Projective transformations are important in imaging because they describe arbitrary plane-onto-plane projections. As such, this transformational class describes the geometric properties of ideal imagers. Projective transformations are described in homogeneous coordinates by the Projective linear group $PL(n)$, and this property ensures closure under composition, the existence of both identity and inverse elements [208]. In more constrained imaging scenarios, it is often useful to consider subgroups of transformations. For example, in the limit as imaging range gets large, projective transformations are restricted to their affine subgroup. Various subgroups are listed below by their respective degrees of freedom

- (1) Rotation w/ implicit origin - $SO(2, \Re)$
- (2) Translation - $E(1) \times E(1)$
- (3) Proper Rotation and Translation - $E^+(2)$
- (4) Proper Rotation, Isotropic Magnification, and Translation
- (5) Proper Rotation, Magnification, and Translation
- (6) Affine group

This set is of particular interest because of the invariants of the members and their relationship to common imaging scenarios. For example, rotation + translation is listed as a 3 degree of freedom subgroup, but isotropic magnification + translation is not. This is due to the fact that stabilizing an imaging system such that one can assume no rotation is present is extremely difficult, while imaging a distant static object while not moving the focal length

is relatively common. Alternative representations that span the same space are possible, but these characterizations given here have been found to be particularly useful.

Consider the mathematical form of these transformations and how one would distinguish between them. Because some of the subclasses contain others, two representations are convenient. The first is the full projective matrix representation $\mathbf{M} \in PL(3)$, and the second is a vector \mathbf{x} of the intrinsic dimension of the transformation. For example, one could represent a rotation using a 3×3 projective matrix, but simply giving the angle would also be sufficient if we 1st agree to use the subclasses given above. More generally, given a predetermined ordering of the dimensions, one can form a mapping between the vector representations and its equivalent vector form. For the aforementioned subgroups, these relationships are summarized below.

$$(1) \mathbf{M} = \begin{bmatrix} \cos(x_1) & -\sin(x_1) & 0 \\ \sin(x_1) & \cos(x_1) & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (5) \mathbf{M} = \begin{bmatrix} x_3 \cos(x_5) & -x_4 \sin(x_5) & x_1 \\ x_3 \sin(x_5) & x_4 \cos(x_5) & x_2 \\ 0 & 0 & 1 \end{bmatrix}$$

$$(2) \mathbf{M} = \begin{bmatrix} 1 & 0 & x_1 \\ 0 & 1 & x_2 \\ 0 & 0 & 1 \end{bmatrix} \quad (6) \mathbf{M} = \begin{bmatrix} x_3 & x_5 & x_1 \\ x_4 & x_6 & x_2 \\ 0 & 0 & 1 \end{bmatrix}$$

$$(3) \mathbf{M} = \begin{bmatrix} \cos(x_3) & -\sin(x_3) & x_1 \\ \sin(x_3) & \cos(x_3) & x_2 \\ 0 & 0 & 1 \end{bmatrix} \quad (8) \mathbf{M} = \begin{bmatrix} x_3 & x_5 & x_1 \\ x_4 & x_6 & x_2 \\ x_7 & x_8 & 1 \end{bmatrix}$$

$$(4) \mathbf{M} = \begin{bmatrix} x_3 \cos(x_4) & -x_3 \sin(x_4) & x_1 \\ x_3 \sin(x_4) & x_3 \cos(x_4) & x_2 \\ 0 & 0 & 1 \end{bmatrix}$$

Note that for these representations to be unique all rotations must be assumed to be proper. Improper rotations are, however, still supported through magnifications. Using this transformation set a roto-inverse requires at least 5 degrees of freedom, with “flipping” and “scaling” treated as part of the same class.

Determining the mapping from the projective transform to the subclass with the fewest degrees of freedom is less straight-forward because the transformation representations are not unique. This stems from the fact magnification within the set $\{\pm 1\}$ and rotation are cyclic subgroups. When composed, these operators have a 2 intrinsic dimensions but are not simply-connected. In other words, we need both dimensions to describe the operator but the topological structure contains cycles that make the representations non-unique. This

same issue is more commonly encountered through the non-uniqueness of the SVD. This ambiguity can be resolved, but doing so is beyond the scope of this work.

APPENDIX D

Warp Partials

This section provides the partial derivatives of the warping operator (camera pose) relative to its parameterization. To begin, recall that the warping operator maps an object \mathbf{f} into an imaging plane through a projective transformation. This transformation is described in terms of the transformation of the coordinates through the 3×3 matrix \mathbf{V} . Depending on the (sub)-class of projective transforms used, this matrix contains at most 8 degrees-of-freedom described by the vector \mathbf{v} . The object representation is not closed under this transformational class, so the continuous object is approximated through an interpolation operator \mathbf{Q} as follows

$$\mathbf{W}(\mathbf{v}) \circ f(x, y, \lambda) \stackrel{\text{def}}{=} f(x(\mathbf{v}), y(\mathbf{v}), \lambda) \tag{D.1}$$

$$\approx \mathbf{Q}(x(\mathbf{v}), y(\mathbf{v})) \mathbf{f}, \tag{D.2}$$

The action of the projective transformation on the coordinates is given in homogenous coordinates as

$$\begin{bmatrix} r'(\mathbf{v}) \\ c'(\mathbf{v}) \\ w'(\mathbf{v}) \end{bmatrix} = \mathbf{L}^{-1} \mathbf{V}^{-1}(\mathbf{v}) \mathbf{L} \begin{bmatrix} r \\ c \\ 1 \end{bmatrix}, \tag{D.3}$$

where the variables “r” and “c” are chosen to suggest row and column, and the similarity matrix \mathbf{L} enables an arbitrary relabeling of this space. The structure of $\mathbf{V}(\mathbf{v})$ depends on the particular transformation class (See Appendix C), and this matrix enters into (D.3) through its inverse to make the transformation description more intuitive. If one wants to describe a transformation of an *image*, then it follows that this transformation is brought about by resampling to a new set of *coordinates* that are transformed by the inverse of the

transformation. Defining transforms without the inverse is not wrong, but one must accept that $\mathbf{V} = \text{diag}([.5 \ 1 \ 1]^T)$ is describing a magnification rather than a demagnification.

Differentiating the homogenous coordinate transformation with respect the warp parameterization one finds

$$\frac{\partial}{\partial \mathbf{v}_i} [(\mathbf{L}^{-1}\mathbf{V}\mathbf{L})^{-1}] = \frac{\partial}{\partial \mathbf{v}_i} [\mathbf{L}^{-1}\mathbf{V}^{-1}\mathbf{L}] \quad (\text{D.4})$$

$$= \mathbf{L}^{-1} \frac{\partial}{\partial \mathbf{v}_i} [\mathbf{V}^{-1}] \mathbf{L} \quad (\text{D.5})$$

$$= -\mathbf{L}^{-1}\mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \mathbf{v}_i} \mathbf{V}^{-1}\mathbf{L} \quad (\text{D.6})$$

$$= -(\mathbf{L}^{-1}\mathbf{V}\mathbf{L})^{-1} \left(\mathbf{L}^{-1} \frac{\partial \mathbf{V}}{\partial \mathbf{v}_i} \mathbf{L} \right) (\mathbf{L}^{-1}\mathbf{V}\mathbf{L})^{-1} \quad (\text{D.7})$$

The codomain of this derivative is still in homogeneous coordinates. Projection back to \mathbb{R}^2 is achieved by applying the equivalence relation for homogeneous representations (Chapter II, Equation 2.1) in conjunction with the quotient rule for differentiation. Let \mathbf{x} be a coordinate in normalized homogeneous form. Then $\mathbf{x}_3 = 1$, and the associated projection into \mathbb{R}^2 is given by $\mathbf{x}_{1:2}$. Let \mathbf{y} be the transformation of this coordinate as given by (D.3). Then the derivative of the transformation w.r.t. \mathbf{v} is given in \mathbb{R}^2 by

$$\frac{\partial}{\partial \mathbf{v}_i} [(\mathbf{L}^{-1}\mathbf{V}\mathbf{L})^{-1}] \mathbf{x} = \frac{\frac{\partial}{\partial \mathbf{v}_i} [(\mathbf{L}^{-1}\mathbf{V}\mathbf{L})^{-1}]_{1:2,:} \mathbf{x} - \left(\frac{\partial}{\partial \mathbf{v}_i} [(\mathbf{L}^{-1}\mathbf{V}\mathbf{L})^{-1}]_{3,:} \mathbf{x} \right) \langle \mathbf{y} \rangle_{1:2}}{\left[(\mathbf{L}^{-1}\mathbf{V}\mathbf{L})^{-1} \right]_{3,:} \mathbf{x}} \quad (\text{D.8})$$

This form of the gradient is particularly useful because it emphasizes ones ability to use traditional linear algebra operations to efficiently compute the derivative over the entire image field. Differentiating the warping operator as defined through (D.2) with respect to \mathbf{v} one finds

$$\frac{\partial}{\partial \mathbf{v}} [\mathbf{Q}(\mathbf{v})] = \frac{\partial \mathbf{Q}}{\partial r}(r(\mathbf{v}), c(\mathbf{v})) \frac{\partial r}{\partial \mathbf{v}} + \frac{\partial \mathbf{Q}}{\partial c}(r(\mathbf{v}), c(\mathbf{v})) \frac{\partial c}{\partial \mathbf{v}}, \quad (\text{D.9})$$

and expression that depends only the derivative of the interpolator with respect to the coordinates, and the coordinates with respect to the warping parameterization. The latter terms are provided in vector form through (D.8).

The Hession of this operator is less straight-forward, and is included here for completeness. As before, we will begin by computing the derivative of the coordinate transformation

with respect to the warp parameterization.

$$\frac{\partial}{\partial \mathbf{v}_j} \left[-\mathbf{L}^{-1} \mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \mathbf{v}_i} \mathbf{V}^{-1} \mathbf{L} \right] = -\mathbf{L}^{-1} \frac{\partial}{\partial \mathbf{v}_j} \left[\mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \mathbf{v}_i} \mathbf{V}^{-1} \right] \mathbf{L} \quad (\text{D.10})$$

$$\frac{\partial}{\partial \mathbf{v}_j} \left[\mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \mathbf{v}_i} \mathbf{V}^{-1} \right] = -\mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \mathbf{v}_j} \mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \mathbf{v}_i} \mathbf{V}^{-1} + \mathbf{V}^{-1} \frac{\partial}{\partial \mathbf{v}_j} \left[\frac{\partial \mathbf{V}}{\partial \mathbf{v}_i} \mathbf{V}^{-1} \right] \quad (\text{D.11})$$

$$\frac{\partial}{\partial \mathbf{v}_j} \left[\frac{\partial \mathbf{V}}{\partial \mathbf{v}_i} \mathbf{V}^{-1} \right] = \frac{\partial^2 \mathbf{V}}{\partial \mathbf{v}_j \partial \mathbf{v}_i} \mathbf{V}^{-1} - \frac{\partial \mathbf{V}}{\partial \mathbf{v}_i} \mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \mathbf{v}_j} \mathbf{V}^{-1} \quad (\text{D.12})$$

Substituting (D.12) into (D.11) fully reduces this term

$$\begin{aligned} \frac{\partial}{\partial \mathbf{v}_j} \left[\mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \mathbf{v}_i} \mathbf{V}^{-1} \right] &= \mathbf{V}^{-1} \frac{\partial^2 \mathbf{V}}{\partial \mathbf{v}_j \partial \mathbf{v}_i} \mathbf{V}^{-1} - \mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \mathbf{v}_i} \mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \mathbf{v}_j} \mathbf{V}^{-1} - \\ &\quad \mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \mathbf{v}_j} \mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \mathbf{v}_i} \mathbf{V}^{-1}, \end{aligned} \quad (\text{D.13})$$

and substituting (D.14) into (D.10) leads to the desired result.

$$\frac{\partial}{\partial \mathbf{v}_j} \left[-\mathbf{L}^{-1} \mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \mathbf{v}_i} \mathbf{V}^{-1} \mathbf{L} \right] = \mathbf{L}^{-1} \mathbf{V}^{-1} \left[\frac{\partial \mathbf{V}}{\partial \mathbf{v}_i} \mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \mathbf{v}_j} + \frac{\partial \mathbf{V}}{\partial \mathbf{v}_j} \mathbf{V}^{-1} \frac{\partial \mathbf{V}}{\partial \mathbf{v}_i} - \frac{\partial^2 \mathbf{V}}{\partial \mathbf{v}_i \partial \mathbf{v}_j} \right] \mathbf{V}^{-1} \mathbf{L} \quad (\text{D.14})$$

This derivative must be transferred from $\mathbb{P}^2 \rightarrow \mathbb{R}^2$ by applying the definition of the equivalence relation. This involves differentiating the ratio of two functions, each of two arguments.

The structure of this transformation is given by

$$\begin{aligned} \frac{\partial^2}{\partial x \partial y} \left[\frac{f(x,y)}{g(x,y)} \right] &= \left[2 \left(\frac{f}{g} \right) \left(\frac{1}{g} \frac{\partial g}{\partial x} \right) - \frac{1}{g} \frac{\partial f}{\partial x} \right] \left(\frac{1}{g} \frac{\partial g}{\partial y} \right) - \left(\frac{1}{g} \frac{\partial g}{\partial x} \right) \left(\frac{1}{g} \frac{\partial f}{\partial y} \right) + \\ &\quad \frac{1}{g} \frac{\partial^2 f}{\partial x \partial y} - \frac{f}{g} \left(\frac{1}{g} \frac{\partial^2 g}{\partial x \partial y} \right), \end{aligned} \quad (\text{D.15})$$

where both the single and mixed partials have been expressed above. Finally, given a interpolator that implements both its first and second partial derivatives with respect to the coordinates, one can easily express the hessian of the warping operator with respect to \mathbf{v}

$$\begin{aligned} \frac{\partial^2}{\partial \mathbf{v}_i \partial \mathbf{v}_j} [\mathbf{Q}(\mathbf{v})] &= \left(\frac{\partial^2 \mathbf{Q}}{\partial r^2} \frac{\partial r}{\partial \mathbf{v}_i} + \frac{\partial^2 \mathbf{Q}}{\partial r \partial c} \frac{\partial c}{\partial \mathbf{v}_i} \right) \frac{\partial r}{\partial \mathbf{v}_j} + \frac{\partial \mathbf{Q}}{\partial r} \frac{\partial^2 r}{\partial \mathbf{v}_i \partial \mathbf{v}_j} + \\ &\quad \left(\frac{\partial^2 \mathbf{Q}}{\partial r \partial c} \frac{\partial r}{\partial \mathbf{v}_i} + \frac{\partial^2 \mathbf{Q}}{\partial c^2} \frac{\partial c}{\partial \mathbf{v}_i} \right) \frac{\partial c}{\partial \mathbf{v}_j} + \frac{\partial \mathbf{Q}}{\partial c} \frac{\partial^2 c}{\partial \mathbf{v}_i \partial \mathbf{v}_j}. \end{aligned} \quad (\text{D.16})$$

APPENDIX E

Geometric Optics for the Asymmetric Generalized Imaging Model

The modeling fidelity needed to compliment the previously discussed wavefront-estimation approach in the context of ray-tracing generally leads to intractable forms that must inexorably be evaluated numerically. See, for example, the excellent texts by Borne [19] and Goodman [77]. On the other hand, photographers and those in the remote sensing community generally apply geometrical approximations that are inadequate for the purposes of wavefront estimation. This section derives the defocus induced by changes in depth-of-field for Goodman's generalized imaging model under a paraxial approximation. The results are compatible with the Fraunhofer diffraction model used throughout this thesis, and make explicit the relationship between this model and the less precise geometric model for the Circle-of-Confusion (CoC).

Figure E.1 illustrates the key elements of the generalized camera model, as well as their relationships to the relevant camera and lens parameters. Let D_{in} and D_{out} represent entrance- and exit-pupil diameters respectively. The entrance-pupil is the image of the aperture stop as seen from the object side of the imaging system. Similarly, the exit-pupil is the image of the aperture stop as seen from the image side of the lens. P and P' represent the entrance- and exit- principal-planes [93]. These planes are the positions where the Gaussian lens formula is valid. As a result one has the relationship

$$\frac{1}{r} + \frac{1}{s} = \frac{1}{f}, \quad (\text{E.1})$$

which implies that with respect the ray-tracing geometry, the two principal-planes behave as though they were collocated. When the entrance- and exit-pupil diameters are of different sizes, the camera system induces a magnification. The pupil magnification is denoted




Figure E.1: Diagram and relevant parameters needed to relate a physical camera to the generalized model

by $m_p = D_{out}/D_{in}$. Notice that when the pupil magnification is not one, the entrance- and exit-pupils do not coincide with principal-planes.

The relationship between the principal-planes and their respective exit-pupils can be determined by considering the system when focused to infinity. When this is the case the cone (beam) of light passing through the entrance-pupil intersects P with diameter D_{in} . Applying (E.1), this light then forms a cone that extends from P' with a diameter of D_{in} , to a focus point a distance f away along the optical axis. This cone also passes through the exit-pupil, and from the similar triangles one finds that the distance between the exit-pupil its principal-plane is $(1 - m_p)f$. Working the opposite direction similarly yields the distance between the entrance-pupil its principle-plane.

Figure E.2 shows the rays corresponding to near and far ranges associated with a range-induced-defocus resulting in a Circle-of-Confusion (CoC) of diameter c . From similar triangles one may write

$$\frac{c}{D_{out}} = \frac{s_n - s}{s_n - (1 - m_p)f} = \frac{s - s_f}{s_f - (1 - m_p)f}, \quad (\text{E.2})$$

where it useful to recall that F-Number is commonly used to describe the ratio between focal




Figure E.2: Rays tracing diagram corresponding to the near- and far-focus positions that yield a fixed Circle-of-Confusion (CoC) of diameter c .

length and entrance-pupil diameter¹.

$$f_{\#} = \frac{f}{D_{in}} = \frac{m_p f}{D_{out}} \quad (\text{E.3})$$

Solving for the distance between the exit principle-plane and the point of focus one finds

$$s_n = \frac{s - f_{\#} c \left(\frac{1}{m_p} - 1 \right)}{1 - \frac{c}{D_{out}}} = \frac{f s - f_{\#} c \left(\frac{1}{m_p} - 1 \right) f}{f - f_{\#} \frac{c}{m_p}}, \quad (\text{E.4})$$

and similarly,

$$s_f = \frac{s + f_{\#} c \left(\frac{1}{m_p} - 1 \right)}{1 + \frac{c}{D_{out}}} = \frac{f s + f_{\#} c \left(\frac{1}{m_p} - 1 \right) f}{f + f_{\#} \frac{c}{m_p}}. \quad (\text{E.5})$$

Applying the thin lens law one may write $s = \frac{r f}{r - f}$, and back-substituting one may express the near and far imaging ranges in terms of the CoC diameter.

$$r_n = \frac{r f - f_{\#} c \left(\frac{1}{m_p} - 1 \right) (r - f)}{(r - f) \left(1 - \frac{c}{D_{out}} \right)} = \frac{r f^2 - f_{\#} c \left(\frac{1}{m_p} - 1 \right) (r - f)}{f^2 + f_{\#} c (r - f)} \quad (\text{E.6})$$

$$r_f = \frac{r f + f_{\#} c \left(\frac{1}{m_p} - 1 \right) (r - f)}{(r - f) \left(1 + \frac{c}{D_{out}} \right)} = \frac{r f^2 + f_{\#} c \left(\frac{1}{m_p} - 1 \right) (r - f)}{f^2 - f_{\#} c (r - f)} \quad (\text{E.7})$$

$$(\text{E.8})$$

¹Smaller $f_{\#}$ optics are often called “faster” because they collect a fixed amount of light in less time than optics with a larger $f_{\#}$. Specifically, the radiance at the detector goes as the square of $f_{\#}$.

Given a (possibly defocused) point object at some range r_d observed through a camera that is in focus at range r , the associated CoC diameter is given by

$$c_d = \frac{m_p f^2}{f_\# (r - f)} \cdot \frac{|r_d - r|}{r_d + (r_d - f)(m_p - 1)}. \quad (\text{E.9})$$

Similarly, for a system focused at some range r , the difference between the near and far ranges yielding the CoC diameter c is given by

$$r_f - r_n = \frac{2f_\# c f^2 \left[r + f \left(\frac{1}{m_p} - 1 \right) \right] (r - f)}{f^4 - f_\#^2 c^2 (r - f)^2} \quad (\text{E.10})$$

The relationship between c_d and peak-to-valley waves of defocus is given by the geometry illustrated in Figure E.3. The difference in sagitta between the wavefronts is given by




Figure E.3: The relationship between Peak-to-Valley (PF) defocus and the Circle-of-Confusion (CoC) is given by the difference in sagitta between the in- and out- of focus converging waves.

$$\delta_{PV} = (s_d - s) + \sqrt{[s - (m_p - 1)f]^2 - \left(\frac{D_{out}}{2} \right)^2} - \sqrt{[s_d - (m_p - 1)f]^2 - \left(\frac{D_{out}}{2} \right)^2} \quad (\text{E.11})$$

Solving in terms of imaging range, one finds

$$\delta_{PV} = f \left[\frac{r_d}{r_d - f} - \frac{r}{r - f} + \sqrt{\left(1 - m_p + \frac{r}{r - f}\right)^2 - \left(\frac{m_p}{2f_{\#}}\right)^2} - \sqrt{\left(1 - m_p + \frac{r_d}{r_d - f}\right)^2 - \left(\frac{m_p}{2f_{\#}}\right)^2} \right] \quad (\text{E.12})$$

and a circular aperture, the relationship between peak-to-valley defocus error and RMS defocus error is $\delta_{PV} = \delta_{RMS}\sqrt{12}$. By understanding the relationship between the CoC and wavefront description of the an imaging system, one is more easily apply physically realistic models to imaging experiments.

APPENDIX F

Numerical PSF Evaluation

High fidelity statistical inverse-techniques, and their associated information-theoretic bounds, are only as accurate as the models used to approximate the physical phenomenon being studied. That is to say, bounds correctly derived for a particular model are only relevant to the extent that the model accurately represents the physical process in question. This section gives insight into the fidelity of the blur model used throughout this composition, and provides support to the implicit claim made in Chapter III that it is sufficient for examining estimation performance on the order of 1×10^{-3} waves RMS. This level of accuracy is in part due to a first-order correction applied to the aperture function which to the author's knowledge has not appeared in the literature. That correction is also derived here for completeness.

Recall that the blur in an optical system induced by wavefront-errors in the exit-pupil relate to the system's PSF as follows

$$h(x, y) = \mathcal{F}^{-1}\{H(\tilde{w}_x, \tilde{w}_y)\} \quad (\text{F.1})$$

$$= c_0 |\mathcal{F}^{-1}\{C(\tilde{w}_x, \tilde{w}_y)\}|^2 \quad (\text{F.2})$$

$$C(\tilde{w}_x, \tilde{w}_y) = A(\tilde{w}_x, \tilde{w}_y) \exp \left[j2\pi \frac{\lambda}{\lambda_{ref}} \Psi(\tilde{w}_x, \tilde{w}_y) \right] \quad (\text{F.3})$$

$$\tilde{w} \stackrel{\text{def}}{=} \frac{\lambda}{\lambda_{ref}} w \quad (\text{F.4})$$

where c_0 is a normalizing constant that ensures the PSF integrates to 1, the 2π in the exponent of the CTF implies that phase errors are being represented in units of waves, and the ratio λ/λ_{ref} accounts for the scaling between the wavelength being used to represent the optical system and the wavelength where the linear forward model is being evaluated.

This expressions are exact, but do not indicate how one should implement them in a finite-precision computer.

After implementing the proceeding equations, two sources of error must be considered; errors in accuracy and errors in precision. These two types of errors are, like the blur operator itself, are implicitly functions of the number of bits used to represent the system. They are analogous to the bias and variance for an estimator in the sense that one can have a very accurate (low-bias) estimator with very high variance, or conversely one can have a very inaccurate (high-bias) estimator with very low variance. In the context of implementing a forward operator, errors in precision (variance) is referred to here as the variation of the operator in the neighborhood of its implicit parameters. As such, a complete description of the errors in any implementation must address both sources¹.

The physics of an optical system leads to parameterizations over Ψ resulting in the most parsimonious representations of the blur operator, however, these parameters related to the PSF in a non-linear way. The Sum of Absolute Error (SAE) is chosen to measure PSF accuracy, and unless otherwise noted errors are measured about an unaberrated system. For example, Figure F.1 illustrates the PSF variation measured in SAE for a 1×10^{-3} wave RMS variation in the wavefront description. This plot is a histogram because not all directions




Figure F.1: PSF variation induced by 1×10^{-3} waves RMS in the exit pupil over 7 Zernike modes

in the parameter space manifest themselves as equal levels of variation in the PSF. Never the less, if one is interested in estimating PSF variations on order of 1×10^{-3} waves, the blur operator must exhibit errors in the PSF less than $\sim 3 \times 10^{-3}$. Given this knowledge, Figure F.2 shows that the use of a binary aperture is extremely inconvenient given our

¹In this case accuracy (bias) errors will dominate precision errors

tight accuracy constraints. This is due to the fact that the computational complexity of computing a PSF goes as $O(k^2 \log k)$, where k is the aperture diameter, and that the accuracy of the blur approximation is rolling off at or near the level of deviation expected from a 1×10^{-3} RMS deviation in the parameters. Taken together, this suggests that the use of a binary aperture essentially prohibits estimation to the desired level. A higher-fidelity blur




Figure F.2: Blur operator error as a function of exit-pupil diameter

operator can be achieved by using shaded aperture. Despite the fact that the true aperture was binary, a shaded aperture can yield a better approximation because one is essentially using numerical quadrature in the calculation of the PSF. Another way of looking at the problem is to note that Equations F.2 and F.3 specify that the peak of the PSF (Strehl ratio), for an aberration-free aperture, is proportional to the square of the area of the aperture function. For a circular aperture, this area is known a priori and may be viewed as a constraint on the aperture approximation. Figure F.3 shows the reduction in PSF error (bias) as a function of exit pupil diameter for using the shaded aperture approximation. At a pupil diameter of only 225, the associated variation in pupil size (variance) is reduced to less than $2E-4$. The details of how one computes this improved approximation is given in Section G, and its necessity when attempting to estimate errors on the order of this described in Chapter III should now be clear.




Figure F.3: Blur operator error as a function of exit-pupil diameter for shaded and binary aperture approximations

APPENDIX G

Shaded Aperture Derivation

The finite numerical approximation to a continuous circular aperture can be improved by allowing each pixel in the approximation to take on the value of the area of intersection between the pixels support and the disc representing the true continuous aperture. For our purposes we'll assume the pixels lie on $\mathcal{Z} \times \mathcal{Z}$ with the disc centered at the origin with radius r . Then the problem is then naturally parameterized by where the disc intersects the square representing the pixel centered about the point P . The area of intersection and the relevant points are illustrated in Figure G.1. Exploiting the symmetry in the problem allows one to




Figure G.1: Area of intersection between (shown in blue) between a disc of radius r and a pixel centered about the point P . The problem is naturally parameterized by the boundary points A and B

consider only those pixels intersecting the first quadrant of the cartesian plane. Furthermore, only “boundary” pixels are of interest, where these pixels are defined as those taking on values in the range $(0, 1)$. Pixels entirely within the circle of radius r are given by

$$\left[x + \frac{1}{2} \operatorname{sgn} \left(x - \frac{1}{2} \right) \right]^2 + \left[y + \frac{1}{2} \operatorname{sgn} \left(y - \frac{1}{2} \right) \right]^2 \leq r^2 \quad (\text{G.1})$$

whereas those entirely outside the circle are given by

$$\left[x + \frac{1}{2} \text{sgn}(x) \right]^2 + \left[y + \frac{1}{2} \text{sgn}(y) \right]^2 > r^2 \quad (\text{G.2})$$

The remaining pixels are “boundary” pixels as illustrated in Figure G.2, and calculating their values is the subject of the remainder of this section.




Figure G.2: Pupil diameter of 9 with boundary pixels (with respect to a circle of radius 4.5) shown in light red

After projecting each of the boundary pixels into quadrant 1, the area of intersection with the disc can be written as a function of the area under the curve defining the boundary of the disc and the half plane

$$\int_a^b \sqrt{r^2 - x^2} dx = \frac{r^2}{2} \left[\arcsin\left(\frac{b}{r}\right) - \arcsin\left(\frac{a}{r}\right) \right] + \frac{b}{2} \sqrt{r^2 - b^2} - \frac{a}{2} \sqrt{r^2 - a^2} \quad (\text{G.3})$$

The means of computing the area of intersection fall into 2 primary cases; (i) “Full left edge” and (ii) “Partial left edge” corresponding to the left edge of the pixel entirely falling within the disc verses only partially falling within the disc. These two primary cases each have two principle sub-cases corresponding to how the bounding circle exits the pixel. Additional exceptions occur when the pixel falls on the principle axes, and the edge cases associated with circles of very small radii. The function `circularAperture.m` implements these ideas and

provides both the typical nearest-neighbor circular aperture, as well as the shaded aperture implementation described in this appendix.

APPENDIX H

OTF Partial

This section provides the partial derivatives of the PSF relative to the phase error parameterization. To begin, recall the PSF relates to the phase errors as follows

$$h(x, y) = \mathcal{F}^{-1}\{H(w_x, w_y)\} \quad (\text{H.1})$$

$$= c_0 |\mathcal{F}^{-1}\{C(w_x, w_y)\}|^2 \quad (\text{H.2})$$

$$C(w_x, w_y) = A(w_x, w_y) \exp[j2\pi\Psi(w_x, w_y, \alpha)] \quad (\text{H.3})$$

where h is the PSF, H is the OTF, C is the CTF, A is the aperture, Ψ is the wavefront phase function, and c_0 is a constant that normalizes the PSF to unit volume. Taking the gradient one may write

$$\nabla_\alpha[h] = 2c_0 \text{Re}\left(\left(\mathcal{F}^{-1}\{C\}\right)^* \mathcal{F}^{-1}\{\nabla_\alpha C\}\right) \quad (\text{H.4})$$

$$= 2c_0 \text{Re}\left(\left(\mathcal{F}^{-1}\{C\}\right)^* (j2\pi\mathcal{F}^{-1}\{C\nabla_\alpha\Psi\})\right) \quad (\text{H.5})$$

$$= -4\pi c_0 \text{Im}\left(\left(\mathcal{F}^{-1}\{C\}\right)^* \mathcal{F}^{-1}\{C\nabla_\alpha\Psi\}\right) \quad (\text{H.6})$$

$$= 4\pi c_0 \text{Im}\left(\mathcal{F}^{-1}\{C\} \left(\mathcal{F}^{-1}\{C\nabla_\alpha\Psi\}\right)^*\right) \quad (\text{H.7})$$

Differentiating a second time results in the Hessian

$$\nabla_\alpha^2[h] = 4\pi \text{Im}\left(j2\pi\mathcal{F}^{-1}\{C\nabla_\alpha\Psi\} \left(\mathcal{F}^{-1}\{C\nabla_\alpha\Psi\}\right)^* + \left(\nabla_\alpha[\mathcal{F}^{-1}\{C\nabla_\alpha\Psi\}]\right)^* \mathcal{F}^{-1}\{C\}\right) \quad (\text{H.8})$$

where the chain rule must again be applied the intermediate term

$$\nabla_{\alpha} [\mathcal{F}^{-1}\{C\nabla_{\alpha}\Psi\}] = \mathcal{F}^{-1}\{j2\pi C\nabla_{\alpha}\Psi\nabla_{\alpha}\Psi + C\nabla_{\alpha}^2\Psi\} \quad (\text{H.9})$$

$$= j2\pi\mathcal{F}^{-1}\{C(\nabla_{\alpha}\Psi\nabla_{\alpha}\Psi + \nabla_{\alpha}^2\Psi)\} \quad (\text{H.10})$$

Plugging this result into Equation H.8 leads to the desired result

$$\nabla_{\alpha}^2[h] = 8\pi^2\text{Re}\left(\mathcal{F}^{-1}\{C\nabla_{\alpha}\Psi\}(\mathcal{F}^{-1}\{C\nabla_{\alpha}\Psi\})^* - \mathcal{F}^{-1}\{C\}(\mathcal{F}^{-1}\{C(\nabla_{\alpha}\Psi\nabla_{\alpha}\Psi + \nabla_{\alpha}^2\Psi)\})^*\right) \quad (\text{H.11})$$

APPENDIX I

Derivations for Known Objects

Consider the derivative of the likelihood function given by Equation 3.5. The differential in general form is

$$\nabla L(\boldsymbol{\theta}) = -\mathbf{a}^T \boldsymbol{\Sigma}^{-1} \nabla \mathbf{a} + \frac{1}{2} \mathbf{a}^T \boldsymbol{\Sigma}^{-1} \nabla \boldsymbol{\Sigma} \boldsymbol{\Sigma}^{-1} \mathbf{a} - \frac{1}{2} \left(\mathbf{b}^{\odot -1} \right)^T \nabla \mathbf{b} \quad \text{s.t.} \quad (\text{I.1})$$

$$\mathbf{a} = \mathbf{d} - c_3 \mathbf{H} (c_1 \mathbf{f} + c_2) - c_4 \quad (\text{I.2})$$

$$\mathbf{b} = c_3^2 \mathbf{H} (c_1 \mathbf{f} + c_2) + \sigma_r^2 \quad (\text{I.3})$$

Next we sequentially evaluate the individual terms

$$\nabla \mathbf{H} = \nabla \left[\text{diag} \left(\left\{ \mathbf{T} \mathbf{S} \mathbf{B}(\boldsymbol{\alpha}) \mathbf{W}(\mathbf{v}_j) \right\}_{j=1}^J \right) (\mathbf{1}_J \otimes \mathbf{I}_{n_f}) \right] \quad (\text{I.4})$$

$$= \text{diag} \left(\left\{ \mathbf{T} \mathbf{S} \left[\mathbf{B}(\boldsymbol{\alpha}) \nabla \mathbf{W}(\mathbf{v}_j) + \nabla \mathbf{B}(\boldsymbol{\alpha}) \mathbf{W}(\mathbf{v}_j) \right] \right\}_{j=1}^J \right) (\mathbf{1}_J \otimes \mathbf{I}_{n_f}) \quad (\text{I.5})$$

Later we will use the fact that $\nabla \mathbf{B}$ and $\nabla \mathbf{W}$ both have $\mathbf{1}$ in their null-space implying so to does $\nabla \mathbf{H}$. This follows immediately from the fact that $\mathbf{H}(\mathbf{1}) = \mathbf{1}$

$$\nabla \mathbf{a} = -c_3 \nabla \mathbf{H} (c_1 \mathbf{f} + c_2) - \mathbf{H} \nabla c_1 c_3 \mathbf{f} + c_2 c_3 - \nabla c_4 \quad (\text{I.6})$$

Note that for vector differentials

$$\nabla(\cdot) = \left[\nabla_{\mathbf{v}}(\cdot) \quad \nabla_{\boldsymbol{\alpha}}(\cdot) \quad \nabla_{\mathbf{c}}(\cdot) \quad \nabla_{\sigma_r}(\cdot) \right] \quad (\text{I.7})$$

Matrix differentials are taken one element at a time, but their actions are similarly defined.

$$[\nabla \mathbf{H} \mathbf{f}]_{(j,:)} = \mathbf{T} \mathbf{S} \left[\mathbf{B}(\boldsymbol{\alpha}) \nabla \mathbf{W}(\mathbf{v}_j) \mathbf{f} + \nabla \mathbf{B}(\boldsymbol{\alpha}) \mathbf{W}(\mathbf{v}_j) \mathbf{f} \right] \quad (\text{I.8})$$

where for clarity we use $[\{\cdot\}]$ to denote the concatenation of the set elements, and $[\cdot]_{(1:J,:)}$ to indicate that the concatenation is occurring along dimension 1 filling elements 1:J (in this case each “element” contains n_f scalars along dimension 1) and “:” indicates that all of the column indices are assigned as necessary.

Explicitly taking the action of the differential we get

$$\nabla \mathbf{a} = -c_1 c_3 \nabla \mathbf{H} \mathbf{f} - \begin{bmatrix} \mathbf{0}_{J n_f \times (n_v + n_a)} & c_3 \mathbf{H} \mathbf{f} & c_3 \mathbf{1}_{J n_f} & c_1 \mathbf{H} \mathbf{f} + c_2 & \mathbf{1}_{J n_f} & \mathbf{0}_{J n_f} \end{bmatrix} \quad (\text{I.9})$$

Similarly

$$\nabla \mathbf{b} = c_1 c_3^2 \nabla \mathbf{H} \mathbf{f} + \begin{bmatrix} \mathbf{0}_{J n_f \times (n_v + n_a)} & c_3^2 \mathbf{H} \mathbf{f} & c_3^2 \mathbf{1}_{J n_f} & 2c_3 (c_1 \mathbf{H} \mathbf{f} + c_2) & \mathbf{0}_{J n_f} & 2\sigma_r \mathbf{1}_{J n_f} \end{bmatrix} \quad (\text{I.10})$$

Turning to the covariance term

$$\boldsymbol{\Sigma}^{-1} = \text{diag} \left([c_3^2 \mathbf{H} (c_1 \mathbf{f} + c_2) + \sigma_r^2]^{\odot -1} \right) \quad (\text{I.11})$$

$$= \text{diag} \left(\left[\text{diag} \left(c_3^2 \{ \mathbf{T} \mathbf{S} \mathbf{B}(\boldsymbol{\alpha}) \mathbf{W}(\mathbf{v}_j) \}_{j=1}^J (\mathbf{I}_J \otimes \mathbf{I}_{n_f}) (c_1 \mathbf{f} + c_2) + \sigma_r^2 \right) \right]^{\odot -1} \right) \quad (\text{I.12})$$

$$= \text{diag} \left(\left[\text{vec} \left(\{ c_3^2 \mathbf{T} \mathbf{S} \mathbf{B}(\boldsymbol{\alpha}) \mathbf{W}(\mathbf{v}_j) (c_1 \mathbf{f} + c_2) + \sigma_r^2 \}_{j=1}^J \right) \right]^{\odot -1} \right) \quad (\text{I.13})$$

$$= \text{diag} \left(\left[c_1 c_3^2 \text{vec} \left(\{ \mathbf{T} \mathbf{S} \mathbf{B}(\boldsymbol{\alpha}) \mathbf{W}(\mathbf{v}_j) \mathbf{f} \}_{j=1}^J \right) + c_2 c_3^2 + \sigma_r^2 \right]^{\odot -1} \right) \quad (\text{I.14})$$

$$= \text{diag} \left(\left[c_1 c_3^2 \text{vec} \left(\{ \mathbf{H}_j \mathbf{f} \}_{j=1}^J \right) + c_2 c_3^2 + \sigma_r^2 \right]^{\odot -1} \right) \quad (\text{I.15})$$

Thus,

$$\mathbf{a}^T \boldsymbol{\Sigma}^{-1} \nabla \mathbf{a} = - \sum_{j=1}^J \left(\frac{\mathbf{d}_j - c_3 (c_1 \mathbf{T} \mathbf{S} \mathbf{B}(\boldsymbol{\alpha}) \mathbf{W}(\mathbf{v}_j) \mathbf{f} + c_2) - c_4}{c_1 c_3^2 \mathbf{T} \mathbf{S} \mathbf{B}(\boldsymbol{\alpha}) \mathbf{W}(\mathbf{v}_j) \mathbf{f} + c_2 c_3^2 + \sigma_r^2} \right)^T \quad (\text{I.16})$$

$$\begin{bmatrix} c_1 c_3 \mathbf{T} \mathbf{S} \mathbf{B}(\boldsymbol{\alpha}) (\nabla \mathbf{W}(\mathbf{v}_j)) \mathbf{f} & c_1 c_3 \mathbf{T} \mathbf{S} (\nabla \mathbf{B}) \mathbf{W}(\mathbf{v}_j) \mathbf{f} & c_3 \mathbf{H}_j \mathbf{f} & c_3 \mathbf{1}_{n_f} \cdots \\ c_1 \mathbf{H}_j \mathbf{f} + c_2 & \mathbf{1}_{n_f} & \mathbf{0}_{n_f} \end{bmatrix}$$

where \mathbf{H}_j identifies the component of the forward operator that generates the j^{th} frame of data. This essentially lays out the individual component gradients of the 1st term in the

likelihood. Similarly

$$\mathbf{a}^T \boldsymbol{\Sigma}^{-1} \nabla \boldsymbol{\Sigma} \boldsymbol{\Sigma}^{-1} \mathbf{a} = \mathbf{a}^T \boldsymbol{\Sigma}^{-1} \nabla \text{diag}(\mathbf{b}) \boldsymbol{\Sigma}^{-1} \mathbf{a} \quad (\text{I.17})$$

$$= \left((\boldsymbol{\Sigma}^{-1} \mathbf{a})^{\odot 2} \right)^T \nabla \mathbf{b} \quad (\text{I.18})$$

$$= \left(\left(\frac{\mathbf{a}}{\mathbf{b}} \right)^{\odot 2} \right)^T \nabla \mathbf{b} \quad (\text{I.19})$$

$$= \sum_{j=1}^J \left(\left(\frac{\mathbf{d}_j - c_3 (c_1 \mathbf{H}_j \mathbf{f} + c_2) - c_4}{c_1 c_3^2 \mathbf{H}_j \mathbf{f} + c_2 c_3^2 + \sigma_r^2} \right)^{\odot 2} \right)^T \quad (\text{I.20})$$

$$\begin{bmatrix} c_1 c_3^2 \mathbf{TSB}(\boldsymbol{\alpha}) (\nabla \mathbf{W}(\mathbf{v}_j)) \mathbf{f} & c_1 c_3^2 \mathbf{TS}(\nabla \mathbf{B}) \mathbf{W}(\mathbf{v}_j) \mathbf{f} & c_3^2 \mathbf{H}_j \mathbf{f} & c_3^2 \mathbf{1}_{n_f} \dots \\ 2c_3 (c_1 \mathbf{H}_j \mathbf{f} + c_2) & \mathbf{0}_{n_f} & 2\sigma_r \mathbf{1}_{n_f} & \end{bmatrix}$$

$$\mathbf{1}^T \mathbf{b}^{\odot -1} \nabla \mathbf{b} = \sum_{j=1}^J \left(\frac{1}{c_3^2 \mathbf{H}_j (c_1 \mathbf{f} + c_2) + \sigma_r^2} \right)^T \quad (\text{I.21})$$

$$\begin{bmatrix} c_1 c_3^2 \mathbf{TSB}(\boldsymbol{\alpha}) (\nabla \mathbf{W}(\mathbf{v}_j)) \mathbf{f} & c_1 c_3^2 \mathbf{TS}(\nabla \mathbf{B}) \mathbf{W}(\mathbf{v}_j) \mathbf{f} & c_3^2 \mathbf{H}_j \mathbf{f} & c_3^2 \mathbf{1}_{n_f} \dots \\ 2c_3 (c_1 \mathbf{H}_j \mathbf{f} + c_2) & \mathbf{0}_{n_f} & 2\sigma_r \mathbf{1}_{n_f} & \end{bmatrix}$$

APPENDIX J

Known Object CRB Terms

To clarify the calculations of the Fisher information matrix, consider a matrix representation of $\nabla \mathbf{a}$ that specifically isolates the components of the individual data frames. Next consider how these components come together to form the individual elements of the Fisher information matrix.

$$\begin{aligned}
 (\nabla \mathbf{a})^T \nabla \mathbf{a} &= \begin{bmatrix} \mathbf{C}_v^T \mathbf{C}_v & \mathbf{C}_v^T \mathbf{C}_\alpha & \mathbf{C}_v^T \mathbf{C}_{c_1} & \mathbf{C}_v^T \mathbf{C}_{c_2} & \mathbf{C}_v^T \mathbf{C}_{c_3} & \mathbf{C}_v^T \mathbf{C}_{c_4} & \\ & \mathbf{C}_\alpha^T \mathbf{C}_\alpha & \mathbf{C}_\alpha^T \mathbf{C}_{c_1} & \mathbf{C}_\alpha^T \mathbf{C}_{c_2} & \mathbf{C}_\alpha^T \mathbf{C}_{c_3} & \mathbf{C}_\alpha^T \mathbf{C}_{c_4} & \\ & & \mathbf{C}_{c_1}^T \mathbf{C}_{c_1} & \mathbf{C}_{c_1}^T \mathbf{C}_{c_2} & \mathbf{C}_{c_1}^T \mathbf{C}_{c_3} & \mathbf{C}_{c_1}^T \mathbf{C}_{c_4} & \vdots \\ & & & \mathbf{C}_{c_2}^T \mathbf{C}_{c_2} & \mathbf{C}_{c_2}^T \mathbf{C}_{c_3} & \mathbf{C}_{c_2}^T \mathbf{C}_{c_4} & 0 \\ & & & & \mathbf{C}_{c_3}^T \mathbf{C}_{c_3} & \mathbf{C}_{c_3}^T \mathbf{C}_{c_4} & \vdots \\ & & & & & \mathbf{C}_{c_4}^T \mathbf{C}_{c_4} & \\ & & \dots & 0 & \dots & & \end{bmatrix} \quad (\text{J.1}) \\
 (\nabla \mathbf{b})^T \nabla \mathbf{b} &= c_3^2 \begin{bmatrix} \mathbf{C}_v^T \mathbf{C}_v & \mathbf{C}_v^T \mathbf{C}_\alpha & \mathbf{C}_v^T \mathbf{C}_{c_1} & \mathbf{C}_v^T \mathbf{C}_{c_2} & 2\mathbf{C}_v^T \mathbf{C}_{c_3} & 2\frac{\sigma_r}{c_3} \mathbf{C}_v^T \mathbf{C}_{c_4} & \\ & \mathbf{C}_\alpha^T \mathbf{C}_\alpha & \mathbf{C}_\alpha^T \mathbf{C}_{c_1} & \mathbf{C}_\alpha^T \mathbf{C}_{c_2} & 2\mathbf{C}_\alpha^T \mathbf{C}_{c_3} & 2\frac{\sigma_r}{c_3} \mathbf{C}_\alpha^T \mathbf{C}_{c_4} & \\ & & \mathbf{C}_{c_1}^T \mathbf{C}_{c_1} & \mathbf{C}_{c_1}^T \mathbf{C}_{c_2} & 2\mathbf{C}_{c_1}^T \mathbf{C}_{c_3} & \vdots & 2\frac{\sigma_r}{c_3} \mathbf{C}_{c_1}^T \mathbf{C}_{c_4} \\ & & & \mathbf{C}_{c_2}^T \mathbf{C}_{c_2} & 2\mathbf{C}_{c_2}^T \mathbf{C}_{c_3} & 0 & 2\frac{\sigma_r}{c_3} \mathbf{C}_{c_2}^T \mathbf{C}_{c_4} \\ & & & & 4\mathbf{C}_{c_3}^T \mathbf{C}_{c_3} & \vdots & 4\frac{\sigma_r}{c_3} \mathbf{C}_{c_3}^T \mathbf{C}_{c_4} \\ & & \dots & 0 & \dots & & 4\left(\frac{\sigma_r}{c_3}\right)^2 \mathbf{C}_{c_4}^T \mathbf{C}_{c_4} \end{bmatrix} \quad (\text{J.2})
 \end{aligned}$$

Only the upper triangle is given because the matrix is clearly symmetric. The individual

terms, which include the inverse covariance matrices, are computed below.

$$\mathbf{C}_v^T \boldsymbol{\Sigma}^{-1} \mathbf{C}_v = \text{diag} \left(\begin{bmatrix} c_1 \left(\frac{\nabla_{v_1}(\mathbf{H}_1) \mathbf{f}}{\mathbf{1}_{n_{v_1}}^T \otimes \left[\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T (\nabla_{v_1}(\mathbf{H}_1) \mathbf{f}) \\ \vdots \\ c_1 \left(\frac{\nabla_{v_J}(\mathbf{H}_J) \mathbf{f}}{\mathbf{1}_{n_{v_J}}^T \otimes \left[\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T (\nabla_{v_J}(\mathbf{H}_J) \mathbf{f}) \end{bmatrix} \right) \quad (\text{J.3})$$

$$\mathbf{C}_v^T \boldsymbol{\Sigma}^{-2} \mathbf{C}_v = \text{diag} \left(\begin{bmatrix} \frac{1}{c_3^2} \left(\frac{\nabla_{v_1}(\mathbf{H}_1) \mathbf{f}}{\mathbf{1}_{n_{v_1}}^T \otimes \left[\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\nabla_{v_1}(\mathbf{H}_1) \mathbf{f}}{\mathbf{1}_{n_{v_1}}^T \otimes \left[\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right) \\ \vdots \\ \frac{1}{c_3^2} \left(\frac{\nabla_{v_J}(\mathbf{H}_J) \mathbf{f}}{\mathbf{1}_{n_{v_J}}^T \otimes \left[\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\nabla_{v_J}(\mathbf{H}_J) \mathbf{f}}{\mathbf{1}_{n_{v_J}}^T \otimes \left[\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right) \end{bmatrix} \right) \quad (\text{J.4})$$

$$\mathbf{C}_v^T \boldsymbol{\Sigma}^{-1} \mathbf{C}_\alpha = \begin{bmatrix} c_1 \left(\frac{\nabla_{v_1}(\mathbf{H}_1) \mathbf{f}}{\mathbf{1}_{n_{v_1}}^T \otimes \left[\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T (\nabla_\alpha(\mathbf{H}_1) \mathbf{f}) \\ \vdots \\ c_1 \left(\frac{\nabla_{v_J}(\mathbf{H}_J) \mathbf{f}}{\mathbf{1}_{n_{v_J}}^T \otimes \left[\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T (\nabla_\alpha(\mathbf{H}_J) \mathbf{f}) \end{bmatrix} \quad (\text{J.5})$$

$$\mathbf{C}_v^T \boldsymbol{\Sigma}^{-2} \mathbf{C}_\alpha = \begin{bmatrix} \frac{1}{c_3^2} \left(\frac{\nabla_{v_1}(\mathbf{H}_1) \mathbf{f}}{\mathbf{1}_{n_{v_1}}^T \otimes \left[\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\nabla_\alpha(\mathbf{H}_1) \mathbf{f}}{\mathbf{1}_{n_\alpha}^T \otimes \left[\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right) \\ \vdots \\ \frac{1}{c_3^2} \left(\frac{\nabla_{v_J}(\mathbf{H}_J) \mathbf{f}}{\mathbf{1}_{n_{v_J}}^T \otimes \left[\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\nabla_\alpha(\mathbf{H}_J) \mathbf{f}}{\mathbf{1}_{n_\alpha}^T \otimes \left[\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right) \end{bmatrix} \quad (\text{J.6})$$

$$\mathbf{C}_v^T \boldsymbol{\Sigma}^{-1} \mathbf{C}_{c_1} = \begin{bmatrix} (\nabla_{v_1}(\mathbf{H}_1) \mathbf{f})^T \left(\frac{\mathbf{H}_1 \mathbf{f}}{\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \\ \vdots \\ (\nabla_{v_J}(\mathbf{H}_J) \mathbf{f})^T \left(\frac{\mathbf{H}_J \mathbf{f}}{\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \end{bmatrix} \quad (\text{J.7})$$

$$\mathbf{C}_v^T \boldsymbol{\Sigma}^{-2} \mathbf{C}_{c_1} = \begin{bmatrix} \frac{1}{c_1 c_3^2} \left(\frac{\nabla_{v_1}(\mathbf{H}_1) \mathbf{f}}{\mathbf{1}_{n_{v_1}}^T \otimes \left[\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\mathbf{H}_1 \mathbf{f}}{\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \\ \vdots \\ \frac{1}{c_1 c_3^2} \left(\frac{\nabla_{v_J}(\mathbf{H}_J) \mathbf{f}}{\mathbf{1}_{n_{v_J}}^T \otimes \left[\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\mathbf{H}_J \mathbf{f}}{\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \end{bmatrix} \quad (\text{J.8})$$

$$\mathbf{C}_v^T \boldsymbol{\Sigma}^{-1} \mathbf{C}_{c_2} = \begin{bmatrix} (\nabla_{v_1}(\mathbf{H}_1) \mathbf{f})^T \left(\frac{\mathbf{1}_{n_f}}{\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \\ \vdots \\ (\nabla_{v_J}(\mathbf{H}_J) \mathbf{f})^T \left(\frac{\mathbf{1}_{n_f}}{\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \end{bmatrix} \quad (\text{J.9})$$

$$\mathbf{C}_v^T \boldsymbol{\Sigma}^{-2} \mathbf{C}_{c_2} = \begin{bmatrix} \frac{1}{c_1 c_3^2} \left(\frac{\nabla_{v_1}(\mathbf{H}_1) \mathbf{f}}{\mathbf{1}_{n_{v_1}}^T \otimes \left[\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\mathbf{1}_{n_f}}{\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \\ \vdots \\ \frac{1}{c_1 c_3^2} \left(\frac{\nabla_{v_J}(\mathbf{H}_J) \mathbf{f}}{\mathbf{1}_{n_{v_J}}^T \otimes \left[\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\mathbf{1}_{n_f}}{\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \end{bmatrix} \quad (\text{J.10})$$

$$\mathbf{C}_v^T \boldsymbol{\Sigma}^{-1} \mathbf{C}_{c_3} = \begin{bmatrix} \frac{c_1}{c_3} (\nabla_{v_1}(\mathbf{H}_1) \mathbf{f})^T \left(\frac{\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1}}{\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \\ \vdots \\ \frac{c_1}{c_3} (\nabla_{v_J}(\mathbf{H}_J) \mathbf{f})^T \left(\frac{\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1}}{\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \end{bmatrix} \quad (\text{J.11})$$

$$\mathbf{C}_v^T \boldsymbol{\Sigma}^{-2} \mathbf{C}_{c_3} = \begin{bmatrix} \frac{1}{c_3^3} \left(\frac{\nabla_{v_1}(\mathbf{H}_1) \mathbf{f}}{\mathbf{1}_{n_{v_1}}^T \otimes \left[\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1}}{\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \\ \vdots \\ \frac{1}{c_3^3} \left(\frac{\nabla_{v_J}(\mathbf{H}_J) \mathbf{f}}{\mathbf{1}_{n_{v_J}}^T \otimes \left[\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1}}{\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \end{bmatrix} \quad (\text{J.12})$$

$$\mathbf{C}_v^T \boldsymbol{\Sigma}^{-1} \mathbf{C}_{c_4} = \begin{bmatrix} \frac{1}{c_3} (\nabla_{v_1}(\mathbf{H}_1) \mathbf{f})^T \left(\frac{\mathbf{1}_{n_f}}{\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \\ \vdots \\ \frac{1}{c_3} (\nabla_{v_J}(\mathbf{H}_J) \mathbf{f})^T \left(\frac{\mathbf{1}_{n_f}}{\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \end{bmatrix} \quad (\text{J.13})$$

$$\mathbf{C}_v^T \boldsymbol{\Sigma}^{-2} \mathbf{C}_{c_4} = \begin{bmatrix} \frac{1}{c_1 c_3^3} \left(\frac{\nabla_{v_1}(\mathbf{H}_1) \mathbf{f}}{\mathbf{1}_{n_{v_1}}^T \otimes \left[\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\mathbf{1}_{n_f}}{\mathbf{H}_1 \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \\ \vdots \\ \frac{1}{c_1 c_3^3} \left(\frac{\nabla_{v_J}(\mathbf{H}_J) \mathbf{f}}{\mathbf{1}_{n_{v_J}}^T \otimes \left[\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\mathbf{1}_{n_f}}{\mathbf{H}_J \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \end{bmatrix} \quad (\text{J.14})$$

$$\mathbf{C}_\alpha^T \boldsymbol{\Sigma}^{-1} \mathbf{C}_\alpha = c_1 \sum_{j=1}^J (\nabla_\alpha(\mathbf{H}_j) \mathbf{f})^T \left(\frac{\nabla_\alpha(\mathbf{H}_j) \mathbf{f}}{\mathbf{1}_{n_\alpha}^T \otimes \left[\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right) \quad (\text{J.15})$$

$$\mathbf{C}_\alpha^T \boldsymbol{\Sigma}^{-2} \mathbf{C}_\alpha = \frac{1}{c_3^2} \sum_{j=1}^J \left(\frac{\nabla_\alpha(\mathbf{H}_j) \mathbf{f}}{\mathbf{1}_{n_\alpha}^T \otimes \left[\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\nabla_\alpha(\mathbf{H}_j) \mathbf{f}}{\mathbf{1}_{n_\alpha}^T \otimes \left[\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right) \quad (\text{J.16})$$

$$\mathbf{C}_\alpha^T \boldsymbol{\Sigma}^{-1} \mathbf{C}_{c_1} = \sum_{j=1}^J (\nabla_\alpha(\mathbf{H}_j) \mathbf{f})^T \left(\frac{\mathbf{H}_j \mathbf{f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.17})$$

$$\mathbf{C}_\alpha^T \boldsymbol{\Sigma}^{-2} \mathbf{C}_{c_1} = \frac{1}{c_1 c_3^2} \sum_{j=1}^J \left(\frac{\nabla_\alpha(\mathbf{H}_j) \mathbf{f}}{\mathbf{1}_{n_\alpha}^T \otimes \left[\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\mathbf{H}_j \mathbf{f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.18})$$

$$\mathbf{C}_\alpha^T \boldsymbol{\Sigma}^{-1} \mathbf{C}_{c_2} = \sum_{j=1}^J (\nabla_\alpha(\mathbf{H}_j) \mathbf{f})^T \left(\frac{\mathbf{1}_{n_f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.19})$$

$$\mathbf{C}_\alpha^T \boldsymbol{\Sigma}^{-2} \mathbf{C}_{c_2} = \frac{1}{c_1 c_3^2} \sum_{j=1}^J \left(\frac{\nabla_\alpha(\mathbf{H}_j) \mathbf{f}}{\mathbf{1}_{n_\alpha}^T \otimes \left[\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\mathbf{1}_{n_f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.20})$$

$$\mathbf{C}_\alpha^T \boldsymbol{\Sigma}^{-1} \mathbf{C}_{c_3} = \frac{c_1}{c_3} \sum_{j=1}^J (\nabla_\alpha(\mathbf{H}_j) \mathbf{f})^T \left(\frac{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.21})$$

$$\mathbf{C}_\alpha^T \boldsymbol{\Sigma}^{-2} \mathbf{C}_{c_3} = \frac{1}{c_3^3} \sum_{j=1}^J \left(\frac{\nabla_\alpha(\mathbf{H}_j) \mathbf{f}}{\mathbf{1}_{n_\alpha}^T \otimes \left[\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.22})$$

$$\mathbf{C}_\alpha^T \boldsymbol{\Sigma}^{-1} \mathbf{C}_{c_4} = \frac{1}{c_3} \sum_{j=1}^J (\nabla_\alpha(\mathbf{H}_j) \mathbf{f})^T \left(\frac{\mathbf{1}_{n_f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.23})$$

$$\mathbf{C}_\alpha^T \boldsymbol{\Sigma}^{-2} \mathbf{C}_{c_4} = \frac{1}{c_1 c_3^3} \sum_{j=1}^J \left(\frac{\nabla_\alpha(\mathbf{H}_j) \mathbf{f}}{\mathbf{1}_{n_\alpha}^T \otimes \left[\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2 \right]} \right)^T \left(\frac{\mathbf{1}_{n_f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.24})$$

$$\mathbf{C}_{c_1}^T \boldsymbol{\Sigma}^{-1} \mathbf{C}_{c_1} = \frac{1}{c_1} \sum_{j=1}^J (\mathbf{H}_j \mathbf{f})^T \left(\frac{\mathbf{H}_j \mathbf{f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.25})$$

$$\mathbf{C}_{c_1}^T \boldsymbol{\Sigma}^{-2} \mathbf{C}_{c_1} = \frac{1}{c_1^2 c_3^2} \sum_{j=1}^J \left(\frac{\mathbf{H}_j \mathbf{f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right)^T \left(\frac{\mathbf{H}_j \mathbf{f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.26})$$

$$\mathbf{C}_{c_1}^T \boldsymbol{\Sigma}^{-1} \mathbf{C}_{c_2} = \frac{1}{c_1} \sum_{j=1}^J \mathbf{1}_{n_f}^T \left(\frac{\mathbf{H}_j \mathbf{f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.27})$$

$$\mathbf{C}_{c_1}^T \boldsymbol{\Sigma}^{-2} \mathbf{C}_{c_2} = \frac{1}{c_1^2 c_3^2} \sum_{j=1}^J \left(\frac{\mathbf{H}_j \mathbf{f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right)^T \left(\frac{\mathbf{1}_{n_f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.28})$$

$$\mathbf{C}_{c_1}^T \boldsymbol{\Sigma}^{-1} \mathbf{C}_{c_3} = \frac{1}{c_3} \sum_{j=1}^J \left(\mathbf{H}_j \mathbf{f} + \frac{c_1}{c_2} \right)^T \left(\frac{\mathbf{H}_j \mathbf{f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.29})$$

$$\mathbf{C}_{c_1}^T \boldsymbol{\Sigma}^{-2} \mathbf{C}_{c_3} = \frac{1}{c_1 c_3^2} \sum_{j=1}^J \left(\frac{\mathbf{H}_j \mathbf{f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right)^T \left(\frac{\mathbf{H}_j \mathbf{f} + \frac{c_1}{c_2}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.30})$$

$$\mathbf{C}_{c_1}^T \boldsymbol{\Sigma}^{-1} \mathbf{C}_{c_4} = \frac{1}{c_1 c_3} \sum_{j=1}^J \mathbf{1}_{n_f}^T \left(\frac{\mathbf{H}_j \mathbf{f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.31})$$

$$\mathbf{C}_{c_4}^T \boldsymbol{\Sigma}^{-1} \mathbf{C}_{c_4} = \frac{1}{c_1 c_3^2} \sum_{j=1}^J \mathbf{1}_{n_f}^T \left(\frac{\mathbf{1}_{n_f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.43})$$

$$\mathbf{C}_{c_4}^T \boldsymbol{\Sigma}^{-2} \mathbf{C}_{c_4} = \frac{1}{c_1^2 c_3^4} \sum_{j=1}^J \left(\frac{\mathbf{1}_{n_f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right)^T \left(\frac{\mathbf{1}_{n_f}}{\mathbf{H}_j \mathbf{f} + \frac{c_2}{c_1} + \frac{1}{c_1} \left(\frac{\sigma_r}{c_3} \right)^2} \right) \quad (\text{J.44})$$

APPENDIX K

Smoothing Regularizer

This section motivates the use of a novel smoothness/roughness based regularizer for continuous objects, and derives a functional with certain desirable invariances. Consider a finite length vector that maps to its continuous representation through a circulant support assumption, and that its Fourier transform is given by the DFT. These assumptions are consistent with the finite-series object model presented in [63] (Section 4.2.1) with a complex exponential basis and inter-sample values defined through sinc interpolation. This is a very commonly applied assumption, but one must be mindful of how Gibb's phenomenon impacts the continuous representation given its circulant support. The remainder of this section addresses this problem for 1D signals, however, the ideas easily generalize to higher dimensions.

Given a finite-series object model, one natural goal is to minimize the bandwidth of the overall signal subject to the observed data. With the uncertainty principle in mind, a natural regularization functional is given by [173]

$$\Lambda(f(x)) = \int (w - w_0)^2 |F(w)|^2 dw \quad (\text{K.1})$$

where w_0 is a free variable related to where one labels the origin of the frequency axis, and $F(w)$ is defined through the Fourier transform $F(w) = \mathcal{F}(f(x))$. The signals of interest in this thesis are always real, implying $|F(w)|^2$ is a symmetric positive function about zero. It follows from symmetry that the minimizer with respect to w_0 is given by $w_0 = 0$, however, this is also a natural selection of w_0 for many applications. By applying Parseval's theorem we see that this objective function is equivalent to minimizing the integrated squared

derivative over the circle.

$$\Lambda(f) = \int |j\omega F(\omega)|^2 d\omega \quad (\text{K.2})$$

$$= \int \left| \mathcal{F}\left(\frac{\partial f(x)}{\partial x}\right) \right|^2 dx \quad (\text{K.3})$$

$$= \int \left| \frac{\partial f(x)}{\partial x} \right|^2 dx \quad (\text{K.4})$$

From this form it is clear that this functional is a member of the more general class of functions

$$R_m(f) = \int \left(\frac{\partial^m f}{\partial x^m} \right)^2 dx, \quad (\text{K.5})$$

and a classic measure of roughness [88, 152]. Thus a constraint on Equation K.1, typically enforced in the form of a penalty term, enforces a degree of smoothness on the object. The fact that the non-commutable induced differential operators associated with the tangent-space of the Heisenberg group are related by first-order differentiation motivates the use of $m = 1$. In other words, we are penalizing the unobserved member of the two incompatible observables.

Integrating (K.4) by parts, and using the fact that f is defined over the circle to address the boundary conditions, also results in the following illuminating relationship

$$\Lambda(f) = \int \left| \frac{\partial f(x)}{\partial x} \right|^2 dx \quad (\text{K.6})$$

$$= - \int \frac{\partial^2 f(x)}{\partial x^2} f dx \quad (\text{K.7})$$

Specifically, when f is constrained to be positive, the regularizer can be interpreted as a weighted inner-product of the laplacian. Equation (K.7) also provides insight into how the circulant model can be used in conjunction with guard banding to approximate non-circulant objects.

The roughness regularization functional given by (K.1) penalizes high-frequency components, but it still lacks some of the invariance properties that are desirable for inverse imaging problems. Specifically, we would like to find a regularizer Λ with the additional properties

1. $\Lambda(f(R \circ x)) = \Lambda(f(x))$ s.t. $R \stackrel{\text{def}}{=} A$ rigid transformation
2. $\Lambda(af(x) + b) = \Lambda(f(x))$ $\forall a, b \in \mathcal{C}$

3. The ability to regularize with respect to another function

Property 1 requires that the roughness measure be invariant to rigid transformations of the coordinates. For a 2D image this implies that circulantly rotating or translating the object must not alter its degree of smoothness. Given that integration occurs over the entire support, and that the determinant of the Jacobian of a rigid transformation is 1, this property holds for (K.1) by applying multivariate substitution of variables [158]. Property 2 requires that an affine shift of the intensity of the object not alter its smoothness property. Finally, property represents the desire to measure smoothness relative to some reference function. As an example, given data from an imaging device, one may wish to penalize non-smooth deviations from some function other than the mean as implied by (K.4). Properties 2 and 3 clearly do not hold for (K.1), so the goal becomes to alter this function to accommodate these new properties without sacrificing either property 1 or the bandwidth interpretation.

Let \mathbf{F} represent the DFT matrix corresponding to a length N signal, then (K.4) may be represented in the form a standard quadratic penalty acting the discrete representation of the band-limited and circulant signal as follows

$$\Lambda(\mathbf{x}) = \frac{1}{n} \mathbf{x}^H \boldsymbol{\Sigma}_x^{-1} \mathbf{x} \quad (\text{K.8})$$

$$= \frac{1}{n} (\mathbf{F}\mathbf{x})^H \mathbf{F}\boldsymbol{\Sigma}_x^{-1} \mathbf{F}^H (\mathbf{F}\mathbf{x}) \quad (\text{K.9})$$

Equating $\mathbf{F}\boldsymbol{\Sigma}_x^{-1} \mathbf{F}$ with w^2 from (K.2) immediately allows one to solve for the penalty matrix. When \mathbf{x} is a lexicographically ordered representation of a 2D signal of size $n_r \times n_c$, and \mathbf{F} the corresponding 2D DFT matrix, then back-substituting one finds that

$$\boldsymbol{\Sigma}_x^{-1} = \mathbf{F}^H \text{diag} \left[\text{vec} \left((\mathbf{w}_r \odot \mathbf{w}_r) \mathbf{1}_{n_c}^T + \mathbf{1}_{n_r} (\mathbf{w}_c \odot \mathbf{w}_c)^T \right) \right] \mathbf{F}, \quad (\text{K.10})$$

where \mathbf{w}_r and \mathbf{w}_c are the frequency axes associated with the DFT in the row and column dimensions respectively. A simple way of accommodating affine scale invariance is to project the object into an affine invariant space. One such space is the space of analytic circulant functions with zero mean and unit-norm. Redefining Λ over this projection with respect to some zero-mean, unit-norm reference function one gets

$$\Lambda(\mathbf{x}) = \frac{1}{n} \left[\mathbf{F} \left(\frac{\mathbf{x} - \bar{\mathbf{x}}}{\|\mathbf{x} - \bar{\mathbf{x}}\|} - \boldsymbol{\mu}_x \right) \right]^H \mathbf{F}\boldsymbol{\Sigma}_x^{-1} \mathbf{F}^H \left[\mathbf{F} \left(\frac{\mathbf{x} - \bar{\mathbf{x}}}{\|\mathbf{x} - \bar{\mathbf{x}}\|} - \boldsymbol{\mu}_x \right) \right], \quad (\text{K.11})$$

where $\bar{\mathbf{x}} \stackrel{\text{def}}{=} 1/n \mathbf{1}^T \mathbf{x} \mathbf{1}$. Due to the structure of the regularization functional it maintains its invariance to rigid transformation while simultaneously picking up the two additional

properties. Furthermore, for zero mean unit-norm signals, with μ_x taken to be zero, (K.11) is precisely the multi-dimensional extension of the classic roughness measure given by (K.5). In general, μ_x is any zero mean reference image, and one often wants to ensure that μ_x is also unit-norm.

The gradient of (K.11) is relatively straight forward once one has first computed a few of the requisite relations

$$\hat{\mathbf{f}} \stackrel{\text{def}}{=} \frac{\mathbf{f}}{\|\mathbf{f}\|} \quad (\text{K.12})$$

$$\frac{\partial \|\mathbf{f}\|}{\partial \mathbf{f}} = \hat{\mathbf{f}}^T \quad (\text{K.13})$$

$$\frac{\partial \mathbf{f}^T \mathbf{Q} \mathbf{f}}{\partial \mathbf{f}} = \mathbf{f}^T (\mathbf{Q} + \mathbf{Q}^T) \quad (\text{K.14})$$

Applying these relationships to the normalized image one finds

$$\frac{\partial}{\partial \mathbf{f}} \left[\frac{\mathbf{f} - \bar{\mathbf{f}}}{\|\mathbf{f} - \bar{\mathbf{f}}\|} \right] = \frac{\|\mathbf{f} - \bar{\mathbf{f}}\| (\mathbf{I} - \frac{1}{n} \mathbf{1} \mathbf{1}^T) - (\mathbf{f} - \bar{\mathbf{f}}) \frac{(\mathbf{f} - \bar{\mathbf{f}})^T}{\|\mathbf{f} - \bar{\mathbf{f}}\|} (\mathbf{I} - \frac{1}{n} \mathbf{1} \mathbf{1}^T)}{\|\mathbf{f} - \bar{\mathbf{f}}\|^2} \quad (\text{K.15})$$

$$= \frac{1}{\|\mathbf{f} - \bar{\mathbf{f}}\|} (\mathbf{I} - \hat{\mathbf{g}} \hat{\mathbf{g}}^T) \left(\mathbf{I} - \frac{1}{n} \mathbf{1} \mathbf{1}^T \right) \quad (\text{K.16})$$

$$\hat{\mathbf{g}} \stackrel{\text{def}}{=} \frac{\mathbf{f} - \bar{\mathbf{f}}}{\|\mathbf{f} - \bar{\mathbf{f}}\|} \quad (\text{K.17})$$

Using these results in conjunction with the chain rule one finds

$$\frac{\partial}{\partial \mathbf{f}} [R_f(\mathbf{f})] = \left[\frac{2}{n \|\mathbf{f} - \bar{\mathbf{f}}\|} \left(\mathbf{I} - \frac{1}{n} \mathbf{1} \mathbf{1}^T \right) (\mathbf{I} - \hat{\mathbf{g}} \hat{\mathbf{g}}^T) \boldsymbol{\Sigma}_f^{-1} (\hat{\mathbf{g}} - \boldsymbol{\mu}_f) \right]^T \quad (\text{K.18})$$

$$= \left[\frac{2}{n \|\mathbf{f} - \bar{\mathbf{f}}\|} (\mathbf{I} - \hat{\mathbf{g}} \hat{\mathbf{g}}^T) \boldsymbol{\Sigma}_f^{-1} (\hat{\mathbf{g}} - \boldsymbol{\mu}_f) \right]^T, \quad (\text{K.19})$$

where (K.19) follows from (K.18) because $\boldsymbol{\Sigma}_f^{-1}$ projects out the mean, \mathbf{F} is orthonormal, and both $\hat{\mathbf{g}}$ and $\boldsymbol{\mu}_f$ are zero mean by construction.

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] Ernst Abbe. Beiträge zur theorie des mikroskops und der mikroskopischen wahrnehmung. *Archiv für mikroskopische Anatomie*, 9(1):413–418, 1873.
- [2] John Aitchison and S David Silvey. Maximum-likelihood estimation procedures and associated tests of significance. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 154–171, 1960.
- [3] John Aitchison and SD Silvey. Maximum-likelihood estimation of parameters subject to restraints. *The Annals of Mathematical Statistics*, pages 813–828, 1958.
- [4] M. Al-Baali. Descent property and global convergence of the fletcher-reeves method with inexact line search. *IMA J. Numer. Anal. (UK)*, 5(1):121 – 4, 1985. global convergence;Fletcher-Reeves method;inexact line search;descent property;.
- [5] Jean-Paul Allouche and Jeffrey Shallit. *Automatic sequences: theory, applications, generalizations*. Cambridge university press, 2003.
- [6] C. Andrieu and Arnaud Doucet. Simulated annealing for maximum a posteriori parameter estimation of hidden markov models. *Information Theory, IEEE Transactions on*, 46(3):994–1004, 2000.
- [7] Christophe Andrieu, Arnaud Doucet, and William J Fitzgerald. An introduction to monte carlo methods for bayesian data analysis. In *Nonlinear Dynamics and Statistics*, pages 169–217. Springer, 2001.
- [8] Sampo Backman, Anssi J. Maekynen, Timo T. Kolehmainen, and Kai M. Ojala. Fast lens testing using random targets. In Thomas J. Glynn, editor, *Optics and Photonics Technologies and Applications*, volume 4876, pages 1100–1109. SPIE, 2003.
- [9] W. J. Bangs. *Array processing with generalized beamformers*. PhD thesis, Yale, New Haven, CT, 1971.
- [10] Johnathan M Bardsley. Laplace-distributed increments, the laplace prior, and edge-preserving regularization. *Inverse Ill-Posed Problems*, 2012.
- [11] Bryce E. Bayer. Color imaging array. US Patent 3971065 A, 1975.
- [12] Dimitri P Bertsekas. *Nonlinear programming*. Athena Scientific, 1999.

- [13] Christophe Biernacki. Testing for a global maximum of the likelihood. *Journal of Computational and Graphical Statistics*, 14(3):657–674, 2005.
- [14] Doron Blatt and Alfred O Hero. On tests for global maximum of the log-likelihood function. *Information Theory, IEEE Transactions on*, 53(7):2510–2525, 2007.
- [15] I Bongartz, AR Conn, Nick Gould, Ph L Toint, and I Bongartz. Constrained and unconstrained testing environment. *Publications Du Département De Mathématique*, 1993.
- [16] E. Borchini and G. Marsiglia. A note on the uncertainty product of bandlimited functions. *Signal Processing*, 9(4):277 – 282, 1985.
- [17] Glenn D Boreman. *Modulation transfer function in optical and electro-optical systems*, volume 4. SPIE press Bellingham, WA, 2001.
- [18] Max Born and Emil Wolf. *Principles of optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. CUP Archive, 1999.
- [19] M. Born and E. Wolf. *Principles of optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Pergamon Press, 1965.
- [20] Ron Bracewell. *The Fourier Transform and IIS Applications*. McGraw-Hill, 1965.
- [21] Johannes Brauers, Claude Seiler, and Til Aach. Direct PSF estimation using a random noise target. In Francisco Imai, Nitin Sampat, and Feng Xiao, editors, *Digital Photography VI*, volume 7537, 2010.
- [22] Aurélie Bugeau, Marcelo Bertalmío, Vicent Caselles, and Guillermo Sapiro. A comprehensive framework for image inpainting. *Image Processing, IEEE Transactions on*, 19(10):2634–2645, 2010.
- [23] Richard H Byrd, Mary E Hribar, and Jorge Nocedal. An interior point algorithm for large-scale nonlinear programming. *SIAM Journal on Optimization*, 9(4):877–900, 1999.
- [24] Richard H Byrd, Jorge Nocedal, and Richard A Waltz. Knitro: An integrated package for nonlinear optimization. In *Large-scale nonlinear optimization*, pages 35–59. Springer, 2006.
- [25] Jian-Feng Cai, Hui Ji, Chaoqiang Liu, and Zuowei Shen. Framelet-based blind motion deblurring from a single image. *Image Processing, IEEE Transactions on*, 21(2):562 –572, feb. 2012.
- [26] Patrizio Campisi and Karen Egiazarian. *Blind Image Deconvolution: Theory and Applications*. CRC Press, 2016.
- [27] Vance Carpenter and Robert E Hopkins. Comparison of three photoelectric methods of image evaluation. *JOSA*, 46(9):764–767, 1956.

- [28] Patrick PK Chan, Bing-Zhong Jing, Wing WY Ng, and Daniel S Yeung. Depth estimation from a single image using defocus cues. In *Machine Learning and Cybernetics (ICMLC), 2011 International Conference on*, volume 4, pages 1732–1738. IEEE, 2011.
- [29] Tony F Chan and Jianhong Jackie Shen. *Image processing and analysis: variational, PDE, wavelet, and stochastic methods*. Siam, 2005.
- [30] Tony F Chan and Chiu-Kwong Wong. Total variation blind deconvolution. *IEEE transactions on Image Processing*, 7(3):370–375, 1998.
- [31] Subhasis Chaudhuri, AN Rajagopalan, and S Chaudhuri. *Depth from defocus: a real aperture imaging approach*, volume 3. Springer New York, 1999.
- [32] Subhasis Chaudhuri, Rajbabu Velmurugan, and RM Rameshan. *Blind image deconvolution*. Springer, 2016.
- [33] Charles K Chui. An mra approach to surface completion and image inpainting. *Applied and Computational Harmonic Analysis*, 26(2):270–276, 2009.
- [34] Charles K Chui and HN Mhaskar. Mra contextual-recovery extension of smooth functions on manifolds. *Applied and Computational Harmonic Analysis*, 28(1):104–113, 2010.
- [35] Roger N. Clark. Digital cameras: Counting photons, photometry, and quantum efficiency, January 2013.
- [36] Christopher D Claxton and Richard C Staunton. Measurement of the point-spread function of a noisy imaging system. *JOSA A*, 25(1):159–170, 2008.
- [37] T.M. Cover and J.A. Thomas. *Elements of information theory*. John Wiley & Sons, Inc., 1991.
- [38] David R Cox. Tests of separate families of hypotheses. In *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 105–123, 1961.
- [39] David R Cox. Further results on tests of separate families of hypotheses. *Journal of the Royal Statistical Society: Series B (Methodological)*, 24(2):406–424, 1962.
- [40] DR Cox and DV Hinkley. *Theoretical Statistics*. Chapman and Hall, London, 1974.
- [41] Harald Cramér. A contribution to the theory of statistical estimation. *Scandinavian Actuarial Journal*, 1946(1):85–94, 1946.
- [42] Harald Cramér. *Mathematical methods of statistics*. Princeton university press, 1946.
- [43] Antonio Criminisi, Patrick Pérez, and Kentaro Toyama. Region filling and object removal by exemplar-based image inpainting. *IEEE Transactions on Image Processing*, 13(9):1200–1212, 2004.

- [44] Mauricio Delbracio. *Two Problems of Digital Image Formation: Recovering the Camera Point Spread Function and Boosting Stochastic Renderers by Auto-similarity Filtering*. PhD thesis, École Normale Supérieure de Cachan, France Universidad de la República, Uruguay, 2013.
- [45] Mauricio Delbracio and Pablo Musé and Andrés Almansa. Non-parametric sub-pixel local point spread function estimation. *Image Processing Online*, 2:8–21, 2012.
- [46] Mauricio Delbracio, Pablo Musé, Andrés Almansa, and Jean-Michel Morel. The non-parametric sub-pixel local point spread function estimation is a well posed problem. *International Journal of Computer Vision*, 96(2):175–194, 2012.
- [47] John D’Errico. Finite difference approximation to pde inpainting. Inpaint_NaNs available at [http://www.mathworks.com/matlabcentral/fileexchange/...](http://www.mathworks.com/matlabcentral/fileexchange/), 2009.
- [48] H. Dym and H McKean. *Fourier Series and Integrals*. Academic Press, 1972.
- [49] M. Elad and A. Feuer. Restoration of a single superresolution image from several blurred, noisy, and undersampled measured images. *Image Processing, IEEE Transactions on*, 6(12):1646–1658, Dec 1997.
- [50] M. Elad and Y. Hel-Or. A fast super-resolution reconstruction algorithm for pure translational motion and common space-invariant blur. *Image Processing, IEEE Transactions on*, 10(8):1187–1193, 2001.
- [51] Michael Elad, J-L Starck, Philippe Querre, and David L Donoho. Simultaneous cartoon and texture image inpainting using morphological component analysis (mca). *Applied and Computational Harmonic Analysis*, 19(3):340–358, 2005.
- [52] Robert F Engle et al. Wald, likelihood ratio, and lagrange multiplier tests in econometrics. *Handbook of econometrics*, 2:775–826, 1984.
- [53] John Ens. *An investigation of methods for determining depth from focus*. PhD thesis, University of British Columbia, 1990.
- [54] Mohamed-Jalal Fadili, J-L Starck, and Fionn Murtagh. Inpainting and zooming using sparse representations. *The Computer Journal*, 52(1):64–79, 2009.
- [55] S. Farsiu, M. Elad, P. Milanfar, et al. Multiframe demosaicing and super-resolution of color images. *IEEE Transactions on Image Processing*, 15(1):141–159, 2006.
- [56] Sina Farsiu, Michael Elad, and Peyman Milanfar. Constrained, globally optimal, multi-frame motion estimation. In *Statistical Signal Processing, 2005 IEEE/SP 13th Workshop on*, pages 1396–1401. IEEE, 2005.
- [57] Sina Farsiu, Dirk Robinson, Michael Elad, and Peyman Milanfar. Fast and robust multi-frame super-resolution. *IEEE Transactions on Image Processing*, 13(10):1327–1344, 2004.

- [58] Paolo Favaro and Stefano Soatto. *3-D Shape Estimation and Image Restoration*. Springer, 2007.
- [59] Charles Fefferman. Fitting a-smooth function to data, iii. *Annals of mathematics*, pages 427–441, 2009.
- [60] Charles Fefferman and Bo’az Klartag. Fitting a-smooth function to data i. *Annals of mathematics*, pages 315–346, 2009.
- [61] Charles Fefferman, Bo’az Klartag, et al. Fitting a $c^\infty m$ -smooth function to data ii. *Revista Matemática Iberoamericana*, 25(1):49–273, 2009.
- [62] Rob Fergus, Barun Singh, Aaron Hertzmann, Sam T. Roweis, and William T. Freeman. Removing camera shake from a single photograph. *ACM Trans. Graph.*, 25:787–794, July 2006.
- [63] J.A. Fessler. Image reconstruction. Technical notes on image reconstruction with an emphasis on tomography, 2013.
- [64] R. Feynman, R. Leighton, and M. Sands. *The Feynman Lectures on Physics*, volume 3. Addison Wesley, 1965.
- [65] James R Fienup. Phase retrieval algorithms: a personal tour [invited]. *Applied Optics*, 52(1):45–56, 2013.
- [66] James R. Fienup, Brian J. Thelen, Richard G. Paxman, and David A. Carrara. Comparison of phase diversity and curvature wavefront sensing. In *Adaptive Optical System Technologies*, volume 3353, pages 930 – 940. Proceedings of SPIE - The International Society for Optical Engineering, 1998.
- [67] D. A. Fish, J. G. Walker, A. M. Brinicombe, and E. R. Pike. Blind deconvolution by means of the Richardson–Lucy algorithm. *J. Opt. Soc. Am. A*, 12(1):58–65, Jan 1995.
- [68] Ronald Aylmer Fisher. Theory of statistical estimation. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 22, pages 700–725. Cambridge University Press, 1925.
- [69] David A Forsyth and Jean Ponce. *Computer vision: a modern approach*. Prentice Hall Professional Technical Reference, 2002.
- [70] Maurice Fréchet. Sur l’extension de certaines évaluations statistiques au cas de petits échantillons. *Revue de l’Institut International de Statistique*, pages 182–205, 1943.
- [71] Li Gan and Jiming Jiang. A test for global maximum. *Journal of the American Statistical Association*, 94(447):847–854, 1999.
- [72] Josep Garcia, Juan María Sanchez, Xavier Orriols, and Xavier Binefa. Chromatic aberration and depth extraction. In *Pattern Recognition, 2000. Proceedings. 15th International Conference on*, volume 1, pages 762–765. IEEE, 2000.

- [73] Joseph M Geary. *Introduction to wavefront sensors*. SPIE Optical Engineering Press London, 1995.
- [74] Mark Gokstorp. Computing depth from out-of-focus blur using a local frequency representation. In *Pattern Recognition, 1994. Vol. 1-Conference A: Computer Vision & Image Processing., Proceedings of the 12th IAPR International Conference on*, volume 1, pages 153–158. IEEE, 1994.
- [75] R. A. Gonsalves and R. Chidlaw. Wavefront sensing by phase retrieval. In A. Tescher, editor, *Applications of Digital Image Processing III*, volume 207 of *Proc. SPIE*, pages 32–39, San Diego, CA, 1979.
- [76] Rafael. C. González and Richard E Woods. *Digital Image Processing*. Addison Wesley, 1992.
- [77] J. W. Goodman. *Introduction to Fourier Optics*. McGraw-Hill, 1996.
- [78] Steven N. Goodman. Toward evidence-based medical statistics: The p value fallacy. *Annals of Internal Medicine*, 130(12):995–1004, 1999.
- [79] Nicholas I.M. Gould, Stefano Lucidi, Massimo Roma, and Philippe L. Toint. Exploiting negative curvature directions in linesearch methods for unconstrained optimization. *Optimization Methods and Software*, 1(2):75–98, 2000.
- [80] Nicholas IM Gould, Dominique Orban, and Philippe L Toint. Cuter and sifdec: A constrained and unconstrained testing environment, revisited. *ACM Transactions on Mathematical Software (TOMS)*, 29(4):373–394, 2003.
- [81] Nicholas IM Gould, Dominique Orban, and Philippe L Toint. Cutest: a constrained and unconstrained testing environment with safe threads for mathematical optimization. *Computational Optimization and Applications*, 60(3):545–557, 2015.
- [82] Loukas Grafakos. *Classical fourier analysis*, volume 2. Springer, 2014.
- [83] E Grisan, F Frassetto, V Da Deppo, G Naletto, and A Ruggeri. Aberration estimation from single point image in a simulated adaptive optics system. In *Engineering in Medicine and Biology Society, 2005. IEEE-EMBS 2005. 27th Annual International Conference of the*, pages 3173–3176, 2006.
- [84] P. Grossmann. Depth from focus. *Pattern Recognition Letters*, 5(1):63 – 69, 1987.
- [85] GSMA Intelligence. Definitive data and analysis for the mobile industry. Technical report, <https://www.gsmaintelligence.com>, January 2016.
- [86] Frédéric Guichard, Hoang-Phi Nguyen, Régis Tessières, Marine Pyanet, Imène Tarchouna, and Frédéric Cao. Extended depth-of-field using sharpness transport across color channels. In *IS&T/SPIE Electronic Imaging*. International Society for Optics and Photonics, 2009.

- [87] Peter Hall and Peihua Qiu. Nonparametric estimation of a point-spread function in multivariate problems. *Annals of Statistics*, 35(4):1512–1534, 2007.
- [88] Bruce Hansen. Bandwidth selection for nonparametric distribution estimation. Technical report, University of Madison, 2004.
- [89] R.C. Hardie, K.J. Barnard, and E.E. Armstrong. Joint map registration and high-resolution image estimation using a sequence of undersampled images. *Image Processing, IEEE Transactions on*, 6(12):1621–1633, 1997.
- [90] Russell C Hardie and Kenneth J Barnard. Fast super-resolution using an adaptive wiener filter with robustness to local motion. *Optics express*, 20(19):21053–21073, 2012.
- [91] Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2003.
- [92] M.H. Hayes. The reconstruction of a multidimensional sequence from the phase or magnitude of its fourier transform. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 30(2):140–154, Apr 1982.
- [93] E. Hecht. *Optics*. Addison Wesley, 1998.
- [94] A.O. Hero. A Cramér-Rao type lower bound for essentially unbiased parameter estimation. Technical Report 890 DTIC AD-A24666, M.I.T. Lincoln Laboratory, Lexington, MA 02173-0073, Jan. 1992.
- [95] A.O. Hero, J.A. Fessler, and M. Usman. Exploring estimator bias-variance tradeoffs using the uniform CR bound. *Signal Processing, IEEE Transactions on*, 44(8):2026–2041, 1996.
- [96] N.S. Hoang and S.B. Damelin. On surface completion and image inpainting by biharmonic functions: numerical aspects. Technical report, University of Minnesota, 2012.
- [97] B.J. Hoenders. On the solution of the phase retrieval problem. *J. Math. Phys. (USA)*, 16(9):1719 – 25, 1975.
- [98] Ian P Howard. *Perceiving in depth, volume 1: basic mechanisms*. Oxford University Press, 2012.
- [99] Wei Hu, Jianru Xue, and Nanning Zheng. PSF estimation via gradient domain correlation. *IEEE Transactions on Image Processing*, 21(1):386 – 392, 2012.
- [100] I. A. Ibragimov and R. Z. Has’ Minskii. *Statistical estimation: asymptotic theory*, volume 16. Springer Science & Business Media, 1981.
- [101] N. Joshi, R. Szeliski, and D.J. Kriegman. PSF estimation using sharp edge prediction. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8, 2008.

- [102] Seung-Won Jung, Byeong-Doo Choi, and Sung-Jea Ko. Iterative PSF estimation and its application to shift invariant and variant blur reduction. *EURASIP Journal on Advances in Signal Processing*, 2009.
- [103] Kenichi Kanatani. *Statistical Optimization for Geometric Computation*. Dover, 1996.
- [104] Eric Kee, Sylvain Paris, Simon Chen, and Jue Wang. Modeling and removing spatially-varying optical blur. In *Computational Photography (ICCP), 2011 IEEE International Conference on*, pages 1–8. IEEE, 2011.
- [105] R. Keys. Cubic convolution interpolation for digital image processing. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 29(6):1153 – 1160, dec. 1981.
- [106] Jan Kotera, Filip Šroubek, and Peyman Milanfar. Blind deconvolution using alternating maximum a posteriori estimation with heavy-tailed priors. In *Computer Analysis of Images and Patterns*, pages 59–66. Springer, 2013.
- [107] Samuel Kotz, Narayanaswamy Balakrishnan, and Norman L Johnson. *Continuous multivariate distributions, Volume 1: Models and applications*, volume 1. John Wiley & Sons, 2004.
- [108] Henry P Kramer. A generalized sampling theorem. *Journal of Mathematics and Physics*, 38(1-4):68–72, 1959.
- [109] Reginald L Lagendijk, A Murat Tekalp, and Jan Biemond. Maximum likelihood image and blur identification: a unifying approach. *Optical Engineering*, 29(5):422–436, 1990.
- [110] Lucien Le Cam. *Asymptotic methods in statistical decision theory*. Springer Science & Business Media, 2012.
- [111] Joel W. LeBlanc, Brian J. Thelen, and Alfred O. Hero. Joint camera blur and pose estimation from aliased data. *J. Opt. Soc. Am. A*, 35(4):639–651, Apr 2018.
- [112] Joel W. LeBlanc, Brian J. Thelen, and Alfred O. Hero. Testing for local minima of the likelihood using reparameterized embeddings. *Mathematical Imaging and Vision (In Review)*, 2019.
- [113] Hsien-Che Lee. Review of image-blur models in a photographic system using the principles of optics. *Optical Engineering*, 29(5):405–421, 1990.
- [114] E. Lehmann and G. Casella. *Theory of Point Estimation*. Springer, 1998.
- [115] Eli Levy, Doron Peles, Michal Opher-Lipson, and Stephen Lipson. Modulation transfer function of a lens measured with a random target. *Applied Optics*, 38(4):679–683, 1999.

- [116] Lerenhan Li, Jinshan Pan, Wei-Sheng Lai, Changxin Gao, Nong Sang, and Ming-Hsuan Yang. Learning a discriminative prior for blind image deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6616–6625, 2018.
- [117] Xin Li. Image recovery via hybrid sparse representations: A deterministic annealing approach. *IEEE Selected Topics in Signal Processing (JSTSP)*, 5:953–962, 2011.
- [118] Jingyu Lin, Xiangyang Ji, Wenli Xu, and Qionghai Dai. Absolute depth estimation from a single defocused image. *IEEE Trans. Image Processing*, 2013.
- [119] Leon B Lucy. An iterative technique for the rectification of observed distributions. *The astronomical journal*, 79:745, 1974.
- [120] Rastislav Lukac. *Single-Sensor Imaging Methods and Applications for Digital Cameras*. CRC Press, 2008.
- [121] Virendra N. Mahajan. Zernike polynomials and aberration balancing. *Proc. SPIE*, 5173:1–17, 2003.
- [122] Aamir Saeed Malik and Tae-Sun Choi. A novel algorithm for estimation of depth map using image focus for 3d shape recovery in the presence of noise. *Pattern Recognition*, 41(7):2200–2225, 2008.
- [123] André Maréchal. Étude des effets combinés de la diffraction et des aberrations géométriques sur l’image d’un point lumineux. *Rev. Opt. Theor. Instrum.*, 26:257–277, 1947.
- [124] Robert J II Marks. *Advanced topics in Shannon sampling and interpolation theory*. Springer Science & Business Media, 1993.
- [125] G. Martial. Strehl ratio and aberration balancing. *J. Opt. Soc. Am. A, Opt. Image Sci. (USA)*, 8(1):164 – 70, 1991.
- [126] Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953.
- [127] P. Milanfar. *Super-Resolution Imaging*. CRC Press, 2010.
- [128] Willard Jr. Miller. *Symmetry Groups and Their Applications*. Academic Press, 1972.
- [129] Jorge J. More and David J. Thuente. Line search algorithms with guaranteed sufficient decrease. *ACM Trans. Math. Software*, 20:286–307, 1992.
- [130] Jorge J. Moré and David J. Thuente. Line search algorithms with guaranteed sufficient decrease. *ACM Trans. Math. Softw.*, 20:286–307, September 1994.

- [131] Ali Mosleh, Paul Green, Emmanuel Onzon, Isabelle Begin, and JM Pierre Langlois. Camera intrinsic blur kernel estimation: A reliable framework. In *Computer Vision and Pattern Recognition (CVPR), 2015 IEEE Conference on*, pages 4961–4968. IEEE, 2015.
- [132] Vinay P. Namboodiri and Subhasis Chaudhuri. Recovery of relative depth from a single observation using an uncalibrated (real-aperture) camera. In *26th IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2008.
- [133] Kamal Nasrollahi and Thomas B Moeslund. Super-resolution: a comprehensive survey. *Machine vision and applications*, 25(6):1423–1468, 2014.
- [134] Rafael Navarro, Ricardo Rivera, and Justiniano Aporta. Representation of wavefronts in free-form transmission pupils with complex zernike polynomials. *Journal of Optometry*, 4, 2011.
- [135] J. Nocedal and S. Wright. *Numerical Optimization*. Springer, 1999.
- [136] Robert J. Noll. Zernike polynomials and atmospheric turbulence. *J. Opt. Soc. Am.*, 66(3):207–211, Mar 1976.
- [137] Soon-Yong Park. An image-based calibration technique of spatial domain depth-from-defocus. *Pattern Recognition Letters*, 27(12):1318 – 1324, 2006.
- [138] M-A Parseval des Chênes. Mémoire sur les séries et sur l’intégration complète d’une équation aux différences partielle linéaire du second ordre, a coefficients constans. *Sciences, mathématiques et Physiques (Savans étrangers)*, 1:638–648, 1806.
- [139] Vorapoj Patanavijit and Somchai Jitapunkul. An iterative super-resolution reconstruction of image sequences using a bayesian approach with btv prior and affine block-based registration. In *Computer and Robot Vision, 2006. The 3rd Canadian Conference on*, pages 45–45, 2006.
- [140] Vorapoj Patanavijit and Somchai Jitapunkul. A lorentzian stochastic estimation for a robust iterative multiframe super-resolution reconstruction with lorentzian-tikhonov regularization. *EURASIP Journal on Advances in Signal Processing*, 2007(2):21–21, 2007.
- [141] G. Pavlovic and A.M. Tekalp. Maximum likelihood parametric blur identification based on a continuous spatial domain model. *Image Processing, IEEE Transactions on*, 1(4):496 –504, oct 1992.
- [142] R. G. Paxman, T. J. Schulz, and J. R. Fienup. Joint estimation of object and aberrations by using phase diversity. *JOSSA*, 9:1072–1085, 1992.
- [143] Alex Paul Pentland. A new sense for depth of field. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 4:523–531, 1987.

- [144] J.R. Price, T.F. Gee, and K.W. Tobin Jr. Blur estimation in limited-control environments. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, volume 3, pages 1669 – 1672, 2001.
- [145] Simon JD Prince. *Computer vision: models, learning, and inference*. Cambridge University Press, 2012.
- [146] John G Proakis. *Digital signal processing: principles algorithms and applications*. Pearson Education India, 2001.
- [147] John G Proakis, Charles M Rader, Fuyun Ling, and Nikias Chrysostomos L. *Advanced Digital signal processing*. Macmillan Publishing, 1992.
- [148] Matan Protter and Michael Elad. Super resolution with probabilistic motion estimation. *Image Processing, IEEE Transactions on*, 18(8):1899–1904, 2009.
- [149] Ambasamudram N Rajagopalan and Subhasis Chaudhuri. A block shift-variant blur model for recovering depth from defocused images. In *Image Processing, 1995. Proceedings., International Conference on*, volume 3, pages 636–639. IEEE, 1995.
- [150] AN Rajagopalan and Subhasis Chaudhuri. Space-variant approaches to recovery of depth from defocused images. *Computer Vision and Image Understanding*, 68(3):309–329, 1997.
- [151] Alexander G.. Ramm. *Inverse problems: mathematical and analytical techniques with applications to engineering*. Springer, 2005.
- [152] J. O. Ramsay and B. W. Silverman. *Functional Data Analysis*. Springer, 2nd edition, 2005.
- [153] C Radhakrishna Rao. Minimum variance and the estimation of several parameters. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 43, pages 280–283. Cambridge University Press, 1947.
- [154] Radhakrishna C Rao. Large sample tests of statistical hypotheses concerning several parameters with applications to problems of estimation. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 44, pages 50–57. Cambridge Univ Press, 1948.
- [155] Stephen E Reichenbach, Stephen K Park, and Ramkumar Narayanswamy. Characterizing digital image acquisition devices. *Optical Engineering*, 30(2):170–177, 1991.
- [156] William Hadley Richardson. Bayesian-based iterative method of image restoration. *JOSA*, 62(1):55–59, 1972.
- [157] Arturo Rosenblueth and Norbert Wiener. The role of models in science. *Philosophy of science*, 12(4):316–321, 1945.

- [158] Walter Rudin. *Real and complex analysis*. Tata McGraw-Hill Education, 1987.
- [159] Walter Rudin. *Fourier analysis on groups*, volume 12. John Wiley & Sons, 1990.
- [160] Jorge LC Sanz. Mathematical considerations for the problem of fourier transform phase retrieval from magnitude. *SIAM Journal on Applied Mathematics*, 45(4):651–664, 1985.
- [161] Jorge L.C Sanz and Thomas S Huang. Phase reconstruction from magnitude of band-limited multidimensional signals. *Journal of Mathematical Analysis and Applications*, 104(1):302 – 308, 1984.
- [162] Ashutosh Saxena, Sung H Chung, and Andrew Y Ng. 3-d depth reconstruction from a single still image. *International Journal of Computer Vision*, 76(1):53–69, 2008.
- [163] Roland V Shack. Characteristics of an image-forming system. *Research of the National Bureau of Standards*, 56:245, 1956.
- [164] R.R. Shannon and J.C. Wyant. *Applied Optics and Optical Engineering*, volume XI. Academic Press, 1992.
- [165] L Shapiro. *Computer Vision*. Prentice Hall, 2001.
- [166] K.C. Sharman and G.D. McClurkin. Genetic algorithms for maximum likelihood parameter estimation. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 2716 – 19, New York, NY, USA, 1989.
- [167] P. Shearer, A.C. Gilbert, and A.O. Hero. Correcting camera shake by incremental sparse approximation. In *IEEE International Conference on Image Processing*, pages 572 – 6, 2013.
- [168] H. Shen, L. Zhang, B. Huang, and P. Li. A map approach for joint motion estimation, segmentation, and super resolution. *Image Processing, IEEE Transactions on*, 16(2):479–490, 2007.
- [169] Yichang Shih, Brian Guenter, and Neel Joshi. Image enhancement using calibrated lens simulations. Technical report, Microsoft Research and Massachusetts Institute of Technology, April 2012.
- [170] S. D. Silvey. The lagrangian multiplier test. *Ann. Math. Statist.*, 30(2):389–407, 06 1959.
- [171] Samuel David Silvey. *Statistical inference*. CRC Press, 1975.
- [172] Yaduvir Singh. *Semiconductor devices*. Wiley, 2009.
- [173] David Slepian. Some comments on fourier analysis, uncertainty and modeling. *SIAM Review*, 25(3):379–393, 1983.

- [174] Christopher G Small, Jinfang Wang, Zejiang Yang, et al. Eliminating multiple root problems in estimation (with comments by john j. hanfelt, cc heyde and bing li, and a rejoinder by the authors). *Statistical Science*, 15(4):313–341, 2000.
- [175] Jean-Luc Starck, Michael Elad, and David Donoho. Redundant multiscale transforms and their application for morphological component separation. *Advances in Imaging and Electron Physics*, 132:287–348, 2004.
- [176] Stephen M Stigler et al. The epic story of maximum likelihood. *Statistical Science*, 22(4):598–620, 2007.
- [177] Petre Stoica and Randolph L Moses. On biased estimators and the unbiased cramer-rao lower bound. *Signal Processing*, 21(4):349–350, 1990.
- [178] David M. Strong and Stephen C. Cain. Maximum a priori estimation of focus aberration in imaging systems. *Optical Engineering*, 46(12):127006, 2007.
- [179] M. Subbarao. Parallel depth recovery by changing camera parameters. In *Computer Vision., Second International Conference on*, pages 149–155, 1988.
- [180] M. Subbarao and N. Gurumoorthy. Depth recovery from blurred edges. In *Proceedings CVPR '88: The Computer Society Conference on Computer Vision and Pattern Recognition (Cat. No.88CH2605-4)*, pages 498 – 503, 1988.
- [181] M. Subbarao and G. Surya. Depth from defocus: a spatial domain approach. *Int. J. Comput. Vis. (Netherlands)*, 13(3):271 – 94, 1994.
- [182] Huadong Sun, Zhijie Zhao, Xuesong Jin, Lianding Niu, and Lizhi Zhang. Depth from defocus and blur for single image. In *Visual Communications and Image Processing (VCIP)*, pages 1–5. IEEE, 2013.
- [183] Libin Sun, Sunghyun Cho, Jue Wang, and James Hays. Edge-based blur kernel estimation using patch priors. In *Computational Photography (ICCP), 2013 IEEE International Conference on*, pages 1–8. IEEE, 2013.
- [184] Richard Szeliski. *Computer Vision: Algorithms and Applications*. Springer, 2011.
- [185] Hiroyuki Takeda, Peyman Milanfar, Matan Protter, and Michael Elad. Super-resolution without explicit subpixel motion estimation. *Image Processing, IEEE Transactions on*, 18(9):1958–1975, 2009.
- [186] B. Thelen, R. Paxman, D. Carrara, and J. Seldin. Maximum a posteriori estimation of fixed aberrations, dynamic aberrations, and the object from phase-diverse speckle data. *J. Opt. Soc. Am. A*, 16:1016–1025, 1999.
- [187] Brian J. Thelen, John R. Valenzuela, and Joel W. LeBlanc. Theoretical performance assessment and empirical analysis of super-resolution under unknown affine sensor motion. *J. Opt. Soc. Am. A*, 33(4):519–526, Apr 2016.

- [188] Richard Tolimieri, Myoung An, Chao Lu, and C Sidney Burrus. *Mathematics of multidimensional Fourier transform algorithms*, volume 2. Springer-Verlag New York, 1993.
- [189] Pauline Trouvé, Frédéric Champagnat, Guy Le Besnerais, Jacques Sabater, Thierry Avignon, and Jérôme Idier. Passive depth estimation using chromatic aberration and a depth from defocus approach. *Applied optics*, 52(29):7152–7164, 2013.
- [190] Pauline Trouvé-Peloux, Frédéric Champagnat, Guy Le Besnerais, and Jérôme Idier. Theoretical performance model for single image depth from defocus. *JOSA A*, 31(12):2650–2662, 2014.
- [191] UN Population Division. World population prospects. Technical report, United Nations, 2015.
- [192] Harry L. Van Trees. *Detection, Estimation, and Modulation Theory - Van Trees, Harry L.* John Wiley & Sons, 2001.
- [193] Patrick Vandewalle, Luciano Sbaiz, Joos Vandewalle, and Martin Vetterli. How to take advantage of aliasing in bandlimited signals. In *Acoustics, Speech, and Signal Processing, 2004. Proceedings.(ICASSP'04). IEEE International Conference on*, volume 3, pages iii–948. IEEE, 2004.
- [194] Patrick Vandewalle, Sabine Süsstrunk, and Martin Vetterli. A frequency domain approach to registration of aliased images with application to super-resolution. *EURASIP Journal on applied signal processing*, 2006:233–233, 2006.
- [195] Klaus Voss. *Discrete images, objects, and functions in Z^n* , volume 273. Springer, 1993.
- [196] Abraham Wald. Tests of statistical hypotheses concerning several parameters when the number of observations is large. *Transactions of the American Mathematical society*, 54(3):426–482, 1943.
- [197] Abraham Wald. Note on the consistency of the maximum likelihood estimate. *The Annals of Mathematical Statistics*, 20(4):595–601, 1949.
- [198] Richard A Waltz, José Luis Morales, Jorge Nocedal, and Dominique Orban. An interior algorithm for nonlinear optimization that combines line search and trust region steps. *Mathematical programming*, 107(3):391–408, 2006.
- [199] Ruichen Wang, Liangpeng Xu, Chunxiao Fan, and Yong Li. Accurately estimating psf with straight lines detected by hough transform. In *Ninth International Conference on Graphic and Image Processing (ICGIP 2017)*, volume 10615, page 106151N. International Society for Optics and Photonics, 2018.
- [200] Masahiro Watanabe and Shree K Nayar. Rational filters for passive depth from defocus. *International Journal of Computer Vision*, 27(3):203–225, 1998.

- [201] Walter Thompson Welford. *Aberrations of optical systems*. Taylor & Francis, 2017.
- [202] Halbert White. Maximum likelihood estimation of misspecified models. *Econometrica: Journal of the Econometric Society*, pages 1–25, 1982.
- [203] S. S. Wilks. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *Ann. Math. Statist.*, 9(1):60–62, 03 1938.
- [204] Edwin B Wilson and Margaret M Hilferty. The distribution of chi-square. *proceedings of the National Academy of Sciences of the United States of America*, 17(12):684, 1931.
- [205] D. Xie and T. Schlick. A more lenient stopping rule for line search algorithms. *Optimization Methods and Software*, 17:683–700, 2000.
- [206] Yalin Xiong and S.A. Shafer. Moment filters for high precision computation of focus and stereo. In *Proceedings. 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems. Human Robot Interaction and Cooperative Robots (Cat. No.95CB35836)*, volume 3, pages 108–113, 1995.
- [207] Li Xu and Jiaya Jia. Two-phase kernel estimation for robust motion deblurring. In *Computer Vision–ECCV 2010*, pages 157–170. Springer, 2010.
- [208] Paul B Yale. *Geometry and symmetry*. Courier Corporation, 1968.
- [209] Junlan Yang and Dan Schonfeld. Virtual focus and depth estimation from defocused video sequences. *Image Processing, IEEE Transactions on*, 19(3):668–679, 2010.
- [210] J Yen. On nonuniform sampling of bandwidth-limited signals. *IRE Transactions on circuit theory*, 3(4):251–257, 1956.
- [211] Y. Yitzhaky and N.S. Kopeika. Identification of blur parameters from motion blurred images. *Graphical Models and Image Processing*, 59(5):310 – 320, 1997.
- [212] Daniel F. Yu and Jeffrey A Fessler. Edge-preserving tomographic reconstruction with nonlocal regularization. *Medical Imaging, IEEE Transactions on*, 21(2):159–173, 2002.
- [213] J. A. Zandhuis, D. Pycock, S. F. Quigley, and P.W. Webb. Sub-pixel non-parametric PSF estimation for image enhancement. *Vision, Image and Signal Processing, IEEE Proceedings -*, 144(5):285–292, 1997.
- [214] von F Zernike. Beugungstheorie des schneidenverfahrens und seiner verbesserten form, der phasenkontrastmethode. *Physica*, 1(7):689–704, 1934.
- [215] Shaojie Zhuo and Terence Sim. Defocus map estimation from a single image. *Pattern Recognition*, 44(9):1852–1858, 2011.
- [216] Djemel Ziou and Francois Deschênes. Depth from defocus estimation in spatial domain. *Computer Vision and Image Understanding*, 81(2):143–165, 2001.