UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title
Quantitative Microscopy using Coded Illumination

Permalink
https://escholarship.org/uc/item/70d9j190

Author
Phillips, Zachary Fitzgerald

Publication Date
2019

Peer reviewed|Thesis/dissertation

Quantitative Microscopy using Coded Illumination

by

Zachary Fitzgerald Phillips

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Applied Science and Technology

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Associate Professor Laura Waller, Chair
Professor Lydia Sohn
Assistant Professor Ren Ng

Spring 2019

Quantitative Microscopy using Coded Illumination

Abstract

Quantitative Optical Microscopy Using Coded Illumination

by

Zachary Fitzgerald Phillips

Doctor of Philosophy in Applied Science and Technology

University of California, Berkeley

Associate Professor Laura Waller, Chair

Quantitative optical microscopy continues to be a powerful tool for biomedical research and the sciences at-large. Building upon centuries of optical theory, computational microscopy leverages large-scale numerical computation to dramatically extend and improve the capabilities of existing optical microscopes through high-speed capture, super-resolution, or opening new application spaces such as quantitative phase imaging. This dissertation describes the theory and reduction to practice of several novel computational microscopy techniques which make use of jointly-designed coded illumination hardware and physics-based reconstruction algorithms. We first introduce label-free quantitative phase imaging using differential phase contrast, and demonstrate a novel single-shot variant using color-multiplexing. Second, we explore coded illumination for high-throughput imaging, and demonstrate a temporal-coding technique which enables significantly higher SNR for high-speed slide scanning and neuropathology applications. For fluorescent imaging, this method can provide up to a $10\times$ improvement in reconstruction SNR compared to conventional high-speed imaging techniques. Next, we explore the design and fabrication of LED illumination devices, including a quasi-dome LED illuminator which enables high-angle illumination for a variety of applications. To address practical calibration concerns for computational microscopy systems, we demonstrate two examples of algorithmic self-calibration. These include aberration recovery using differential phase contrast as well as source calibration for the quasi-dome LED array, using both image-based calibration and an online method based on Fourier ptychography. These works demonstrate the benefits and practical challenges related to computational optical microscopy with coded illumination.

For Mandie

Contents

Acknowledgments

First and foremost, I would like to thank my advisor Laura Waller for providing careful guidance throughout my time at Berkeley, and for putting faith and resources behind myself and fellow lab members so that we can freely explore, create, and publish in the field of computational imaging. I would also like to thank my committee, Assistant Professor Ren Ng and Professor Lydia Sohn, for many useful conversations throughout my time at UC Berkeley. The Waller Lab has been a fantastic place to learn about the field of computational imaging, and I am grateful to each of my fellow group members. I would like to specifically thank my co-authors and collaborators within the Waller lab, including Regina Eckert, Gautam Gunjala, Michael Kellman, Nick Antipa, and David Ren. Specifically, I would like to thank Michael Chen for being an excellent collaborator in both research and entrepreneurship, and my closest neighbor, Li-Hao Yeh, with whom I have enjoyed thousands of conversations during our nearly 5 years in 558 Cory Hall. I have genuinely enjoyed my many whiteboard discussions with Sarah Dean, and feel very fortunate to have worked extensively with Postdocs Lei Tian, Volker Jaedicke, Shwetadwip Chowdhury, and Emrah Bostan. I would also like to acknowledge those who took the time to proofread this dissertation through its many iterations who have not already been mentioned: Kristina Monokhova, Michael Kellman, Henry Pinkard, and Emma Alexander.

During my time at Berkeley, I found that being productive and focused during the week requires taking time to unwind during the weekend (and occasionally during the evening). I would first like to thank my girlfriend, Chelsea, for always being a steadfast adventure partner and sounding board through the many seasons of graduate school. Since my first days in Berkeley, I've enjoyed many waves and adventures with my close friend and fellow AS&T student Brandon Wood. I also want to thank my friends Chris Lalau-Keraly, Alex Forse, Kyle Ginthner, JB Chapman, (Day-leader) Mike Chapman, Jon Morris, Alix Charles, Mike Francis, Dan Wooten, and many others for our many good times.

Attending Berkeley has been my proudest and must fulfilling accomplishment, and it would have never been possible without my Mom, Dad, and sisters Allie and Mandie, who have always been supporting through our shared challenges as well as my own. I credit my dad, together with my grandfather, for teaching me how to break down a complicated problem into a series of small problems, and the importance of tenacity. My mother taught me the importance of a sense of humor and being personable. I would also like to acknowledge all of my grandparents, especially Phil and Bertha Phillips, who continue to teach importance of family, travel, and spending time outdoors.

I also want to thank those who helped and encouraged me prior to coming to Berkeley, including Professor Amy Oldenburg, Daniel Marks, and Raghav Chhetri.

Finally, I want to thank the many undergraduate students I've worked with, including Patrick Oare, Joel Whang, Nitin Sadras, Aditya Gande, Gautam Gunjala, Sachin Deyoung, Taehyung Kim, Paroma Varma, Jared Rulison, Sudarshan Seshadri, and Kuan Chang for their generous assistance despite their many obligations.

Chapter 1

Introduction

1.1 Optical Microscopy

The optical microscope is one of the oldest scientific instruments, and continues to be an essential tool for researchers, medical professionals, and engineers across many disciplines. Microscopes are typically defined as having two or more refractive surfaces to provide magnification between the object of interest and the imaging plane, enabling the user to see things much smaller than the resolving power of the human eye. Credit for the invention of the compound microscope is generally attributed to Hans and Zacharias Jansen [1], although the first published work on microscope design wasn't released until 1665 (Hooke and van Leeuwenhoek) [2,3]. The term "microscope" is generally used to describe optical microscopes - those which are designed for use with light within the optical band of the electro-magnetic spectrum ($390nm \leq \lambda \leq 700nm$), which is approximately the electromagnetic spectrum detectable by the human eye.

Imaging and Resolution

Light interacts with our world in many ways, including diffraction, refraction, reflection, and absorption. At optical wavelengths, many common materials (such as glass) have favorable properties for refractive optics (providing significant phase delay with little absorption), facilitating precise control of an optical wavefront using these elements. Imaging is the process of creating a copy of a particular optical signal at a different position in space, typically with the goal of making a measurement using a film or electronic detector. In the simplest case, a single lens may be used to form a magnified image of an optical signal by placing the lens and detector at a particular distance from the sample A single-lens imaging system has many practical issues, however, including lacking telecentricity (consistency of magnification across the field) and aberrations, both geometric and chromatic. Including multiple optical elements into a compound microscope can dramatically improve image quality by providing aberration compensation and enabling telecentricity. Typically, the exact number and design of these components is abstracted to the end-user and can be defined by a relatively low

number of descriptive quantities. Magnification and numerical aperture (NA) are the most important of these; the magnification of an objective sets the field of view, while the NA sets a minimum bound on the diffraction-limited resolution. The numerical aperture is defined by the formula $NA = n\sin(\theta)$, where n is the refractive index of the medium, and θ is the maximum half-angle at which light may pass through the objective relative to the radial (optical) axis. The angular dependence of numerical aperture is completely described by interference effects which arise from the wave-optics model of light propagation. As multiple off-axis sources of the same wavelength converge to a point, the wavefronts of these sources will cause constructive and destructive interference. The minimum distance between two peaks formed by constructive interference is proportional to both the wavelength of the illumination and the angular separation between the two beams (which is set by the maximum NA of the illumination source and imaging optics). Practically, the size of this spot defines the resolution of the optical system. By the Rayleigh criteria, the resolution of an optical system x_{min} is defined as:

$$\Delta x_{min} = \frac{1.22\lambda}{(NA_{objective} + NA_{illumination})} \tag{1.1}$$

This quantity defines the minimum separation between two points which can be detected by a system with a circular aperture and is defined by the distance between the center of the point spread function (PSF) and its first null. Note that Eq. 1.1 is dependent on both detection side NA ($NA_{objective}$) and illumination side NA ($NA_{illumination}$). This equation holds true while $NA_{illumination} \leq NA_{objective}$. High-angle illumination ($NA_{illumination} > NA_{objective}$) provides no direct improvement to the resolution of the measurement, as the signal from these illumination angles is not able to interfere with the background (DC) term. This sets a lower limit on the resolution of a brightfield microscope. When a sample is illuminated from high angles only ($NA_{illumination} > NA_{objective}$), high-resolution features are revealed as diffraction-limited edge contrast but are not individually resolved with more than $2NA_{objective}$ resolution. This is the working principle of darkfield microscopy.

1.2 Fourier Optics

Fourier Optics provides an important bridge between Fourier theory and optical systems by modeling common optical phenomena (such as propagation and lensing) using the Fourier transform. The seminal text on imaging using Fourier theory to analyze imaging systems was published in 1968 [4] which presents the framework underpinning many common computational techniques such as deconvolution, holography, and free-space propagation. The Fourier optics description is especially useful for an optical system configured as a telecentric ($4f$) system, which is a typical model for an optical microscope:
In this system, the lenses in this system operate as forward and inverse Fourier transform operators on the input optical field at (P1) as it propagates through the pupil plane (P2) to the image plane (P3). The electric field at position P2 can be modeled as the Fourier

Figure 1.1: Schematic of a $4f$ optical system

transform of P1, which is often occupied by an aperture stop (circular low-pass filter) to limit the NA of the objective. This stop defines the resolution of the optical system and can be used to reduce aberrations and prevent aliasing of the optical signal at the camera plane. Most modern infinity-corrected microscopes may be approximated to a high degree using these systems, assuming minimal aberrations and misalignment.

1.3 Computational Microscopy

The concept of using computational tools to simulate and invert optical imaging systems was developed soon after the emergence of large-scale computation in the 1960s; the theory of image formation and the propagation of light was developed prior to this time [5, 6]. Recently, the field of computational imaging has expanded considerably due to increasing availability of computing power and digital sensing hardware. In modern microscopes, digital cameras allow the detection of the intensity of an optical field using a grid of photodetectors, which digitize the optical field and facilitate computational imaging reconstructions on a host computer. As graphical processing units (GPUs) have become faster and more widely available, computational algorithms have likewise accelerated both in speed and scale.

An early example of computational imaging was the application of a cubic phase plate at the microscope pupil, which provides significantly increased depth of field but produces a highly distorted image. With knowledge of these distortions the original image with extended depth of field can be deconvolved using knowledge of the system's point spread function (PSF) [7]. Since these early works, the field of computational microscopy has expanded considerably due to the widespread availability of computing hardware and software tools for simulating optical systems and performing quantitative analysis. Prominent examples include super-resolution methods such as structured illumination [8, 9], which enhances resolution by projecting a pattern onto the sample, localization microscopy [10, 11], which employs statistical analysis to localize sparse fluorophores using temporal dynamics, and both conventional [12] and Fourier [13] ptychography. Three-dimensional imaging has likewise become a powerful tool for imaging 3D biological quantities, and becomes abso-

lutely necessary for high-NA imaging of thick samples which encounter multiple scattering. Various approaches have been proposed, including deconvolving focal stacks [14], light-field microscopy [15, 16], PSF engineering [17], and diffraction tomography [18–20]. In addition, computational imaging has been widely used for quantitative phase imaging, using interferometry [19, 21, 22], off-axis holography [23], or commercial add-ons [24, 25]. Another add-on option uses two cameras to capture defocused images which can then be used to solve the Transport of Intensity Equation (TIE) [26]. Alternatively, if chromatic aberrations are large enough, they can enable single-shot color TIE [27] without any hardware changes.

In most cases the propagation and refraction of light can be modeled using a small number of linear and non-linear operations and may be accurately simulated using linear algebra software packages such as `numpy` or `MATLAB`. With knowledge of this forward model, an inverse problem may be formed to recover the object without any distortions imposed by the imaging system (provided the information is still present in the measurement). In computational imaging these distortions are carefully designed to reveal contrast in ways a conventional system cannot, enabling the recovery of high-dimensional or high-resolution information using computation after an acquisition in performed. The system operator $A\{\cdot\}$ describes and models this entire acquisition process, including a model for the propagation of light, the design of the system, and any calibration or mis-alignment which may be present (Fig.1.3). To facilitate the recovery of an unknown object from measurements made under this system y, an inverse problem formulation is used to invert the system operator $A\{\cdot\}$, often having the common standard form:

$$\hat{\boldsymbol{x}} = \operatorname*{argmin}_{\boldsymbol{x}} \quad ||\mathcal{A}\{\boldsymbol{x}\} - \boldsymbol{y}||_2^2 \tag{1.2}$$

where $\boldsymbol{x}$ represents the variable of interest (generally the object), $\mathcal{A}\{\cdot\}$ is the mathematical operation describing the optical system, $\boldsymbol{y}$ is the measured intensity, and $\hat{\boldsymbol{x}}$ is an estimate of the object $\boldsymbol{x}$. The forward operator $\mathcal{A}\{\cdot\}$ is normally formed based on the physics of the optical systems, need not be linear or represented by a matrix.

The goal of a computational imaging system is to invert the forward operator $\mathcal{A}\{\cdot\}$ in a way which minimizes the distance between the object estimate distortions of the forward-inversion process ($||\hat{\boldsymbol{x}} - \boldsymbol{x}||_2^2$). If $\mathcal{A}\{\cdot\}$ is linear, is can be inverted in a closed-form solution using the Moore-Penrose Pseudo-Inverse [28], or with an iterative method such as gradient descent. If $\mathcal{A}\{\cdot\}$ is non-linear but smooth, it must be inverted iteratively using analytic expressions for each regularization. If $\mathcal{A}\{\cdot\}$ is non-smooth, it can, in some cases, still be inverted using iterative soft-thresholding methods such as FISTA [29].

In general, linear problems (characterized by satisfying the relationship $\mathcal{A}\{\boldsymbol{a} + \boldsymbol{b}\} = \mathcal{A}\{\boldsymbol{a}\} + \mathcal{A}\{\boldsymbol{b}\}$) are much easier to invert and solve, having lower memory requirements and complexity as well as a direct inverse. Linear convolution operators are particularly common for telecentric imaging systems. When a convolution is well-posed, it may be efficiently inverted using a FFT-based deconvolution algorithm [30], which has complexity $N \log(N)$ as opposed to N^2 for normal operators.
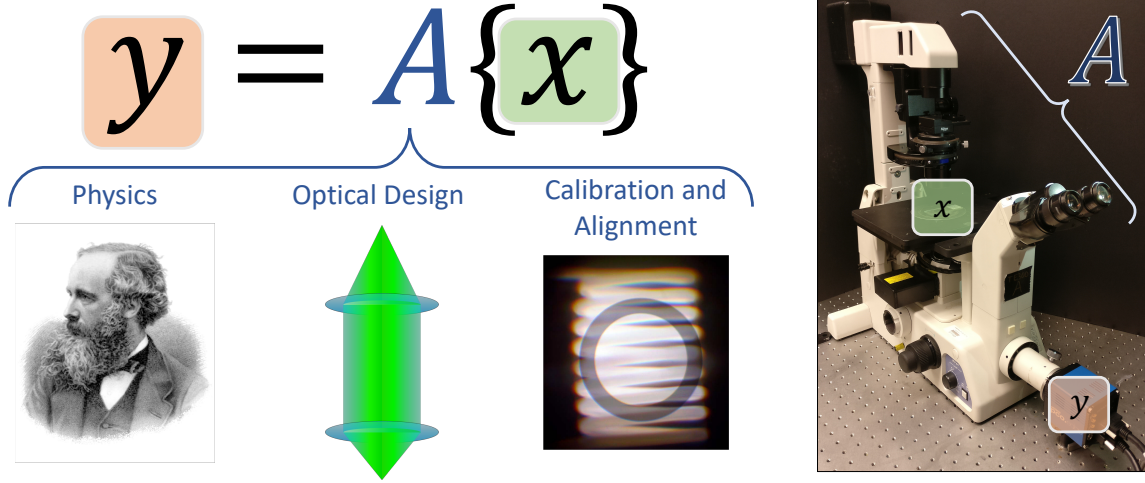
Figure 1.2: Overview of Computational Imaging. The forward model $\mathcal{A}$ is a function of the physical properties of light, the optical system design, and (mis)calibration of the system. An image of a Nikon TE300 microscope used in this work is provided for context.

The performance of inversion processes may be improved by adding regularization term to penalize certain undesirable characteristics of the signal, such as noise. The most commonly used regularization method is Tikhonov (or ℓ_2) regularization [31] which enforces a prior on the total energy of a system. Tikhonov regularization is equivalent to adding an additional ℓ_2 term to Eq. 1.2:

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x}}{\operatorname{argmin}} \quad ||A\{\boldsymbol{x}\} - \boldsymbol{y}||_2^2 + \alpha||\boldsymbol{x}||_2^2 \tag{1.3}$$

where α is a tuning parameter which represents the weight of the Tikhonov prior (normally set to $\frac{1}{SNR}$). If $\mathcal{A}\{\cdot\}$ is linear and can be represented as a matrix, Eq. 1.3 can be directly inverted using a closed-form expression:

$$\hat{\boldsymbol{x}} = ((\mathbf{A}^H\mathbf{A})^{-1} + \alpha\mathbf{I})\mathbf{A}^H\boldsymbol{y} \tag{1.4}$$

where $\mathbf{A}$ is the matrix form of $\mathbf{A}\{\cdot\}$ and $\mathbf{I}$ is the identity matrix with the same dimensions as $\mathbf{A}^H\mathbf{A}$. This closed-form solution makes Tikhonov regularization popular for many inverse problems, although the total energy prior may not be accurate in all cases.

A second common class of priors enforce sparsity of the object in some domain. Mathematically the ℓ_0 "norm" returns the number of non-zero elements of the input. This norm is non-convex, however, requiring a large combinatorial search which is intractable for most problems [32]. As a proxy, the ℓ_1 norm is conventionally employed as a convex, though non-smooth alternative [33]. When coupled with a generalized sparsifying operator $\mathcal{W}\{\cdot\}$ and a differential forward model $\mathcal{A}\{\cdot\}$, this problem is convex, and can be written as:

$$\hat{\boldsymbol{x}} = \underset{\boldsymbol{x}}{\operatorname{argmin}} \quad ||\mathcal{A}\{x\} - \boldsymbol{y}||_2^2 + \alpha||\mathcal{W}\{\boldsymbol{x}\}||_1. \tag{1.5}$$

Because the regularization term is non-smooth, iterative solvers must be used to recover the optimal $\hat{\boldsymbol{x}}$. When $\mathcal{W}$ is a unitary function $\mathbf{w}$ (or the identity matrix), a solver implementing proximal gradient descent using soft-thresholding may be used to minimize this objective function, such as FISTA [29], ADMM [34], or TwIST [35]. In general, $\mathbf{W}$ can be any unitary transform, including the Fourier transform or Wavelet Transform, or a learned unitary operator which is optimized using a machine-learning framework [36]. In all of these cases, the optimal $\hat{\boldsymbol{x}}$ may be recovered by performing many iterations of proximal gradient descent. In the case where $\mathbf{W}$ is not unitary, the above relationship does not hold, and other proximal methods must be used. One prominent example is total-variation regularization (TV), which enforces sparsity of the image gradients [37]. TV regularization can be implemented using ADMM [38], FISTA [29], or using soft-thresholding on wavelet coefficients [39].

1.4 Noise in Computational Microscopy Systems

All measurements contain noise from various sources, including photon quantization or camera electronics. In general, these noise sources can be additive or multiplicative, and may take on a variety of statistical profiles including Gaussian and Poisson distributions. Analyzing the propagation of noise through computational imaging systems is extremely important for practical implementation of the algorithms presented in this work, as they help motivate when computational imaging makes sense. Analyses previously introduced in computational photography [40] have been adopted for microscopy applications here. In this section, we provide a derivation which applies to all linear systems presented in this dissertation, particularly Chapter 2 and Chapter 3. Here, we generally assume the presence of an additive, Gaussian noise term $\boldsymbol{\eta}$ with zero-mean, and variance $\sigma_{\boldsymbol{\eta}}$, which is added to each measurement made under a general forward operator $\mathcal{A}\{\cdot\}$:

$$\boldsymbol{y} = \mathcal{A}\{\boldsymbol{x}\} + \boldsymbol{\eta} \tag{1.6}$$

This approximation is normally valid for measurements made with more than 10 photon counts, which includes every case presented in this dissertation (including fluorescence imaging). The effect of this noise on image quality is generally represented as the signal-to-noise ratio (SNR). Here, we use the common imaging SNR definition:

$$SNR = \frac{\bar{\boldsymbol{y}}}{\sigma_{\boldsymbol{\eta}}} \tag{1.7}$$

where $\sigma_{\boldsymbol{\eta}}$ is the standard deviation of the noise term and $\bar{y}$ is the mean signal (DC term) of the measurement y. When inverting the forward operator $\mathcal{A}\{\}$ to recover $\boldsymbol{x}$, the presence of $\boldsymbol{\eta}$ will lead to error in the measurements compared to the ground truth. Take, for example,

if $\mathcal{A}$ can be represented as a matrix $\mathbf{A}$, the Moore-Penrose pseudoinverse of the object is given by:

$$\hat{\boldsymbol{x}} = \boldsymbol{x} + (\mathbf{A}^H \mathbf{A})^{-1} \mathbf{A}^H \boldsymbol{\eta}. \tag{1.8}$$

Clearly, the recovered object $\hat{\boldsymbol{x}}$ will be corrupted by an additive term, which is essentially the inversion process applied to the noise term $\hat{\boldsymbol{\eta}}$. Based on Eq. 1.8, the root-mean-squared error (RMSE) between $\hat{\boldsymbol{x}}$ and the true $\boldsymbol{x}$ is $(\mathbf{A}^H \mathbf{A})^{-1} \mathbf{A}^H \boldsymbol{\eta}$. Taking the covariance of this term, we can find an expression for the covariance of the error term in the reconstruction:

$$\boldsymbol{\Sigma}_{\mathbf{A}^\dagger \boldsymbol{\eta}} = \sigma_{\boldsymbol{\eta}}^2 (\mathbf{A}^H \mathbf{A})^{-H} \tag{1.9}$$

The main result of Eq. 1.9 is that the inversion process re-weights the spectrum of the original Gaussian white noise $\boldsymbol{\eta}$. When $\mathbf{A}$ has an ℓ_2 operator norm of 1, the minimum singular value of $\mathbf{A}$ defines the maximum noise amplification, while the sum of the inverse singular values defines the total RMSE:

$$\sigma_{\mathbf{A}^\dagger \boldsymbol{\eta}} = \sigma_{\boldsymbol{\eta}} \sqrt{\boldsymbol{\Sigma}_{i=0}^N \frac{1}{\sigma_i^2 \{\mathbf{A}\}}}, \tag{1.10}$$

where N is the length of $\boldsymbol{x}$ and $\sigma_i^2\{\mathbf{A}\}$ represents the i^{th} singular value of $\mathbf{A}$. This definition is consistent with [41]. This relationship between the singular values of A and the amplification of $\boldsymbol{\eta}$ enables a closed-form relationship to the reconstruction signal-to-noise ratio (SNR_{recon}) of a measurement $\boldsymbol{y}$, defined as:

$$SNR_{recon} = \frac{\bar{\boldsymbol{y}}}{\sigma_{\boldsymbol{\eta}} \sqrt{\boldsymbol{\Sigma}_{i=0}^N \frac{1}{\sigma_i^2 \{\mathbf{A}\}}}}, \tag{1.11}$$

From this relationship, it becomes clear that the deconvolution process will reduce the SNR by a factor:

$$f = \sqrt{\boldsymbol{\Sigma}_{i=0}^N \frac{1}{\sigma_i^2 \{\mathbf{A}\}}}, \tag{1.12}$$

where f is the deconvolution noise factor (DNF).

The main result here is that the deconvolution process will always amplify measurement noise, and that for linear systems this amplification can be computed algebraically so long as the singular values of the forward operator are known. This result is particularly useful for convolutional forward operators, where the singular values of A can be computed quickly and efficiently from the Fourier coefficients of the convolution kernel h due to the circulant structure of A. This analysis of RMSE amplification is equivalent to A-optimal design [42], which is used here due to strict compatibility with the definition of imaging SNR (Eq. 1.7).

As a demonstration of this relationship, we simulated a convolutional forward model with additive Gaussian noise, and perform deconvolution of the noisy measurement. These
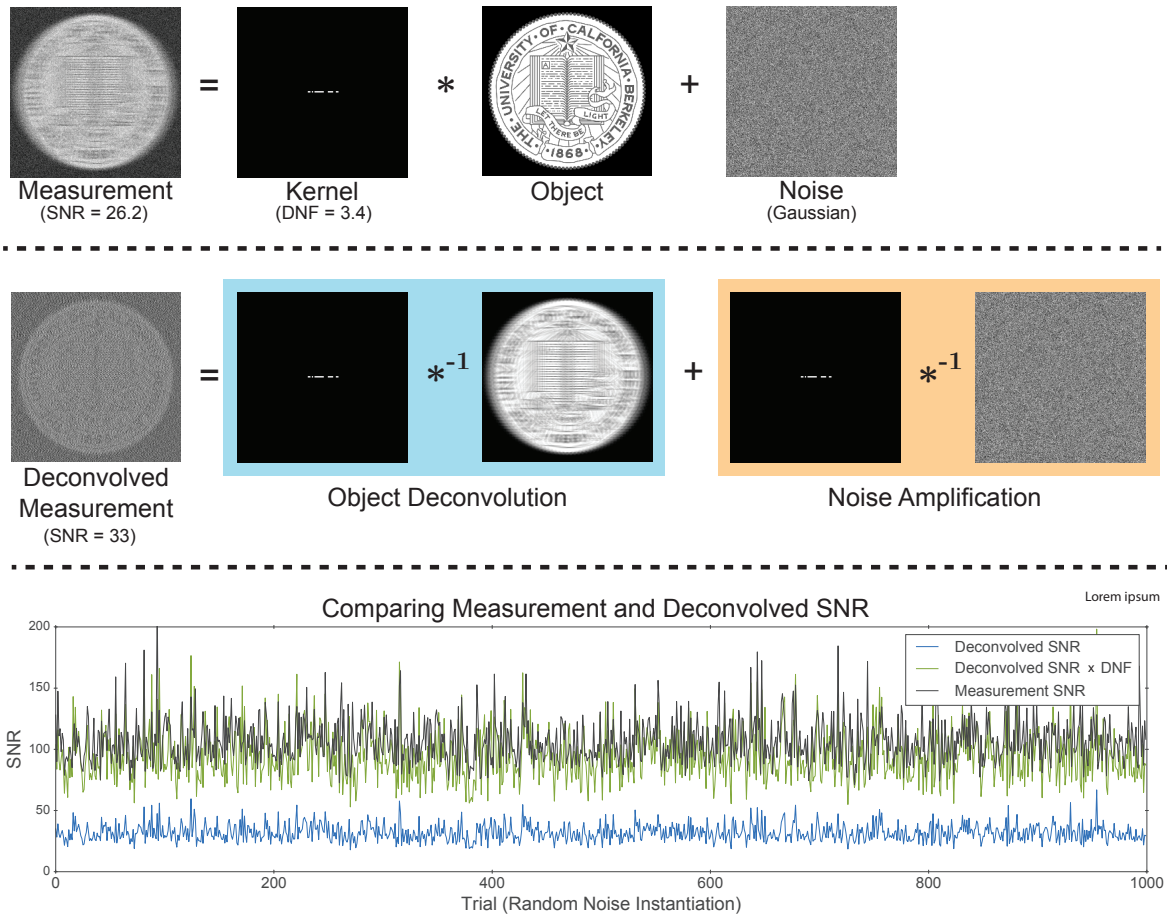
Figure 1.3: Simulation of a convolutional forward model and verification of DNF calculations. Top row shows forward modal with additive Gaussian noise, while the middle row shows the deconvolution of the same measurement, separated into object deconvolution and noise amplification terms. Bottom row shows the measurement and reconstruction SNR across 1000 random generations of the Gaussian white noise term, overlaid with the multiplication of the reconstruction SNR multiplied by the DNF as a verification. The operators $*$ and $*^{-1}$ represent 2D convolution and deconvolution, respectively.

results show that the DNF provides a very close scalar relationship between measurement SNR and deconvolved SNR, which is illustrated by comparing the estimated measurement SNR (deconvolved SNR multiplied by the DNF) to the original measurement SNR. Small discrepancies between the expected and predicted SNR calculations are likely due to sampling error, since we evaluate the noise standard deviation across a 20-pixel square in the top left corner of each image, rather than the full image (to avoid including the standard deviation of the object in our SNR calculation). These relationships are used in both Chapter 2 and

Chapter 3 for analyzing the noise propagation of linear forward models.

1.5   Coded Illumination for Optical Microscopy

Since the 17^{th} century, microscopes have employed some sort of light source to illuminate a specimen, from candles to semiconductor light sources. Illumination affects image contrast in variety of ways - for purely absorptive samples, brightfield microscopy images the light that is attenuated by the sample by illuminating within the range of angles defined by the numerical aperture (NA) of the objective. Conversely, darkfield microscopy uses illumination from angles outside of the illumination NA, imaging only light that is scattered or "bent" by a refractive medium such as water. Phase contrast microscopy [43] uses annular illumination to reveal qualitative phase contrast, while differential interference contrast (DIC) [44] uses coherent, polarized illumination to reveal the gradient of an object's phase in a single direction. Recently, programmable light sources such as LED arrays [45, 46] have enabled a variety of qualitative contrast modalities at high speed [45, 47], as well as the capability to perform quantitative reconstructions of the complete complex field of the sample [13, 48, 49], enabling the measurement of the dry mass of many aqueous biological samples [50, 51].

In the following chapters, I will describe several novel applications and of coded illumination in optical microscopy, fabrication methods for coded illumination devices, and self-calibration techniques for the aforementioned methods, Fig. 1.5 shows the system which was used in most experiments presented in this dissertation.

Chapter 2 describes methods for qualitative and quantitative phase recovery, including a single-shot quantitative phase imaging method which uses partially coherent color-multiplexed illumination to recover the complete optical field of an object from a single measurement. Chapter 3 describes a novel method for recovering a large field-of view with high SNR by using a coded illumination sequence to introduce motion blur during each capture, which is then computationally removed using a motion deblurring algorithm to recover the static object. We demonstrate the performance of this technique for both brightfield imaging and fluorescence imaging and provide an analysis of optimal acquisition strategy in terms of common system parameters such as camera noise level and illumination power. In Chapter 4, we describe the fabrication of several prototypes used for coded illumination, including a programmable domed LED array, LED sources for high-throughput imaging, and Computational CellScope, a prototype device which uses a programmable domed LED illumination to perform quantitative phase imaging, digital refocusing, and multi-contrast imaging in a portable (smartphone-based) form factor. Chapter 5 describes self-calibration techniques for quantitative phase imaging, including LED position recovery for LED domes and aberration recovery using a linearized model. These methods are essential for practical implementation of many quantitative phase imaging techniques. Chapter 6 concludes this dissertation on quantitative microscopy using coded illumination and provides future extensions of the work presented in the previous chapters.

Figure 1.4: Nikon TE300-based system used for most experiments presented in this dissertation. This system uses a digital camera, programmable LED illumination source (Quasidome), mechanical motion stage, and acquisition PC for controlling the acquisition and processing the data.

Chapter 2

Phase Imaging using Coded Illumination

Light is a propagating wave, having a velocity, wavelength (color), amplitude, and phase. As a single ray of light reaches an interface, it can be absorbed (reduced in amplitude) or undergo a change in phase through reflection, refraction, or diffraction based on the material properties and geometry of the object. Measuring this electromagnetic field at the sample plane is the primary goal of an imaging system; however, the current camera technology can only measure the amplitude of this complex field due to the relatively slow speed of silicon-based computing systems compared to the frequency of electromagnetic waves at optical wavelengths. For example, a coherent wavefront with wavelength $530nm$ traveling in air will have an oscillation period of approximately 1.76 femtoseconds ($1.76 \times 10^{-15}s$), which would need to be sampled at a rate of 1.13 petahertz ($1.13 \times 10^{15}Hz$) to sample the waveform at the Nyquist rate. Practical limitations of silicon-based computers limit their clock speeds of around 10 GHz ($10^9 Hz$), making it impossible to measure an optical wavefront by directly sampling it electronically. Instead, conventional cameras measure hundreds of thousands of periods of a waveform, averaging these periods to capture the intensity of the wavefront, which contains no direct phase information (although the phase of a wavefront can be indirectly inferred using camera hardware employing interferometric [24, 52] or geometric [16] methods).

Mathematically, the process of measuring the intensity of a wavefront described as taking the magnitude of the complex field $I = |\boldsymbol{E}|^2 = |\boldsymbol{A}e^{i\phi}|^2 = \boldsymbol{A}^2$ where $\boldsymbol{A}$ is the amplitude of a wave and $\boldsymbol{\phi}$ is the phase of the wavefront in this phasor notation. Phase is related to the mechanical geometry of the cell by the relationship $\boldsymbol{\phi} = \frac{2\pi}{\lambda}\boldsymbol{nd}$, where λ is the system wavelength, $\boldsymbol{n}$ is the refractive index change, and $\boldsymbol{d}$ is the thickness of the object.

This loss of phase information is particularly problematic for imaging aqueous samples, such as biological cells. To counter this, biologists often apply chemical stains to add absorption contrast artificially, but this process is cumbersome and can modify the micro-environment in significant ways. Optical phase imaging methods such as Differential Interference Contrast (DIC) [44] and Zernike Phase Contrast (PhC) [53] were developed to

provide phase contrast, which is a mix phase and amplitude of the wavefront and requires no chemical staining. These methods have since become widely adopted due to their simplicity and ability to reveal phase information in aqueous samples without adding additional contrast agents (stains).

2.1 Quantitative Phase Imaging

Quantitative Phase Imaging (QPI) involves recovering the complete amplitude and phase, or complex field, of a sample. In contrast to *qualitative* phase imaging methods, such as Zernike phase contrast and DIC, *quantitative* methods recover the phase delay caused by the sample, decoupled from absorption information. Modifications of PhC [54] and DIC [55] can make these setups quantitative, at a cost of requiring multiple images. More commonly, QPI methods use interferometry with coherent illumination and a reference beam [21, 22, 56], making them expensive and sensitive to misalignment and vibrations.

In optical microscopy, low cost and low-complexity methods which require little hardware modification are often preferred to more complex interferometric methods for practical reasons such as cost and calibration limitations. Through-focus phase retrieval [57–59] recovers the complex field using defocused intensity measurements of the sample by solving the transport of intensity (TIE) equation. TIE-based phase imaging has also been extended to use partially-coherent sources, which provides higher resolution than under coherent illumination [60]. The TIE equation is non-linear, however, and requires a second-order solver and phase unwrapping algorithm to accurately recover the full complex field.

Most quantitative phase imagine methods require the user to capture multiple images to facilitate phase recovery. Amongst the wide array of existing QPI methods only a small number are single-shot, since disambiguating amplitude and phase from a single-measurement is an ill-posed problem. Off-axis holography interferes the sample beam with a tilted reference beam, then recovers phase by Fourier filtering [23]. Parallel phase-shifting can spatially multiplex several holograms within a single exposure via an array of polarizers [61]. And single-shot QPI add-ons based on amplitude gratings work with commercial microscopes, replacing the traditional camera module [24, 25]. Another add-on option uses two cameras to capture defocused images which can then be used to solve the transport of intensity equation (TIE) [26]. Alternatively, if chromatic aberrations are large enough, they can enable single-shot color TIE [27, 62] without any hardware changes.

In the remaining sections of this chapter, we describe a quantitative phase imaging method which uses a programmable partially-coherent light source to perform quantitative phase imaigng, and demonstrate a single-shot variant which uses a color-multiplexing and inexpensive 3D-printed inserts to enable high-speed recovery of the complex field of a sample.

2.2 Differential Phase Contrast

Differential Phase Contrast (DPC) [49,63–65] is a partially coherent QPI technique that uses asymmetric illumination to shift the sample's spectrum in Fourier space, revealing phase information in a weak object. While raw DPC measurements provide phase contrast which is similar to DIC contrast, the DPC method provides a linearization of the image formation model which can be inverted using a single-step deconvolution process (Quantitative DPC) [49,64]. Quantitative DPC recovers both amplitude and phase with resolution up to the incoherent resolution limit (2× better than coherent methods). Practically, the illumination switching can be done quickly and at low cost with an LED array [47,49,65]. At least two complementary source patterns are required, but generally 3-4 patterns (top, bottom, left, right half-circles) are used to avoid missing spatial frequencies. The DPC method was recently extended to color multiplexing [66], where the 4 source patterns were encoded into two images by using a color camera in combination with a color LED array. Similarly, color photometric stereo has been used for retrographic surface profiling of large objects using off-axis color illumination in reflection mode [67].
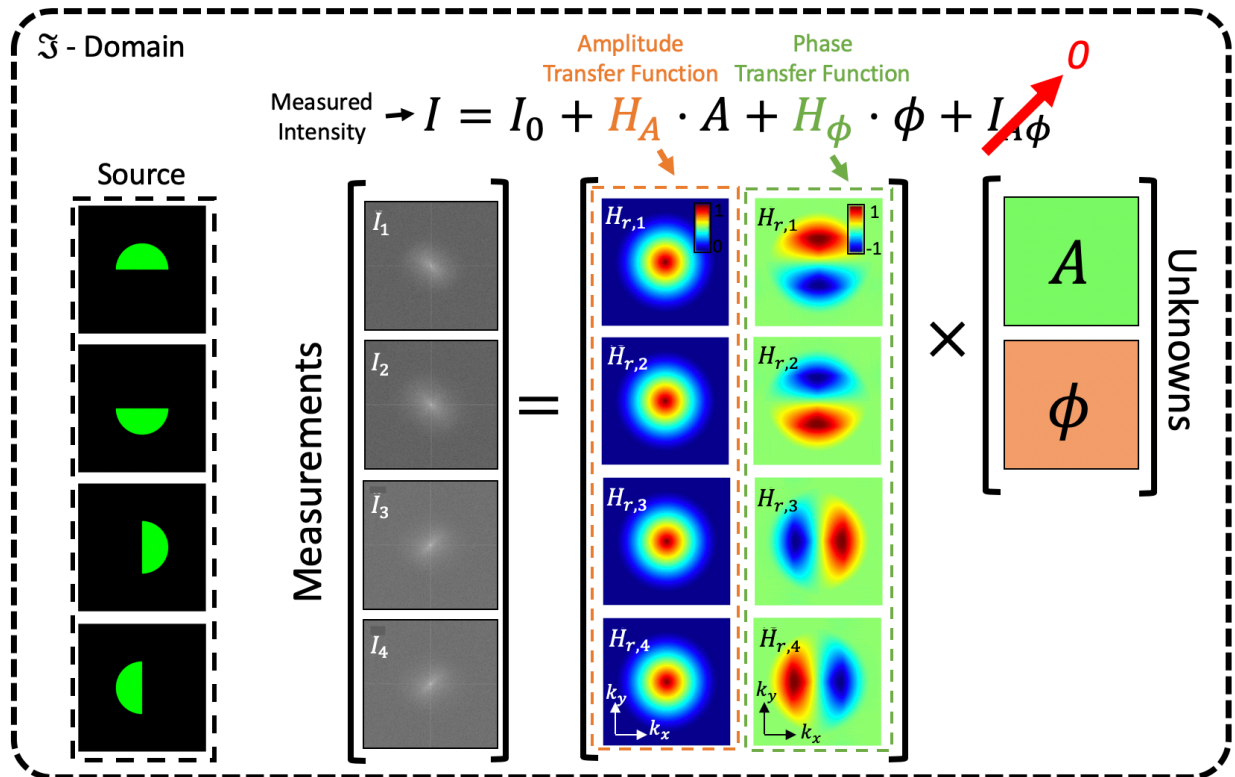
Figure 2.1: Example transfer functions for typical half-circle illumination patterns.

Forward model

The general image formation model for a sample illuminated under a partially-coherent source is defined as the sum of many coherent forward models [64]:

$$I = \sum_i^N |\boldsymbol{P} * \boldsymbol{O} e^{-i2\pi \frac{\sin \theta_i}{\lambda} \boldsymbol{r}}|^2 \tag{2.1}$$

where N sources illuminating from angles θ_i with wavelength λ illuminate an object $\boldsymbol{O}$ which is then filtered by a pupil $\boldsymbol{P}$, having spatial coordinates $\boldsymbol{r}$. In general, this model is nonlinear and difficult to invert directly. DPC employs a linearized forward model to describe the intensity images that result from a given complex-field. This linearization is achieved by making a weak phase assumption on the sample's complex-field and ignoring higher order (nonlinear) terms in the Taylor expansion of the complex-field, $\boldsymbol{E} = e^{i\phi - \boldsymbol{\mu}} \approx 1 - \boldsymbol{\mu} + i\boldsymbol{\phi}$. Under the same approximation, the Weak Object Transfer Functions (WOTFs) for absorption ($\tilde{\boldsymbol{H}}_\mu$) and phase ($\tilde{\boldsymbol{H}}_\phi$) can be derived [49, 63, 68] to linearly related to intensity measurements:

$$\boldsymbol{I}(r) = \boldsymbol{I}_0 + \boldsymbol{H}_\mu(r) * \boldsymbol{\mu}(r) + i \cdot \boldsymbol{H}_\phi(r) * \boldsymbol{\phi}(r) + \boldsymbol{I}_{ss} \tag{2.2}$$

Here, r represents 2D real-space coordinates, $\boldsymbol{I}$ is the intensity measurement, $\boldsymbol{I}_0$ is the background signal, and $*$ denotes convolution. $\boldsymbol{I}_{ss}$ is the 2^{nd} order scatter-scatter term, which is assumed to be small and is ignored in the DPC algorithm. Eq. 2.2 can also be represented in the Fourier domain:

$$\tilde{\boldsymbol{I}}(r) = \tilde{\boldsymbol{I}}_0 + \tilde{\boldsymbol{H}}_\mu(r) \cdot \tilde{\boldsymbol{\mu}}(r) + i \cdot \tilde{\boldsymbol{H}}_\phi(r) \cdot \tilde{\boldsymbol{\phi}}(r) + \tilde{\boldsymbol{I}}_{ss} \tag{2.3}$$

Here, $\tilde{\ }$ represents the Fourier transform of the measurements and sample amplitude ($\boldsymbol{\mu}$) and phase ($\boldsymbol{\phi}$). Given a known source ($\boldsymbol{S}$), and pupil function ($\boldsymbol{P}$) whose bandwidth is set by the objective numerical aperture (NA) and wavelength (λ), the WOTFs are [49, 68]:

$$\tilde{\boldsymbol{H}}_\mu(k) = \tilde{\boldsymbol{P}}(k) \star (\tilde{\boldsymbol{P}}(k) \cdot \tilde{\boldsymbol{S}}(k)) + (\tilde{\boldsymbol{P}}(k) \cdot \tilde{\boldsymbol{S}}(k)) \star \tilde{\boldsymbol{P}}(k) \tag{2.4}$$

$$\tilde{\boldsymbol{H}}_\phi(k) = \tilde{\boldsymbol{P}}(k) \star (\tilde{\boldsymbol{P}}(k) \cdot \tilde{\boldsymbol{S}}(k)) - (\tilde{\boldsymbol{P}}(k) \cdot \tilde{\boldsymbol{S}}(k)) \star \tilde{\boldsymbol{P}}(k), \tag{2.5}$$

where k represents Fourier-domain coordinates. Generally, multiple DPC measurements are acquired together, resulting in a series of measurements $\{\boldsymbol{y}_1, \boldsymbol{y}_2, \cdots \boldsymbol{y}_n\}$ and a stacked forward model based on one pair of $\tilde{\boldsymbol{H}}_\phi$ and $\tilde{\boldsymbol{H}}_\mu$ for each source:

$$\begin{bmatrix} \tilde{\boldsymbol{y}}_1 \\ \tilde{\boldsymbol{y}}_2 \\ \tilde{\boldsymbol{y}}_3 \end{bmatrix} = \begin{bmatrix} \tilde{\boldsymbol{H}}_{\mu,1} & \tilde{\boldsymbol{H}}_{\phi,1} \\ \tilde{\boldsymbol{H}}_{\mu,2} & \tilde{\boldsymbol{H}}_{\phi,2} \\ \tilde{\boldsymbol{H}}_{\mu,3} & \tilde{\boldsymbol{H}}_{\phi,3} \end{bmatrix} \times \begin{bmatrix} \tilde{\boldsymbol{\mu}} \\ \tilde{\boldsymbol{\phi}} \end{bmatrix} \tag{2.6}$$

Phase Reconstruction

Eq. 2.6 is linear and can be inverted using linear least squares [49] to recover phase and absorption of the complex field. Due to nulls in the WOTFs, it is often prudent to add regularization (such as Tikhonov) to ensure noise is not amplified unnecessarily during the reconstruction process. With m measurements and Tikhonov regularization parameters γ_μ and γ_ϕ, the linearized complex field absorption ($\boldsymbol{\mu}$) and phase ($\boldsymbol{\phi}$) can be recovered using the following equations:

$$\mu = F^{-1} \left\{ \frac{\left(\sum_m |\tilde{\boldsymbol{H}}_{\phi,m}|^2 + \gamma_\phi\right) \cdot \sum_m \left(\tilde{\boldsymbol{H}}_{\mu,m}^* \cdot \tilde{\boldsymbol{I}}_m'\right) - \sum_m \left(\tilde{\boldsymbol{H}}_{\mu,m}^* \cdot \tilde{\boldsymbol{H}}_{\phi,m}\right) \cdot \sum_m \left(\tilde{\boldsymbol{H}}_{\phi,m}^* \cdot \tilde{\boldsymbol{I}}_m'\right)}{\left(\sum_m |\tilde{\boldsymbol{H}}_{\mu,m}|^2 + \gamma_\mu\right) \cdot \left(\sum_m |\tilde{\boldsymbol{H}}_{\phi,m}|^2 + \gamma_\phi\right) - \sum_m \left(\tilde{\boldsymbol{H}}_{\mu,m} \cdot \tilde{\boldsymbol{H}}_{\phi,m}^*\right) \cdot \sum_m \left(\tilde{\boldsymbol{H}}_{\mu,m}^* \cdot \tilde{\boldsymbol{H}}_{\phi,m}\right)} \right\}$$

$$(2.7)$$

$$\phi = F^{-1} \left\{ \frac{-\mathrm{i} \cdot \left[\left(\sum_m |\tilde{\boldsymbol{H}}_{\mu,m}|^2 + \gamma_\mu\right) \cdot \sum_m \left(\tilde{\boldsymbol{H}}_{\phi,m}^* \cdot \tilde{\boldsymbol{I}}_m'\right) - \sum_m \left(\tilde{\boldsymbol{H}}_{\mu,m} \cdot \tilde{\boldsymbol{H}}_{\phi,m}^*\right) \cdot \sum_m \left(\tilde{\boldsymbol{H}}_{\mu,m}^* \cdot \tilde{\boldsymbol{I}}_m'\right) \right]}{\left(\sum_m |\tilde{\boldsymbol{H}}_{\mu,m}|^2 + \gamma_\mu\right) \cdot \left(\sum_m |\tilde{\boldsymbol{H}}_{\phi,m}|^2 + \gamma_\phi\right) - \sum_m \left(\tilde{\boldsymbol{H}}_{\mu,m} \cdot \tilde{\boldsymbol{H}}_{\phi,m}^*\right) \cdot \sum_m \left(\tilde{\boldsymbol{H}}_{\mu,m}^* \cdot \tilde{\boldsymbol{H}}_{\phi,m}\right)} \right\},$$

$$(2.8)$$

For improved performance, other regularizers such as Total Variation (TV) regularization [69], although many of these methods require iterative solvers such as FISTA [29] which is slower than direct inversion.
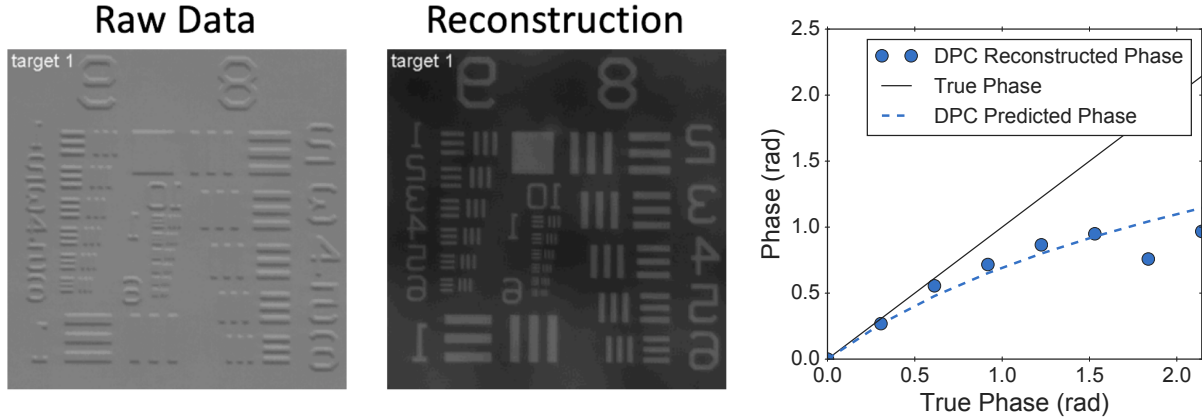
Figure 2.2: Differential Phase Contrast Reconstruction of a USAF 1951 resolution target printed as a phase object. The DPC linearization becomes less accurate for objects with strong phase (> 1 radian) or a strong phase gradient.

Limitations

The linearization made by DPC is valid when absorption and phase gradient are small. Practically, DPC will fail to recovery objects which have sharp edges and strong gradients but will successfully recover large phase values which vary slowly (such as a micro-lens array in Fig. 2.3) [68,70]. To provide experimental validation of the DPC approximation, we acquired test images for a series of phase targets across a range of physical heights - each having a phase which increases linearly across 9 targets [71]. These targets were manufactured on a single test slide (Benchmark Technologies), having USAF1951, star targets, a "wedding-cake" structure with three discrete heights. The results are shown in Fig. 2.2, illustrating the decreasing validity of the DPC approximation with increasing target height (and target phase). The WOTF prediction can be derived by taking the natural logarithm of the weak-object approximation ($i\phi_{true} \approx 1 + i\phi_{dpc}$), which reveals that the true phase ϕ_{dpc} is related to the DPC-recovered phase by the relationship $i\phi_{true} = \texttt{ln}(1 + i\phi_{dpc})$, although this approximation is only valid for a single-scattering model. To place this in context, a typical biological cell monolayer immersed in water will have less than 0.2 radians of phase deviation, making DPC well-suited for these sorts of imaging tasks.

2.3   cDPC: Single-Shot Quantitative Phase Imaging

To improve temporal resolution of a DPC system without compromising spatial resolution, we propose color Differential Phase Contrast (*cDPC*), which requires only a *single* color image using color-multiplexed source patterns[1]. In this method, the source is discretized into three color channels which are used to display three different half-circle source patterns. It is important to note that while the original DPC algorithm presented in [49] requires 4 images, the proposed method only requires 3, since the 4th source configuration can be synthesized by taking the sum of two images acquired with opposite half-circle illuminations (a synthetic brightfield image) and subtracting that of a 90 degree rotated half-circle source. In early prototypes the color source pattern was implemented in an LED array microscope, which offers many imaging modalities in one platform [13, 45, 47, 49, 65, 72, 73]. More recent work proposed the use of a motion-compensation algorithm to solve for motion between DPC frames, effectively providing a single-shot phase imaging using computation [74]. Unlike these methods, however, our proposed configuration does not require a dynamic source, making it possible to use a static multi-color filter placed in the condenser back focal plane, assuming Köhler illumination. Both configurations simplify hardware and reduce costs significantly as compared to phase contrast or DIC, while providing quantitative phase, which is more general and can be used to synthesize both of the aforementioned methods digitally [75].

---

[1]This work was performed in close collaboration with Michael Chen (Waller Lab, EECS, UC Berkeley).
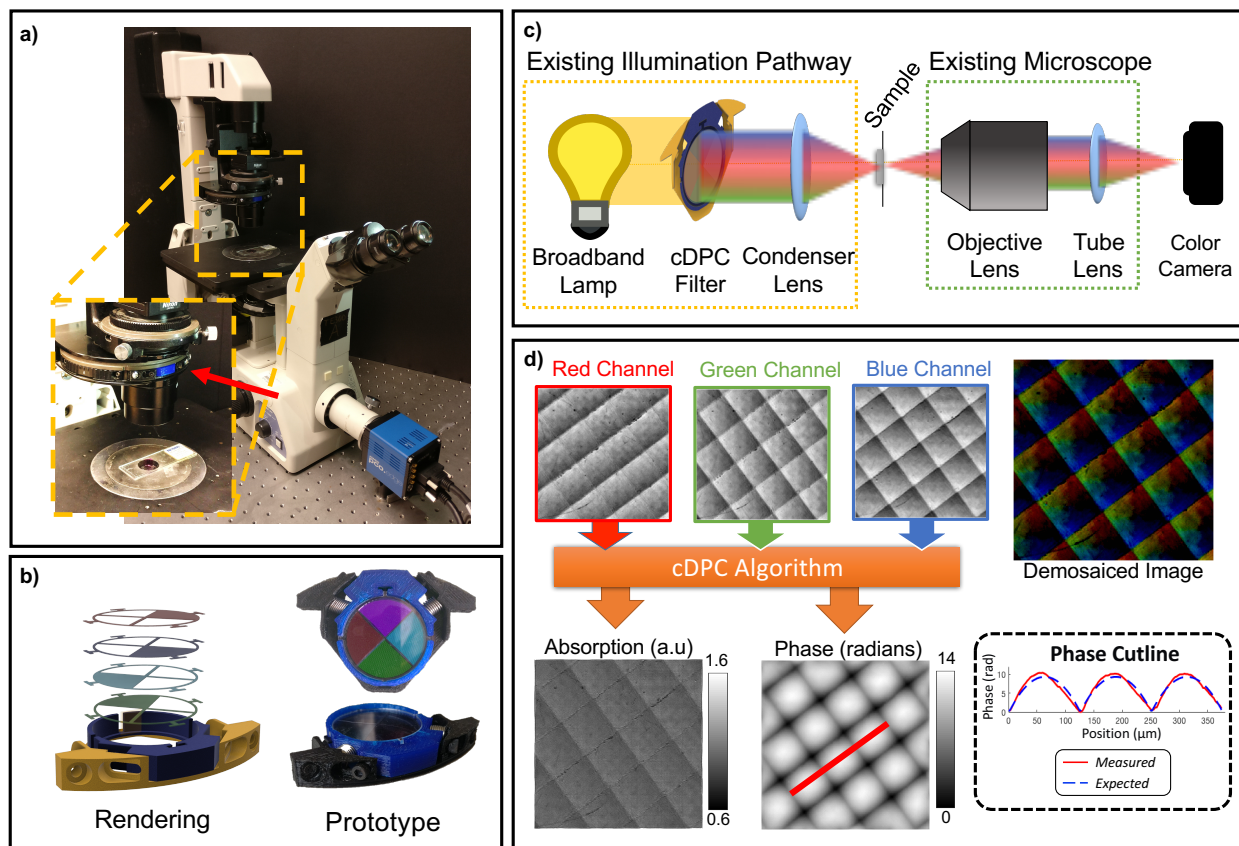
Figure 2.3: Single-shot color Differential Phase Contrast (*cDPC*) microscopy. a) Installation in Nikon TE300 microscope condenser turret. b) CAD model and image of fabricated *cDPC* insert.c) Optical schematic of a brightfield microscope with a *cDPC* color filter placed at the back focal plane of the condenser in Köhler configuration. d) Reconstruction: the captured color image is separated into its RGB components, which are then used to recover two unknowns (amplitude and phase) via a well-posed linear deconvolution. The sample is a micro-lens array (Fresnel Technologies 605).

Hardware Design

As in conventional DPC, this method requires measurements of the sample illuminated by known asymmetric sources. In *cDPC*, however, we make use of the microscope's existing condenser unit, which has a turret commonly used for phase contrast inserts or DIC prisms. This intermediate plane can usually be accessed easily by removing the mechanical inserts. Taking advantage of this configuration, a simple 3D printed color filter was designed and fabricated that can be placed in the condenser turret of a Nikon TE300 microscope (Figure 2.3a).

The filter prototype consists of Polyethylene Terephthalate (PET) color filters (Lee Filter, Inc.) laser cut to size and installed into a 3D printed insert designed to fit our microscope. Narrow bandwidth illumination filters (e.g. multi-layer coated glass) would separate colors better but suffer from low light throughput and high cost. Therefore, inexpensive and easy-to-cut PET film filters were used; the resulting cross-talk between color channels will be accounted for in post-processing, described below.

The total cost of raw materials is approximately \$30 and filters were produced quickly with a 3D printer and laser cutter. One filter is shown in Figure 2.3b; it was installed in the condenser turret of an inverted microscope (Figure 2.3a), replacing one of the removable phase contrast (Ph1, Ph2 or Ph3) inserts.

Calibration

Ideally, the color filters would provide perfect separation of the three source patterns into the three-color channels. In reality, both the illumination and camera color channels have cross-talk between the desired wavelengths. To account for this, system calibration is separated into two separate steps: detection-side and illumination-side.

Illumination-side calibration corrects for the relative spectral transmittance of each of the source color filters. The illumination pattern simultaneously encodes three half-circle sources, one each for the RGB color channels. Red and green are opposite half-circles, and blue is rotated by 90 degrees relative to the others. Where the blue and green patterns overlap, a cyan filter (blue + green) was used. Where the blue and red patterns overlap, a purple filter (blue + red) was used. Hence, the final filter design actually contains four quadrants having red, green, cyan and purple filters (see Fig. 2.4).

When filtered by the sensor Bayer pattern, the filter spectrum bases are not orthogonal. This can be seen in the spectra of each PET film after capture with a color camera (left column of Fig. 2.4). The result is an undesirable loss of asymmetry in the source that reduces phase SNR. However, it is possible to account for the asymmetry during reconstruction by modeling the source patterns as in Fig. 2.4.

Detection-side calibration accounts for spectral cross-talk of the camera color channels. Standard RGB Bayer filters do not provide perfect discrimination between RGB wavelengths but coupling artifacts can be removed by calibration. Given the pixel values from the raw color image with an RGBG Bayer filter ($\boldsymbol{I}_r$, $\boldsymbol{I}_{g1}$, $\boldsymbol{I}_{g2}$, $\boldsymbol{I}_b$), it is possible to solve for the decoupled color image ($\boldsymbol{I}_R$, $\boldsymbol{I}_G$, $\boldsymbol{I}_B$) that would be obtained if the sample were illuminated with a single color, according to the following equation,

$$\begin{bmatrix} \boldsymbol{I}_r \\ \boldsymbol{I}_{g1} \\ \boldsymbol{I}_{g2} \\ \boldsymbol{I}_b \end{bmatrix} = \mathbf{C} \begin{bmatrix} \boldsymbol{I}_R \\ \boldsymbol{I}_G \\ \boldsymbol{I}_B \end{bmatrix} . \tag{2.9}$$

The matrix $\mathbf{C}$ is a 4×3 calibration matrix describing the coupling between each color channel. It is generated by filtering the broadband source with each filter independently, then measur-
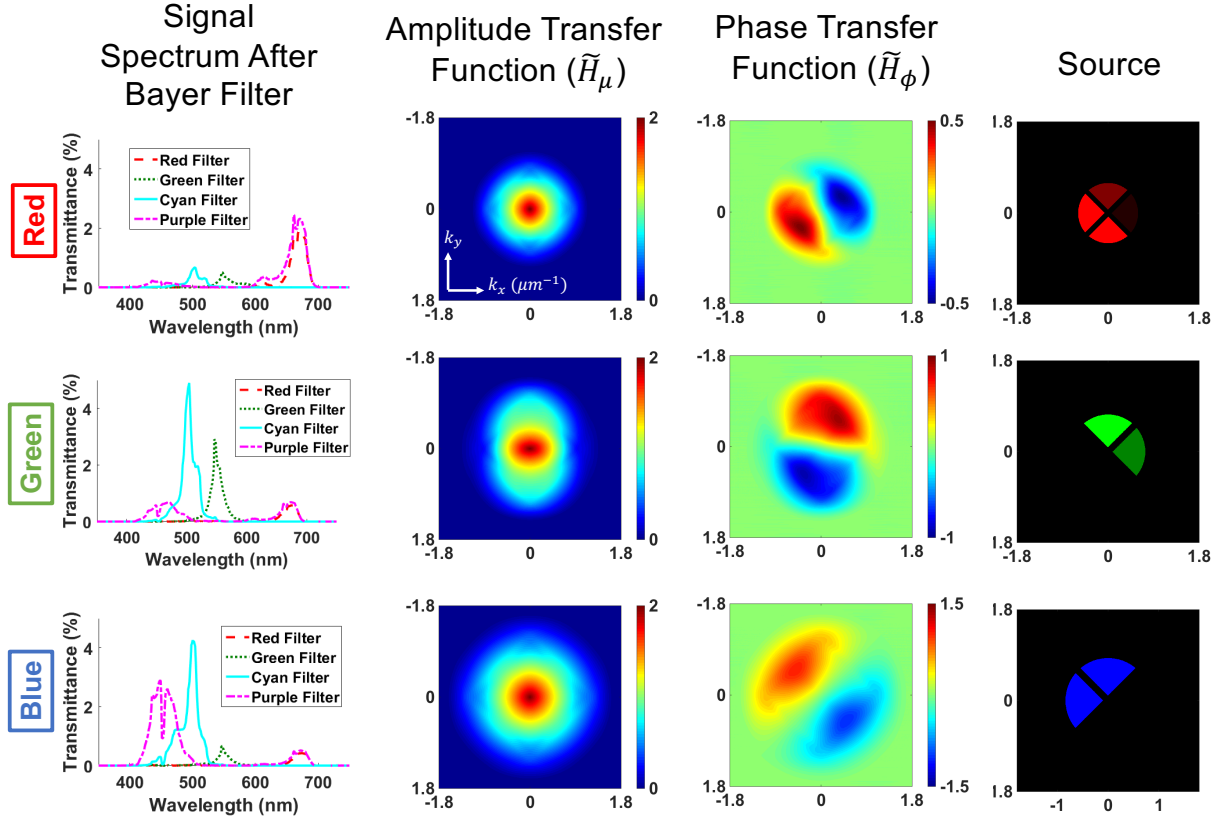
Figure 2.4: Transfer functions for amplitude and phase contrast in each *cDPC* color channel. Left: Spectral contribution of each illumination filter as captured by the camera's Bayer pattern. The following columns show the components of the amplitude and phase transfer functions in the spatial frequency domain and the source represented in each image. Bottom row: sum of each column, representing the calibrated and scaled source and the total coverage of amplitude and phase transfer functions, respectively.

ing the relative red (I_R), green (I_G) and blue (I_B) read-outs to populate the corresponding column vectors of the **C** matrix. The ratio between the intensities of each flat-field image at each detection channel provides a linear weighting of the contribution of each source to the color measurement. Once **C** has been measured once, it can be used to pre-process all later measurements by solving Eq. (2.9). This step is important for reducing artifacts in the phase results.

Another important step for *cDPC* is to account for wavelength-dependent changes in phase and spatial frequency. DPC recovers absorption (μ) and phase (ϕ) information from intensity measurements. These quantities are defined as:

$$\mu = \frac{2\pi}{\lambda_0}\boldsymbol{\alpha d}, \ \ \phi = \frac{2\pi}{\lambda_0}\boldsymbol{nd}, \tag{2.10}$$

where λ_0 is a reference wavelength, $\boldsymbol{d}$ is the thickness of the sample, $\boldsymbol{n}$ represents refractive index and $\boldsymbol{\alpha}$ indicates absorption coefficient. Absorption and phase transfer functions are determined by illumination numerical aperture (NA), objective NA and illumination wavelength [49]. In the proposed color-multiplexed DPC method, the transfer functions must also consider the change in wavelength of each color channel. Phase ($\boldsymbol{\phi}$) depends on which wavelength is used. By assuming no dispersion in the sample, it is possible to use Eq. (2.10) to synthesize phase for any wavelength by simply multiplying the optical path length ($\boldsymbol{nd}$) by the wave number ($\frac{2\pi}{\lambda_0}$) of a desired reference wavelength λ_0.

Examining Fig. 2.4, it is clear that the absorption transfer functions for each color channel are symmetric low-pass filters. The phase transfer functions, on the other hand, are asymmetric band-pass-like filters with a line of missing frequencies along the axis of asymmetry. By rotating the blue half-circle by 90 degrees relative to the red and green ones, the missing line is filled. The overall amplitude and phase transfer functions for *cDPC* are shown in the last row of Fig. 2.4, calculated by summing the absolute values of each color transfer function. As with previous DPC implementations, absorption information loses contrast at high spatial frequencies. Phase has a similar drop-off at high frequencies, but also loses contrast in the low spatial frequency regions. Hence, SNR will be important for accurately recovering low-frequency phase information. The maximum spatial frequency range captured is 2× the NA of the blue color channel. However, the final resolution using *cDPC* is set by the diffraction limit of green light, since total frequency coverage is set by the maximum spatial frequency which is measured by *two or more* color channels. This comes as an implication of trying to recover two unknowns, amplitude and phase, thus requiring at least two measurements.

Inverse problem

Using the forward model developed in Section 2.2, the *cDPC* inverse problem aims to minimize the difference between the measured color image and that which would be measured, given the estimate of the sample's amplitude and phase:

$$\min_{\mu,\phi} \sum_{m=1}^{3} \frac{1}{2} \parallel \tilde{\boldsymbol{I}}'(\lambda_m) - \tilde{\boldsymbol{H}}_\mu(\lambda_m) \cdot \tilde{\boldsymbol{\mu}} - \mathrm{i} \cdot \tilde{\boldsymbol{H}}_\phi(\lambda_m) \cdot \tilde{\boldsymbol{\phi}} \parallel_2^2 + \mathcal{R}(\tilde{\boldsymbol{\mu}}, \tilde{\boldsymbol{\phi}}), \tag{2.11}$$

where $\tilde{\boldsymbol{I}}'$ is the spatial frequency spectrum of the background-subtracted intensity, m is the wavelength index and $\mathcal{R}(\tilde{\boldsymbol{\mu}}, \tilde{\boldsymbol{\phi}})$ is a regularization term (typically on the order of 10^{-3}). This problem is linear and can be solved with a one-step least-square solution (e.g. Wiener deconvolution [76]) or by an iterative algorithm (e.g. gradient descent). The ideal choice of regularizer $\mathcal{R}(\tilde{\boldsymbol{\mu}}, \tilde{\boldsymbol{\phi}})$ depends on the sample and noise.

$$\min_{\boldsymbol{\mu},\boldsymbol{\phi}} \sum_{m=1}^{3} \frac{1}{2} \parallel \tilde{\boldsymbol{I}}'(\lambda_m) - \tilde{\boldsymbol{H}}_\mu(\lambda_m) \cdot \tilde{\boldsymbol{\mu}} - \mathrm{i} \cdot \tilde{\boldsymbol{H}}_\phi(\lambda_m) \cdot \tilde{\boldsymbol{\phi}} \parallel_2^2 + \gamma_{\boldsymbol{\mu}} \cdot \parallel \boldsymbol{\mu} \parallel_2^2 + \gamma_{\boldsymbol{\phi}} \cdot \parallel \boldsymbol{\phi} \parallel_2^2, \quad (2.12)$$

which remains differentiable and allows us to find the global minimum solution for absorption and phase with a single matrix inversion step. The reconstructed amplitude and phase are obtained using Eq. 2.7 and Eq. 2.8 as in conventional DPC
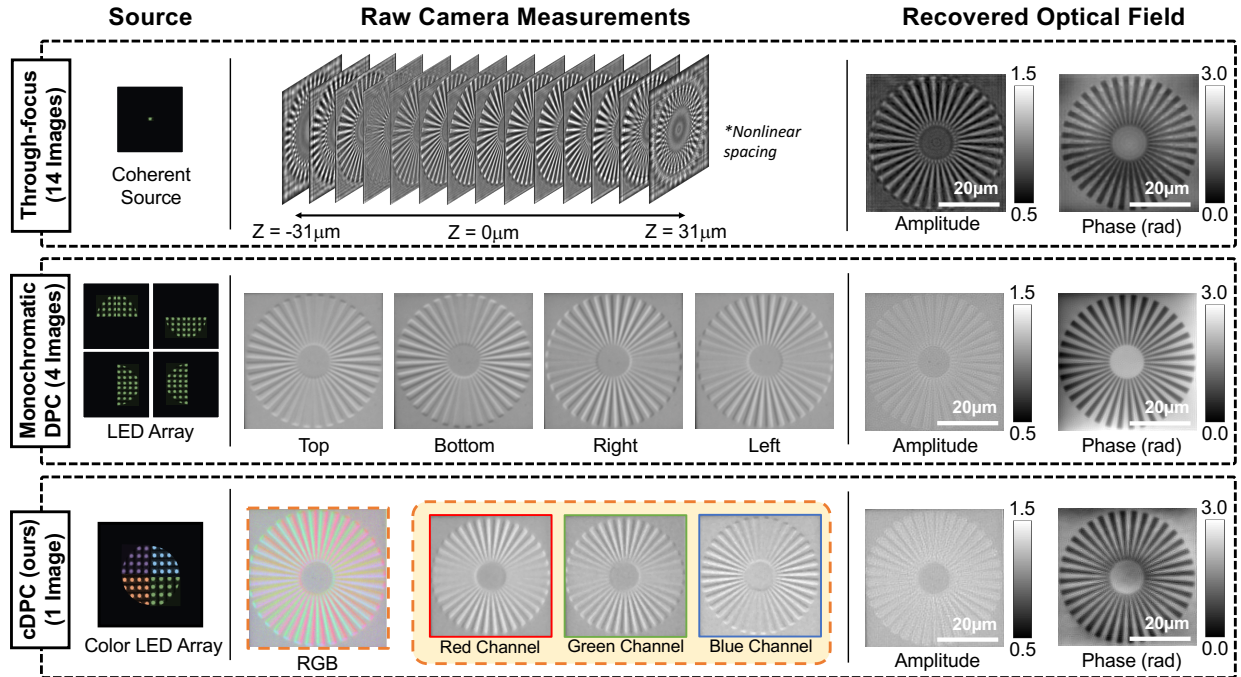
Figure 2.5: Experimental comparison of single-shot *cDPC* with monochromatic DPC and through-focus phase retrieval methods. (Left) Source patterns. (Middle) Raw camera measurements. (Right) Recovered optical field. DPC methods (partially coherent) were acquired using a 20× 0.4 NA objective lens, while through-focus images (spatially coherent) were captured using 60× 0.8 NA, in order to ensure equal resolution in all cases.

Validation

To experimentally validate the proposed *cDPC* method, results were compared with two established QPI methods: monochromatic DPC and through-focus phase retrieval (Fig. 2.5). For fair comparison, all are implemented on the same Nikon TE300 microscope using illumination generated by an RGB LED array (Adafruit). Each *cDPC* experiment uses a discretized version of the *cDPC* color filter design displayed on the LED array. Monochromatic DPC

uses 4 images captured with each of 4 asymmetric source patterns [47]. Through-focus phase imaging uses only the central green LED (for temporal and spatial coherence) while capturing 14 images at different focus depths; phase is then recovered by a nonlinear optimization phase retrieval method [60].

Because of the coherent illumination, through-focus phase imaging has 2× worse resolution than DPC methods. Thus, a 20× 0.4 NA objective lens was used for DPC methods, but switched to a 60× 0.8 NA objective for through-focus phase, in order keep resolution equal for all three. Spatial resolution is quantified using a spoke-pattern phase target [77].

As can be seen in Fig. 2.5, the RGB color channel images have similar contrast to the left, right and top images of the monochromatic DPC, as expected. The phase results are also similar, with equivalent spatial resolution. Because the *cDPC* image is captured in one shot with color filters, it has lower SNR than monochromatic DPC and deviates in its low-frequency fluctuations, which have weaker transfer function values. Overall, however, single-shot *cDPC* performs comparably to multi-shot DPC.

Next, the LED array was removed and replaced with the existing illumination pathway. For illumination, a broadband arc lamp light source was used. Alternatively, a high-power blue-phosphor static LED source could be used. The color filter insert shown in Fig. 2.3b was then installed into the condenser turret. Figure 2.6 shows amplitude and phase reconstructions from the proposed *cDPC* method with objectives of various magnification. The *cDPC* method is compatible with any standard objective having $NA_{objective} \leq NA_{illumination}$. If an objective with larger NA than the condenser NA is used, the low frequencies of phase will not be transmitted during image formation (see the phase transfer function in Fig. 2.4), since phase contrast comes primarily from high-angle illumination. The spatial coherence factor σ is often defined as:

$$\sigma = \frac{NA_{illumination}}{NA_{objective}}. \tag{2.13}$$

In other words, $\sigma < 1$ will result in reduced phase contrast as compared to the $\sigma \geq 1$ case. The Nikon TE300 microscope used in this study was configured with a 0.53 NA condenser lens. Imaging with a higher objective NA would require high-NA illumination (e.g. by using a domed LED array [73]). Temporal coherence is set by the bandwidth of the color filters, since these have narrower bandwidth than the camera filters. The full-width-half-maximum (FWHM) bandwidth for the filters used in this study was approximately 50nm, which is similar to the emission spectrum of the LED array used previously [49].

Temporal Resolution

Since *cDPC* is single-shot, temporal resolution is set by the camera's frame rate, giving a factor of 4 improvement over conventional DPC. Single-shot methods reduce artifacts due to motion blur and image registration. This can be seen in Fig. 2.7, where the performance *cDPC* and conventional DPC are compared when imaging a live *C. elegans* culture. Motion
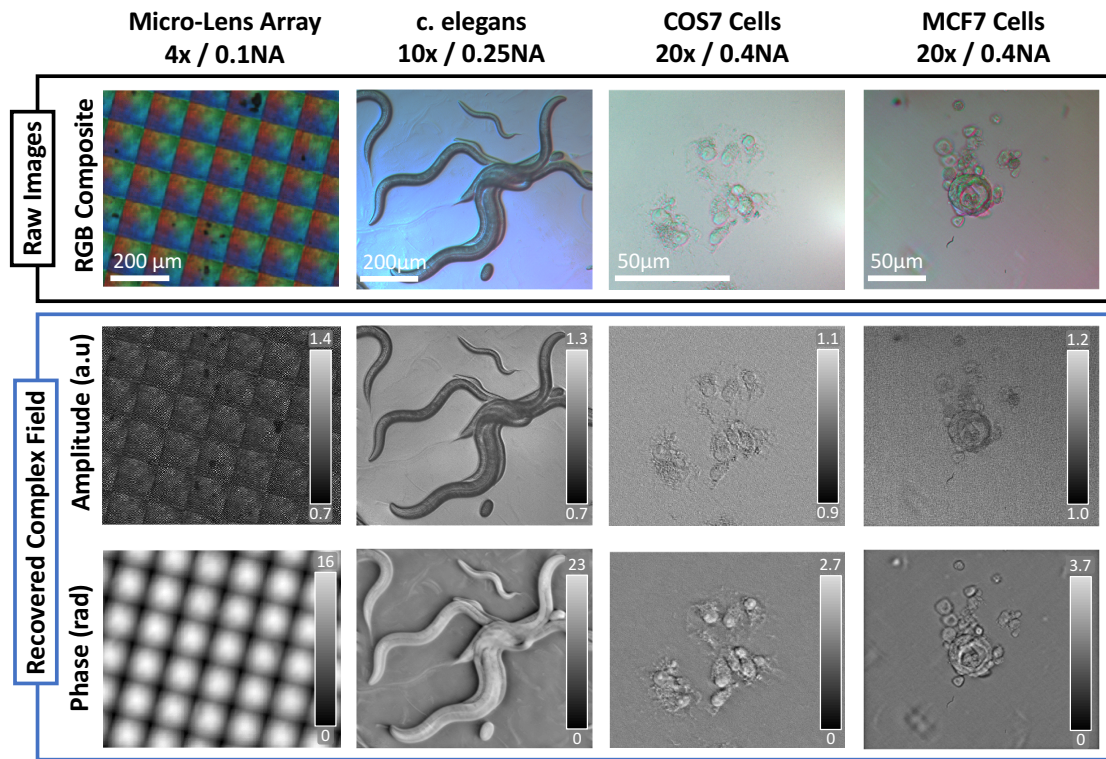
Figure 2.6:   Phase and amplitude reconstructions for various samples and magnifications. (First column) Micro-lens array, 4x 0.1 NA. (Second column) Wild-type *C. elegans*, 10x 0.25 NA. (Third column) HEK 293T cells, 20× 0.4 NA). (Fourth column) MCF7 cells, 20× 0.4 NA.

blur is significantly reduced with *cDPC*, since the sample changes rapidly between frames, even at 12.5 frames per second.

Synthesized PhC and DIC Images

Differential Interference Contrast (DIC) and conventional Phase Contrast (PhC) microscopy are examples of the widespread adoption of phase imaging methods in medicine and biomedical research. Though both methods have gained widespread adoption, optical components required for their implementation remain expensive, and alignment by an experienced user is required for acceptable performance. Both DIC and phase contrast can be described by forward models which produce a qualitative mixture of amplitude and phase images [43, 44]. Since the forward models of these systems are well known, quantitative phase imaging methods can be used to form these images digitally, mimicking the physical optical system through numerical simulation. Synthesized images from *cDPC*, as well as ground truth DIC and PhC images, are shown in Figure 2.8 to be comparable.
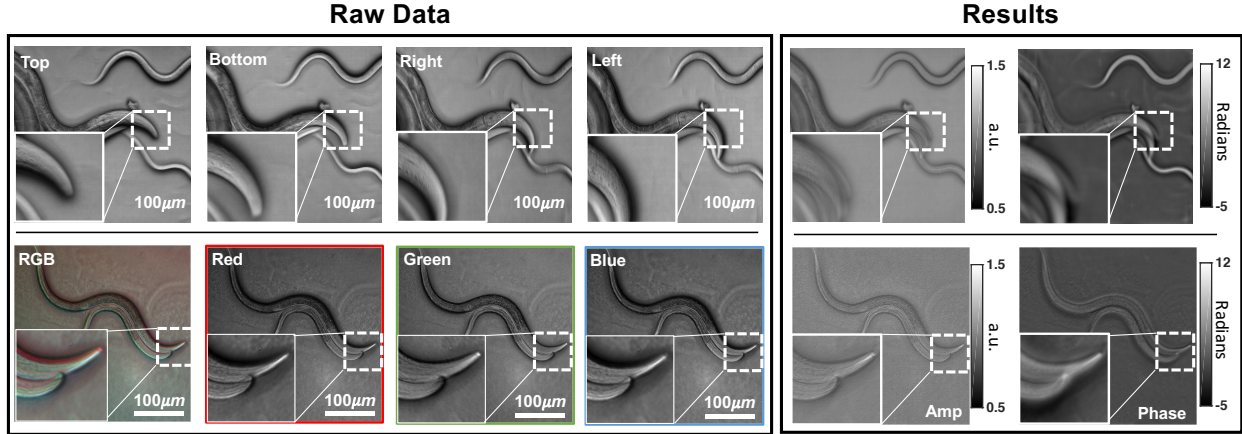
Figure 2.7: Experimental demonstration of motion blur reduction with *cDPC* vs. conventional DPC. The *cDPC* method results in significantly reduced motion blur artifacts due to its single-shot acquisition.

Synthesizing DIC and PhC is of particular use for clinicians and researchers who have been trained to make diagnoses or decisions based on these images. While all QPI methods can be used to synthesize these images, the *cDPC* method is particularly well-suited since it is single-shot, allowing for real-time digital synthesis. In addition, *cDPC* is much cheaper to implement than either DIC or PhC, since it requires only the addition of an inexpensive color filter insert and no specialized objectives. In contrast, DIC prisms and phase contrast objectives (specific to a given NA) can drive up the cost of a microscope significantly.

Compatibility with Stained and Dispersive Samples

The *cDPC* method uses color multiplexing to recover complex-field, making an inherent assumption that the sample is both non-dispersive and colorless. Non-dispersive means that the refractive index does not change appreciably with wavelength:

$$\phi(\boldsymbol{n}(\lambda), \boldsymbol{d}, \lambda) \approx \phi(\boldsymbol{n}_0, \boldsymbol{d}, \lambda). \tag{2.14}$$

This assumption implies that the optical path length ($OPL = nd$) will remain constant for all measurements. The relative phase delay will always vary with λ (Eq. 2.10), but this is accounted for in the *cDPC* algorithm by scaling the transfer functions based on the relative wavelength of each color channel. Unless the dispersion curve is known and the material is assumed to be uniform, one cannot account for dispersive effects in the sample using the proposed algorithm. However, in practice these effects do not corrupt our phase reconstructions results significantly due to relatively small dispersive effects of water across optical wavelengths.
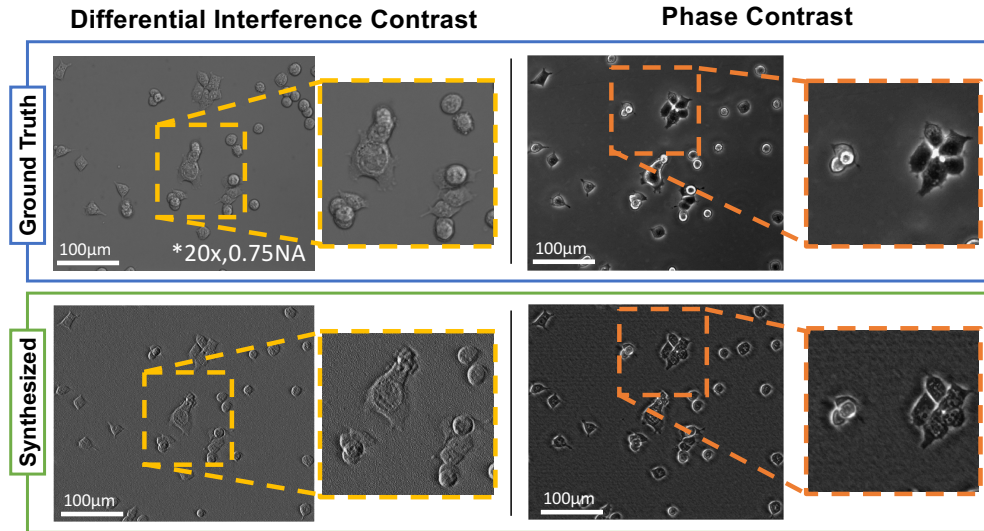
Figure 2.8: Comparison of standard DIC and PhC images to their synthesized counterparts from *cDPC*. Ground truth DIC images were acquired using a 20x 0.75 NA objective and phase contrast images using a 20x 0.4 NA PhC objective. *cDPC* images were acquired using a 20x 0.4 NA objective and the filter insert.

The second assumption is that the sample is colorless, meaning that the absorption does not have chromatic dependence:

$$\boldsymbol{\mu}(\lambda) \approx \boldsymbol{\mu}_0. \tag{2.15}$$

This is generally valid for unstained biological samples, which are transparent. Color variations due to filter transmission coefficients at different wavelengths are present, but can be removed by the calibration procedure described in Section 1.2. Color-dependent absorption, such as that created by stained samples, cannot be recovered and will cause errors in the phase result. In practice, these assumptions limit the applicability of the *cDPC* method to unstained uncolored samples. However, quantitative phase reveals the mechanical structure of the microenvironment with high contrast, which may eliminate the need for staining in many applications.

2.4 SNR Analysis of DPC Phase Recovery Systems

Differential Phase Contrast linearizes the partially coherent image formation model, enabling the analysis of this system using the noise analysis methods described in Chapter 1, Section 1.4. This section will analyze the choice of source pattern and pupil function in the context of our propsed linear noise model, and examine the practical benefit of using a variable number of measurements, variable source diameter, and adding a defocus aberration to the pupil, in terms of reconstruction SNR. Recent works have used learning techniques to
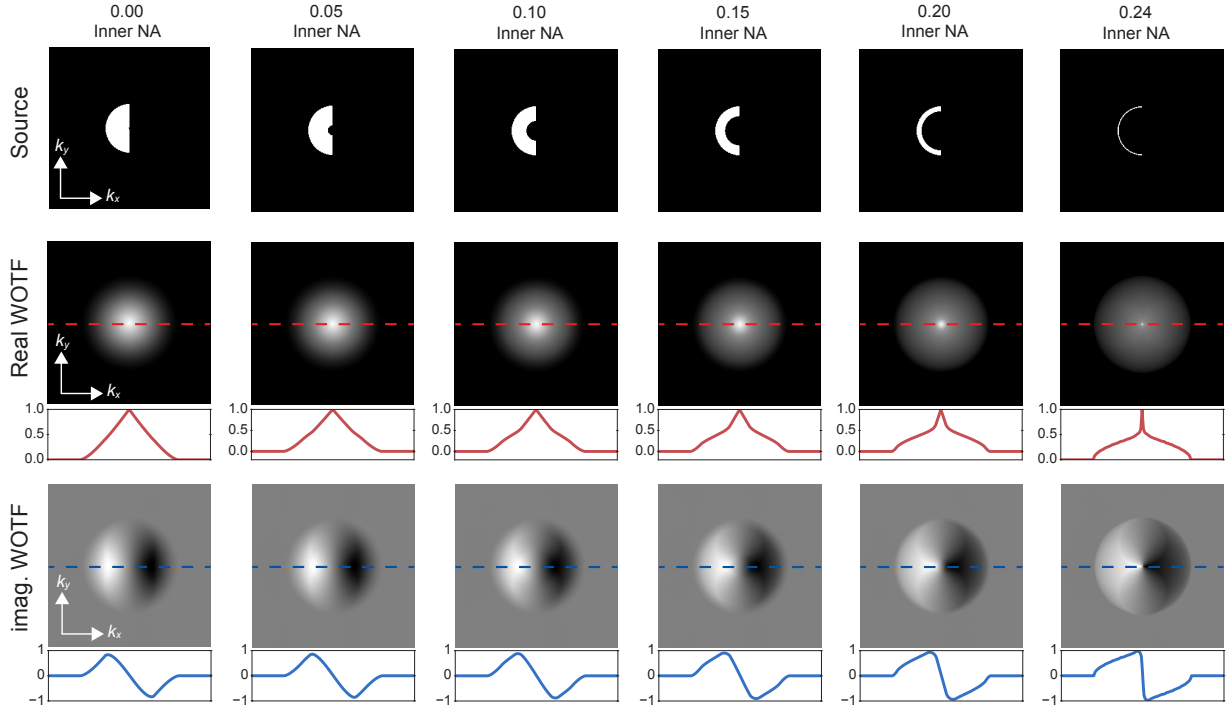
Figure 2.9: DPC transfer functions for amplitude and phase for a range of source cavitation values. As illumination is removed from the center of the source, the overall spectrum becomes more normalized, leading to higher relative values across the mid-range spatial frequencies. However, signal is simultaneously reduced due to lower light throughput. The trade-off between better conditioning and light throughput depends on the illumination power per pixel (or per led).

explore optimal source designs using non-linear recovery methods (such as an unrolled convolutional network) [78]) in terms of root-mean-square error (RMSE) - these techniques are non-linear and object dependent, but are compatible with non-linear regluarizers, offering better performance for a specific class of samples with slightly less generality. Our proposed method for optimizing DPC source patterns assumes nothing about the spectrum of the sample, but does not take advantage of non-linear regularization methods.

In the presence of measurement noise (valid for most practical systems) the multi-image DPC forward model described in Eq. 2.6 can be modeled as having an additional additive noise term $\boldsymbol{\eta}$:

$$\begin{bmatrix} \tilde{\boldsymbol{y}}_1 \\ \tilde{\boldsymbol{y}}_2 \\ \tilde{\boldsymbol{y}}_3 \end{bmatrix} = \begin{bmatrix} \tilde{\boldsymbol{H}}_{\mu,1} & \tilde{\boldsymbol{H}}_{\phi,1} \\ \tilde{\boldsymbol{H}}_{\mu,2} & \tilde{\boldsymbol{H}}_{\phi,2} \\ \tilde{\boldsymbol{H}}_{\mu,3} & \tilde{\boldsymbol{H}}_{\phi,3} \end{bmatrix} \times \begin{bmatrix} \tilde{\boldsymbol{\mu}} \\ \tilde{\boldsymbol{\phi}} \end{bmatrix} + \boldsymbol{\eta} \tag{2.16}$$

Source Cavitation

Figure 2.10: Expected SNR of DPC reconstructions as a function of source cavitation. Source patterns with no cavitation generally provide nearly-optimal SNR, although a source with cavitation NA of approximately $NA2$ is optimal for high illuminance values.

Here $\boldsymbol{\eta}$ is modeled as white zero-mean Gaussian noise with variance σ_η^2. The goal of the DPC linearization is to facilitate the inversion of the block WOTF matrix; when $\boldsymbol{\eta}$ is included in the forward model, this inversion will amplify η based on the singular values of the block matrix $\mathbf{H}$, which are defined by the illumination source $\tilde{\boldsymbol{S}}$ and system pupil $\tilde{\boldsymbol{P}}$. Here, we explore the design of DPC source patterns $\tilde{S}$ in the presence of additive noise, in terms of reconstruction SNR. The use of a programmable light source such as a LED illuminator enables the easy and efficient programming of $\tilde{\boldsymbol{S}}$ simply by changing illumination pattern electronically, making the design space of each source quite large. In this analysis, we explore the use of continuous sources as an approximation of discrete LED sources, as in previous works [49, 79].

The noise amplification of any linear system can be related to the singular values of the forward model by the deconvolution noise factor (DNF, see Chapter 1, Section 1.4). In the DPC forward model described in Eq. 2.16, these singular values may be found by the following relationship [80]:

$$\boldsymbol{\sigma} = \left| \frac{tr(\mathbf{A}^H\mathbf{A})}{2} \pm \sqrt{\frac{tr(\mathbf{A}^H\mathbf{A})^2}{4} - det(\mathbf{A}^H\mathbf{A})} \right| \qquad (2.17)$$

Where $\pm$ indicates that both the positive and negative combinations should be included in singular value calculations, $tr(\cdot)$ is the trace and $\det(\cdot)$ is the determinant. Using Eq. 1.12 to calculate the deconvolution noise factor (DNF), the noise amplification by the DPC inversion process may be estimated using these singular values. Practically, it is necessary to evaluate the DNF within a range of support in the frequency domain, corresponding to the optical bandwidth of the optical system. In addition, the inherent structure of the phase WOTF (H_ϕ) provides very little support near the DC term, which places a lower bound on the support region of the DNF calculation. In this analysis, we evaluate the DNF at all frequencies between 3% and 97 % of the optical bandwidth ($2NA_{objective}$), which corresponds to the region where the phase WOTF has more than 5% of it's maximum value for all angles of a half-circle source and circular pupil.

With this analysis framework, explore the use of a cavitated DPC source in order to improve the conditioning. The cavitated ("C"-shaped) source was previously proposed [49] as an alternative to a filled source ("D"-shaped) due to better conditioning of the deconvolution, particularly for low-frequencies of phase. This poor conditioning leads to low-frequency artifacts that often lead to a "halo"-like effect around the sample. Figure 2.9 illustrates the improvement in the phase transfer function at low frequencies for sources with high cavitation. These sources have a overall spectrum which is more flat, corresponding to lower expected noise amplification by the deconvolution process. However, we also expect the additional illumination power provided by center LEDs will also improve SNR, especially in the presence of signal-independent noise (readout or pattern noise). Therefore, we analyze our system across a two-dimensional space of illumination power (defined by illuminance (Lux) at the sample) and source cavitation (defined by the cavitation NA). The goal of this analysis is to determine the optimal cavitation (cavitation NA), given a particular LED source power.

To characterize the relative benefit to using a cavitated source, we simulated measurements performed under various illumination power values (illuminance), which correspond to different measurement SNR values, and evaluated the SNR in the presence of both signal-dependent (shot) noise as well as signal-independent (readout or fixed-pattern) noise sources corresponding to common parameters (The full list of system parameters is provided in Appendix 7.1, Section 7.1). These illuminance values can be compared to measured source illuminance for LEDs used in this work (Appendix 7.1, Section 7.1). While most DPC sources are discrete, most DPC inversion implementations model the source as continuous to avoid (or average out) LED position calibration error. In this analysis, we adopt the same convention so that our analysis is smooth in terms of cavitation NA.

The results of this analysis are shown in Fig. 2.4. For per-source LED power of less than 1000 lux, a D-shaped source is optimal due to the additional signal which is provided by the inner LEDs. For brighter sources, an inner NA corresponding to approximately 40-50% of the objective NA is optimal in terms of SNR, although the SNR improvement is generally less than 5% compared to a filled-in (non-cavitated) source.

In addition to source cavitation, we can use the same methods to explore other DPC
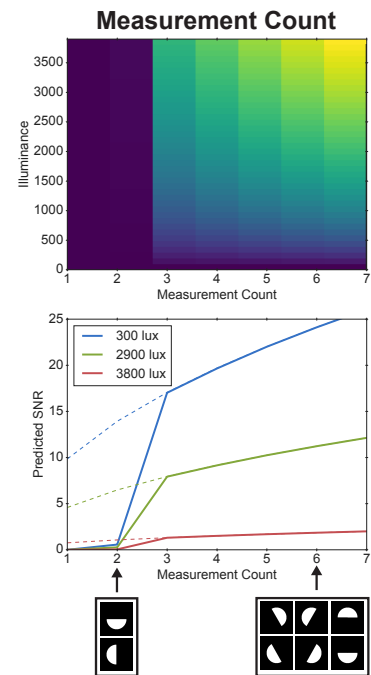
Figure 2.11: Expected SNR of DPC reconstructions as a function of measurement count (equidistant angles). SNR increases with the square root of measurement count for more than three DPC measurements, which is the equivalent of averaging measurements under the same noise conditions. For two or less measurements, expected SNR decreases significantly, since there is not enough information to disambiguate phase and amplitude for all frequencies.

system parameters, including measurement count. The original DPC implementation used 4 measurements [49], but it was later shown that three measurements is sufficient for well-conditioned recovery [81]. To explore this space, we compared using different DPC measurement counts across a range of illumination powers (Fig. 2.4). This analysis confirms that reconstructions will fail for less than three measurements, and that using greater than 3 measurements provides improvements which are characteristic of signal averaging ($SNR \propto \sqrt{N}$). This analysis confirms these previous claims and motivates using three DPC measurements in most circumstances (unless signal averaging is needed to improve SNR).

This section provides two examples of optimizing DPC acquisitions using a noise-based analysis. In general, these data indicate that a D-shaped source is still nearly-optimal in most circumstances, with a C-shaped source providing only marginal SNR improvement when using a powerful source. These data also motivate using 3 DPC measurements instead of 4 (spaced equally in angle). Future work could explore the optimal angular distribution of DPC sources, as well as the relative benefit of using diverse DPC sources such as monopole, dipole, or other non-semicircular patterns.

2.5 Summary

Phase imaging enables label-free imaging of samples by revealing variations in refractive index which are naturally present in nearly all biological samples. In this chapter, we described a method for performing quantitative phase imaging using coded illumination (differential phase contrast), which can be implemented easily on a conventional optical microscope. DPC makes use of the weak-object approximation, which linearizes the image formation process, but is less valid for samples having a strong phase gradient (steep or sharp features). Conventional DPC requires 3 or 4 images to recover the full optical field from intensity measurements; Here, we have proposed novel method for performing DPC using only single measurement using a wavelength-multiplexing technique. This method can be implemented using a simple 3D-printed insert (costing less than $30 to fabricate), or with a programmable light source such as a LED array. While our method does not match the SNR of conventional DPC using 4 images, it enables significantly higher acquisition frame rates, enabling quantitative phase imaging of fast-moving samples such as *C. elegans* or cells moving through a microfluidic channel. Our method, which we call *cDPC*, also enables the synthesis of conventional phase images, like DIC or phase contrast, since quantitative phase is more general than these qualitative phase measurements. Finally, we proposed a framework for performing DPC source optimization in terms of SNR and explore two areas of source design which reveal optimal illumination patterns across a range of source power. Together, these contributions aim to bring quantitative phase imaging to a wider audience through their simple and inexpensive implementations.

Chapter 3

High-Throughput Imaging using Coded Illumination

High-throughput wide-field microscopy enables the collection of large amounts of image data at high-speed, using optimized hardware and computational techniques to push system throughput beyond conventional limits. These systems play a critical role in many fields, including drug-discovery [82–84], functional protein analysis [85,86] and neuropathology [87–89], enabling the rapid acquisition of large volumes of data. Many of these imaging methods involve both resolving and localizing biological structures using fluorescence imaging, which uses proteins conjugated with a fluorescent dye to tag a particular binding site of the sample. When these proteins are illuminated with a beam of light within their excitation spectra, they emit light at a longer wavelength (lower energy). This emitted light can be isolated using spectral filters which remove the excitation wavelength. The collected signal contains only localized proteins within the sample, but has low signal compared compared to brightfield and phase images due to multiple spectral filters and the efficiency of the fluorophores used.

In all wide-field microscopes, the choice of objective lens defines both the resolution and field of view (FOV) of the system, requiring the user to allocate optical throughput to either high-resolution features or a wide FOV (but not both). Addressing this trade-off has been the subject of a large number of computational imaging techniques [8, 10–13, 90], which have, through various mechanisms, demonstrated the ability to enhance the resolution of an imaging system beyond the wide-field diffraction limit while maintaining the same FOV. Similarly, system throughput may also be enhanced by increasing the FOV directly through mechanical scanning and image-stitching while maintaining the resolution of the imaging optics. This direct approach has been widely employed for commercial slide-scanning systems [91], which, when coupled with state-of-the-art analysis tools such as CellProfiler [92], have enabled statistical analysis of the cellular micro-environment at larger scales than ever before.

Despite their wide adoption for a large variety of imaging tasks, the performance of slide-scanning systems is often limited by the mechanical parameters of the motion stage rather than optical parameters of the camera. This leads to lower information throughput

than comparable computational imaging systems which employ coded illumination. The information throughput of an imaging system can be quantified by the space-bandwidth product (SBP), which is the dimensionless product of the spatial coverage (FOV) and Fourier coverage (resolution) of a system [93], as well as the Space-bandwidth rate (SBR), which is SBP per unit time. Improving the SBP and SBR has been the subject of several seminal works in the field of computational imaging, including structured illumination [8], localization microscopy [10, 11], and both conventional [12] and Fourier [13, 90, 94] ptychography. While these methods are diverse in their approaches and application spaces, they share a common theme of acquiring multiple wide-field measurements under diverse imaging conditions to improve the throughput of an imaging system.

Quantifying the SBR of high-throughput imaging systems reveals bottlenecks in their acquisition strategy. For example, conventional slide-scanning systems are often SBR-limited by the time required for a motion stage to move between scan positions and stabilize. These mechanical motions can lead to long acquisition times, especially when imaging very large samples such as coronal sections of the human brain [95] at cellular resolution. Conversely, a Fourier-domain super-resolution technique such as Fourier ptychography only requires electronic scanning of LED illumination, so is more likely to be SBR-limited by photon counts or camera readout, since the time to change LED patterns is on the order of microseconds. However, most super-resolution methods have a fixed FOV, requiring additional mechanical scanning to capture extended samples.

Conventional slide-scanning microscopes employ one of two imaging strategies. The first, commonly referred to as "stop-and-stare" involves moving the sample to each scan position serially, halting the stage motion before each exposure and resuming motion only after the exposure has finished. While this method produces high-quality images due to long exposures, it is slow due to the time required to stop and start motion between exposures. A second approach, often referred to as "strobed illumination", involves illuminating the sample with a very short, bright pulse as it moves continuously, such that the motion blur which would otherwise be introduced by an extended pulse is avoided. While fast, acquisitions performed under strobed illumination will generally produce images with much lower SNR than a comparable stop-and-stare acquisition due to short pulse times (often on the order of micro-seconds). The choice between these two acquisition strategies requires user to trade SNR for acquisition rate, often in ways which make large-scale imagery impractical due to extremely long acquisition times.

In the following sections, we will describe a technique for improving the speed of real-space scanning techniques by multiplexing using motion deblurring and coded illumination. Multiplexing was previous applied to Fourier ptychography [90] by illuminating LEDs from different angles during each frame, reducing acquisition times and increasing measurement SNR. This work presents a similar technique, where LED intensities are coded in time as the sample is moved continuously, and the resulting image is deconvolved using the known illumination sequence to recover the static image. Further, we show situations where our method is sub-optimal compared to the current state-of-the-art techniques in both theory

and experiment[1].

3.1 Quantifying Throughput in Microscopy Systems

The information throughput of an imaging system can be quantified by the space-bandwidth product (SBP), which is the dimensionless product of the spatial coverage (FOV) and Fourier coverage (resolution) of a system [93], as well as the space-bandwidth rate (SBR), which is SBP per unit time. Improving the SBP and SBR has been the subject of several seminal works in the field of computational imaging, including structured illumination [8], localization microscopy [10,11], and both conventional [12] and Fourier [13,90,94] ptychography. While these methods are diverse in their approaches and application spaces, they share a common theme of acquiring multiple wide-field measurements under diverse imaging conditions to improve the throughput of an imaging system.

Evaluating the space-bandwidth product in conventional imaging systems involves comparing the minimum resolution enabled by the optical system with the aberration-free FOV shape. If a camera is perfectly matched to the resolution and FOV supported by the imaging optics, this is equivalent to counting the number of pixels on the camera. The nominal optical FOV in a microscope is normally set based on the quality of the optical designs used in the objective lens; calculating this is beyond the scope of this work (although Chapter 5, Section 5.2 illustrates the graduation degradation of resolution across the FOV). Practically, the camera is what defines the FOV, as most objectives are designed to have small aberrations within the dimensions of most imaging sensors. For incoherent imaging, the minimum resolution is readily calculated from the system numerical aperture NA using the Rayleigh criterion [96]:

$$\Delta x = \frac{1.22\lambda}{2NA} \tag{3.1}$$

where λ is the illumination wavelength. This resolution limit defines the maximum spatial frequency in the optical system which is related to the numerical aperture of the camera by:

$$k_{max} = \frac{NA}{\lambda} \tag{3.2}$$

The camera sampling must be chosen such that this maximum spatial frequency is sampled as the Nyquist rate, which is $2k_{max}$. Practically, this is accomplished by introducing magnification to the optical system to shrink the pixel shape to match the required spatial frequency. Mathematically, the goal is to determine a magnification factor M which satisfies the following bound, given a camera pixel size Δ:

$$\frac{\Delta}{M} \leq \frac{1}{2k_{max}} \tag{3.3}$$

---

[1]This work was performed in collaboration with Sarah Dean and Benjamin Recht (ADEPT Lab, EECS, UC Berkeley).

Figure 3.1: Illustration of ideal space-bandwidth product allocation for various common objective lenses as plotted in phase-space. While the area of each rectangle remains constant, the aspect ratio illustrates the allocation of bandwidth between high-resolution and a wide FOV.

For color imaging system, note that the factor Δ must be doubled, since measurements at each color channel occur at twice the distance between each measurement.

Assuming these criteria are met, the total bandwidth of the system is set by the system resolution (Eq. 3.1) and the sensor shape scaled by the magnification. A phase-space representation provides a geometric description of the aspect ratio of the system coverage in the spatial and spectral dimensions as a function of the prescribed resolution and FOV as set by the system magnification. For example, the SBP of a $10\times$ / 0.25NA and a $40\times$ / 1.0NA objective may be the same, but the former will allocate more throughput to resolution than FOV, while the latter will provide a wide FOV at lower resolution. This well-known trade-off is illustrated in Figure 3.1.

The illustration in Fig. 3.1 reveals the regions which are most amenable to scanning for common objective types. For example, when using an objective with a high magnification and NA, it would be more optimal to scan a sample in the real domain, since each measurement will already have significant spectral coverage (resolution). Conversely, low-magnification optics are more amenable to scanning in the frequency domain to build resolution, since these systems already have a very large FOV. This intuition explains the practical benefits of performing Fourier ptychography using a low-magnification objective [13], since scanning is performed in the frequency domain - there would be little practical benefit to performing
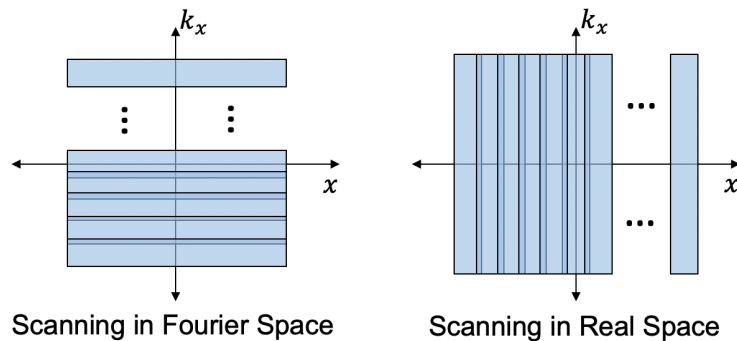
Figure 3.2: Comparison of spectral (Fourier-domain) and spatial (real-domain) scanning techniques in high-throughput systems in one dimension (x). A phase-space representation provides a visual representation of the scanning geometry of each approach.

Fourier ptychography using a high-magnification objective and the maximum illumination NA would be roughly the same as the imaging NA, making any resolution improvement marginal compared to brightfield or DPC phase imaging.

On the other hand, mechanical scanning provides a means for augmenting the SBP in real-space, making it amenable for high-magnification imaging systems. A basic example of this technique is slide-scanning systems, which move a sample while serially capturing images between each movement. A second example is conventional ptychography [12], which probes a sample in real-space while imaging the diffracted field in the Fourier domain (usually using free-space propagation onto a detector array). The choice of which scanning method to use is a general question and is beyond the scope of this work but would certainly depend on the relative acquisition bottlenecks present in each system and would be best compared using the SBR. Figure 3.2 illustrates these two scanning techniques.

A final parameter of interest is the signal-to-noise ratio (SNR) of the recovered images. In imaging, SNR is generally defined as ratio of the mean signal to the standard deviation of the background (As described in Eq. 1.7). The absolute SNR requirements vary based on applications; in many cases, a $SNR > 10$ is considered a threshold where images have good SNR, although images below this threshold are also useful for many applications. The SNR is inherently tied to space-bandwidth rate because SNR will always increase for larger exposure times, which would conversely lead to a lower SBR due to longer dwell times for each measurement. Therefore, it is important to consider a minimum SNR as an additional constraint on the speed of the optical system. The remaining sections of this chapter present a method for increasing the SNR of measurements while maintaining acquisition speed, which could also be interpreted as increasing the maximum SBR to maintain a necessary minimum signal-to-noise ratio.
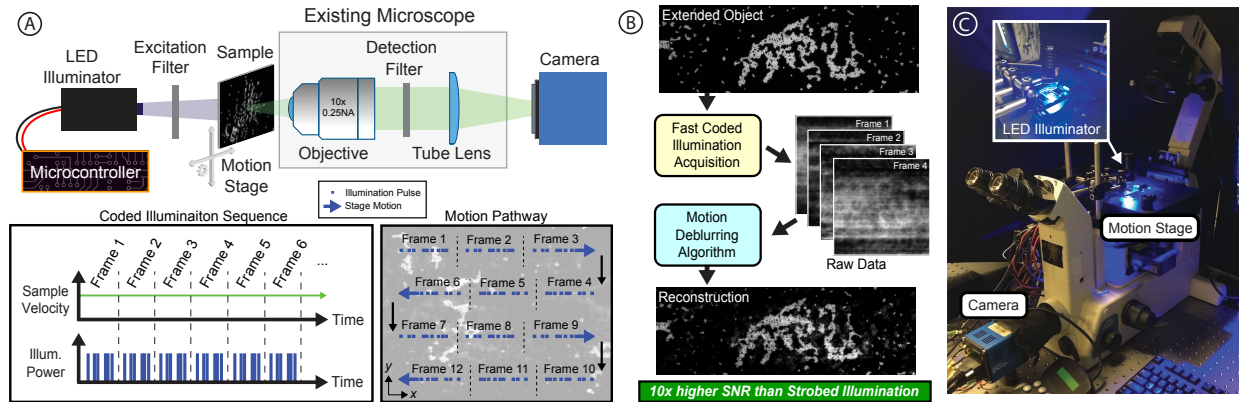
Figure 3.3: High throughput imaging system with coded illumination. A.) Our system consists of an inverted optical fluorescence microscope with a 2-axis motion stage, illuminated using a programmable LED illumination source. In our proposed method, the sample is illuminated with many discrete pulses while being scanned at constant velocity. B.) Comparison between conventional high-throughput imaging techniques (stop-and-stare, strobed illumination) and our proposed coded illumination technique. Coded illumination provides a trade-off between SNR and acquisition speed, particularly for low-light situations such as fluorescence imaging. C.) Image of our system, which is a Nikon TE300 microscope configured with a Prior motion stage and LED illuminator.

3.2 High-Throughput Microscopy with Motion Deblurring

In this section, we propose a novel computational imaging technique which employs a coded-illumination acquisition and deconvolution to meet the demands of high-throughput applications which require a minimum SNR. Our method involves illumination the sample with multiple pulses during each acquisition in order to improve the SNR compared to strobed illumination, while maintaining a high acquisition rate. With knowledge of this pulse sequence and motion trajectory, the blurred image can be used to perform a reconstruction of the static image, with higher SNR than a comparable strobed acquisition. Our method employs a motion-multiplexing technique to enhance the measurement SNR of our system by illumination with a sequence of short, bright pulses every frame. The captured images contain motion-blur artifacts, which must be removed computationally through a motion deblurring algorithm [97]. The overall gain in SNR is proportional to the number of pulses as well as the conditioning of the motion deblurring process [41], necessitating careful design of pulse sequences to produce the highest-quality image. In the following sections, we detail the joint design of the hardware and algorithms to enable high-speed gigapixel-scale fluorescence imaging with SBR > 50 million pixels per second, and compare the performance of
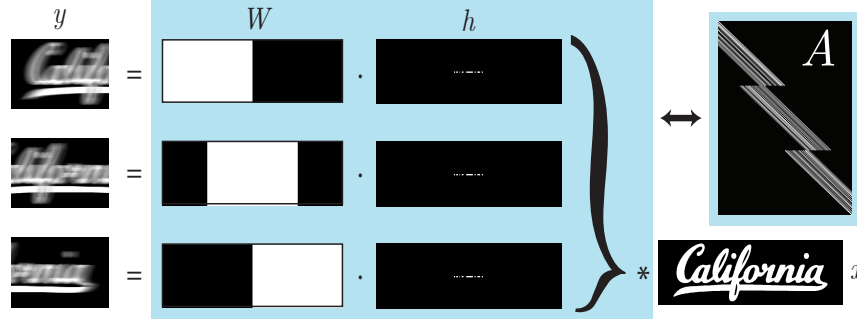
Figure 3.4: Multi-frame forward model. Here $[*]$ represents 2D convolution and $[\cdot]$ represents the element-wise product. Top right inset is a an example of a full 1D multi-frame convolution matrix.

our proposed framework against traditional methods and provide an experimental demonstration of situations where coded illumination is both optimal (e.g. fluorescence imaging) and sub-optimal (e.g. brightfield imaging) as a function of common system parameters such as illumination power and camera noise levels. Thus, our contribution is both the proposal of a new high-throughput imaging technique as well as an analysis of when it is practically useful for relevant applications.

Methods

Motion Blur Forward Models

With knowledge of both sample trajectory and illumination sequence, the physical process of capturing a single measurement of a sample illuminated by a sequence of pulses while in motion can be mathematically described as a convolution:

$$\boldsymbol{y} = \boldsymbol{h} * \boldsymbol{x} + \eta, \tag{3.4}$$

where $\boldsymbol{y}$ is the blurred measurement, $\boldsymbol{x}$ is the static object to be recovered, η represents additive noise, $*$ represents 2D convolution, and $\boldsymbol{h}$ represents the blur kernel, the mapping of the temporal illumination intensity to positions in the imaging coordinate system using kinematic motion equations. With appropriate padding, this convolution can be computed efficiently using the Fourier Transform, and can be similarly inverted using FFT-based deconvolution.

To use coded illumination for large FOV imaging, we propose an extension of the single-frame forward model to the multi-frame case. Mathematically, we model the serial acquisition of motion-blurred frames as the vertical concatenation of many single-frame forward models, which are individually convolutional (Fig. 3.4). In our model, each captured image has an associated blur operator $\mathbf{B}_j$ defined by each blur kernel $\boldsymbol{h}_j$ such that $\boldsymbol{h}_j * \boldsymbol{x} = \mathbf{B}_j \boldsymbol{x}$. Additionally, we prepend each convolutional sub-unit with a crop operator $\mathbf{W}_j$, which selects
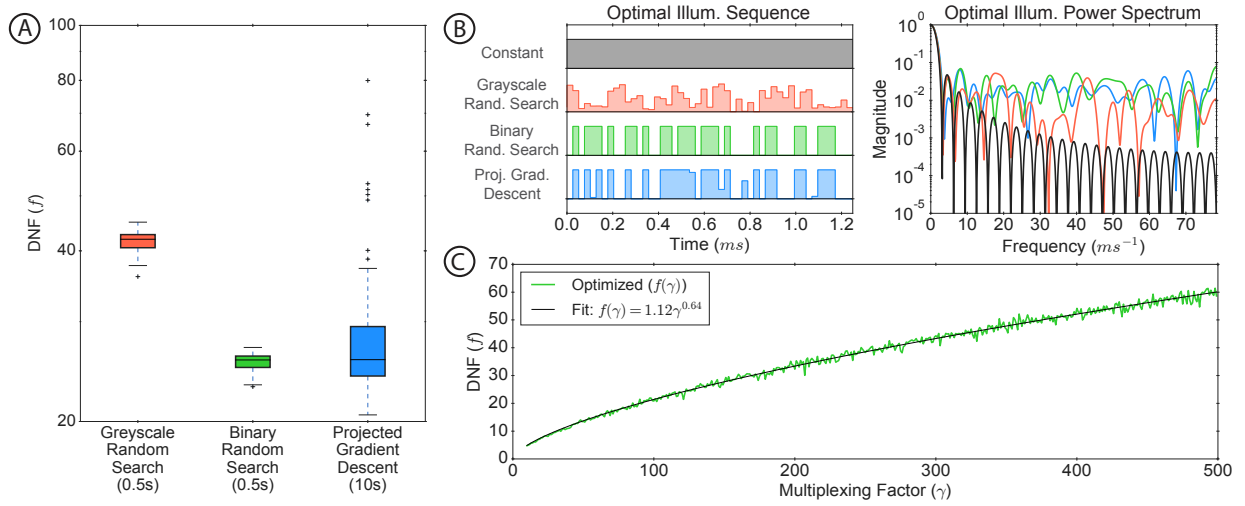
Figure 3.5: A.) Optimal DNF values for illumination patterns generated by different optimization methods show that random search over binary patterns is of comparable effectiveness to the projected gradient descent method. B.) Example illumination sequences show differing power spectra between constant illumination and patterns generated by random search over greyscale, random search over binary, and projected gradient descent. C.) Optimized DNF (binary random search method) is plotted over different values of γ. A power law fit of $f(\gamma) = 1.12\gamma^{0.641}$ has standard error of 1.075.

an area of the object based on the field-of-view of the camera. Together, these operators encode both spatial coverage and the local blurring of each measurement, and are concatenated to form the complete multi-frame forward operator:

$$
\begin{bmatrix} \boldsymbol{y}_1 \\ \vdots \\ \boldsymbol{y}_n \end{bmatrix} = \begin{bmatrix} \mathbf{W}_1\mathbf{B}_1 \\ \vdots \\ \mathbf{W}_n\mathbf{B}_n \end{bmatrix} \boldsymbol{x} + \eta = \mathbf{A}\boldsymbol{x} + \eta \tag{3.5}
$$

This forward operator $\mathbf{A}$ is no longer generally convolutional, taking the form of a spatially variant convolution based on the coverage of each individual $\mathbf{W}_j$. A one-dimensional illustration of the multi-frame smear matrix $\mathbf{A}$ is displayed in Fig. 3.4. With the addition of appropriate windowing logic, the forward operation and its adjoint can be computed efficiently for use in iterative reconstruction.

Reconstruction Algorithm

To invert our forward model ((3.5)), we employ Nesterov accelerated gradient descent [98] algorithm to minimize the error between our measurements $\boldsymbol{y}$ and estimated object $\widehat{\boldsymbol{x}}$ passed through forward model $\mathbf{A}$ in the ℓ_2 metric. For the purposes of this dissertation, we seek to

minimize an unregularized cost function, $O(\boldsymbol{x}) = \frac{1}{2}\|\mathbf{A}\boldsymbol{x} - \boldsymbol{y}\|_2^2$ using the following update equations at each k^{th} iteration:

$$
\begin{aligned}
\boldsymbol{z}^{k+1} &= \boldsymbol{x}^k - \alpha\mathbf{A}^{\mathsf{H}}(\mathbf{A}\boldsymbol{x}^k - \boldsymbol{y}) \\
\boldsymbol{x}^{k+1} &= \beta^k\boldsymbol{z}^k + (1 - \beta^k)\boldsymbol{z}^{k+1}
\end{aligned}
\tag{3.6}
$$

Here, α is a fixed step size and β^k is set each iteration by the Nesterov update equation.

In all reconstructions we performed 30 iterations of (3.6), which we found was a favorable balance of reconstruction quality and reconstruction time. While adding a regularization term such to enforce signal priors could improve reconstruction quality (and was previously analyzed in the context of motion deblurring [99]), we chose not to incorporate regularization to provide a more fair and straightforward comparison between the proposed coded illumination acquisition and conventional strobed and stop-and-stare acquisitions. It should be noted that running our algorithm for a pre-defined number of iterations may provide some regularization from early-stopping [100].

Reconstructions were performed in Python using the Arrayfire GPU computation library [101]. Due to the structure of our scans (Figure 3.3A), the reconstruction algorithm is highly parallelizable – we separate our reconstruction into strips along the major translation axis to perform 1D convolutions and stitch these strips together after computation. Using this parallelization, we are able to reconstruct an approximately 1 gigapixel FOV in approximately 2 minutes. Further details of the reconstruction implementation are discussion in Section 3.2.

Reconstruction SNR

To minimize the error between the reconstructed $\widehat{\boldsymbol{x}}$ and the true object $\boldsymbol{x}$, the blur kernels and scanning pattern should be chosen such that the noise is minimally amplified by the inversion process. This amplification is controlled by the singular values of the forward model $\mathbf{A}$.

In the case of single-frame blur, the singular values are controlled by the length and coding of the blur kernel $\boldsymbol{h}$. While early works [72, 97] used a non-linear optimization routine (i.e. the `fmincon` function in MATLAB (Mathworks)) to minimize the condition number of $\mathbf{A}$, more recent work proposed maximizing the reconstruction signal-to-noise ratio directly, using camera noise parameters, source brightness, and the well-posedness of the deconvolution [40, 41].

We extend this work to the multiframe setting. In our analysis, we adopt the convention of imaging SNR, which is the ratio of the mean signal to the signal variance (due to photon shot noise, camera readout noise, fixed pattern noise, and other camera-dependent factors). Under a simplified model, the noise variance will be the addition in quadrature of the camera read noise variance σ_r^2 plus a signal dependent term $\bar{s}$:

Figure 3.6: (1.18 gigapixel, 23mm × 20mm full-field reconstruction of $4.7\mu m$ fluorescent microspheres. Inset scale bars are $50\mu m$. While coded reconstructions have lower SNR than stop-and-stare (S&S) measurements, measurements acquired using coded illumination (Coded) were more than $5.5\times$ faster while maintaining enough signal to distinguish individual microspheres compared to strobed illumination (Strobed)

$$SNR = \frac{\bar{s}}{\sqrt{\bar{s} + \sigma_r^2}} \tag{3.7}$$

Here, we ignore exposure-dependent noise parameters such as dark current and fixed-pattern noise, since these are usually small for short exposure times relative to read noise σ^2. Note that the denominator of (3.7) is equivalent to the standard deviation of η in (3.4). This definition is valid for both strobed illumination and stop-and-stare acquisitions. For an acquisition performed using coded illumination, it is necessary to consider the noise amplification that results from inverting the forward model. As defined in previous work [41], this amplification is controlled by the deconvolution noise factor (DNF), which for single-frame blurring is defined as:

$$f = \sqrt{\frac{1}{m} \sum_{i=0}^{m} \frac{\max_j |\tilde{\boldsymbol{h}}|_j^2}{|\tilde{\boldsymbol{h}}|_i^2}} \tag{3.8}$$

where f is the DNF, m is the size of the blur kernel $\boldsymbol{h}$, and $\tilde{\boldsymbol{h}}$ represents the Fourier transform of $\boldsymbol{h}$.

To write an expression for the singe-frame SNR under coded illumination, we first define a multiplexing factor $\gamma = \sum_i h_i$, the total amount of illumination imparted during exposure.

If $\boldsymbol{h}$ is constrained to be binary, γ will be equal to the number of pulses in $\boldsymbol{h}$. (3.7) can the be modified using both f and γ:

$$SNR = \frac{\gamma \bar{s}_0}{f \sqrt{\gamma \bar{s}_0 + \sigma_r^2}} \tag{3.9}$$

In (3.9), $\bar{s}_0$ is the mean signal imparted by a single illumination pulse. We include the complete derivation of (3.9) in Appendix 7.3 for completeness.

The derivation of this quantity relies on the convolutional nature of single-frame motion blur. As shown in Fig. 3.4, each column of this matrix defines a localized convolutional kernel; near the horizontal boundaries between the different frames, the vertical combinations of kernels are non-trivial. While methods for analyzing spatially variant convolution matrices exist [102], we focus our analysis on a practical simplifying assumption: blur path and illumination patterns are fixed to be the same across all frames, i.e. $\boldsymbol{h}_j = \boldsymbol{h}$ for all j.

In this special case, the resulting SNR of the proposed multi-frame model is governed by both the power spectrum of the blur kernel and the spatial coverage of the crop operators. We define c_i to be the *coverage* at pixel i, i.e. the number of times pixel i is included in the windows $\{\mathbf{W}_1, ..., \mathbf{W}_n\}$. The SNR for a multi-frame acquisition with coded illumination is bounded as

$$SNR \geq \sqrt{\min_i c_i} \cdot \frac{\gamma \bar{s}_0}{f \cdot \sqrt{\gamma \bar{s}_0 + \sigma_r^2}}$$

Thus, the SNR improves by at least a factor of square root the minimum coverage compared with the single-frame coded illumination case in (3.9). The derivation is included in Appendix 7.3.

Notably, this bound decouples the spatial coverage from the spectral quality of the blur, allowing for blur kernel optimization independent of the multi-frame aspect. A good motion path ensures even spatial coverage through c_i while a good illumination sequence ensures spectral coverage through traditional single-frame methods. In what follows, we focus system design on the maximization of this decoupled lower bound.

The decision to use the same blur kernel in every frame has several practical implications as well: the micro-controller memory would be saturated when storing more than one thousand kernels and post-processing registration is much easier since all measurements have been distorted by the same blurring operator. Additionally, the fact that this reduction requires the motion to be along single motion axis is not limiting, since in practice horizontal strips are reconstructed independently to accommodate computer memory.

Illumination Optimization

Previous work [41, 97] showed that reconstructions performed using constant (non-coded) illumination will have very poor quality (in terms of SNR) compared to using optimized pulse sequences or a short, single pulse. Here, we explore several approaches for generating illumination pulse sequences which maximize the reconstruction SNR ((3.9)). We first

consider the problem of minimizing the DNF (f) with respect to $\boldsymbol{h}$:

$$f(\gamma) := \min_{\boldsymbol{h}} \ f$$
$$s.t. \ \ 0 \le h_i \le 1 \ \forall \ i, \quad \sum_i h_i = \gamma \,, \tag{3.10}$$

where inequality constraint on $\boldsymbol{h}$ represents the finite optical throughput of the system. This optimization problem is non-convex, and it resembles those used in previous work [41, 72, 97] (Our multiplexing factor γ is related to the kernel length N and throughput coefficient β in previous work by $\gamma = N\beta$). This definition enables a layered approach to maximize the SNR: after solving (3.10) for each multiplexing factor, it is possible to find the one which optimizes (3.9) in the context of camera noise parameters.

To simplify the optimization task, the positions encoded in $\boldsymbol{h}$ may be restricted a priori, e.g. to a centered horizontal line with fixed length as in Fig. 3.4. In the following, we constrain the positions to a straight line with length $N = 2\gamma$. This provides a sufficiently large sample space for kernel optimization and is supported as optimal by analysis in [41].

We consider several methods for solving the DNF optimization in (3.10): random search over greyscale kernels, random search over binary kernels, and a projected gradient descent (PGD) approach. Our random search approach is simple: a fixed number of candidate kernels are randomly generated, and the one with the lowest DNF is chosen. The grayscale candidates were generated by sampling uniform random variables, while the binary was generated by sampling indices without replacement. In our PGD approach, the kernel optimization problem ((3.10) is reformulated as the minimization of a smooth objective $g(\boldsymbol{h})$ subject to convex constraints $\mathcal{S}$.

Starting from an initial $\boldsymbol{h}_0$, the update rule includes a gradient step followed by a projection:

$$\tilde{\boldsymbol{h}}^{k+1} = \text{Proj}_{\mathcal{S}}(\boldsymbol{h}^k - \alpha^k \nabla g(\boldsymbol{h}^k)) \,. \tag{3.11}$$

Details of the reformulation and optimization approach are presented in Appendix 7.3.

The box plots in Fig. 3.5A show the distribution of optimization results for 100 trials of each approach, where the random search methods sample 1000 candidates and PGD runs with step-size determined by backtracking line search until convergence from a random binary initialization. Example illumination sequences from each method and their corresponding power spectra are displayed in Fig. 3.5B.

Though the kernels with the lowest DNF were generated through PGD, binary random search results in comparable DNF values and is up to $20\times$ faster than PGD. Further, we note that a random binary search resulted in significantly lower DNF calculations than grayscale random search. A binary restriction also achieves fast illumination updates, since grayscale illumination (as in [72]) would require a pulse-width-modulation (PWM) cycle spread across multiple clock cycles.

Plotting the DNF ($f(\gamma)$) generated through binary random search in Fig. 3.5C reveals a concave curve. Fitting the curve with a power-law, a closed-form approximation for the

DNF as is $f(\gamma) = 1.12\gamma^{0.64}$. This analytic relationship allows for a direct optimization of SNR: substituting any $f(\gamma) \propto \gamma^p$ into (3.9) and differentiating with respect to γ, we can determine the approximately optimal multiplexing factor γ^* as a function of mean strobed signal $\bar{s}_0$ and camera readout noise σ_r^2:

$$\gamma^* = \frac{2 - 2p}{2p - 1} \cdot \frac{\sigma_r^2}{\bar{s}_0} \ . \tag{3.12}$$

We note that for smaller p, i.e. slower DNF growth with γ, the optimal multiplexing factor will be larger. With $p = 0.5$, the expression for SNR in (3.9) only increases with increasing multiplexing factors, meaning that the optimal γ^* would be as large as possible given hardware constraints. We show in Appendix 7.3 the lower bound $f(\gamma) \geq \gamma^{0.5}$ regardless of optimization method or illumination sequence. The experimental $p = 0.64$ accurately reflects the practical relationship, given that we experience limitations from imperfect optimization methods. Part of the observed relationship may come from the increasing difficultly of optimization as the decision space grows with γ.

The expression for optimal γ^* also solidifies the intuition that a higher multiplexing should be used for systems with high noise in order to increase detection SNR, while a lower multiplexing is appropriate for less noisy systems. This result is in agreement with [41], in which it was demonstrated that the choice of multiplexing factor depends on the relative magnitude of the acquisition noise.

System Hardware and Software

Our system is built around an inverted microscope (TE300 Nikon) using a lateral motion stage (Prior, H117). Images were acquired using a SCMOS camera (PCO.edge 5.5, PCO) through hardware triggering, and illumination was provided by a high-power LED (M470L3, Thorlabs) which was controlled by a micro-controller (Teensy 3.2, PJRC). Brightfield measurements were illuminated using one of two sources: a custom LED illuminator with 40 blue-phosphor LEDs (VAOL-3LWY4, VCC), or a single, high-power LED source (Thorlabs M470L3), both modulated using a simple single-transistor circuit through the same micro-controller. The first illuminator was designed to have a broad spectrum for brightield imaging, while the second was intended for fluorescence imaging, having a narrow spectral bandwidth. For this project, we found it was necessary to adopt very simple LED circuitry to avoid electronic speed limitations associated with dimming (PWM) and serial control of LED driver chips. A schematic and layout of our system is shown in Fig. 3.3A and a photograph in Fig. 3.3C.

The micro-controller firmware used in this device was developed as part of a broader open-source LED array firmware project [103]. Images were captured through the python bindings of the Micro-Manager software [104], which were controlled through a Jupyter notebook [105], enabling fast prototyping of both acquisition and reconstruction pipelines in the same application. With the exception of our custom illumination device, everything in our optical system is commercially available in combinations commonly available at many

imaging centers. Therefore, our method is amenable to most traditional optical setups, and can be implemented cheaply through the simple addition of our temporally coded light source and open-source software tools.

All acquisitions in this work were performed using a $10\times$, 0.25NA objective, which had a depth of field of approximately 8.5 μm. Because this depth-of-field was relatively large compared to our sample, we were able to level the sample manually prior to acquisition using adjustment screws present on the motion stage to ensure the sample remained in focus across an approximately $20mm$ movement range. For larger areas or shallower depth of field, we anticipate a more sophisticated leveling technique will be necessary, such as active autofocus [106, 107]. In practice, the focal adjustment process required only a few minutes for our $10\times$ objective, and was stable between samples (requiring infrequent adjustment).

Results

Gigapixel Reconstruction

To demonstrate the performance benefit of our system, we performed a 1.18 gigapixel reconstruction of a microscope slide plated with $4.7\mu m$ polystyrene fluorescent beads (Thermo-Fisher) using the high-throughput coded illumination microscope shown in Fig. 3.3. Multiple images were captured at spatial offsets, enabling a reconstruction with extend field of view (FOV). Each scan pathway consisted of multiple 1D continuous scans structured in a raster-scanning pattern (Fig 3.4B) to enable fast 2D scanning of the sample. For a complete comparison, we acquired stop-and-stare, strobed illumination, and coded illumination datasets serially, and performed image stitching and registration of the three datasets for a direct comparison of image quality. The total acquisition time for the coded reconstructions of this size was 31.6 seconds, while a comparable stop-and-stare acquisition required 210.9 seconds. The computation time for coded-illumination reconstructions using (3.6) with step-size $\alpha = 0.5$ was approximately 30 minutes on a MacBook Pro (Apple) with attached RX580 external GPU (Advanced Micro-Devices), or approximately 2 minutes when parallelized across 18 EC2 p2.xlarge instances with Nvidia Titan GPUs (Amazon Web Services), excluding data transfer to and from our local machine.

Method Comparison

To provide a complete comparison of our coded illumination method with existing high-throughput imaging techniques, we sought to quantify the expected SNR for each method based on relevant system parameters such as source illuminance, camera noise parameters, and desired acquisition frame-rate. In conventional high-throughput imaging it is well-understood that stop-and-stare strategy will provide higher SNR than imaging using strobed illumination but is only feasible for low-frame rates due to mechanical limitations of the motion stage. Therefore, at high frame-rates we restrict our comparison to strobed illumination and coded illumination. Later, we perform a comprehensive analysis of these
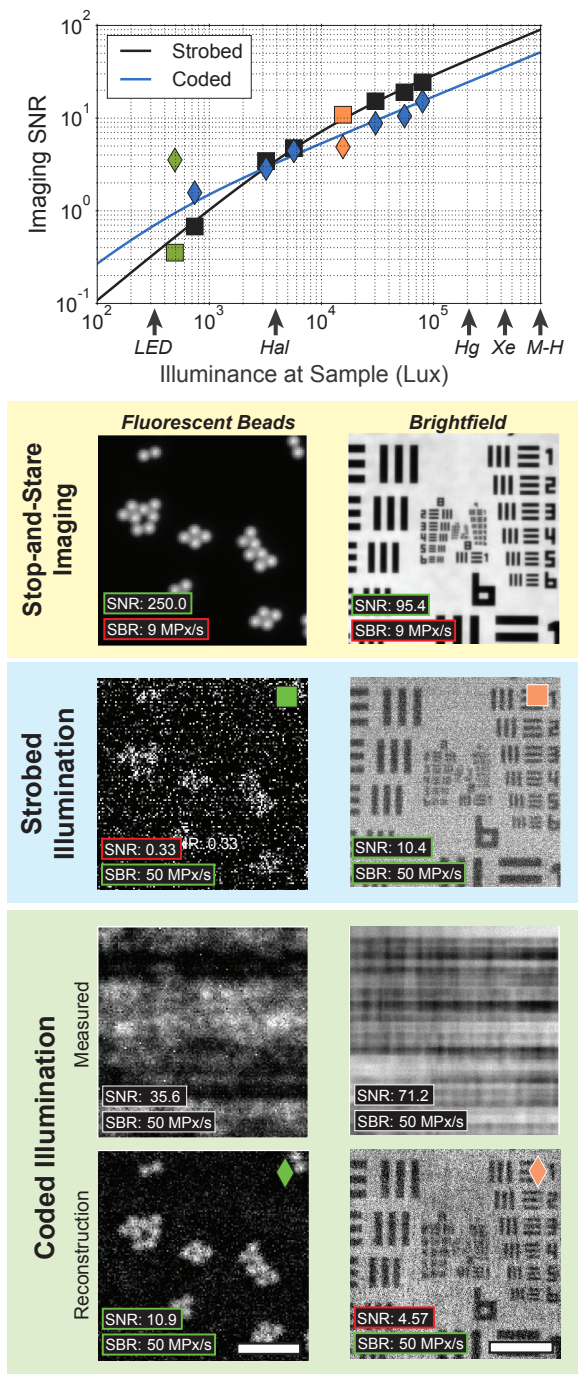
Figure 3.7: (Top) Experimental SNR values for a USAF 1951 resolution illuminated across a range of illumination values under strobed (square) and coded illumination (diamond). Solid lines illustrate predicted SNR based on known system parameters. Experimental SNR values are the average of 3 SNR measurements performed across the field. Green and orange data points represent inset data for fluorescent beads and resolution targets respectively. Characteristic illuminance values for LED sources, Halogen-Tungsten Lamps (Hal.), Mercury Lamps (Hg), Xenon Lamps (Xe), and Metal-Halide Lamps (M-H) are shown for reference. (Bottom) Selection of measurements used to generate the above plot. Scale bar is 25 μm.

trade-offs; Fig. 3.8A provides a visual representation of where stop-and-stare is possible in terms of mechanical parameters of the motion stage.

Using the experimental setup described in Fig 3.3 we compare acquisitions performed under brightfield and fluorescence configurations while sweeping the output power of the LED source. In each case, we measured the illuminance at the camera plane using an optical power meter (Thorlabs). For each image and reconstruction, we computed the average imaging SNR across three different regions of interest to produce an experimental estimate of the overall imaging SNR of the scene. To ensure a fair comparison, we performed no prepossessing on the data except for subtracting a known, constant offset from each measurement which was characterized before acquisitions were performed and verified with the camera datasheet. Reconstructions of measurements made using coded illumination were performed using the reconstruction algorithm described in Section 3.2.

Fig. 3.7 shows that coded illumination can provide up to $10\times$ higher SNR in low-illumination situations, which are most relevant for fluorescence microscopy (illuminance < 1000 lux). For brightfield microscopy, strobed illumination provides a higher SNR without additional computational complexity associated with our coded illumination method. The solid lines in Fig. 3.7 are theoretical predictions of reconstruction SNR for each method based on our system parameters, which are generally in agreement with our experimental data. The methods for computing these curves are described in the following section.

Component Analysis

In Section 3.2, we showed that the choice to use strobed or coded illumination depends heavily on the illumination power of the source; however, other system parameters may also affect this trade-off. In this section, we consider other parameters such as camera noise and motion velocity and include an analysis of when stop-and-stare should be used as opposed to continuous-scan methods (strobed illumination and coded illumination). As a first step, we derive the expected photons per pixel, per second (J) we expect to measure in a transmission microscope, incorporating system magnification (M), numerical aperture (NA), camera pixel size (Δ), mean wavelength ($\bar{\lambda}$, and the photometric look-up table $K(\bar{\lambda})$:

$$J = K(\bar{\lambda})\bar{\lambda}\hbar c \cdot I_{lux} * (NA)^2 * (\frac{\Delta}{M})^2 \qquad (3.13)$$

Here, $\hbar$ is Planck's constant, c is the speed of light, I_{lux} is the source illuminance in lux, and J is the photon flux per pixel-second. Given J, the mean signal $\bar{s}$ is a function of the illumination time t_{illum} and the camera quantum efficiency Q:

$$\bar{s} = JQt_{illum} \qquad (3.14)$$

Substituting (3.14) into (3.9), we can define the expected SNR as a function of these parameters as well as the blur kernel h and camera readout noise σ_r:
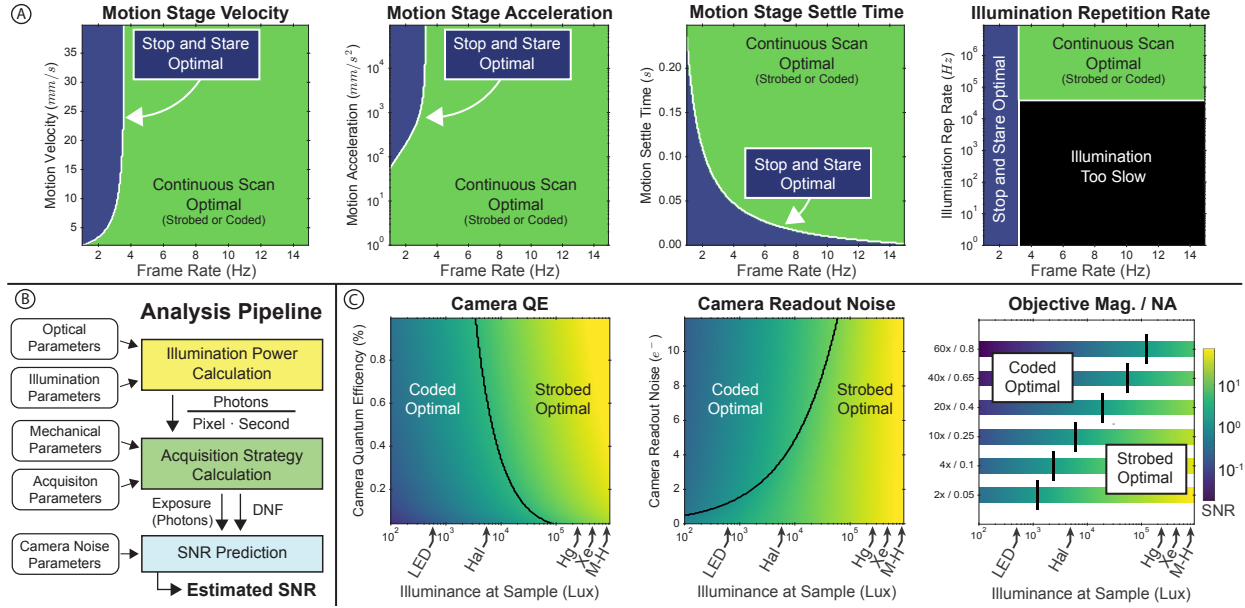
Figure 3.8: Limiting Analysis of Imaging System. A.) Analysis pipeline for predicting SNR from system parameters, including illumination power, mechanical parameters of the motion stage, and camera noise parameters. B.) The stop-and-stare acquisition strategy is only possible for some configurations of mechanical system parameters and frame rates. C.) Different combinations of optical system parameters and system illuminance determine the best possible SNR and whether strobed or coded illumination is preferable. Characteristic illuminance values for LED sources, Halogen-Tungsten Lamps (Hal), Mercury Lamps (Hg), Xenon Lamp (Xe), and Metal-Halide Lamps (M-H) are shown for reference.

$$SNR = \frac{JQt_{illum}}{f\sqrt{JQt_{illum} + \sigma_r^2}} \tag{3.15}$$

(3.15) is used in the analysis of Fig. 3.7 and Fig. 3.8.

The parameters t_{illum} and f are functions which change based on acquisition strategy. For stop-and-stare and strobed acquisitions, we set $f = 1$, since no deconvolution is being performed, while for coded acquisitions f depends on the parameter γ as derived in Section 3.2. Similarly, t_{illum} is set based on acquisition strategy and motion stage parameters. For stop-and-stare illumination, t_{illum} is proportional to the residual time after stage movement, including stage maximum velocity (v_{stage}), stage acceleration (a_{stage}), the field of view of a single frame along the blur axis (FOV), mechanical settle time (t_{stage}), and desired acquisition frame rate r_{frame}:

$$t_{illum} = \frac{1}{r_{frame}} - 2\frac{v_{stage}}{a_{stage}} - \frac{FOV - 0.5 * \frac{v_{stage}^2}{a_{stage}}}{v_{stage}} \; . \tag{3.16}$$

For strobed and coded illumination, the minimum pulse duration t_{illum} is set by the overlap fraction (O) between frames, the number of pixels along the blur direction, N_{px}, and the multiplexing by γ (with strobed encoded by $\gamma = 1$):

$$t_{illum} = \frac{\gamma}{r_{frame}(1 - O)N_{px}} \; . \tag{3.17}$$

In (3.17), we implicitly calculate the velocity as the fastest speed where two frames may overlap with O within a time set by r_{frame}.

Derivations for the above relationships are provided in Appendix 7.3. With these theoretical solutions for t_{illum}, we are able to derive closed-form solutions for expected SNR as a function of acquisition rate with knowledge of our system parameters (Listed in Appendix 7.3), using the pipeline shown in Fig. 3.8B.

Our system analysis is divided into two parts; When stop-and-stare is possible given a desired acquisition frame rate, it will always provide higher SNR than strobed or coded illumination due to high photon counts compared to strobed illumination and no deconvolution noise. Fig. 3.8A analyzes where stop-and-stare is both possible and optimal compared to a continuous acquisition technique as a function of frame rate. If acquisition frame rate is low or limited by other factors (such as sample stability), stop-and-stare will always be optimal in terms of SNR. If a high-frame rate is desired, however, a continuous acquisition strategy is optimal, so long as the illumination repetition rate of the source is fast enough to accommodate sample velocity.

Fig. 3.8C describes the optimal continuous imaging technique as a function of illuminance, imaging objectives, camera readout noise (σ_r, and camera quantum efficiency (QE). Generally speaking, higher illuminance values favor strobed illumination (being shot-noise limited), while lower illuminance values favor coded illumination (being read-noise limited). Conversely, as read noise (σ_r) increases or camera QE decreases, coded illumination becomes more beneficial. Practically, a camera with high read-noise with a low QE will favor coded illumination more strongly (at higher source illuminance) than a high-end camera (such as the PCO.edge 5.5 used in this study), where $\sigma_r \approx 3.7e^-$ and $QE \approx 0.6$). In addition, objectives with a higher magnification and NA will generally favor coded illumination more strongly due to the decreasing $\frac{NA}{Mag}$ ratio, although this is not necessarily true for all NA/magnification combinations due to differing optical designs. It should be noted, however, that higher NA values will require more sophisticated autofocusing methods than those presented in this work. Example illuminance values for common microscope sources were calculated based on estimated source power at 550nm [108].
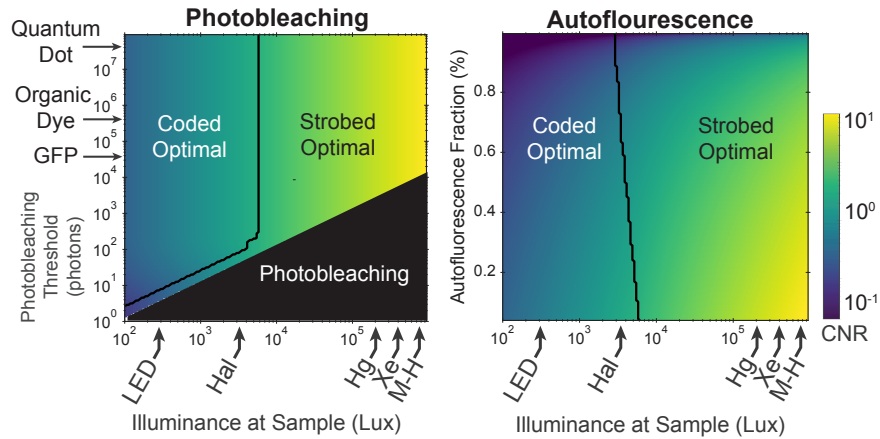
Figure 3.9: Limiting analysis of constraints imposed by the chemical fluorescence process, using contrast-to-noise ratio as the figure of merit. A.) Photobleaching influences the choice between coded and strobed illumination only when introducing a coding scheme would cause photobleaching, corresponding to a thin area of strobed optimality near the photobleaching limit. This plot assumes no background autofluorescence, so contrast-to-noise and SNR are equivalent. B.) The amount of autofluorescence relative to the signal mean has a slight effect on the optimality of strobed and coded illumination, but the effect is not strong relative to the other parameters studied here. Generally, the presence of autofluorescence degrades contrast-to-noise ratio significantly for all methods, for all illumination levels.

Biological Limitations

In fluorescence imaging, autofluorescence [109] and photobleaching [110] are primary considerations when assessing system throughput. Photobleaching, the result of chemical interactions of activated fluorophores with the surrounding medium, is of particular concern for motion deblurring applications due to higher captured SNR compared to strobed acquisitions. Practically, photobleaching can limit the maximum number of pulses which a sample may tolerate before exhibiting a non-linear response, causing strobed illumination to become a more favorable option, even at low-light. However, the region where this condition occurs is small (for the proposed system), and near the photobleaching limit (Fig. 3.9A). Autofluorescence is also a well-studied process which can further degrade the quality of fluorescence images. Autofluorescence affects contrast, which is best quantified using the contrast-to-noise ratio:

$$CNR = \frac{\gamma \bar{s}_0 - \bar{b}}{f \sqrt{\gamma \bar{s}_0 + \bar{b} + \sigma_r^2}} \tag{3.18}$$

where $\bar{b}$ is the mean background signal (autofluorescence) and all other variables are the same as in (3.9). Note that in the absence of a background ($\bar{b} = 0$), CNR is equivalent to SNR. Fig. 3.9B illustrates the relative optimality of strobed and coded illumination in the

presence of background autofluorescence, as expressed as a fraction of the primary signal.

While significant, the lifetime of various fluorophores is not limiting at illumination speeds presented in this work. Most endogenous fluorophores and fluorescent proteins have a lifetime of less than $10ns$, while organic dyes may have lifetimes of less than $100ns$ [111]. The fastest illumination source used in this work had a repetition period of approximately $4\mu s$, which is $40\times$ faster than fluorophore-limited update rate for organic dyes. Still, for future systems using motion stages moving at high velocities (such as $> 100\frac{mm}{s}$ and high magnifications (greater than $40\times$), it will become more important to consider the lifetime of the dyes used. These same constraints would also apply to strobed imaging, but not stop-and-stare imaging (which does not require high-speed signal modulation).

3.3 Summary

In this chapter, we have demonstrated a novel high-throughput imaging framework which employs multi-frame motion deblurring using coded illumination. Through both experiment and theoretical analysis, we have shown the applicability of our method for fluorescence microscopy and performed a comprehensive analysis of when our method makes sense in terms of source power and other system parameters. These results indicate that coded illumination provides up to $10\times$ higher SNR than conventional strobed illumination methods in low-light situations while maintaining acquisition rates of 50 Megapixels per second, making our method particularly well-suited for applications in drug-discovery and whole-slide imaging. Our analysis of optimal kernel selection indicates that efficient illumination sequences can be calculated quickly and cheaply using a simple random search, and our analysis of optimal pulse length provides an approximate relationship between the length of a pulse sequence and source illuminance. Further, our proposed multi-frame reconstruction algorithm produces good results using simple accelerated gradient descent with no regularization and can be scaled to multiple cloud instances for fast data processing. Future work should address improvements such as regularization, more complicated motion pathways, self-calibration, and reconstructions using under-sampled data.

Chapter 4

Fabrication of Coded Illumination Devices

In this chapter, we explore the design and fabrication of microscopes which employ coded illumination and describe the design process for coded illumination devices used in this dissertation. Illumination sources vary significantly between microscope setups based on applications, capabilities, availability, and cost. Early microscope sources involved using a candle or oil lamp for illumination. As early as 1665, Hooke used an oil lamp and water-filled globe to focus an oil lamp onto a sample to be imaged using his early microscope [3]. These methods (along with gas lamps) were used for centuries, until the carbon arc-lamp enabled electronic illumination in the late 1800s. Around this time, Köhler provided a more complete understanding of proper illumination collection to avoid imaging the source directly onto the sample [112], which was widely adopted and continues to be used today (including this work). As new lamp sources were invented (such as the Xenon arc lamp [113], tungsten lamp [114]) they found use in microscopy under the Köhler configuration.

The rise of the fluorescent microscope led to the development of light sources with engineered spectral properties, which were designed to have sharp (temporally coherent), non-overlapping excitation and emission spectra and high source power. Spatially coherent sources were used to provide phase contrast [43, 44] by filtering spatially incoherent lamps using a pinhole aperture. With the invention of the laser [115], spatially and temporally coherent illumination could be directly created and used at high power, improving the SNR of both coherent methods and florescence imaging. This high degree of coherence also leads to artifacts arising from random diffraction off of dust and edges in the optical system, limiting its use for wide-field brightfield microscopy.

Illumination using light-emitting diodes (LEDs) has emerged as an energy-efficient and low-cost alternative to lamp-based sources. As more diverse and efficient semiconductor materials have been discovered, LED light sources have found greater usage for both fluorescence and brightfield imaging. In addition, arrays of LEDs have been used to achieve even greater light throughput [46], improve sectioning in confocal microscopy [116], and provide brightfield and darkfield contrast as well as 3D digital refocusing [45]. When paired

with computation, LED arrays have found significant usage for quantitative phase imaging [49, 70, 73], super-resolution imaging [13, 90], and high-throughput imaging [72]. These devices have found widespread usage within the computational imaging community due to their simple and low-cost implementation, fast switching, and relatively high light output. However, conventional LED arrays suffer from low light-throughput at high angles due to the angular emission profile of the sources and the large distances between the LED and sample. In addition, high-throughput imaging systems (such as those used in Chapter 3) require specific electornic designs to accomidate illumination repitation rates on the order of megahertz.

In this chapter, we detail the design process and manufacture of several LED array designs and discuss their performance in relevant use-cases and applications described in the previous chapters of this dissertation.

4.1 Domed Illumination Devices for Coded Illumination

Designing the ideal LED array source for optical microscopy is challenging, both in design and manufacture. Several early implementations of the LED array microscope use a planar off-the-shelf LED array such as the 32×32 LED unit available from Adafruit [13, 45]. While these arrays are inexpensive and widely available, they have limited angular illumination range, low light throughput, and introduce flickering artifacts due to their electronic design. For high-angle high-speed FPM, a more specialized LED array becomes necessary. As we showed previously in [73, 117], a domed illuminator provides significantly higher intensity at high-NA compared to a planar array. For a planar array, intensity at the sample I can be related to the angle of illumination θ of the emitter by the equation:

$$I_\theta \propto \cos(\theta)^4 \tag{4.1}$$

where θ is the angle between the illumination vector and the optical axis. When the LEDs are arranged in a domed shape, reducing the intensity falloff to:

$$I_\theta \propto \cos(\theta)^1 \tag{4.2}$$

The illumination throughput benefits are a result of two phenomena, shown in Fig. 4.2e. The first is that off-axis LEDs in a planar array will have a larger LED-to-sample distance and thus decreased intensity at the sample. For example, if we assume that each LED is a point emitter, the intensity falloff due to increased distance can be expressed as $I(\theta) = I_0 \cos^2(\theta)$, where I_0 is the intensity at the sample from the on-axis LED and θ is illumination angle. The second improvement in light efficiency comes from the fact that LEDs have significant angular variation in intensity (typically emitting more light in the forward direction). In a planar array, the LEDs at higher angles provide less effective illumination, a problem corrected by the dome geometry, where all LEDs are radially oriented. In both the domed

Figure 4.1: Overview of dome designs from the Adafruit 32×32 planar LED array to the proposed quasi-domed device. Our first iteration used a true domed geometry, increasing illumination power at high angles, but was very difficult to fabricate. A second iteration led to the quasi-domed array, which uses standard printed circuit board fabrication processes to make the dome easy to fabricate while maintaining the benefits of a domed geometry.

and planar geometries we note that intensity further decreases with a final factor of $\cos(\theta)$ due to the smaller profile of objective window when viewed off-axis; combining these factors and assuming a Lambertian ($\sim \cos(\theta)$) angular dependence for physical (non-point-source) LEDs results in an expected intensity falloff of $\sim \cos^4(\theta)$ for the planar geometry but only $\sim \cos(\theta)$ for the domed geometry, a vast improvement at high incidence angles. Thus, the difference between geometries is proportional to $\cos^3(\theta)$, or a factor of $> 50\%$ at 40 degree and 99% at 77-degree incidence, having a substantial impact on illumination throughput, and therefore required exposure times to achieve a good SNR.

In most planar arrays (such as the Adafruit 32×32 array used in early work [13, 45]), LEDs are uniformly spaced in Cartesian coordinates across a planar circuit board. When the LED positions are projected onto numerical aperture (NA) coordinates this spacing becomes tighter at high-angles, leading to an uneven sampling across the full range of illumination angles. While having a tighter spacing is usually not a problem for computational imaging applications such as differential phase contrast or Fourier ptychography, it is inefficient, and may lead to unnecessarily long acquisition times. Using a domed geometry is more amenable to creating a uniform LED pattern in angle (NA coordinates) due to geometry

and will maintain the LED directionality at high LED angles.

The minimum LED spacing depends on the application. In Fourier Ptychography, for example, early work [13,65,118] indicated that at least a 60% pupil overlap ($0.4 \times NA_{objective}$) was necessary in order to properly recover the complex field at high resolution. This quantity depends on the diameter of the pupil (and therefore numerical aperture of the objective); systems with smaller, low-NA objectives will have tighter requirements on LED spacing. The 60% overlap requirement is developed to provide at least $2\times$ redundancy of information to reconstruct both amplitude and phase from intensity measurements. Increased overlap beyond this requirement may provide a benefit in terms of SNR, but this effect is difficult to quantify due to the non-linear algorithm used for Fourier ptychographic reconstructions. For DPC-based methods (See chapter 2) LED spacing is naturally restricted to less than the objective NA, since only brightfield LEDs are used. The primary consideration for DPC sources is having LEDs which illuminate at $NA \approx NA_{objective}$ for maximum phase contrast at low-frequencies.

While an LED dome is clearly necessary for obtaining high illumination intensity at high angles, practical limitations on manufacturing and assembly of non-planar circuit boards make building these domes at low cost and large scale non-trivial. In the following sections, we detail the design of a domed LED illuminator through several iterations of the design and prototyping process.

3D-Printed Approach

The first iteration of domed illuminators was heavily inspired by the AWARE gigapixel camera series [119–121], where hundreds of individual cameras were manually inserted into an aluminum dome which enforced opto-mechanical constraints on directionality and position through a hexagonal camera packing. With this inspiration, we designed and fabricated a 3D-printed domed illuminator consisting of 508 individually addressable broad spectrum (white) LEDs uniformly distributed across a domed surface. This first-generation device consists of 4 major components: a hemispherical dome frame for mounting the LEDs, the LEDs themselves, controller circuit boards and the sample stage. The dome mounting structure is a rigid hemisphere designed to constrain the individual LEDs within an array of computationally positioned bores, aligning the LED with the radius vector to the sample center. This hemisphere was designed with a $60mm$ radius in order to provide maximum intensity at the sample, given our desired number of LEDs and a minimum distance between neighboring LEDs. To pack these LEDs, we used a hexagonal packing pattern across a 77-degree cone of angles with maximum $NA_{illumination} = 0.62$.

The dome hemisphere scaffold was 3D printed (InterPro Models) to achieve the necessary $100\mu m$ printing resolution for accurate LED positioning. The LED angular positions were computed algorithmically to ensure uniform spacing across the dome, constrained by a minimum $150\mu m$ distance between bores for mechanical rigidity and a maximum angular separation of 3.85 degrees allowing for sufficient coherence area at the sample. This angular spacing means that 38 LEDs make up the brightfield region for a common low-NA objective
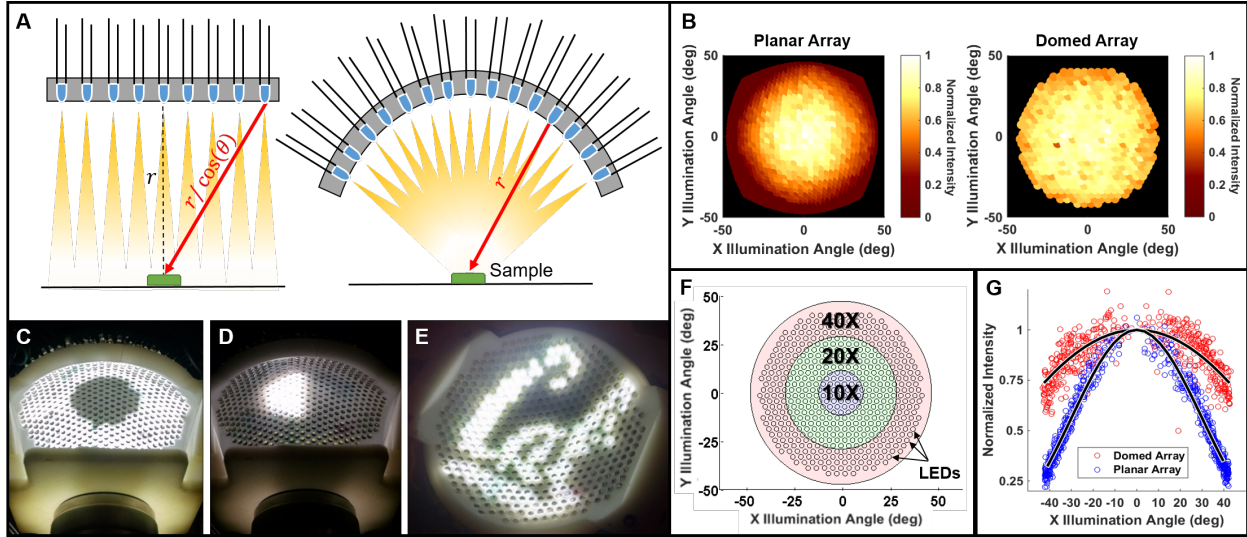
Figure 4.2: Domed LED Illuminator. (a) Visual comparison of a planar LED array with a domed array. Since the intensity of a spherical wave drops as a function of the inverse square of radius, the illumination at the sample depends on the distance between the LEDs and the sample. In the planar case (left), LED distance r increases as a function of illumination angle, causing weaker illumination at higher angles. A domed LED array (right) eliminates this variation (r is constant). (b) Normalized mean pixel intensities measured at the sensor for the planar and domed arrays. Intensity decreases as a function of angle in both cases, but much more strongly in the case of the planar geometry. Values were normalized to the central LED's brightness in both cases. (c) Illumination pattern used to acquire darkfield images with a 0.25 NA objective. (d) Illumination pattern used to generate differential phase contrast images with a 0.25 NA objective. (e) Illustration of the arbitrary illumination patterning capabilities of the device. (f) Plot illustrating the relative objective NA for several common magnifications, as compared to our dome's LED placement (small black circles). (g) Normalized measured intensity falloff as a function of angle relative to the optical axis for the domed and planar LED arrays. Falloff is proportional to $\cos(\theta)$ for the domed geometry and $\sim \cos^4(\theta)$ for the planar geometry. Black lines are $\cos(\theta)$ and $\cos^4(\theta)$ fits for the domed and planar geometries, respectively.

($4\times$ / $0.1NA$), with even more for larger NA objectives, ensuring high quality digital refocusing results across a large range of depth slices for all objectives. The 3mm through-hole, white LEDs (Mouser 593-VAOL-3LWY4) were press fit into the dome and a rigid lateral constraint was provided by acrylic retaining inserts behind each individual LED. 508 of these LEDs were soldered directly to controller boards arranged above the array, with insulated leads to prevent electrical shorting.

Accounting for mechanical tolerances of the 3D printed dome and the LED epoxy lenses,

manufacturing tolerances suggest that the maximum angular pointing error will be no greater than ± 4.8 degrees. This corresponds to a maximum intensity attenuation of only 1.2% due to assembly variation and tolerances across all illumination angles. Our illumination is also quite uniform across the field of view. The maximum field of view of our optical system has a radius of 1.25mm, set by the eyepiece field-stop diameter of 10mm and assuming a 4× objective. Given the 60mm radius of curvature of our dome, this corresponds to illumination variation due to mechanical tolerances being less than 1% across the field of view for each single LED illumination. While this result is quite good, the spread of intensities between different LEDs is significantly larger (see Fig. 4.2g), as a result of combined mechanical, electrical, and parts tolerances. Conveniently, a one-time calibration sweep of illumination angles, taken with no sample present, is sufficient to allow computational removal of this variation for all practical purposes.

The device used nine identical printed circuit boards placed in a fanned arrangement above the dome, each containing four LED controller chips (Texas Instruments TLC5926) serving up to 64 LEDs. These were controlled by a single Arduino Micro micro-controller, which calculates the appropriate bit pattern based on serial commands from an included Bluetooth transceiver. The array is fully addressable through a standard Bluetooth serial link or USB interface. We operate the serial output at 115K baud and note that we can update the entire pattern with approximately 100ms latency, although we predefine some of the more complex LED illumination patterns and store them in the Arduino flash memory to further improve acquisition time. Thus, our acquisition times are primarily limited by the camera rather than the LED array control.

The dome's power control board is tolerant of voltages between 7 and 20 VDC to allow compatibility with a large range of power sources, including a standard 12V automotive battery and a 100W wall-plug variable output power supply, as well as many commercially available portable power supplies for consumer electronics. During regular usage, the device requires no more than 2A of current, though it could potentially draw up to 4.8A of current when all LEDs are illuminated. This is not a typical use case, however, since simultaneous illumination inside and outside the objective NA amounts to an undesirable mixing of darkfield and brightfield contrast. Noting that for 4×, 10×, and 20× objective configurations there are more darkfield than brightfield LEDs, to reduce power consumption we perform darkfield illumination by default using an annulus with a width equivalent to 0.15NA rather than using all of the darkfield LEDs. This moderately reduces the contrast and resolution of darkfield images but significantly reduces power use during the darkfield illumination cycle. We note that the device can operate indefinitely without overheating issues for both multi-contrast and digital refocusing under normal use.

The completed prototype is shown in Fig. 4.2. Overall, this prototype was very difficult to manufacture because LEDs had to be inserted by hand, and the leads of each LED had to be manually soldered to one of the 9 circuit boards above the dome. This 3D reconstruction process required significant time and effort to ensure all of the LEDs were functional. The designs of this illuminator were published publicly, but to our knowledge there were no successful reproductions, owing to the difficult fabrication process.
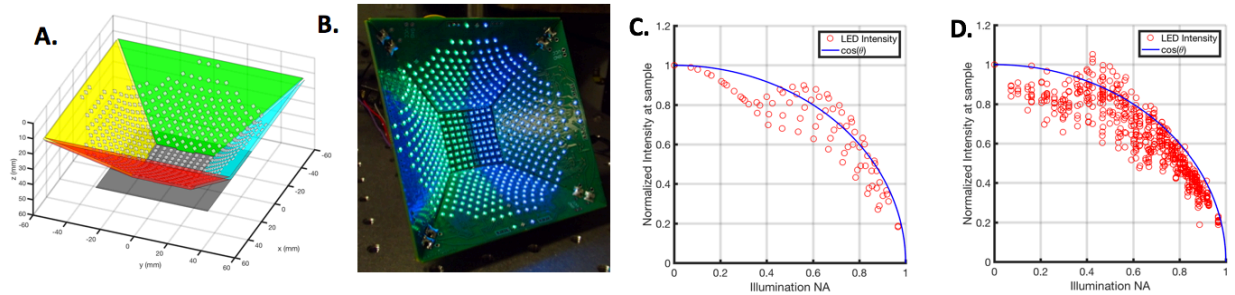
Figure 4.3: Quasi-dome programmable LED illuminator. (a) CAD model of LED flange positions. (b) Assembled LED array quasi-dome displaying two half circles (center line of LEDs are turned off). (c) Simulated intensity falloff of LED array considering all factors (normalized). (d) Experimental normalized intensity falloff.

Quasi-Dome

In response to difficulties experienced in fabricating the 3D-printed dome, we developed a quasi-domed LED array consisting of 5 panels (see Fig. 4.3a,b), each of which are standard printed circuit boards (PCBs) with all components pre-assembled using standard manufacturing techniques. This domed illuminator used multi-channel LEDs (Knightbright APTF1616 series) which have individual channel control without multiplexing using on-board LED controllers (Texas Instruments TLC5955), connected in a serial daisy-chain (Fig. 4.4). This electrical configuration enables 16-bit LED intensity control of each channel (using pulse-width modulation) with fast pattern updates (10ms) due to high-speed serial control via a Teensy 3.2 Microcontroller (PJRC). The software used to control this device was released as the open-source Illuminate LED array firmware [103].

The LED flange geometry was selected to provide sufficient overlap of sample spectrum areas for Fourier Ptychography. We selected 0.1 NA as our minimum objective numerical aperture, commonly corresponding to a 4× objective. This spacing was used to select LED positions with even spacing in Fourier space, lending to the large LED spacings on the circuit boards at high angles. Fig 4.3d shows the experimental intensity fall-off of radially projected LEDs as measured at the sensor, showing good agreement between theoretical illumination falloff due to the $\cos(\theta)$ term and experimental measurements. The use of color LEDs enables color multiplexing, which was used to perform single-shot quantitative phase imaging (cDPC) [81].

In general, this illumination source solves most of the problems with the illuminator presented in the previous section. It provides more LEDs and simple and scalable fabrication using conventional PCB fabrication techniques while maintaining the many benefits of a domed geometry. For these reasons, the Quasi-dome geometry is practically superior to both planar geometries and the dome presented in Section 4.1.

Figure 4.4: Circuit schematic for Quasi-dome device. A chain of up to 40 TLC5955 chips are connected in a daisy-chain configuration, having common Power (VCC), ground (GND), Grayscale-clock (GS), serial clock (SCLK), and latch (LAT) pins. These are controlled by a micro-controller upstream, which enables the control of up to $16 \times N_{chips}$ LEDs for N_{chips} TLC5955 chips.

4.2 Coded Illumination Devices for High-Throughput Microscopy

High throughput imaging using both strobed illumination and motion deblurring [97] requires light sources to be extremely fast compared to the domed devices presented in the previous section. In these systems, the sample is imaged and illuminated while being moved continuously by a mechanical motion stage, without stopping. A key parameter of a high-throughput illumination source is repetition rate, or the minimum pulse duration the source can provide. This time must be less than the time required for the sample to move one effective pixel size (pixel size divided by magnification) while being scanned. Mathematically, the source repetition period T is related to the motion velocity v_{motion}, camera pixel size Δ, and system magnification M by the following relationship:

$$T \leq \frac{\Delta}{M v_{motion}} \tag{4.3}$$

For practical systems (such as those described in Chapter 3), this threshold is on the order of micro-seconds, which is 3-4 orders of magnitude faster than the LED domes described in previous sections. For these applications, it was necessary to adopt very simple LED circuitry to avoid electronic speed limitations associated with dimming (using pulse-width modulation) and serial control of LED driver chips. As such, we adopt an extremely simple circuit, consisting of only a micro-controller and transistor, to control an arbitrary number of LED sources at very high speed, limited only by the clock-speed of the microcontroller. The drawback of this configuration is that only binary illumination is possible, but binary illumination is often optimal for both strobed and coded illumination (See Chapter 3).

Figure 4.5: Circuit schematic for fast LED driver circuit. A transistor (NPN type) is used to modulate a large current source for controlling many LEDs at once, which are connected in serial. Resistor R0 is normally set to 1KΩ, and R1 is set such that the current is not too large based on the VCC voltage. This circuit enables micro-controller limited illumination updates, although it does not allow per-channel dimming and supports only binary illumination patterns.

For these applications, we developed two high-speed sources based on the same platform, one for high-throughput brightfield and quantitative phase imaging, and one for fluorescence imaging. The first device used 40 white (blue-phosphor) LED emitters (VAOL-3LWY4, VCC) controlled by four transistor circuits to modulate 4 quadrants of a circle. The intention of this design was to enable extremely fast brightfield (color) imaging, as well as quantitative phase imaging using DPC. The second used a single, high-power LED source (Thorlabs M470L3) modulated using a simple single-transistor circuit through the same micro-controller. Chromatic filters could be added to the illumination pathway and detection pathway to enable fluorescence imaging, and this device was compatible with a wide range of LED sources due to its modular design. Both of these devices used the same firmware as the micro-controller firmware, which was designed to be modular to accommodate various LED arrays [103]. A schematic of this high-speed circuit is shown in Fig. 4.5.

4.3 Coded Illumination Devices for Portable Microscopy

Optical microscopy is an important tool for disease screening and diagnosis throughout the world. Significant resources have been devoted to developing portable and affordable compact microscopes for remote clinical applications [122–133]. Compact microscopes based on mobile phones, including CellScope [134,135], have demonstrated that microscopy can be effectively performed outside of hospitals and diagnostic laboratories by minimally trained healthcare workers, that images can be transmitted for confirmation of diagnosis, and that phone-based computational analysis can be used to provide automated diagnosis.

Computational CellScope

Here, a new variation of the CellScope microscope is demonstrated which incorporates recently developed techniques of computational illumination [45, 47, 65] using the LED dome shown in Fig. 4.2 to enable new imaging modalities, including darkfield, phase imaging and digital refocusing[1]. Using only the on-board processing power of a smartphone, our device implements quantitative phase recovery using DPC, as well as lightfield digital refocusing, so that a sample focus can be changed after the fact (without mechanically changing focus).

The flexibility and speed of the programmable LED array illuminator, as well as the lack of moving parts and low cost, make the hardware very amenable to modification as a CellScope attachment. In order for our device to be practically useful in the field, we have here enforced the requirement that all of our processing and control be performed on the smartphone, without use of a PC. Thus, the device can be field-deployable as a simple add-on to CellScope, requiring only a 7V-20V DC power source which could be adapted to many power sources, such as batteries, car batteries, or wall-outlets using commonly available adapters. In the following sections we detail the design and performance of the hardware and software of our new Computational CellScope device.

While our addition involves custom LED drive circuitry and a 3D printed structure, complexity was kept low to preserve the low-cost nature of CellScope. Part counts, cost and especially size may be further reduced in design-for-manufacture. The size of the illuminator could be reduced to essentially the dimensions of the dome itself, and cost could be comparable to the price of a modern smartphone, matching and improving upon the functionality of a full-size microscope at a fraction of the cost.

---

[1]This work was developed in close collaboration with Daniel Fletcher, Mike D'Ambrosio, and Neil Switz (Fletcher Lab, Bioengineering, UC Berkeley), as well as Lei Tian, Jared Rulison, Hurshal Patel, Nitin Sadras and Aditya Gande (Waller Lab, EECS, UC Berkeley).
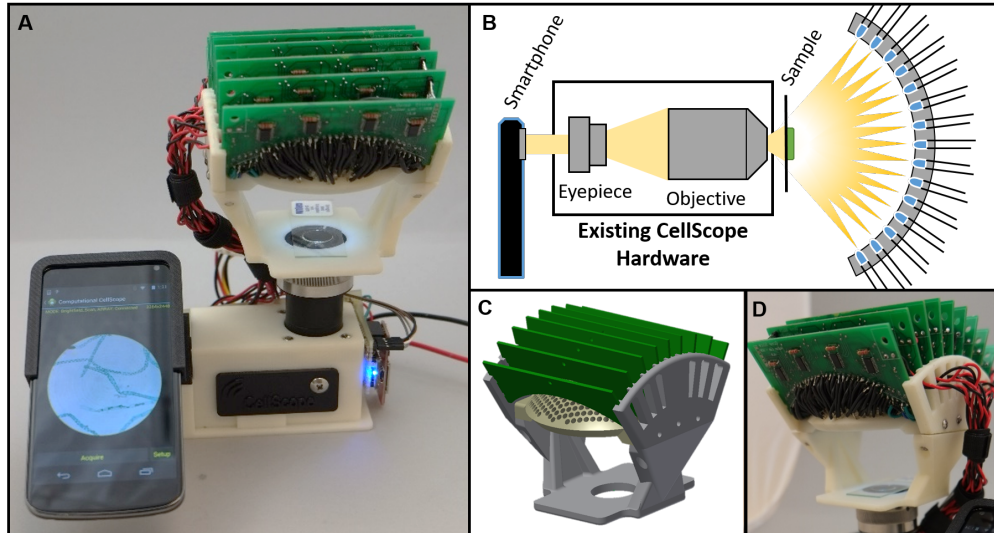
Figure 4.6: Computational CellScope. a). Device observing a sample using a Nexus 4 smartphone. b). Optical schematic of the CellScope device with our custom-made domed LED illuminator. c). CAD assembly of the dome. d). Assembled dome and control circuitry.

Multi-Contrast Imaging

To achieve brightfield, darkfield and phase contrast simultaneously, we time-multiplex images taken with different LED patterns and post-process them on the smartphone to synthesize pseudo-real-time multi-contrast imaging, as in [47]. Brightfield images correspond to illumination by LEDs that lie within the cone of angles described by the objective numerical aperture (NA). Darkfield images are obtained by illuminating the sample from angles beyond the angular acceptance of the objective (Fig. 4.2a) [45]. Since different objectives have different NA, one must specify in the software which objective is being used, with larger NA corresponding to a larger brightfield region of LEDs. Our dome is designed to enable darkfield contrast for any objective of NA< 0.62, roughly corresponding to a typical 40× objective.

Phase contrast can be achieved in a single-shot image by any asymmetric illumination pattern [136,137]. Here, we choose to employ a differential phase contrast (DPC) scheme [63–65, 138], which requires two images having complementary illumination patterns, because it gives good phase contrast at all spatial frequencies and can be quantitatively interpreted as the gradient of sample phase. The method involves sequentially illuminating the sample with the two opposite halves of the brightfield circle while capturing an intensity image for each. For example, one may first take an image, I_R, with only the right half of the LEDs on and then a second image, I_L, with only the left half of LEDs on (see Fig. 4.2b). The two images are processed as follows to obtain brightfield and phase contrast:
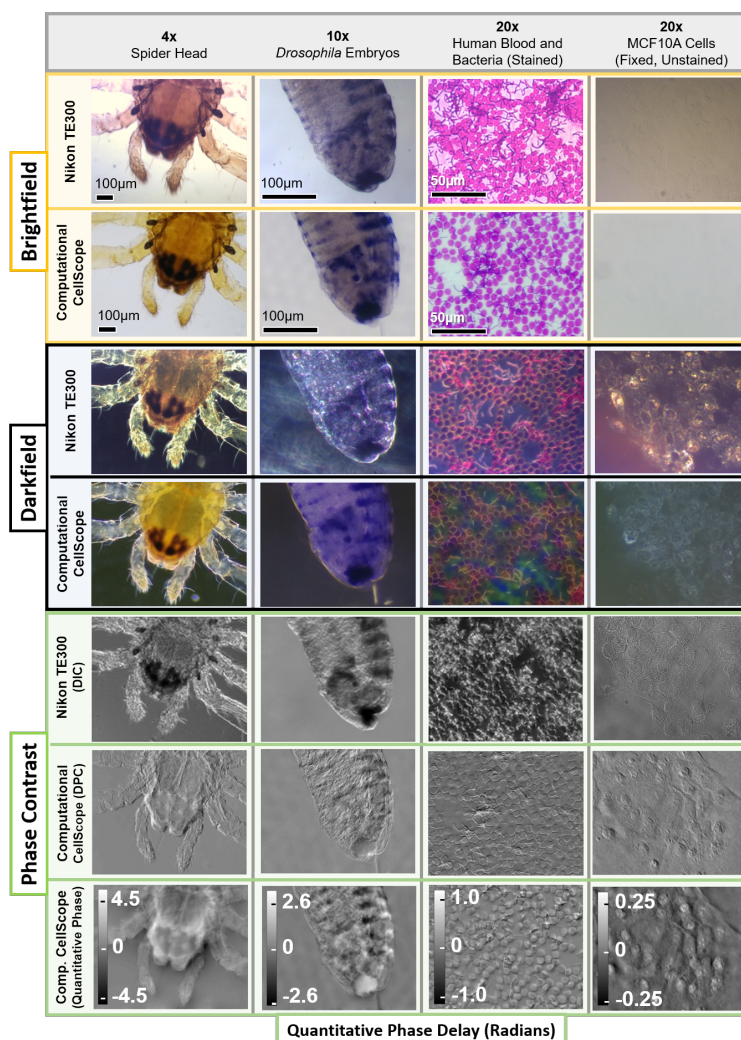
Figure 4.7: Image results compared to a standard microscope. Computational CellScope acquires brightfield and darkfield images of similar quality to a standard upright microscope (Nikon TE300) without the use of hardware inserts. Additionally, it enables phase imaging using Differential Phase Contrast (DPC), which contains similar information to standard phase contrast imaging, and can be inverted to obtain quantitative phase of the sample (bottom row). Differences in color shades are caused by the relative differences in hue of the halogen lamp and the white LEDs. Note the additional dark features in DIC results, as compared to DPC, illustrating mixing of phase and absorption information in DIC. In the right-most column, we show images for an unstained transparent sample, illustrating the utility of phase imaging methods for label-free imaging.

$$I_{\mathrm{BF}} = I_{\mathrm{L}} + I_{\mathrm{R}}, \qquad I_{\mathrm{DPC}} = \frac{I_{\mathrm{L}} - I_{\mathrm{R}}}{I_{\mathrm{L}} + I_{\mathrm{R}}}, \tag{4.4}$$

where I_{BF} is the brightfield image and I_{DPC} is the phase contrast image. Since the LEDs are mutually incoherent, adding the two images gives an equivalent brightfield image and subtracting them produces phase contrast, due to asymmetric clipping in Fourier space. The intensity of the DPC image can be shown to be approximately proportional to the first derivative of phase along the direction of illumination asymmetry [63], and different axes of rotation can be programmed by changing the LED array pattern accordingly. Typically, we capture an additional two images in order to compute both the left-right and top-bottom phase derivative results representing both orthogonal directions. DPC images are qualitatively similar to Differential Interference Contrast (DIC); however, the latter is not a quantitative method. To obtain quantitative phase from DPC images, we solve the inverse problem [64,139] using a simple deconvolution in Fourier space. Results for all of the contrast modalities are shown in Fig. 4.7.

Thus, by acquiring four half-brightfield images and a single darkfield image for each time point, we can synthesize brightfield, darkfield, and phase contrast modes in near real-time. Users have the option of saving and post-processing time-multiplexed frames or viewing a live multi-contrast display of the sample, though display speed is significantly faster in the latter case due to limited file write speeds on the smartphone. We developed an application to stream these four contrast modes size-by-side while updating each frame sequentially as the illumination pattern cycles through the different patterns (Fig. 4.9b). The user may touch any of the four images for a live full-screen display of that contrast mode only, and the illumination pattern cycle will update to reflect this.

Some image results for each of the contrast modes are shown in Fig. 4.7, using different objective magnifications and samples. For comparison, we show the same samples imaged in a commercial inverted microscope with traditional hardware. Darkfield was obtained by using a Ph3 condenser aperture in combination with objectives having NA smaller than the sine of the half-angle of the Ph3 annulus inner diameter. Since DPC is not currently commercially available, we instead compare our DPC phase contrast images to (similar-appearing) DIC. Both provide images whose contrast is related to the first derivative of phase along a single direction; however, DIC mixes absorption and birefringence information with phase, so that dark features in the image may result from either absorption of the sample or phase contrast interferences. In the DPC images, on the other hand, the image is related purely to the sample phase distribution (see Fig. 4.7), which can be inverted to reveal quantitative phase, as shown in the bottom row. Provided in a portable package, these multi-contrast video and streaming methods have the potential to allow clinicians to view a sample with three separate contrast methods at once, enhancing the information available for diagnosis and disease discrimination.

Digital Refocusing

For thick samples, our system can capture a different sequence of images in order to recover 3D images and enable digital refocusing. In this case, we sequentially capture images for each of the LEDs in the brightfield region. The resulting dataset is similar to limited angle tomography with many angles, which provides depth sectioning from angular information [140]. For simpler processing more amenable to mobile phone programming, we use a lightfield approach here [16, 45]. This involves a simple shift-and-add algorithm to digitally refocus the image to different axial (z) planes. We calculate the digitally refocused intensity image at a distance Δz away from the physical focus plane as:

$$I^{\Delta z} = \sum_{\text{all brightfield LEDs}} I_i(x + \Delta z \tan \theta_x, y + \Delta z \tan \theta_y), \qquad (4.5)$$

where I_i denotes the intensity image for the i^{th} LED, shifted according to its angle of illumination at the sample (θ_x, θ_y) and the desired refocus distance Δz.

The number of individual LEDs making up the brightfield region roughly determines the number of depth planes that can be accurately reconstructed, and the range of illumination angles determines the axial resolution of the 3D result. Since a separate image is taken for each illumination angle, both acquisition and processing time are a function of the numerical aperture of the objective, as illustrated in Fig. 4.9. Acquisition speed was primarily limited by the time required to save an image to the smartphone's flash memory at full resolution (8 Megapixels on the Nexus 4). This is important because data acquisition remains fast, while processing can occur in the background. Using the same dataset, we can also calculate 3D phase contrast images by digitally refocusing the two halves of the brightfield region separately [65]. It is expected that this mode of imaging intensity or phase in 3D with no moving parts will give better diagnostic information for thick samples. Alternatively, it could be used for correcting misfocus, obviating the need for automatic axial translation or automated focus adjustment in long time-lapse studies.

Results are shown in Fig. 4.8 for digitally refocused images as compared to physically refocused images on an inverted microscope (Nikon TE300), both with a $10\times$ objective (0.25 NA). The phase contrast images show the first derivative of phase along both the vertical and horizontal directions, calculated from the same dataset using only the green color channel. The algorithm successfully refocused features across 400μm depth of field, limited by object thickness. Our refocusing achieves an axial resolution of approximately 5μm within $\pm50\mu$m of physical focus position but degrades approximately linearly with increasing refocus distance [65]. Processing time is approximately 1.5 minutes per depth slice for a $10\times$ objective.

Acquisition Software and Processing

It has previously been shown that using a smartphone as a microscope poses unique challenges intrinsic to the phone software [135]. Smartphone cameras may only allow minimal
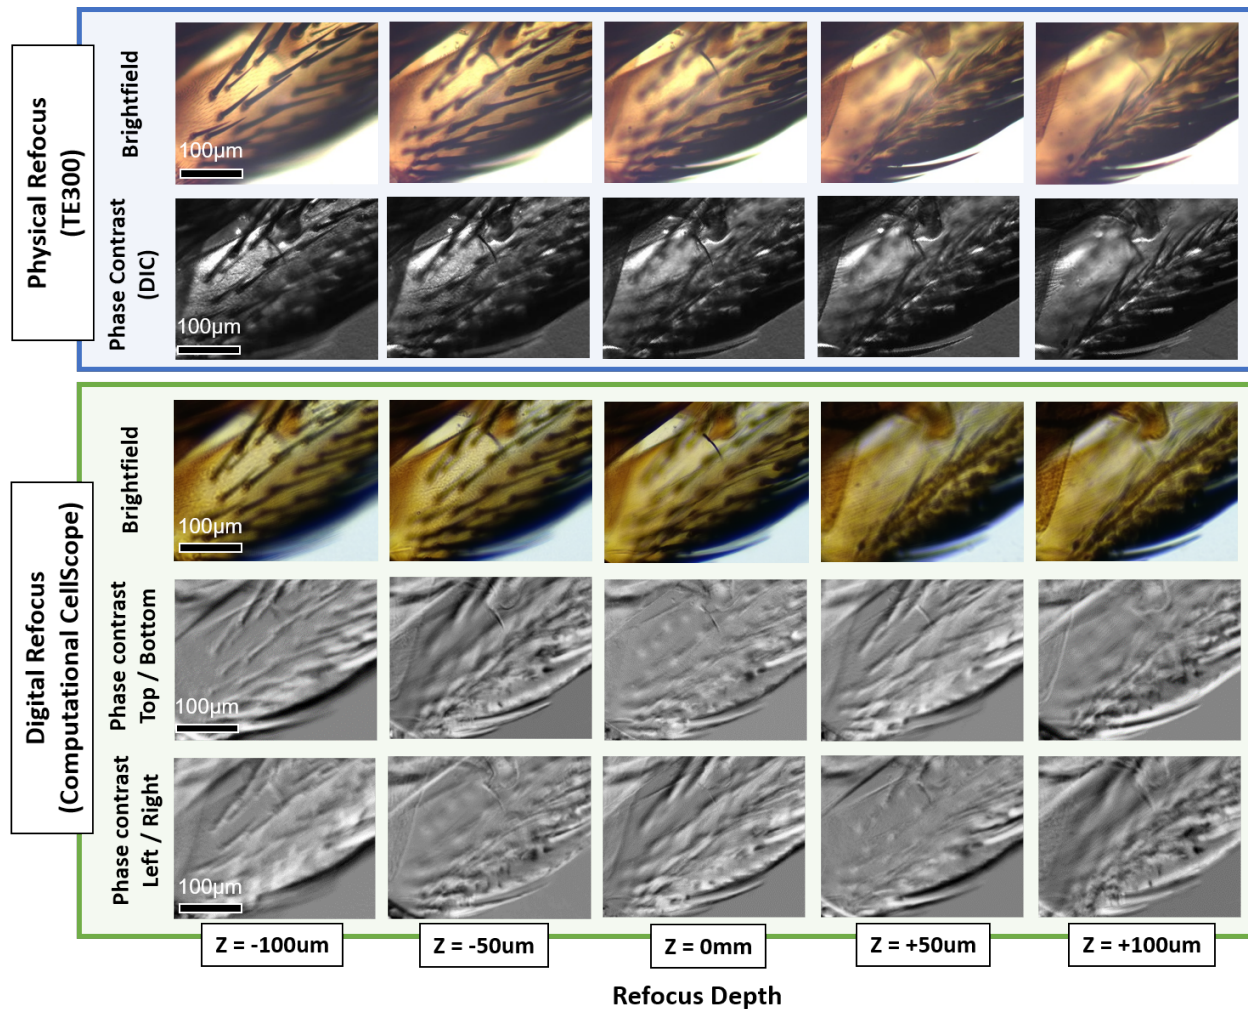
Figure 4.8: Digital refocusing on the Computational CellScope. Comparison of digital refocusing to physical refocusing on a commercial microscope (Nikon TE300) of a house fly wing sample (AmScope PS200) with a 10× objective. Digitally refocused phase contrast images are also computed for both vertical and horizontal phase derivatives at different focus depths.
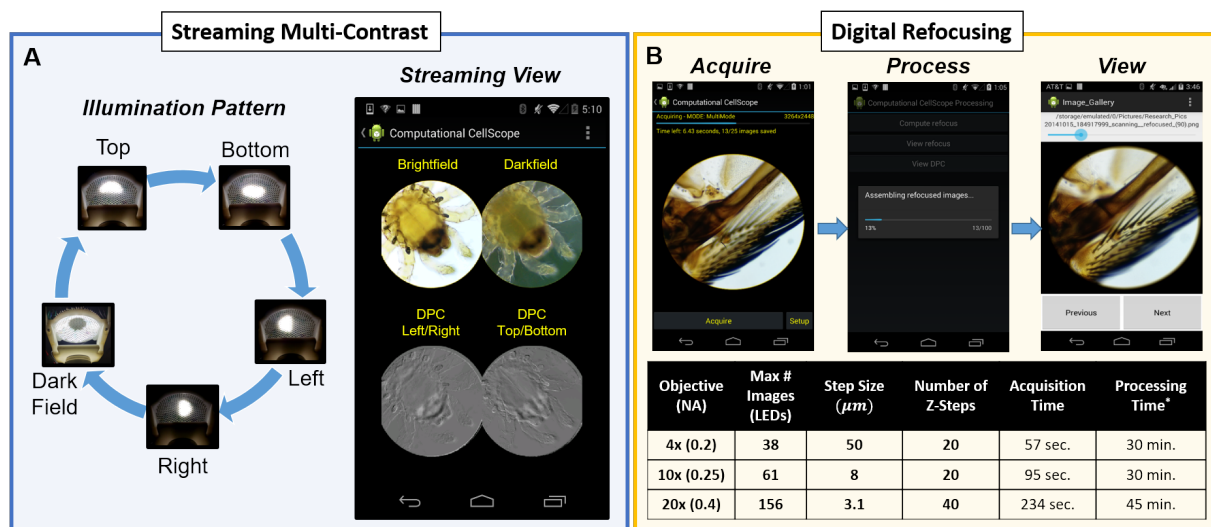
Figure 4.9: Android Application Workflow. a). Schematic of streaming multi-contrast LED patterns. Here we vary the LED pattern in time and acquire and process images on the smartphone, producing a streaming multi-contrast display of a sample without any further post-processing. The user can touch any image to zoom in and stream an individual image. Total cycle time is 2.3 seconds. b). Overview of workflow for digital refocusing mode. Table shows example processing and acquisition times for a typical dataset reconstruction. Axial Resolution is determined by the range of illumination angles sampled (defined by the objective NA). The number of z-steps were chosen such that refocus blur does not exceed 20 pixels. Processing and acquisition time can be reduced by selecting fewer refocus planes or by sparsely sampling LEDs, trading axial resolution for faster acquisition time.

quantitative control over standard imaging parameters (e.g. focus, exposure, gain), opting for opaque automatic algorithms that simplify the user experience. To circumvent these restrictions, we wrote a custom Android application that attempts to achieve the optimum imaging characteristics with our coded illumination configuration. The software application also initiates and handles the Bluetooth connection to the domed array, enabling synchronized acquisition and array control through the standard Android API. Array control is thus transparent to the end-user, requiring them to simply pair the phone with the illuminator and press a connect button to initiate a Bluetooth connection. Our application was developed specifically for the Android platform and will be compatible with any phone running Android OS version 4.0 or later. However, our algorithms were developed using the OpenCV Library, which is cross-platform for iOS (Apple, Inc.) and other operating systems. Thus, most of our application code is portable to other smartphone platforms with moderate development effort. A screenshot of the acquisition and processing are shown in Fig. 4.9. The user may choose to collect and synthesize any or all of darkfield, refocused brightfield and

DPC images.

In our application, images were acquired using the standard Android API, which does not provide an interface to set explicit exposure times, but supports only continuous auto-exposure with predefined exposure offset values. To circumvent this issue, we included a short pre-illumination sequence before each dataset acquisition to lock exposure at the appropriate value and enforce the requirement that equal numbers of LEDs be chosen for each half-circle. Finally, we incur significant latency between the camera shutter and the availability of the frame to our application within the API, due to post-processing algorithms integral to the phone and performed in the background (e.g. white balance and demosaicing). This severely limits our acquisition speeds, which will likely be improved in newer versions of Android that allow finer camera control through the API. Apple iOS offers a different camera API that may also offer improvements in acquisition speed.

Data post-processing was performed in a standalone Android app, where image stacks were loaded and processed on the phone. We employ a number of functions of the OpenCV Imaging library for Android to perform most of our computation. Individual DPC images are computed in less than a second, as demonstrated in our multi-contrast view mode. Digitally refocusing an image into 21 depth planes (± 100 μm range with 10 μm sectioning) requires approximately 30 minutes of processing time, but the resulting 3D image stack can be interacted with in real-time; all other computational imaging results are much faster ($\sim$0.43 frames/sec). The long processing time is attributed to frequent loading and saving to the smartphone's internal storage. We note that significant improvement in processing speeds for all of our algorithms is possible through implementation using the Android NDK and is also expected as phone computational power increases with each product generation. These performance metrics were calculated on a Nexus 5 smartphone (LG Electronics) and may vary on other devices.

4.4 Summary

In this chapter, we presented several examples of the fabrication and design process for coded illumination devices. These devices are each the result of many design iterations which incorporate the joint design of hardware and software - a key paradigm of computational imaging. The recent development and commercialization of 3D-printing has enabled rapid prototyping of optical devices, enabling the fabrication and iteration of domed LED illuminators (As show in in Section 4.1). However, due to practical fabrication limitations, we found that reverting to more standardized printed circuit board (PCB) approach (Section 4.1) enabled significantly improved manufacturability without significant performance degradation, enabling more rapid dissemination of domed LED illuminators around the world. We also developed coded illumination prototypes with much higher temporal illumination resolution for high-throughput imaging (Section 4.2), which required different design decisions to be made to accommodate rapid temporal coding of LEDs. Finally, in Section 4.3 we described a complete coded illumination solution for portable microscopy using a 3D-printed domed

illuminator. This device enabled portable implementations of qualitative microscopy (darkfield and brighfield), quantitative phase imaging using DPC, and 3D digital refocusing of thick samples, all using only the existing smartphone and the addition of our domed illuminator. This device demonstrates how flexible and amenable coded illumination devices can be through example; future work could demonstrate an improved prototype which is fully field-ready in terms of robustness and flexibility and could enable new applications such as high-throughput imaging and fluorescence imaging.

Chapter 5

Self-Calibration of Coded Illumination Systems

In Chapter 1, we described a general computational imaging framework which consisted of both a forward model generation step as well as a computational inversion. In all imaging systems, the forward model is synthesized from the fundamental laws of light propagation (e.g. Maxwell's equations), the optical design of the system, and the accumulated error from optical mis-alignment, dust, manufacturing imperfections, and other sources which can be collectively referred to mis-calibration. While some of these imperfections can be removed using simple processes (such as background subtraction), others, such as system aberrations, are relatively complex, and need to be tolerated or removed using a more sophisticated procedure.

5.1   Algorithmic Self-Calibration

The concept of algorithmic self-calibration (solving jointly for the reconstructed image and the calibration parameters) has proven particularly useful in coherent computational imaging. Examples include probe retrieval in Ptychography [141–144], source recovery for through-focus phase imaging [145, 146], pupil and source recovery for Fourier Ptychography Microscopy (FPM) [147–149], and calibration-invariant inverse scattering models [150]. The standard approach to self-calibration uses alternating projections (AP), which optimizes multiple variables serially, keeping other parameters fixed during each sub-iteration. The non-convexity of AP provides no guarantee of global convergence, but in practice it works with sufficiently diverse data. Self-calibration of aberrations has been demonstrated previously in an LED array microscope for the cases of through-focus phase [151] and FPM [94, 147, 152–154], but required a large number of images to be captured. For example, a typical FPM setup [155] uses approximately 5× as many images as our system to achieve the equivalent resolution.

   Mathematically, a forward model $\mathcal{A}\{\cdot\}$ can be made a function of any variable, including

an arbitrary calibration vector $\boldsymbol{c}$, with a simple modification of Eq. 1.6:

$$y = \boldsymbol{A}\{\boldsymbol{x}; \boldsymbol{c}\} + \boldsymbol{\eta} \tag{5.1}$$

In the case where $\mathcal{A}$ is differentiable with respect to $\boldsymbol{c}$, it becomes possible to recover $\boldsymbol{c}$ with knowledge of $\boldsymbol{x}$ through direct inversion or gradient methods, though this is only very rarely the case. When $\boldsymbol{x}$ is unknown, it becomes necessary to solve for both $\boldsymbol{x}$ and $\boldsymbol{c}$ using an alternating minimization. When $\boldsymbol{c}$ is the point spread function of the system (convolution kernel), this becomes a blind deconvolution problem, which has been extensively studied in both microscopy [156, 157] and photography [158–161]. However, $\boldsymbol{c}$ can also represent system aberrations [153], illumination parameters such as LED positions [148], or other parameters.

When $\mathcal{A}$ is differentiable with respect to both $\boldsymbol{x}$ (the object) and $\boldsymbol{c}$ (the PSF), it can be solved by taking the gradient with respect to $\boldsymbol{c}$ and $\boldsymbol{x}$ in alternating steps. So long as each gradient step is not too large (or is optimized locally using a line search), this alternating minimization technique will not diverge, and will only improve the initial estimate of $\boldsymbol{x}$ and $\boldsymbol{c}$ until convergence. However, this alternating approach is non-convex, and therefore is very sensitive to initialization since there are many local minima.

5.2 Aberration Self-Calibration using Differential Phase Contrast

System aberrations are nearly always present in optical systems and may be field-dependent in some cases (such as low-magnification objectives). When performing DPC, these aberrations can corrupt reconstructions due to model mis-match, especially at high frequencies. It is well-known that aberrations arising from optical mis-alignments can be mostly classified by a small number of Zernike Polynomials [162], which makes these functions particularly well-suited for self-calibration.

In this section, we propose a method of algorithmic self-calibration for Differential Phase Contrast (DPC) microscopy [49, 64, 68, 81, 136, 163], where we avoid the need for pre-calibration by jointly recovering both the sample's complex-field *and* the spatially-varying aberrations of the system, directly from raw images[1]. The method is an extension of illumination-based DPC microscopy [49, 64], where images are captured with different source patterns, then a reconstruction algorithm recovers the complex-field. Experiments are implemented in a commercial brightfield microscope with a low-cost programmable LED array light source, such that patterns can be switched quickly with no moving parts [13, 49, 94, 164]. Among the wide variety of QPI methods, those which use partially coherent illumination, like DPC, are advantageous since they provide 2× better resolution, higher light throughput and reduced speckle [49, 56, 145, 165], as compared to coherent methods.

---

[1]This work was developed in close collaboration with fellow Ph.D. student Michael Chen (Waller Lab, EECS, UC Berkeley).

To efficiently recover both the complex field and system aberrations, we employ an AP framework that uses 4 captured images for simultaneous phase retrieval and digital correction of spatially-varying aberrations (Fig. 5.1). Three of the measurements are partially-coherent conventional DPC images (with rotated half-circle sources); these provide good phase contrast, but poor aberration contrast. The fourth image uses single-LED (spatially coherent) illumination; this provides aberration contrast, but alone cannot resolve the ambiguity between complex-field and pupil aberration [166]. Using both partially-coherent and coherent images together improves sensitivity to aberrations without sacrificing the benefits of partial coherence. We model the aberrations parametrically (with a Zernike basis [151, 162]) to dramatically reduce the number of unknowns. In addition, by segmenting the field-of-view (FOV), we are able to recover and digitally correct for spatially-varying aberrations across the FOV.
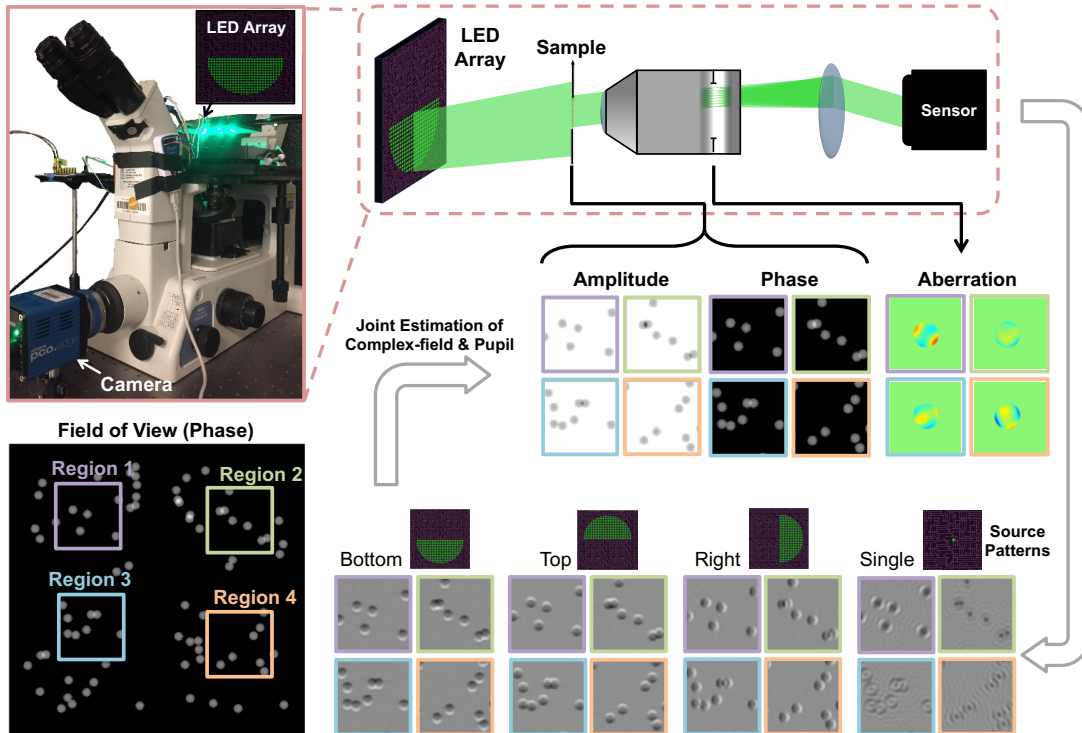
Figure 5.1:   Our LED array microscope captures 4 images with different illumination source patterns (three half-circles and one single LED). The intensity images are used to simultaneously reconstruct both amplitude and phase of the sample, and to estimate the pupil aberrations at each spatial location, which are then digitally corrected for. We show reconstructions for 4 regions with different spatially-varying aberrations.

Joint Estimation of Complex-Field and Aberrations

Our LED array microscope and data capture scheme are shown in Fig. 5.1. Quantitative DPC (without aberration correction) requires a minimum of 3 intensity measurements to reconstruct phase [81], though only two quantities (amplitude and phase) are reconstructed at each pixel. Hence, there is significant redundancy in the data which may, in principle, be used to solve for aberrations. Unfortunately, intensity images formed by partially-coherent illumination do not exhibit significant aberration contrast. Hence, we modify the capture scheme to add one additional measurement with spatially-coherent illumination (a single on-axis LED). The example in Fig. 5.1 has a different aberration in each quadrant of the FOV, and only the single-LED image displays visible differences in contrast for each region. An off-axis LED would also provide the necessary coherent contrast, but the on-axis LED provides higher intensity and better signal-to-noise ratio (SNR). Having achieved both phase and aberration contrast with 1 on-axis and 3 half-circle sources, we use this dataset to jointly recover both spatially-varying aberrations and the complex-field of the sample (with resolution set by the incoherent diffraction limit).

The process of estimating the sample's complex-field and the system's aberrations simultaneously is a joint estimation large-scale nonlinear non-convex problem. To simplify, DPC typically makes a weak scattering approximation which linearizes the forward model. This approximation is generally valid for optically thin or index-matched samples, like biological cells. Intensity measurements can then be related to absorption ($\boldsymbol{\mu}$) and phase ($\boldsymbol{\phi}$) of the sample [49, 81] by simple convolutions. As derived in Appendix A, the DPC forward model in matrix form is:

$$\boldsymbol{I}_{\mathrm{n}} = \mathbf{F}^{-1}\left(\mathrm{diag}(\boldsymbol{H}_{\mu})\mathbf{F}\boldsymbol{\mu} + \mathrm{i} \cdot \mathrm{diag}(\boldsymbol{H}_{\phi})\mathbf{F}\boldsymbol{\phi}\right) . \tag{5.2}$$

Here, $\mathrm{diag}(\boldsymbol{x})$ denotes a diagonal matrix with diagonal values $\boldsymbol{x}$, $\mathbf{F}$ and $\mathbf{F}^{-1}$ represent the DFT and inverse DFT matrices, $\boldsymbol{I}_{\mathrm{n}}$ is the vectorized normalized measured intensity (with the DC term subtracted), and $\boldsymbol{H}_{\mu}$ and $\boldsymbol{H}_{\phi}$ are vectorized transfer functions for absorption and phase. If the source satisfies the Köhler illumination configuration, the transfer functions can be numerically evaluated based on the cross-correlation property of Fourier transforms:

$$\boldsymbol{H}_{\mu} = \frac{1}{\boldsymbol{I}_{\mathrm{o}}}\left[\mathbf{F}^{-1}\mathrm{diag}(\boldsymbol{O}^{*})\mathbf{F} + \mathbf{F}^{-1}\mathrm{diag}(\mathbf{F}^{*}\boldsymbol{P}^{*})\mathbf{F}\mathrm{diag}(\boldsymbol{S})\right]\boldsymbol{P} \tag{5.3}$$

$$\boldsymbol{H}_{\phi} = \frac{1}{\boldsymbol{I}_{\mathrm{o}}}\left[\mathbf{F}^{-1}\mathrm{diag}(\boldsymbol{O}^{*})\mathbf{F} - \mathbf{F}^{-1}\mathrm{diag}(\mathbf{F}^{*}\boldsymbol{P}^{*})\mathbf{F}\mathrm{diag}(\boldsymbol{S})\right]\boldsymbol{P} , \tag{5.4}$$

where $*$ is the complex conjugate operation, $\boldsymbol{I}_{\mathrm{o}}$ is the total intensity of the source passing through the system, $\boldsymbol{S}$ and $\boldsymbol{P}$ are the vectorized source and pupil, and $\boldsymbol{O} = \mathbf{F}\mathrm{diag}(\boldsymbol{S})\boldsymbol{P}$. Typically, the space-invariant exit pupil $\boldsymbol{P}$ is a circular function with its radius determined by numerical aperture (NA), and wavelength, λ. The phase of $\boldsymbol{P}$ is the pupil aberration we wish to recover, modeled as a weighted sum of Zernike modes on spatial frequency coordinate ($\boldsymbol{u}$) [162]:

$$\boldsymbol{P}(\boldsymbol{c}) = \mathrm{Circ}\Big(\frac{\lambda \boldsymbol{u}}{NA}\Big) \prod_{m=0}^{M} e^{\mathrm{i}c_m \boldsymbol{Z}_m} \ , \tag{5.5}$$

where M is the total number of Zernike modes and $\boldsymbol{c}$ contains the coefficients, c_m, of each orthogonal mode $\boldsymbol{Z}_m$. To recover spatially-varying aberrations, we solve for individual pupil aberrations at different spatial regions across the FOV and assume the aberrations are locally space-invariant within each region. Given Eqs. (5.2)-(5.5), an objective function for the joint optimization of absorption, phase and pupil aberrations can be formulated as:

$$\min_{\boldsymbol{\mu},\boldsymbol{\phi},\boldsymbol{c}} \sum_{s=1}^{N_s} \left\| \mathbf{F}\boldsymbol{I}_{\mathrm{n},s} - \mathrm{diag}\left(\boldsymbol{H}_{\mu,s}(\boldsymbol{c})\right) \mathbf{F}\boldsymbol{\mu} - \mathrm{i} \cdot \mathrm{diag}\left(\boldsymbol{H}_{\phi,s}(\boldsymbol{c})\right) \mathbf{F}\boldsymbol{\phi} \right\|_2^2 + \tau \mathcal{R}(\boldsymbol{\mu},\boldsymbol{\phi}) \ , \tag{5.6}$$

where s is the measurement index of each corresponding source pattern, N_s is the total number of measurements, $\|\cdot\|_2$ represents an ℓ_2 norm, τ is a regularization parameter and R is a regularization term that helps mitigate noise artifacts. It is inferred from Eq. (5.6) that the aberration coefficients $\boldsymbol{c}$ are coupled with both $\boldsymbol{\mu}$ and $\boldsymbol{\phi}$, so simultaneously optimizing all variables does not guarantee convergence. An alternating projection update strategy instead provides a non-divergence guarantee, as was previously used for phase-from-focus joint source recovery [146]. Similarly, we iteratively solve for both the complex object and system aberrations as two sub-problems.

Our alternating projections algorithm initializes $\boldsymbol{c}$ with zero. At the start of one iteration, the Zernike coefficients $\boldsymbol{c}$ are fixed and a DPC deconvolution sub-procedure (Section 5.2) is performed in order to update the estimates of amplitude, $\exp(\boldsymbol{\mu})$, and phase, $\boldsymbol{\phi}$. This new complex-field estimate is then held fixed while an aberration estimation procedure (Section 5.2) is performed. After updating the aberration estimate based on this procedure, a new iteration begins. Eventually, the objective function converges to a stationary point, giving the final estimates of amplitude, phase, and aberration coefficients. In general, this optimization strategy works as long as there exist enough diversity and redundancy in the measurements.

DPC phase retrieval sub-procedure

The phase retrieval sub-procedure amounts to solving the conventional DPC inverse problem [49,81], except that it incorporates the current estimate of the Zernike coefficients $\boldsymbol{c}_k$ at iteration k:

$$\boldsymbol{\mu}_{k+1}, \boldsymbol{\phi}_{k+1} = \arg\min_{\boldsymbol{\mu},\boldsymbol{\phi}} \sum_{s=1}^{N_s} \left\| \mathbf{F}\boldsymbol{I}_{\mathrm{n},s} - \mathrm{diag}\left(\boldsymbol{H}_{\mu,s}(\boldsymbol{c}_k)\right) \mathbf{F}\boldsymbol{\mu} - \mathrm{i} \cdot \mathrm{diag}\left(\boldsymbol{H}_{\phi,s}(\boldsymbol{c}_k)\right) \mathbf{F}\boldsymbol{\phi} \right\|_2^2 + \tau \mathcal{R}(\boldsymbol{\mu},\boldsymbol{\phi}) \ . \tag{5.7}$$

The regularization term, R, should be chosen based on *a priori* information about the sample. For instance, Tikhonov regularization can mitigate noise, and the solution of Eq. (5.7) can then be found using a non-iterative deconvolution [81]. If the gradients of the object are relatively sparse, the Total Variation (TV) regularizer can be used to reduce noise without degrading edges. Since the regularization term for TV is not differentiable, iterative algorithms for solving Eq. (5.7) are needed. In this paper, we use the Alternating Direction Method of Multipliers (ADMM) to implement TV regularization [34], typically requiring about 20 iterations.

Aberration recovery sub-procedure

The aberration recovery sub-procedure uses a nonlinear optimization algorithm to update the aberration estimate based on the newly updated complex-field, which is held fixed. The sub-procedure is initialized with the Zernike coefficients estimate from the previous iteration, $\boldsymbol{c}_k$. Mathematically, the sub-procedure problem is written as:

$$\boldsymbol{c}_{k+1} = \arg\min_{\boldsymbol{c}} \sum_{s=1}^{N_s} \left\| \mathbf{F}\boldsymbol{I}_{\mathrm{n},s} - \mathrm{diag}\left(\boldsymbol{H}_{\mu,s}(\boldsymbol{c})\right)\mathbf{F}\boldsymbol{\mu}_{k+1} - \mathrm{i}\cdot\mathrm{diag}\left(\boldsymbol{H}_{\phi,s}(\boldsymbol{c})\right)\mathbf{F}\boldsymbol{\phi}_{k+1} \right\|_2^2. \qquad (5.8)$$

Equation (5.8) may be solved by a gradient descent (first-order optimization) approach, or more sophisticated second-order optimization routines (*e.g.* Newton's method [146]). All of these require computation of the gradient of the objective function with respect to $\boldsymbol{c}$. If we define the cost function as $f = \sum_{s=1}^{N_s} \|\varepsilon_s\|_2^2$, in which $\varepsilon_s = \mathbf{F}\boldsymbol{I}_{\mathrm{n},s} - \mathrm{diag}(\boldsymbol{H}_{\mu,s})\mathbf{F}\boldsymbol{\mu}_{k+1} - \mathrm{i}\cdot\mathrm{diag}(\boldsymbol{H}_{\phi,s})\mathbf{F}\boldsymbol{\phi}_{k+1}$ is the residual vector, the gradient becomes $\nabla_{\boldsymbol{c}}f = \sum_{s=1}^{N_s} [\partial\varepsilon_s/\partial\boldsymbol{c}]^{\mathrm{H}}\varepsilon_s$, where H denotes Hermitian transpose. Using Eqs. (5.3)-(5.5), the gradient can be calculated analytically as:

$$\nabla_{\boldsymbol{c}}f = \frac{\mathrm{i}}{\boldsymbol{I}_{\mathrm{o}}}\boldsymbol{Z}^{\mathrm{T}}\mathrm{diag}(P^*)\sum_{s=1}^{N_s}\left[\boldsymbol{\mathcal{F}}^{-1}\mathrm{diag}(\boldsymbol{O}_s)\mathbf{F}\left[\mathrm{diag}(\mathbf{F}^*\boldsymbol{\mu}_{k+1}^*) - \mathrm{i}\,\mathrm{diag}(\mathbf{F}^*\boldsymbol{\phi}_{k+1}^*)\right] + \right.$$
$$\left. \mathrm{diag}(\boldsymbol{S}_s)\boldsymbol{\mathcal{F}}^{-1}\mathrm{diag}(\mathbf{F}\boldsymbol{P})\mathbf{F}\left[\mathrm{diag}(\mathbf{F}^*\boldsymbol{\mu}_{k+1}^*) + \mathrm{i}\,\mathrm{diag}(\mathbf{F}^*\boldsymbol{\phi}_{k+1}^*)\right]\right]\varepsilon_s.$$
$$(5.9)$$

In this gradient, T denotes transpose and the Zernike basis, $\mathbf{Z}^{\mathrm{T}} = [\boldsymbol{Z}_0, \boldsymbol{Z}_1, \ldots, \boldsymbol{Z}_M]^{\mathrm{T}}$, contains a finite number of modes where $\boldsymbol{Z}_0 \ldots \boldsymbol{Z}_M$ are the vectorized Zernike modes. For efficient computation, we adopt the L-BFGS algorithm [167] and use the gradient in Eq. (5.9) to solve this nonlinear optimization problem, which generally takes ~ 10 iterations to converge.

Simulation results

To verify the performance of our joint estimation framework, we show simulation results in Fig. 5.2. The system parameters were chosen to match our experimental setup (0.4NA,

wavelength $514nm$, 177 source LEDs), with the LED array placed sufficiently far away from the sample such that the illumination from each LED is effectively spatially coherent (plane wave) [13, 94, 153].

We compare our results with joint phase and aberration recovery FPM in Fig. 5.2. FPM captures a separate image for each of the 177 LEDs, whereas DPC requires only 4 images to reconstruct the same quantities. FPM intensity images are simulated by using different tilted plane wave illuminations corresponding to each single LED. All the intensity images contain the same pupil aberration, which is a weighted sum of the first 21 Zernike modes. Our simulated DPC measurements are the sum of the intensity images from half-circle blocks of LEDs on the top, bottom, or right regions of the LED array [49]. For all measurements, we added synthesized noise using a Poisson distribution with a mean of $\sim$3000 photons per pixel. Equation (5.6) was then solved with ℓ_2 regularization using the 4 DPC images, while we implemented the same algorithm in [90] to recover complex-field and aberrations using FPM with the full 177 image dataset. In this setup, FPM could use as few as 32 images (illuminations from the outer-most annular LEDs only) to achieve the same spatial frequency coverage as our DPC method. Therefore, we also include the results of FPM with 32 measurements for comparison.

Reconstructions from both FPM and our DPC algorithm match the ground truth (See Fig. 5.2(b)); however, our method only requires 4 measurements, reducing acquisition time and memory requirements. Figure 5.2(c) plots the normalized root-mean-square error ($\|\boldsymbol{x} - \boldsymbol{x}_{\text{true}}\|_2/\|\boldsymbol{x}_{\text{true}}\|_2$) at each iteration for both complex-field and pupil aberrations. FPM incurs lower complex-field error than DPC, likely due to both the weak scattering approximation and the larger dataset. As for the pupil aberration error, FPM performance varies significantly with dataset size. While FPM with the full dataset outperforms the proposed DPC framework, our method provides a better result than FPM with 32 images. This is because FPM with fewer measurements has lower effective SNR, adding noise to the recovered pupil aberration (inset of Fig. 5.2(b)). Our method requires significant computation since it solves two optimization problems at each iteration, but the computation time is comparable to a sequential FPM reconstruction. Both methods were implemented in MATLAB on a desktop computer (Intel Core i7 CPU, Nvidia Tesla C2075 GPU). With a 650×584 pixel object, each iteration took 2.2s for FPM and 2.5s for our DPC algorithm.

Experimental results

Experimentally, we use an LED array microscope with the illumination module replaced with a custom-built LED array ($\lambda = 0.514\mu m$) [49,94]. A phase target (Benchmark Technologies), which contains periodic patterns of continuous spatial frequencies, is imaged by a $20\times$ 0.4NA objective lens (Nikon, CFI Plan Achro) in a Nikon TE300 microscope and images are recorded by a PCO.edge 5.5 sCMOS camera on the front port of the microscope (which adds $2\times$ magnification). To test the weak phase gradient assumption, phase images of 6 resolution targets of different heights are recovered using DPC. After validating the reconstructed phase values against theoretical ones, we find that DPC provides accurate results when the
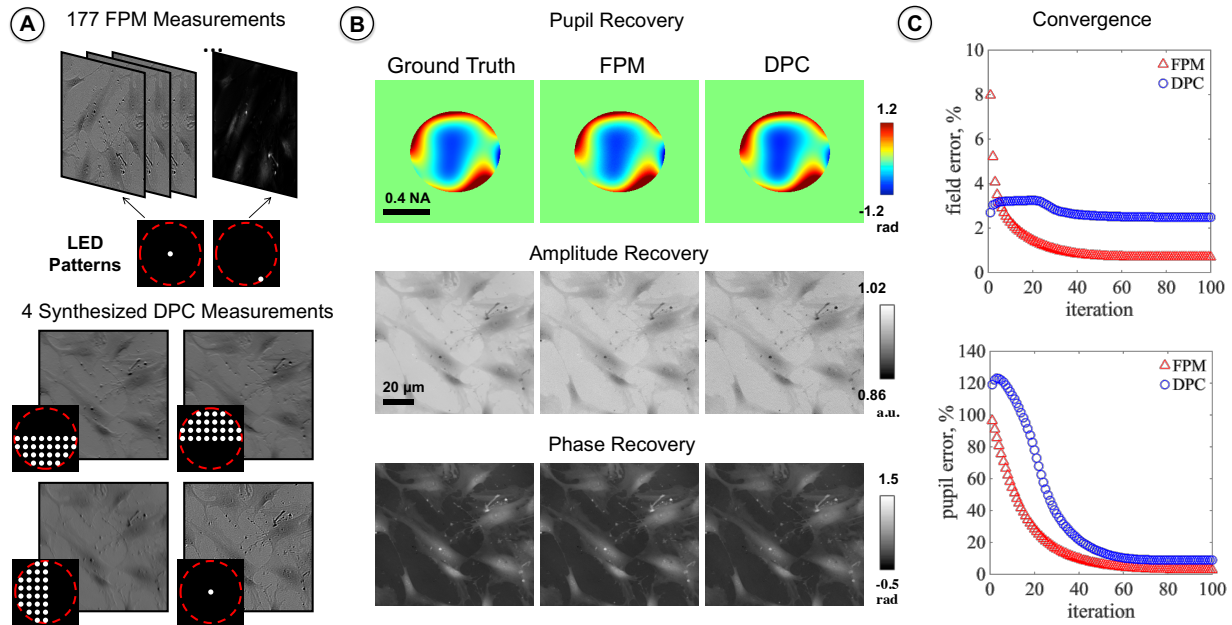
Figure 5.2: Performance of joint phase and aberrations estimation on a simulated dataset. (a) Simulated FPM and DPC measurements. Red dashed circles indicate the NA of the objective lens. (b) Joint estimation of optical field and pupil aberrations, comparing ground truth, FPM and DPC measurements. (c) Errors for complex-field and aberrations at each iteration.

phase of the sample is below 0.64 radians, then underestimates the phase values due to breakdown of the approximation (see Visualization 1). We collect 177 measurements by scanning individual LEDs within a maximum 0.4NA illumination angle. To provide a fair comparison, we use the same measurements to synthesize the 4 images for our method. As shown in Fig. 5.3(a), DPC measurements with half-circle source patterns have high resolution and qualitatively reveal the phase gradients of the sample, while measurements with single-LED illumination have lower resolution. Two single-LED image zoom-ins that contain the same structure at different orientations are shown in Fig. 5.3(a). One has high contrast, while the other does not, because of directional aberrations. By processing the images using the FPM algorithm and the proposed method with TV regularization, we recover the phase of the sample as shown in Fig. 5.3(b). The FPM and DPC reconstructions are similar and can resolve features with period as small as $\lambda/\left(2 \times NA\right) = 0.643\mu m$. In addition, both results provide reliable quantitative phase of the sample. The refractive index of the binary phase target is 1.52, with height of $100nm$, resulting in ~ 0.64 radians peak-to-valley. Looking at the 1D cut-lines (taken along dashed lines) in Fig. 5.3(b), after subtracting the mean of each, the reconstructions show good agreement with the ideal height.

Pupil aberrations recovered by both FPM and DPC algorithms are in Fig. 5.3(c). While we have no ground truth, aberrations estimated from FPM and DPC match well within the 4^{th} radial degree of Zernike modes. The dominant aberration in the objective lens is the 8^{th} Zernike mode (horizontal coma); this agrees with the evidence of directional aberration mentioned above. One reason causing a difference between the reconstructed pupils is the high-frequency fluctuation shown in the pupil aberration from FPM, which does not exist in the low-order Zernike modes. Although high-order aberrations might be estimated with more measurements to avoid overfitting, it's usually enough to improve image quality by correcting the low-order aberrations.
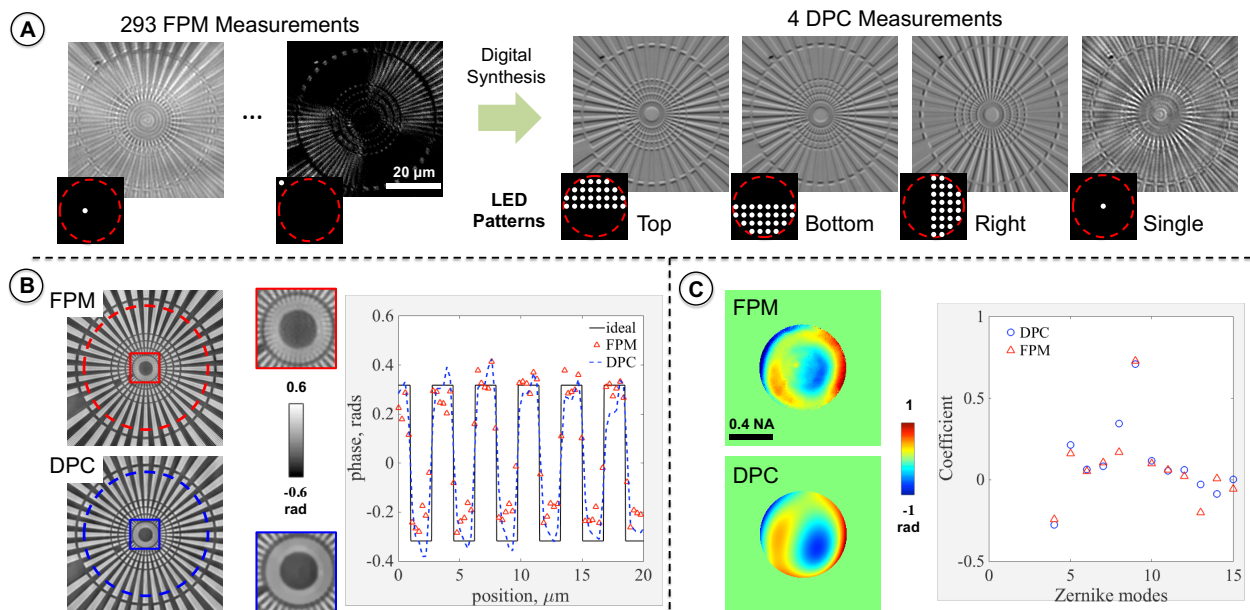
Figure 5.3: (a) Experimental FPM and DPC measurements for different LED source patterns. Zoomed regions at different orientations for coherent illumination are marked in cyan and pink boxes, respectively. (b) Quantitative phase of a star target using FPM and DPC, along with 1D cutlines for FPM (red) and DPC (blue) along the dashed lines. (c) Reconstructed wavefront error function and the weights of each Zernike mode up to the 4^{th} radial degree.

To quantify the performance of our method, we introduced known defocus aberrations using an axial motion stage (Thorlabs, MZS500-E). We translated the phase test target across a range of known axial steps over a total range of $40\mu m$. Each $2\mu m$ of translation, the 4 images for our method were acquired using a $10\times$ 0.25NA objective lens (Nikon, CFI Plan Achro) at an acquisition rate of 20Hz. Quantitative phase reconstructions assuming zero aberration, aberration-corrected phase images and the recovered aberrations are shown

in Fig. 5.4(a). With increased defocus, the uncorrected phase results degrade. Though our method also suffers from resolution degradation at large defocus distances ($\pm 20\mu m$), it provides better resolution than the uncorrected phase retrieval. For example, in Fig. 5.4(b) the numbers in group 9 are resolved with pupil correction, while they are not readable in the uncorrected result.

We can also recover the focus distances from the defocus term of the Zernike modes and compare the estimated defocus values to known values (Fig. 5.4(c)). Since the Zernike basis is normalized in a range from -1 to 1, the defocus value, d, at each time point, t, can be evaluated as $d(t) = c_4(t) \times \lambda/\pi/(1 - (1 - \text{NA}^2)^{0.5})$. As expected, the predicted pupil aberrations have quadratic form, which indicates that defocus dominates. Experimentally, our method overestimates defocus values by a factor of 1.16. This discrepancy might originate from mis-calibration of the experiment or parameters used in computation (*e.g.* wavelength, NA of the objective lens, precision of the motion stage). This linear relationship between the predicted positions and the true focus holds when the magnitude of defocus is less than 16 μm, $\sim 2\times$ Depth of Field (DoF). However, the accuracy of aberration correction drops as the defocus value exceeds $2\times$ DoF, when the maximum phase difference in the pupil becomes larger than 2π and the algorithm becomes more likely to converge to a local minimum. From these experiments, we conclude that our method accurately estimates time-varying aberrations within a total range of $4\times$ DoF.

Our method is easily extended to account for spatially-varying aberrations, simply by solving the joint estimation problem separately over different patches of the FOV. In order to visualize spatially-varying aberrations, we prepared oil (Cargille, $n_D = 1.58$) immersed $10\mu m$ polystyrene beads (Sigma-Aldrich) as our sample, which is assumed to be nearly spatially invariant. The sample was imaged by a $4\times$ 0.2NA objective lens (Nikon, CFI Plan Apo Lambda) that has a FOV of 1.7×2.1mm. Four images were captured at 12.5Hz using the same source patterns as in Fig. 5.3(a). Field-dependent aberrations are primarily caused by two types of system imperfections. First, the incident angle of individual LEDs changes from one FOV to another, since the relative position between the LEDs and the sample at each FOV varies. Second, there exist spatially-varying aberrations native to the optical system. To accommodate this spatial variance, we break the full FOV into small patches, and apply the joint estimation algorithm on each patch independently. In this case, the incident angle of individual LEDs and the aberrations are assumed invariant within the patch, which is usually valid when the region is much smaller than the total FOV. In Fig. 5.5(a), the full FOV is divided into 100 patches. Within each patch, the absorption, phase and pupil aberration were recovered independently before being stitched together to form the full FOV images. In practice, optimization for each patch converges in $20 - 30$ iterations.

Looking at the aberrations recovered, the center of the FOV is essentially aberration-free, which is consistent with the fact that optical systems are usually optimized there. Therefore, the absorption and phase images with pupil recovery are similar to that without aberration correction. Aberrations are much stronger along the edges of the FOV, where the field curvature blurs the image. Consequently, obvious differences occur between the normal DPC and aberration-corrected DPC reconstructions at those patches. The recovered amplitude
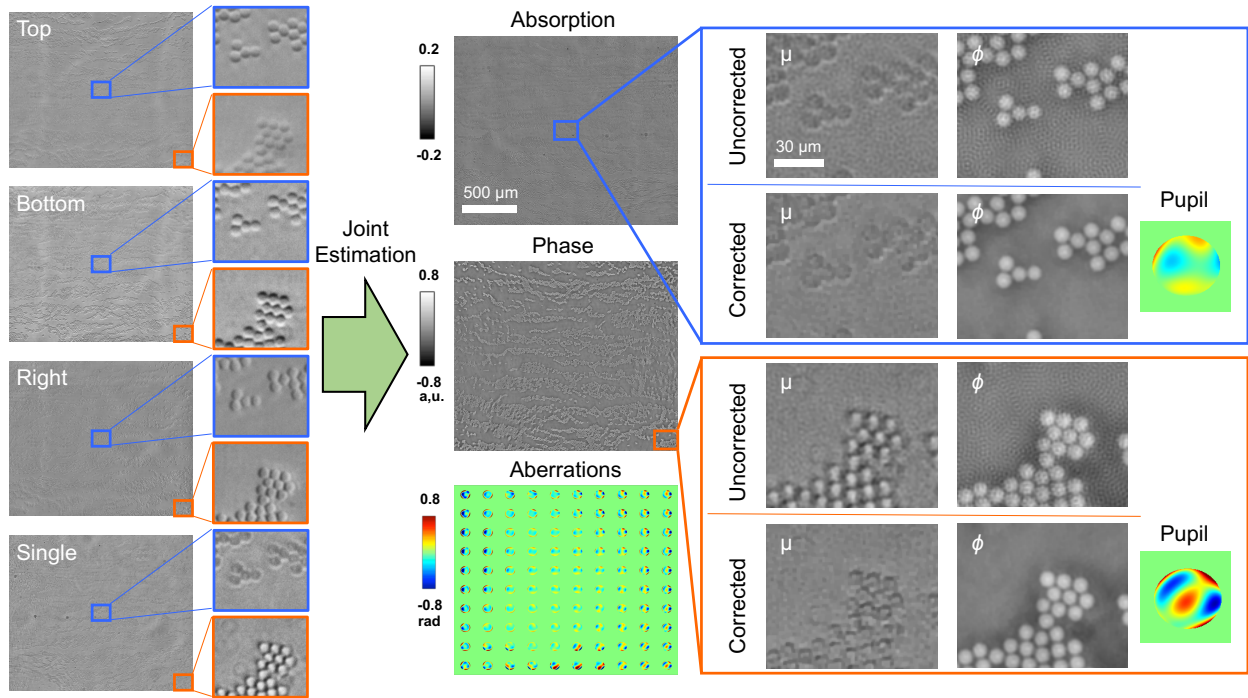
Figure 5.4: (a) Quantitative phase reconstructions of a USAF 1951 resolution target at various defocus distances with and without aberration correction, and the corresponding recovered aberrations (see Visualization 2). (b) Zoomed-in reconstructions at defocus of $10\mu m$. (c) Known and experimentally-estimated defocus values from the 4^{th} Zernike mode over time.

(absorption) at the bottom right corner is dramatically changed after applying aberration correction. Qualitatively, the uncorrected absorption at the edge of the field appears to contain some phase information, leading to a shadow-like appearance; these artifacts are removed with pupil recovery. In addition, while aberrated phase evaluated with Tikhonov deconvolution suffered from high-frequency noise, phase with pupil correction shows reduced noise due to the use of TV regularization.

In computational imaging systems, system aberrations can significantly degrade quantitative phase and absorption reconstructions if not properly compensated. In order to correct the aberration without conducting additional measurements to calibrate the system, joint estimation of the sample and the pupil function is needed. While Fourier ptychography or coded aperture imaging provides pupil recovery, these methods often require many measurements. In this paper, we demonstrated a DPC-based phase retrieval technique that simultaneously recovers the system pupil function and quantitative phase and absorption of a sample with only 4 intensity images. By combining 3 DPC images as well as 1 measurement made
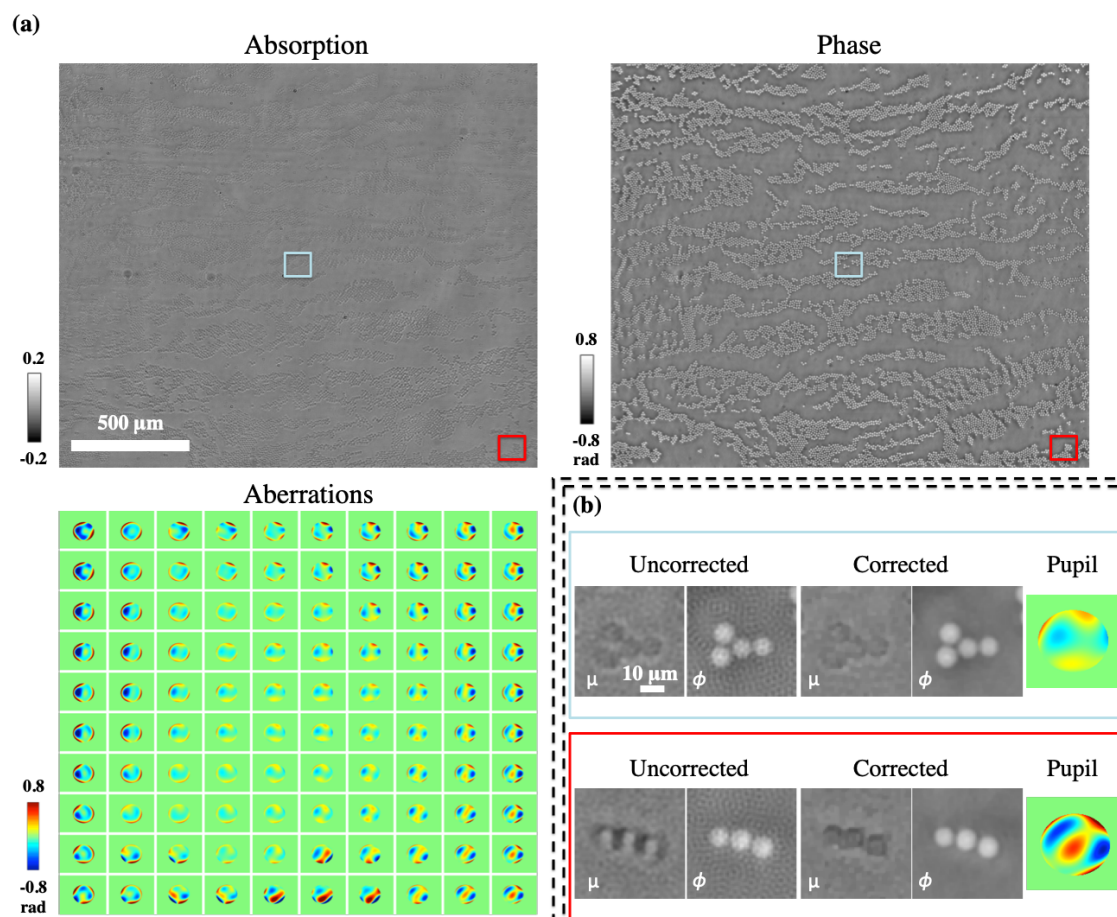
Figure 5.5: (a) Reconstructed absorption, phase and spatially-varying aberrations (recovered pupil wavefronts for different regions of the field-of-view). (b) Comparison of results with and without pupil estimation for central and edge regions of the field-of-view.

under coherent illumination, diffraction-limited field and aberration with arbitrary Zernike orders were acquired through an alternating non-convex optimization. This method not only reduces the data acquisition time, but also increases signal-to-noise ratio due to higher light throughput compared to single-LED acquisition. The recovered system aberrations are general to the system and may be used for image correction in other imaging modalities, such as deconvolution of fluorescence images [168].

5.3   Source Self-Calibration using Fourier Ptychography

Computational imaging leverages the power of both optical hardware and computational algorithms to reconstruct images from indirect measurements. In optical microscopy, programmable sources have been used for computational illumination techniques including multi-contrast [45, 164], quantitative phase [13, 49, 65, 169] and super-resolution [13, 90, 170, 171]. Implementation is simple, requiring only an inexpensive source attachment for a commercial microscope. However, these methods are also sensitive to experimental misalignment errors and can suffer severe artifacts due to model mismatch. Extensive system calibration is needed to ensure that the inverse algorithm is consistent with the experimental setup, which can be time- and labor-intensive. This often requires significant user expertise, making the setup less accessible to reproduction by non-experts and undermining the simplicity of the scheme. Further, pre-calibration methods are not robust to changes in the system (e.g. bumping the setup, changing objectives, sample-induced aberrations) and require precise knowledge of a ground-truth test object.

Algorithmic self-calibration methods [147–149, 152–154, 172–180] eliminate the need for pre-calibration and precise test objects by making calibration part of the inverse problem. These methods jointly solve two inverse problems: one for the reconstructed image of the object and the other for the calibration parameters. By recovering system calibration information directly from captured data, the system becomes robust to dynamic changes in the system.

Here, we focus on *illumination angle* self-calibration for Fourier Ptychographic Microscopy (FPM) [13]. FPM is a coherent computational imaging method that reconstructs high-resolution amplitude and phase across a wide field-of-view (FOV) from intensity images captured with a low-resolution objective lens and a dynamically-coded illumination source. Images captured with different illumination angles are combined computationally in an iterative phase retrieval algorithm that constrains the measured intensity in the image domain and pupil support in the Fourier domain. This algorithm can be described as stitching together different sections of Fourier space (synthetic aperture imaging [181, 182]) coupled with iterative phase retrieval. FPM has enabled fast *in vitro* capture via multiplexing [49, 90], fluorescence imaging [154], and 3D microscopy [139, 183]. It requires significant redundancy (pupil overlap) in the dataset [171, 184], making it suitable for joint estimation self-calibration.

Self-calibration routines have previously been developed to solve for pupil aberrations [152, 153, 172], illumination angles [148, 173–176], LED intensity [147], sample motion [149], and auto-focusing [177] in FPM. The state-of-the-art self-calibration method for illumination angles is simulated annealing [148, 173], a joint estimation solution which (under proper initialization) removes LED misalignment artifacts that usually manifest as low-frequency noise. Unfortunately, because the simulated annealing procedure operates inside the FPM algorithm iterative loop, it slows the run-time of the solver by an order of magnitude or more. For 3D FPM (which is particularly sensitive to angle calibration [139]), the computational
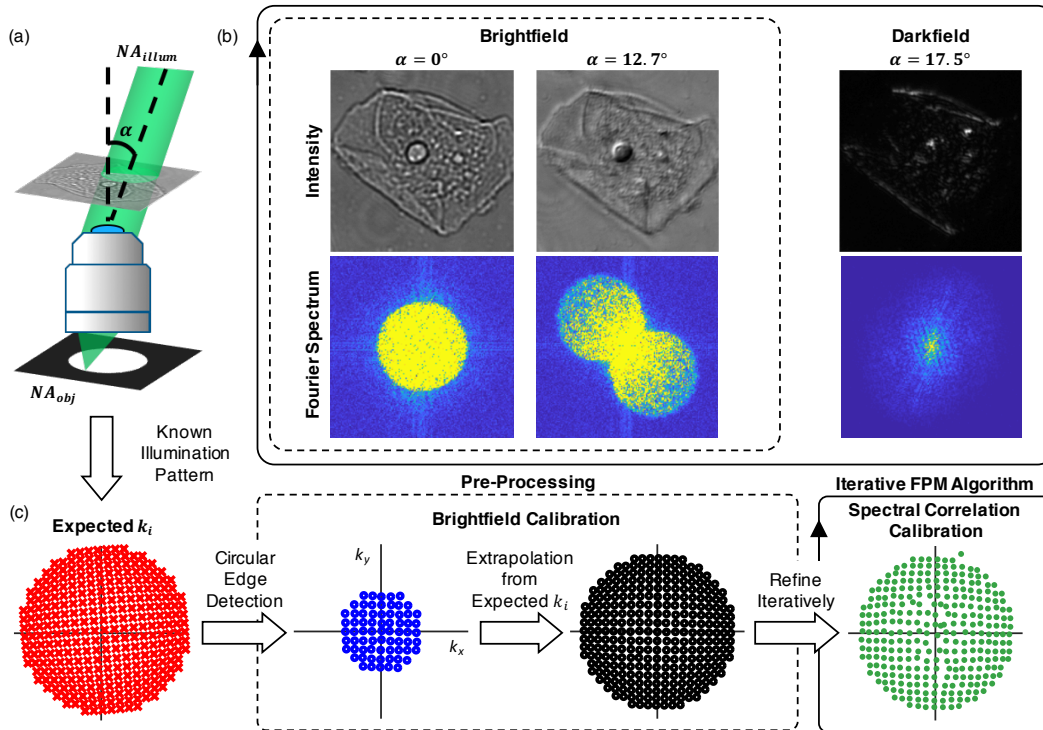
Figure 5.6: Illumination angles are calibrated by analyzing Fourier spectra. (a) A cheek cell is illuminated at angle α and imaged with NA_{obj}. (b) Brightfield images contain overlapping circles in their Fourier spectra; darkfield images do not. (c) We perform a fast and efficient brightfield calibration in pre-processing, then extrapolate the correction to darkfield images and, finally, iteratively calibrate angles inside the FPM algorithm using a spectral correlation calibration.

costs become infeasible.

Moreover, most self-calibration algorithms require a relatively close initial guess for the calibration parameters. This is especially true when the problem is non-convex or if multiple calibration variables are to be solved for (*e.g.* object, pupil, and angles of illumination). Of the relevant calibration variables for FPM, illumination angles are the most prone to error, due to shifts or rotations of the LED array [118], source instabilities [178, 185], non-planar illuminator arrangements [73, 79, 168, 186], or sample-induced aberrations [187, 188]. Sample-induced aberrations can also change the effective illumination angles dynamically, such as when the sample is in a moving aqueous solution.

We propose here a two-pronged angle self-calibration method that uses both pre-processing (*brightfield calibration*) and iterative joint estimation (*spectral correlation calibration*) that is quicker and more robust to system changes than state-of-the-art calibration methods[2]. A

---

[2]This work was developed in close collaboration with fellow Ph.D. student Regina Eckert (Waller Lab, EECS, UC Berkeley)

circle-finding step prior to the FPM solver accurately identifies the angles of illumination in the brightfield (BF) region. A transformation between the expected and BF calibrated angles extrapolates the correction to illuminations in the darkfield (DF) region. Then, a local grid-search-based algorithm inside the FPM solver further refines the angle estimates, with an optional prior based on the illuminator geometry (Fig. 5.6). Our method is object-independent, robust to coherent noise, and time-efficient, adding only seconds to the processing time. We demonstrate on-line angle calibration for 2D and 3D FPM with 3 different source types: an LED array, a galvanometer-steered laser, and a high-NA (max $NA_{illum} = 0.98$) quasi-dome illuminator [79].

Methods

The image formation process for a thin sample under off-axis spatially coherent plane wave illumination can be described by:

$$I_i(r) = |O(r)e^{-i2\pi k_i r} * P(r)|^2 = |\mathbf{F}^{-1}(\tilde{O}(k - k_i)\tilde{P}(k))|^2, \tag{5.10}$$

where k_i is the spatial frequency of the incident light, $\tilde{P}(k)$ is the system pupil function, $\tilde{O}(k)$ is the object Fourier spectrum, and $\mathbf{F}$ is the 2D Fourier transformation operation, valid for shift-invariant systems. Intensity images are captured at the sensor plane, corresponding to auto-correlation in the Fourier domain:

$$\begin{aligned}
\tilde{I}_i(k) &= \mathbf{F}(|O(r)e^{-i2\pi k_i r} * P(r)|^2) \\
&= \tilde{O}(k - k_i)\tilde{P}(k) \star \tilde{O}(k - k_i)\tilde{P}(k),
\end{aligned} \tag{5.11}$$

where $*$ denotes convolution and $\star$ denotes auto-correlation. $\tilde{O}(k - k_i)\tilde{P}(k)$ corresponds to the shifted spectrum of the object within the circle $|k| \leq \frac{NA_{obj}}{\lambda}$ and 0 everywhere else. The auto-correlation operation essentially scans two copies of $\tilde{O}(k - k_i)\tilde{P}(k)$ across each other, coherently summing at each shift to give $\tilde{I}_i(k)$. Typically, the object spectrum has a large zero-order (DC) term that decays sharply towards higher frequencies. In the brightfield region, when the DC term at k_i is within the pupil's passband, the auto-correlation effectively scans this DC term across the conjugate spectrum, giving high values where the DC term overlaps with the conjugate pupil and negligible signal elsewhere. This interference between the DC term and pupil in the auto-correlation creates two distinct circles centered at k_i and $-k_i$ in the intensity spectrum amplitude (Fig. 5.6). Hence, we can calibrate the illumination angle by finding these circle centers. For darkfield images, the DC term is outside $\frac{NA_{obj}}{\lambda}$ and so we do not observe clearly defined circles in $|\tilde{I}_i|$ (Fig.5.6b), making calibration more complicated.

Our algorithm relies on analysis of the raw intensity Fourier transform to recover illumination angles. Fourier domain analysis of intensity images has been used previously to deduce aberrations [189] and determine the center of diffraction patterns [190, 191] for system calibration. We show here that the individual Fourier spectra can be used to accurately determine illumination angles in both the brightfield and darkfield regime.

Figure 5.7: Circular edge detection on brightfield images finds circle centers, giving illumination angle calibration. (a,b) Comparison of uncalibrated (red) and calibrated (black) illumination $\boldsymbol{k}_i$. The blue box in (b) indicates the search range for $\boldsymbol{k}_i$. (c,d) $\tilde{\boldsymbol{I}}_i$ along radial lines, $f(r, \phi_n)$, and derivatives with respect to r. (e,f) $\boldsymbol{E}_1$ and $\boldsymbol{E}_2$, sums of the derivatives at known radii R and $R + \sigma$, peak near the correct center. Boxes show uncalibrated (red) and calibrated (black) $\boldsymbol{k}_i$ centers.

Brightfield Calibration

Locating the center of the circles in the amplitude of a Fourier spectrum is an image processing problem. Previous work in finding circles in images uses the Hough transform, which relies on an accurate edge detector as an initial step [192, 193]. In practice, however, we find that edge detectors do not function well on our datasets due to speckle noise, making the Hough transform an unreliable tool for our purpose.

Intuitively, circular edge detection can be understood as performing edge detection (*i.e.*

calculating image gradients) along a circular arc around a candidate center point in k-space (the Fourier domain). To first approximation, we assume $|\tilde{\boldsymbol{I}}_i|$ is a binary function that is 1 inside the two circles and 0 everywhere else. Our goal is to find the strong binary edge in order to locate the circle center. We need only consider one of the circles, since the intensity image is real-valued and so its Fourier transform is symmetric. Based on information we have about our illumination set-up, we *expect* the illumination spatial frequency (and therefore circle center) for spectrum $\tilde{\boldsymbol{I}}_i$ to be at $\boldsymbol{k}_{i,0} = (k_{x,i,0}, k_{y,i,0})$ (polar coordinates $\boldsymbol{k}_{i,0} = (d_{i,0}, \theta_{i,0})$) (Fig. 5.7a). If this is the *correct* center $\boldsymbol{k}'_i$, we expect there to be a sharp drop in $|\tilde{\boldsymbol{I}}_i|$ at radius R along any radial line $f(r, \phi_n)$ out from $\boldsymbol{k}'_i$ (Fig. 5.7b). This amplitude edge will appear as a peak at $r = R$ in the first derivative of each radial line with respect to r, $f'(r, \phi_n)$ (Fig. 5.7d). Here (r, ϕ_n) are the polar coordinates of the radial line with respect to the center $\boldsymbol{k}_i$, considering the n^{th} of N radial lines.

We identify the correct $\boldsymbol{k}'_i$ by evaluating the summation of the first derivative around the circular arc at $r = R$ from several candidate $\boldsymbol{k}_i = (d_i, \theta_i)$ positions:

$$\boldsymbol{E}_1(R, d_i, \theta_i) = \sum_{n=1}^{N} f'(r = R, \phi_n, d_i, \theta_i). \tag{5.12}$$

When $\boldsymbol{k}_i$ is incorrect, the edges do not align and the derivative peaks do not add constructively at R (Fig. 5.7c). The derivatives at R are all maximized *only* at the correct center $\boldsymbol{k}'_i$ (Fig. 5.7d), creating a peak in $\boldsymbol{E}_1$ (Fig. 5.7e). This is analogous to applying a classic edge filter in the radial direction from a candidate center and accumulating the gradient values at radius R.

In order to bring our data closer to our binary image approximation, we divide out the average spectrum $\text{mean}_i(|\tilde{\boldsymbol{I}}_i|)$ across all i spectra. Because the object remains constant across images while the angles of illumination change, the average spectrum is similar in form to the object's auto-correlation spectrum, with a sharp central peak decaying towards higher frequencies. The resulting normalized spectra contain near-constant circles on top of background from higher-order terms. We then convolve with a Gaussian blur kernel with standard deviation σ to remove speckle noise (Alg. 1.1-2). Experimentally, we choose $\sigma = 2$ pixels, which balances blurring speckle noise and maintaining the circular edge. Under this model, the radial line $f(r, \phi_n)$ from our correct center $\boldsymbol{k}'_i$ can be modeled near the circular edge as a binary step function convolved with a Gaussian:

$$f(r, \phi_n, d'_i, \theta'_i) = \text{rect}(\frac{r}{2R}) * \frac{1}{\sqrt{2\pi}\sigma} e^{\frac{-r^2}{2\sigma^2}}. \tag{5.13}$$

By differentiating through $f'''(r, \phi_n)$ and setting equal to zero, we find the peak of $f'(r, \phi_n)$ still occurs at $r = R$. Additionally, we find that the second derivative $f''(r, \phi_n)$ has a maximum at $r = R + \sigma$. Experimentally, we have found that considering both the first *and* second derivatives increases our accuracy and robustness to noise across a wide variety of datasets. We therefore calculate a second derivative metric,

$$\boldsymbol{E}_2(R + \sigma, d_i, \theta_i) = \sum_{n=1}^{N} f''(r = R + \sigma, \phi_n, d_i, \theta_i), \qquad (5.14)$$

which is jointly considered with Eq. 5.12. We identify candidate centers $\boldsymbol{k}_i$ that occur near the peak of *both* $\boldsymbol{E}_1$ and $\boldsymbol{E}_2$ (Fig. 5.7e-f), then use a least-squares error metric to determine the final calibrated $\boldsymbol{k}_i'$ (Alg. 1.5-9). In practice, we also only consider the non-overlapping portion of the circle's edge, bounding $\boldsymbol{\phi}$.

Until now, we have assumed that the precise radius R of the pupil is known. However, in pixel units, R is dependent on the pixel size of the sensor, p_s, and the system magnification, mag:

$$R = \frac{NA_{obj}}{\lambda} \frac{p_s * M}{mag}, \qquad (5.15)$$

as well as NA_{obj} and λ, where $\tilde{\boldsymbol{I}}_i$ is dimension $M \times M$. Given that mag and NA_{obj} are often imprecisely known but are unchanged across all images, we calibrate the radius by finding the R' which gives the maximum gradient peak $\boldsymbol{E}_1$ across multiple images before calibrating $\boldsymbol{k}_i'$ (Alg. 1.3). A random subset of images may be used to decrease computation time.

Algorithm 1 Brightfield Calibration

1: $\tilde{\boldsymbol{I}}_f \leftarrow |\tilde{\boldsymbol{I}}|/\text{mean}_i(|\tilde{\boldsymbol{I}}_i|)$ ▷ Divide out mean spectrum
2: $\tilde{\boldsymbol{I}}_f \leftarrow \text{gauss}(\tilde{\boldsymbol{I}}_f, \sigma)$ ▷ Smooth speckle
3: $R' \leftarrow \text{argmax}_R \boldsymbol{E}_1(R, d_i, \theta_i), \text{subset }(\tilde{\boldsymbol{I}}_{f,i})$ ▷ Calibrate radius
4: **for** i^{th} image **do** ▷ Circular edge detection
5:      $\boldsymbol{k}_{i,1} \leftarrow (d_i, \theta_i)$ where $\boldsymbol{E}_1$ near max (within 0.1 std)
6:      $\boldsymbol{k}_{i,2} \leftarrow (d_i, \theta_i)$ where $\boldsymbol{E}_2$ near max
7:      $\boldsymbol{k}_i \leftarrow \boldsymbol{k}_{i,1} \cap \boldsymbol{k}_{i,2}$ ▷ Consider both metrics
8:      $\boldsymbol{k}_i' \leftarrow \text{argmin}_{\boldsymbol{k}_i} ||\boldsymbol{I}_i - \mathbf{F}(\tilde{\boldsymbol{I}}_i \cdot \tilde{\boldsymbol{P}}(\boldsymbol{k} - \boldsymbol{k_i}))||_2$ ▷ Evaluate $\boldsymbol{k}_i$
9: **end for**
10: $\mathbf{A}, i_{outliers} \leftarrow \text{RANSAC}(\mathbf{A} = \boldsymbol{k}_i'/\boldsymbol{k}_{i,0})$ ▷ Identify outliers
11: $\boldsymbol{k}_{inliers}^{(0)} \leftarrow \boldsymbol{k}_{inliers}'$ ▷ Initialize for FPM
12: $\boldsymbol{k}_{outliers}^{(0)} \leftarrow \mathbf{A}\boldsymbol{k}_{outliers,0}$
13: $\boldsymbol{k}_{darkfield}^{(0)} \leftarrow \mathbf{A}\boldsymbol{k}_{darkfield,0}$

Spectral Correlation Calibration

While the brightfield (BF) calibration method localizes illumination angles using intrinsic contrast from each measurement, this contrast is not present in high-angle (darkfield) measurements (Fig. 5.6b). Therefore, we additionally solve a more general joint estimation problem to refine the initialization provided by BF calibration, where the object $\boldsymbol{O}(\boldsymbol{r})$, pupil

$P(\boldsymbol{k})$, and illumination angles $\boldsymbol{k_i}$ are optimized within the FPM algorithm. At each inner iteration, we estimate the i^{th} illumination angle by minimizing the FPM objective function with respect to illumination angle (Fig. 5.8a). This step finds the relative k-space location of the current spectrum $\tilde{\boldsymbol{I}}_i$ relative to the overall object, providing an estimate $\boldsymbol{k}_i^{(m)}$ relative to the other illuminator angles $\boldsymbol{k}_j^{(m)}$, $j \neq i$. We call this the spectral correlation method because this optimization implicitly finds $\boldsymbol{k}_i^{(m)}$ which best aligns the i^{th} spectrum with the estimated object spectrum $\tilde{\boldsymbol{O}}(\boldsymbol{k})^{(m)}$.
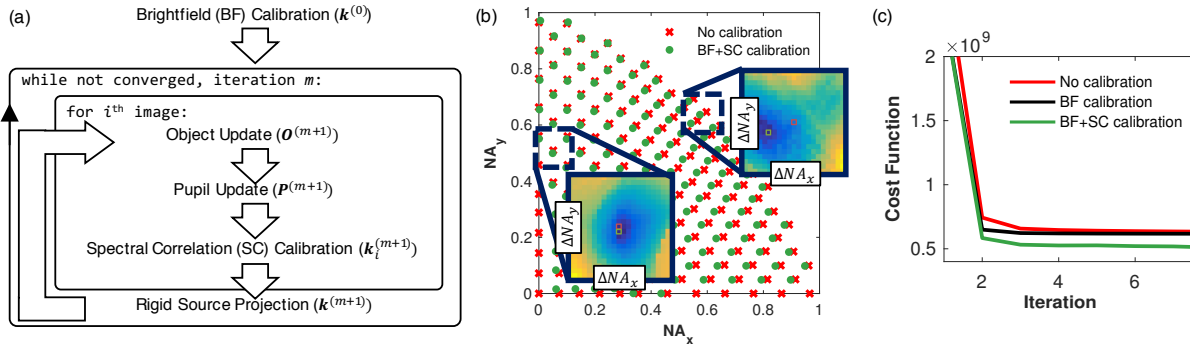
Figure 5.8: BF calibration uses a fast pre-processing step to estimate illumination angles, then SC calibration iteratively refines them within the FPM solver. (a) Algorithm block diagram, (b) uncalibrated (red) and BF + SC calibrated (green) illumination angle map. Insets are example search spaces, showing local convexity. (d) FPM convergence plot for different methods.

Unlike previous joint estimation methods [148, 173], we constrain $\boldsymbol{k}_i$ to exist on the k-space grid defined by our image sampling. Our k-space resolution is band-limited by the size of the image patch, $\boldsymbol{s} = (s_x, s_y)$, across which the illumination can be assumed coherent. This coherent area size is determined by the van Cittert-Zernike theorem, which can be simplified [194] to show that the coherence length l_c of illumination with mean source wavelength $\bar{\lambda}$ produced by a source of size ρ at a distance R is determined by:

$$l_c = \frac{0.61R\bar{\lambda}}{\rho}. \tag{5.16}$$

For example, a $300\mu m$ wide LED placed $50mm$ above the sample with $\bar{\lambda} = 530nm$ gives $l_c = 53.8\mu m$, which provides an upper bound on the size of image patch used in the FPM reconstruction, $(s_x, s_y) \leq l_c$. This limitation imposes a minimum resolvable discretization of illumination angles $\Delta\boldsymbol{k} = \frac{2}{\boldsymbol{s}}$ due to the Nyquist criterion. Since we cannot resolve finer angle changes, we need only perform a local grid search over integer multiples of $\Delta\boldsymbol{k}$, which makes our joint estimation SC method much faster than previous methods.

SC calibration is cast as an iterative optimization of discrete perturbations of the estimated angle using a local grid search. At each FPM iteration, we solve for the optimal pertur-

bation of illumination angle $\boldsymbol{k}_i^{(m)}$ over integer multiples $\boldsymbol{n} = (n_x, n_y)$ of k-space resolution-limited steps $\Delta\boldsymbol{k}$ such that the updated illumination position $\boldsymbol{k}_i^{(m+1)} = \boldsymbol{k}_i^{(m)} + \boldsymbol{n} \cdot \Delta\boldsymbol{k}$ minimizes the $\ell 2$ distance between the object and illumination angle estimates and measurements,

$$\operatorname*{argmin}_{\boldsymbol{n}} \quad ||\boldsymbol{I}_i - |\boldsymbol{O}^{(m+1)}e^{-i2\pi(\boldsymbol{k}_i^{(m)}+\boldsymbol{n}\Delta\boldsymbol{k})\boldsymbol{r}} * \boldsymbol{P}^{(m+1)}|^2||_2^2$$
$$\text{subject to} \quad \boldsymbol{n} = (n_x, n_y), \qquad (n_x, n_y) \in [-1, 0, 1]. \tag{5.17}$$

This grid search is performed iteratively within each sequential iteration of an FPM reconstruction until $\boldsymbol{k}_i$ converges, giving a lower reconstruction cost than BF calibration alone (Fig. 5.8b-c).

The choice of $\boldsymbol{n} = (n_x, n_y)$ to search can be tuned to match the problem. In most experimental cases, we find that a search of the immediate locality of the current estimate ($(n_x, n_y) \in [-1, 0, 1]$) gives a good balance between speed and gradient performance when paired with the close initialization from our BF calibration. A larger search range (e.g. $(n_x, n_y) \in [-2, -1, 0, 1, 2]$) may be required in the presence of noise or without a close initialization, but the number of points searched will increase with the square of the search range, causing the algorithm to slow considerably.

Including prior information about the design of the illumination source can make our calibration problem more well-posed. For example, we can include knowledge that an LED array is a rigid, planar illuminator in our initial guess of the illumination angle map, $\boldsymbol{k}_{i,0}$. By forcing the current estimates $\boldsymbol{k}_i^{(m)}$ to fit a transformation of this initial angle map at the end of each FPM sub-iteration, we can use this knowledge to regularize our optimization (Fig. 5.8a). The transformation model used depends on the specific illuminator. For example, our quasi-dome LED array is composed of five circuit boards with precise LED positioning within each board but variable board position relative to each other. Thus, imposing an affine transformation from the angle map of each board to the current estimates $\boldsymbol{k}_i^{(m)}$ significantly reduces the problem dimensionality and mitigates noise across LEDs, making the reconstruction more stable.

Results

Planar LED Array

We first show experimental results from a conventional LED array illumination system with a $10\times$, 0.25 NA and a $4\times$, 0.1 NA objective lens at $\lambda = 514nm$ and $NA_{illum} \leq 0.455$ (Fig. 5.9). We compare reconstructions with simulated annealing, our BF pre-processing alone, and our combined BF+SC calibration method. All methods were run in conjunction with EPRY pupil reconstruction [153]. We include results with and without the SC calibration to illustrate that the BF calibration is sufficient to correct for most misalignment of the LED array since we can accurately extrapolate LED positions to the darkfield region when the LEDs fall on a planar grid. However, when using a low NA objective ($NA_{obj} \leq 0.1$), as in Fig. 5.9d,
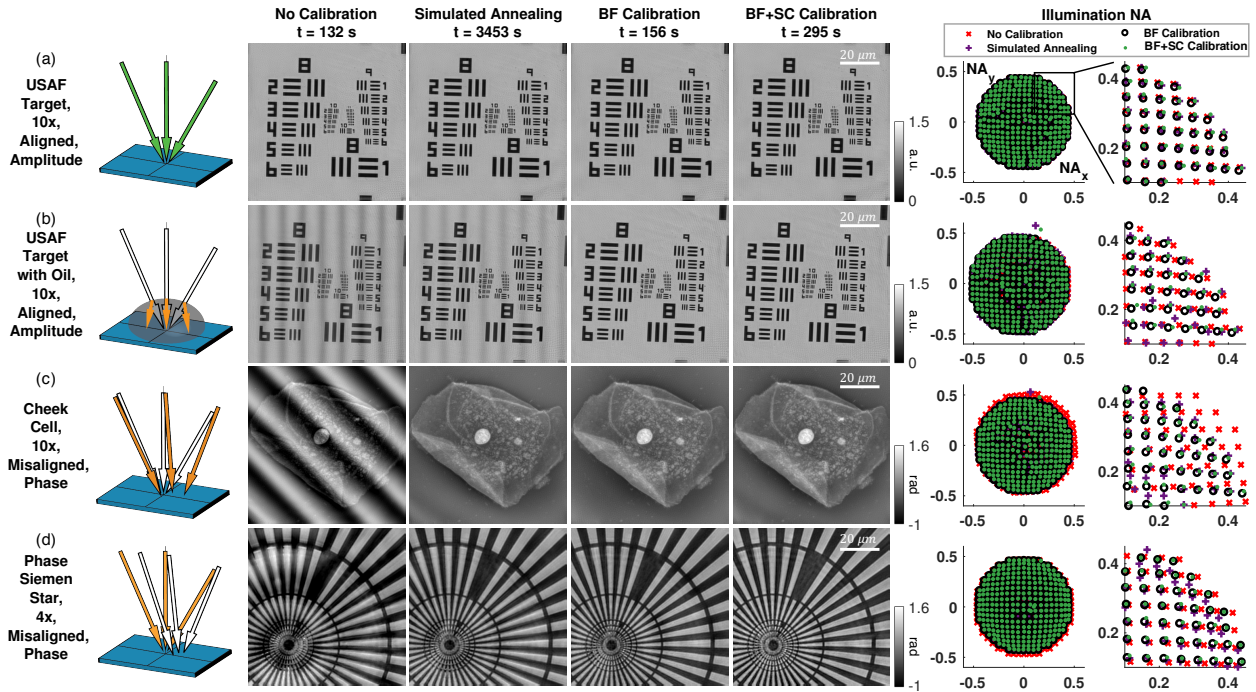
Figure 5.9: Experimental results with an LED array microscope, comparing reconstructions with no calibration (average reconstruction time 132 seconds), simulated annealing (3453 s), our BF calibration (156 s), and our BF + SC calibration (295 s). (a) Amplitude reconstructions of a USAF target in a well-aligned system. (b) Amplitude reconstructions of the same USAF target with a drop of oil placed on top of the sample to simulate sample-induced aberrations. (c) Phase reconstructions of a human cheek cell with computationally misaligned illumination, and (d) a Siemens star phase target with experimentally misaligned illumination.

the SC method becomes necessary because the BF calibration is only able to use 9 images (compared to 69 brightfield images with a $10\times$, 0.25 NA objective, as in Fig. 5.9a-c).

Our method is object-independent, so can be used for phase and amplitude targets as well as biological samples. All methods reconstruct similar quality results for the well-aligned LED array with the USAF resolution target (Fig. 5.9a). To simulate an aqueous sample, we place a drop of oil on top of the resolution target. The drop causes uneven changes in the illumination, giving low-frequency artifacts in the uncalibrated and simulated annealing cases which are corrected by our method (Fig. 5.9b). Our method is also able to recover a $5°$ rotation, 0.02 NA shift, and $1.1\times$ scaled computationally-imposed misalignment on well-aligned LED array data for a cheek cell (Fig. 5.9c), and gives a good reconstruction of an experimentally misaligned LED array for a phase Siemens star (Benchmark Technologies, Inc.) (Fig. 5.9d). In contrast to simulated annealing, which on average takes $26\times$ as long to
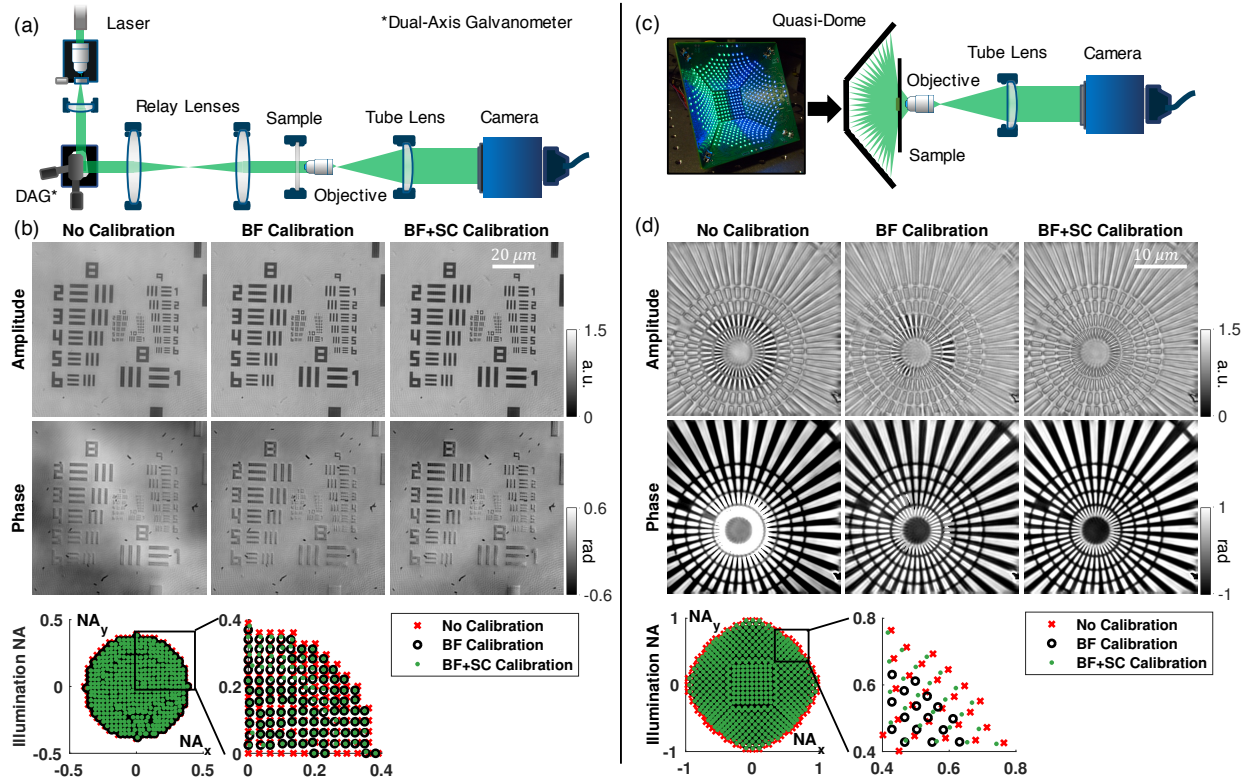
Figure 5.10: Experimental angle calibration in laser and high-NA quasi-dome illumination systems. (a) Laser illumination is steered by a dual-axis galvanometer. The angled beam is relayed to the sample by 4", 80 mm focal length lenses. (b) Our calibration method removes low-frequency reconstruction artifacts. (c) The quasi-dome illuminator enables up to 0.98 NA_{illum} using programmable LEDs. (d) Our 1.23 NA reconstruction provides isotropic 425 nm resolution with BF + SC calibration.

process as FPM without calibration, our brightfield calibration only takes an additional 24 seconds of processing time and the combined calibration takes roughly only 2.25× as long as no calibration.

Steered Laser

Laser illumination can be used instead of LED arrays to increase the coherence and light efficiency of FPM [168,185]. In practice, laser systems are typically less rigidly aligned than LED arrays, making them more difficult to calibrate. To verify the performance of our method, we constructed a laser-based FPM system using a dual-axis galvanometer to steer a 532 nm, 5 mW laser, which is focused on the sample by large condenser lenses (Fig. 5.10a). This laser illumination system allows finer, more agile illumination control than an LED

array, as well as higher light throughput. However, the laser illumination angle varies from the expected value due to offsets in the dual-axis galvonometer mirrors, relay lens aberrations, and mirror position mis-estimations when run at high speeds. Our method can correct for these problems in a fraction of the time of previous methods (Fig. 5.10b).

Quasi-Dome

Since the FPM resolution limit is set by $NA_{obj} + NA_{illum}$, high-NA illuminators are needed for large space-bandwidth product imaging [79, 195]. To achieve high-angle illumination with sufficient signal-to-noise ratio in the darkfield region, the illuminators must become more dome-like, rather than planar [73]. We previously developed a novel programmable quasi-dome array made of five separate planar LED arrays that can illuminate up to 0.98 NA [79]. This device uses discrete LED control with RGB emitters ($\bar{\lambda} = [475nm, 530nm, 630nm]$) and can be easily attached to most commercial inverted microscopes (Fig. 5.10c).

As with conventional LED arrays, we assume that the LEDs on each board are rigidly placed as designed. However, each circuit board may have some relative shift, tilt, or rotation since the final mating of the 5 boards is performed by hand. LEDs with high-angle incidence are both harder to calibrate and more likely to suffer from mis-estimation due to the dome geometry, so the theoretical reconstruction NA would be nearly impossible to reach without self-calibration. Using our method, we obtain the theoretical resolution limit available to the quasi-dome (Fig. 5.10d). The SC calibration is especially important in the quasi-dome case since it usually has many darkfield LEDs.

Discussion

Our calibration method offers significant gains in speed and robustness as compared to previous methods. BF calibration enables these capabilities by obtaining a good calibration that needs to be calculated only once in pre-processing, reducing computation. Since an estimation of a global shift in the illuminator based only on the brightfield images provides such a close initialization for the rest of the illumination angles, we can use a quicker, easier joint estimation computation in our SC calibration than would be otherwise possible. Jointly, these two methods work together to create fast and accurate reconstructions.

We analyze the robustness of our method to illumination changes by simulating an object illuminated by a grid of LEDs with $NA_{illum} < 0.41$, with LEDs spaced at $0.041NA$ intervals. We define the system to have $\lambda = 532nm$, with a $10\times$, 0.25 NA objective, a $2\times$ system magnification, and a camera with $6.5\mu m$ pixels. While the actual illumination angles in the simulated data remain fixed, we perturb the expected angle of illumination in typical misalignment patterns for LED arrays: rotation, shift, and scale (analogous to LED array distance from sample). We then calibrate the unperturbed data with the perturbed expected angles of illumination as our initial guess.

Our method recovers the actual illumination angles with error less than 0.005 NA for rotations of $-45°$ to $45°$ (Fig. 5.11a); shifts of -0.1 to 0.1 NA, or approximately a displacement
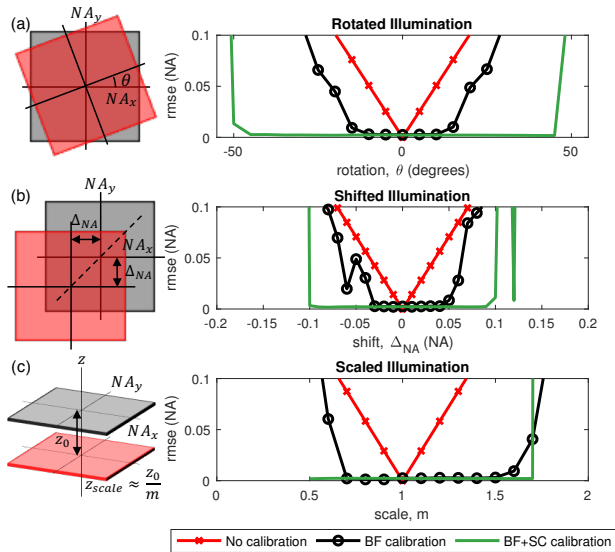
Figure 5.11: Our calibration methods are robust to large mismatches between estimated and actual LED array position. Simulation of misaligned illumination by (a) rotation, (b) shift, and (c) scale. Our calibration recovers the illumination with <0.005 NA error for rotations of $-45°$ to $45°$, shifts of -0.1 to 0.1 NA, and scalings of $0.5\times$ to $1.75\times$ before diverging.

of +/- 2 LEDs (Fig. 5.11b); and scalings of $0.5\times$ to $1.75\times$ (or LED array height between 40-140 *cm* if the actual LED array height is 70 *cm*) (Fig. 5.11c). In these ranges, the average error is 0.0024 NA, less than the k-space resolution of 0.0032 NA. Hence, our calibrated angles are very close to the actual angles even when the input expected angles are extremely far off. This result demonstrates that our method is robust to most mis-alignments in the illumination scheme.

In summary, we have presented a novel two-part calibration method for recovering the illumination angles of a computational illumination system for Fourier ptychography. We have demonstrated how this self-calibrating method makes Fourier ptychographic microscopes more robust to system changes and aberrations introduced by the sample. The method also makes it possible to use high-angle illuminators, such as the quasi-dome, and non-rigid illuminators, such as laser-based systems, to their full potential. Our pre-processing brightfield calibration further enables 3D multislice Fourier ptychography to reconstruct high-resolution features across larger volumes than previously possible. These gains were all made with minimal additional computation, especially when compared to current state-of-the-art methods. Efficient self-calibrating methods such as these are important to make computational imaging methods more robust and available for broad use in the future.

5.4 Summary

In computational imaging systems, mis-calibrations can significantly degrade reconstructions if not properly compensated for. Here, we have presented two frameworks for algorithmic self-calibration. In Section 5.2, we presented a novel method for performing self-calibration of system aberrations using just four measurements - three half-circle (DPC) measurements and one coherent (single-LED) measurement. These four measurements enable the recovery of *both* the complex field of the object as well as spatially-variant system aberrations. These aberrations can be acquired quickly and reconstructed using standard reconstruction techniques such as ADMM. In Section 5.3, we presented a novel method for recovering source illumination angles of planar, quasi-domed, and steered laser sources using an offline (image-based) self-calibration algorithm, which is then refined using an online (FPM-based) self-calibration algorithm. For devices subject to manufacturing variations (such as the quasi-dome), these techniques are absolutely essential for high-resolution FPM reconstructions.

Owing to their non-convex formulations, self-calibration will never be as good as proper physical calibration of the systems, when such calibration is possible. However, when mis-calibration is suspected, these methods will always improve reconstruction quality compared to the uncorrected case, and therefore should be used whenever possible. Practically, the performance Quasi-dome presented in Section 4.1 is significantly improved using self-calibration due to positioning error between printed circuit boards it is composed of. The open-source code used in these sections is provided in the Appendix 7, Section 7.2.

Chapter 6

Conclusion

Driving the performance limits of optical microscopy systems requires a unique blend of both hardware and software innovations. Among a wide variety of techniques for probing a sample, coded illumination using a programmable light source is advantageous due to its simplicity and wide utility for various computational imaging techniques. In this work, we proposed a variety of techniques which take advantage of programmable illumination to enable label-free contrast, quantitative imaging, and high-throughput imaging at relatively low cost and complexity on existing optical systems. We have leveraged novel hardware and computational techniques to push performance boundaries in critical areas, such as label-free imaging and neuropathology. Further, we have put significant effort towards quantifying when computational techniques do not provide benefit, acknowledging limiting cases where conventional methods are still state-of-the-art. Taken together, the methods presented in this work illustrate the capabilities of a computational imaging system, and offer evidence of the impact these systems could have on the broader microscopy community.

In Chapter 2 we described differential phase contrast (DPC), a quantitative phase imaging method which uses partially coherent illumination to recover the complex field of a thin (weak) object. We then introduced a single-shot variant of DPC which uses color-multiplexed illumination to recover the linearized optical field from a single measurement. Our hardware requirements are simple, inexpensive, and compatible with most commercial microscopes through of the use of a color camera and a simple color filter insert placed at the back focal plane of the condenser lens, the same position as many removable phase contrast annulus rings. Unlike phase contrast and DIC, our method does not require special objectives or prisms, which reduces our hardware costs to that of the 3D printed filter itself. In addition, we can use our quantitative phase and amplitude methods to synthesize phase con-

trast and DIC images digitally, matching the functionality of existing phase imaging systems at a fraction of the cost. Finally, we presented a framework for analyzing the source design for DPC-based system in terms of SNR and explored source calibration and measurement count as examples of this analysis.

In Chapter 3, we introduced the concept of high-throughput imaging, and described conventional methods for obtaining images with a wide field-of-view and high resolution. After discussing the limitations of existing methods, we proposed a novel coded illumination technique where measurements are acquired while the sample is in motion, which is synchronized with multiple illumination pulses during each exposure to introduce a known motion blur. These images, which have higher SNR than images captured at the same speed under strobed illumination, are then computationally deblurred to recover the static object at high speed while maintaining the SNR of much slower acquisition methods. We compare our technique to existing methods, stop-and-stare and strobed illumination, which require a significantly longer acquisition time or produce images with significantly lower SNR, respectively. Using a generalized framework for predicting the reconstruction SNR of measurements acquired using each method, we show, through both theory and experiment, that our coded illumination technique can produce images with up to $10\times$ the SNR of strobed acquisitions, at significantly faster acquisition rates than images captured under strobed illumination.

of this benefit is not constant across system parameters such as camera readout noise and illumination power. To quantify how our method compares to existing methods, we perform an analysis of the relative performance of our method compared to conventional high-throughput imaging techniques such as strobed illumination or stop-and-stare. For low-light situations such as fluorescence imaging, our method is uniquely suited to provide faster imaging at higher SNR than existing methods.

In Chapter 4, we detail the design and fabrication of coded illumination devices for the computational microscopy applications presented in this dissertation. We begin by defining the common design requirements of programmable illumination devices, focusing on programmable LED arrays as our primary application due to their low cost and wide availability. We then describe several design iterations of a programmable LED dome which enables high-angle coded illumination. In the first design iteration, we designed and fabricated a 3D-printed LED dome which was carefully assembled by hand, requiring several months of fabrication. This device was intended to be used with CellScope, a portable microscope

platform which uses a smartphone to capture and process images for telemedicine applications. In the second design iteration, we designed and developed a LED quasi-dome which uses 5 printed-circuit boards arranged in a dome-like structure. This device requires significantly less time and effort to assemble and provides RGB illumination across 581 LEDs up to 0.9NA. The ease of manufacturing has enabled the wide distribution of these devices to collaborators around the world. Finally, we describe the Computational CellScope platform in greater detail, which implements coded illumination on a portable platform to perform digital refocusing using light-field methods, quantitative phase imaging, and multi-contrast imaging.

Recovered Aberrations

In Chapter 5, two methods for self-calibration of computational imaging systems are described. Computational imaging methods are uniquely susceptible to mis-calibration since the inferred forward model is often assumed to be ideally calibrated. In reality, the forward model is a function of the optical physics and optical design (which can be inferred), as well as mis-calibration, which cannot be assumed to be negligible in many systems. We first discuss a technique for performing aberration self-calibration using differential phase contrast. Our method employs an alternating-minimization approach to solve for both the complex field of the object as well as the aberrations, which are projected into a Zernike basis. This method is non-convex, but is differentiable, and guaranteed not to diverge so long as an appropriate step size is used. Further, we use a patch-wise solver to solve for aberrations across the field, revealing spatially variant aberration functions. We verify our method by solving for the defocus aberration when introducing a known defocus term. Our method requires only a single coherent measurement, in addition to three DPC measurements for recovering the object's complex field. Next, we presented a method for performing self-calibration of LED positions in both the brightfield and darkfield regions for Fourier ptychography. After performing a brightfield calibration method, we employ a novel gradient-based technique for solving for the homography between printed circuit boards in a quasi-dome device. This technique makes the alternating minimization process much more stable and enables high-NA Fourier ptychography with resolution below 450nm. While a good hardware system alignment is most effective, these self-calibration techniques can mitigate mis-calibration when this is not possible.

6.1 Entrepreneurship in Computational Microscopy

During my time at Berkeley I was fortunate to benefit from the many resources for entrepreneurship on campus. Along with fellow lab member Michael Chen and my advisor Laura Waller, I founded Spectral Coded Illumination, Inc. in early 2017 with the purpose of bridging the gap between the scientific publication of our prototypes and the large-scale distribution of LED dome prototypes. The initial motivation to form this venture came from the early Quasi-dome prototype presented in Chapter **??**, Section **??**, which brought significant interest from both collaborators and visitors from industry. After quickly distributing our initial production run of 10 Quasi-domes through the Waller Lab, we saw a need in both industry and academia for functional LED dome prototypes and sought to use our knowledge in this area to drive progress in our research field at a larger scale. We began by getting involved on through campus programs, including the NSF I-Corps program, the Citris Foundry, Berkeley Skydeck, and the Bakar Innovation Fellows program, which allowed us to meet other entrepreneurs in the sciences and better understand how to build a successful scientific venture. For the past 2.5 years we have continued to sell LED domes and provide consulting services in areas related to our research for the past two and a half years, selling over 15 LED arrays across the United States and in Europe, to customers in both industry and academia.

I view this experience as a critical component of my education at Berkeley, as it allowed me to view my research with a broader view, and to understand the effort and technical challenges which come with commercializing and scaling up a research proposal. While the direct motivations for forming a corporation can be quite different from academic research, I have found that my underlying motivation for both ventures has been the same: to accelerate the pace of scientific discovery through the development of novel imaging methods and hardware. Having a start-up motivated me to align my research to be more useful for the field at-large, which has been very useful in the later stages of my education.

6.2 Proposed Future Work

As my time at Berkeley comes to a close, there are still many open questions which I find interesting and would like to pursue, given the opportunity. In the future, I hope that myself or someone else will have the time and a motivation to explore a few of the following proposals.

The first, and most general, is the application of a task-based imaging paradigm to the applications discussed in this work. In each of the previous chapters, we have quantified the quality of our reconstructions using system parameters such as signal-to-noise ratio, acquisition throughput (space-bandwidth rate), resolving power and field of view. Rarely is emphasis placed on the quality of these images towards a specific task, such as cell counting, pathogen recognition, or disease diagnosis. Conventionally the field of computational imaging has tended to focus on imaging system performance, but I believe that defining a

differentiable performance metric (cost function) which is relevant to the imaging task would be extremely interesting. As a first step, I would propose cell-counting as a metric. Recent work from the machine learning community has proposed optimizing network layers for the counting and classification of cells [196, 197]. Including programmable hardware elements (such as LED arrays which could introduce different contrast) into the learning pipeline could provide significant improvements in counting accuracy and performance.

A second, more specific future direction, is the application of compressed sensing to solve under-determined motion deblurring problems for slide-scanning and digital pathology applications. Currently, the volume of data required for large-scale imaging (e.g. neuropathology) can make acquisition times and data storage difficult or infeasible. If data acquisition processing requirements were lower, acquisitions could be made much faster, and large-scale imaging analysis could be more readily performed. For a compressed-sensing acquisition to perform well, the aliasing introduced by undersampling should be incoherent with a domain in which the sample is sparse. An acquisition with a random forward operator is one example of such a forward model, because this matrix is incoherent with every other matrix. While it is very difficult to design a truly random forward operator in this problem, we have control over the structure of our forward operator through the illumination sequence as well as the motion pathway. So long as this acquisition introduces incoherent aliasing, we have some hope of reconstructing a sparse object from an under-determined system. To accomplish this, I would propose optimizing the entire pipeline using an unrolled network, which allows us to differentiate with simple parameters such as step size as well as more complicated parameters, such as the sparsifying operator or acquisition strategy. Using large amounts of training data, we could then generate a principled acquisition and reconstruction strategy together which would enhance the capabilities of high-throughput acquisitions well-beyond those presented in this work.

A third future direction is active imaging, or the active determination of acquisition trajectories based on current image data. These trajectories need not span Cartesian space, but could also include LED angle, defocused positions, or other metrics. A key component to these methods being advantageous would be that the sample should be localized, meaning that most of the densely sampled data is unnecessary, and that the learning process should be fast, so that the compute + scan time does not exceed the time to densely scan a sample. An obvious example of this would be multi-well plates, where a microscope may perform a fast scan over many wells which aliases them together, then uses this blurred data to determine which wells are worth imaging more closely at high-resolution for quantification. Such a technique could speed up acquisitions significantly for many applications.

In terms of fabrication of LED devices, I believe that engineering improvements in the power supply of the devices as well as a more careful layout of the high-speed traces (such as serial clock and pulse-width-modulation clock) would enable imaging with shorter exposure times. In addition, a more compact portable microscope device incorporating a PCB-based LED array would be much more robust than the current Computational CellScope device, enabling wider distribution to collaborators around the world.

I have made much of my existing code available under the BSD open-source license at

repositories listed in Appendix, Section 7.2. I hope that these examples and this dissertation will inspire others to iterate and advance the field of computational imaging.

Chapter 7

Appendix

7.1 Quantitative Phase Imaging

DPC System Parameters

Parameter	Value
Objective Mag / NA	$10 \times /0.25NA$
Illumination Power	See table 7.1
Camera Readout Noise	$3.7\ e^-$
Camera Quantum Efficiency	60%
Camera Pixel Size	$6.5\mu m$
Number of Brightfield LEDs	72

LED Illuminance Values for DPC Analysis

This section describes measured LED power values for a range of common LEDs. Each of these measurements was made at the sample plane using an optical power-meter (Thorlabs PM100D) at a distance of $50mm$ from the LED.

LED Mfg. Part Number	LED Type	Illuminance	Comment
VCC / VAOL-3LWY4	Through-hole	1000.42 Lux	Used in Computational CellScope [73]
Knightbright / APTF1616SEEZ	SMD	(Red) 2738 Lux (Green) 4666 Lux (Blue) 3500 Lux	Used for Quasi-Dome [198]
VCC / VAOL-3LWY4	Through-hole	1000.42 Lux	Used for High-Throughput Fluorescence Imaging (Chapter 3)

7.2 Open-Source Code

Most of the source code used for these methods is available publicly through the following repositories. Each project is licensed under the BSD 3-clause license.

Repository	Description	URL
llops	GPU-acceleration library with numpy-like syntax and interchangeable backends. Also supports auto-differentiation and latex rendering through the operators submodule.	https://github.com/zfphil/llops
comptic	Computational microscopy library for python, including common functions for simulating optical systems. Also included scripts for acquisition using micro-manager.	https://github.com/zfphil/comptic
pydpc	Python implementation of DPC algorithm, including simulation, processing real data, and aberration recovery. Also included analysis tools for analyzing SNR in DPC systems.	https://github.com/zfphil/pydpc
pyfpm	Python implementation of FPM algorithm, including simulation, processing real data, aberration recovery, and source recovery.	https://github.com/zfphil/pyfpm
htdeblur	Python implementation of motion deblur algorithm. Also includes system analysis tools for convolutional forward models.	https://github.com/zfphil/htdeblur
illuminate	LED array firmware with standardized API	https://github.com/zfphil/illuminate
illuminate controller	High-level interfaces for controlling LED arrays, including MATLAB, python, micro-manager, and GUI interfaces.	https://github.com/zfphil/illuminate_controller

7.3 High-Throughput System Analysis

Derivation of SNR Expression

In this section we derive the expression for the SNR of a recovered image. Considering the additive noise acquisition model, $\boldsymbol{y} = \mathbf{A}\boldsymbol{x} + \eta$, the recovered image $\widehat{\boldsymbol{x}}$ is given by:

$$\widehat{\boldsymbol{x}} = \mathbf{A}^\dagger \boldsymbol{y} = \boldsymbol{x} + \mathbf{A}^\dagger \eta \,.$$

In what follows, we assume only that η is zero mean with covariance $\sigma_\eta^2 \mathbf{I}$,

Defining the mean of the recovered object $\mu = \mathbb{E}[\widehat{\boldsymbol{x}}]$, as well as the covariance $\boldsymbol{\Sigma} = \mathbb{E}[(\widehat{\boldsymbol{x}} - \boldsymbol{x})(\widehat{\boldsymbol{x}} - \boldsymbol{x})^\mathsf{H}]$, we calculate the imaging SNR using the root mean squared error (RMSE):

$$SNR = \frac{\frac{1}{m}\sum_{i=1}^m \mu_i}{\sqrt{\frac{1}{m}\mathrm{Tr}(\boldsymbol{\Sigma})}} \,.$$

Assuming zero-mean noise, the numerator is simply the average object signal $\bar{s}$. Expanding the covariance term in the denominator,

$$\Sigma = \mathbb{E}[\mathbf{A}^\dagger \eta (\mathbf{A}^\dagger \eta)^\mathsf{H}] = \sigma_\eta^2 \mathbf{A}^\dagger (\mathbf{A}^\dagger)^\mathsf{H} \,,$$

where we apply the assumption that the covariance of η is $\sigma_\eta^2 \mathbf{I}$. Then,

$$\mathrm{Tr}(\mathbf{A}^\dagger (\mathbf{A}^\dagger)^\mathsf{H}) = \sum_{i=1}^m \frac{1}{\sigma_i(\mathbf{A})^2} = \frac{1}{\sigma_1(\mathbf{A})^2} \sum_{i=1}^m \frac{\sigma_1(\mathbf{A})^2}{\sigma_i(\mathbf{A})^2} \,.$$

Thus we have that

$$SNR = \frac{\bar{s}}{\frac{1}{\sigma_1(\mathbf{A})}\sqrt{\frac{1}{m}\sum_{i=1}^m \frac{\sigma_1(\mathbf{A})}{\sigma_i(\mathbf{A})}} \cdot \sigma_\eta} := \frac{\sigma_1(\mathbf{A})\bar{s}}{f\sigma_\eta} \,,$$

where f is the general definition of the deconvolution noise factor. This expression is consistent with the definition in (3.8) for convolutional operators, where we note that the singular values are given by the power spectrum of the kernel $\boldsymbol{h}$. Further, we note that in this case $\sigma_1(\mathbf{A}) = \gamma$ since that is the DC component of a non-negative signal.

Multi-frame Decomposition

We consider the case of a multiframe operator with the same blur kernel $\boldsymbol{h}$ used in every frame. In this case, the forward operator has the form

$$\mathbf{A} = \begin{bmatrix} \mathbf{W}_1 \\ \vdots \\ \mathbf{W}_n \end{bmatrix} \mathbf{B} := \mathbf{W}\mathbf{B} \,.$$

Following the derivation of SNR from the previous section, we compute $\text{Tr}(\mathbf{A}^\dagger (\mathbf{A}^\dagger)^{\mathsf{H}})$. First,

$$\mathbf{A}^\dagger = (\mathbf{B}^{\mathsf{H}} \mathbf{W}^{\mathsf{H}} \mathbf{W} \mathbf{B})^{-1} \mathbf{B}^{\mathsf{H}} \mathbf{W}^{\mathsf{H}} = \mathbf{B}^{-1} (\mathbf{W}^{\mathsf{H}} \mathbf{W})^{-1} \mathbf{W}^{\mathsf{H}} \, ,$$

assuming that $\mathbf{B}$ and $\mathbf{W}^{\mathsf{H}} \mathbf{W}$ are invertible. Then we have that

$$\text{Tr}(\mathbf{A}^\dagger (\mathbf{A}^\dagger)^{\mathsf{H}}) = \text{Tr}(\mathbf{B}^{-1} (\mathbf{W}^{\mathsf{H}} \mathbf{W})^{-1} \mathbf{W}^{\mathsf{H}} \mathbf{W} (\mathbf{W}^{\mathsf{H}} \mathbf{W})^{-\mathsf{H}} \mathbf{B}^{-\mathsf{H}})$$
$$= \text{Tr}(\mathbf{B}^{-1} (\mathbf{W}^{\mathsf{H}} \mathbf{W})^{-1} \mathbf{B}^{-1}) = \text{Tr}((\mathbf{W}^{\mathsf{H}} \mathbf{W})^{-1} \mathbf{B}^{-2}) \, .$$

We now consider the form of $\mathbf{W}^{\mathsf{H}} \mathbf{W} = \sum_{j=1}^{n} \mathbf{W}_j^{\mathsf{H}} \mathbf{W}_j$. Each $\mathbf{W}_j^{\mathsf{H}} \mathbf{W}_j$ is a square diagonal matrix with either a 0 or 1 for each diagonal entry, depending on whether the corresponding pixel is included in the window. Thus the sum $\mathbf{W}^{\mathsf{H}} \mathbf{W}$ is a diagonal matrix with the ith diagonal value given by the number of times pixel i is included in the windows $\{\mathbf{W}_1, ..., \mathbf{W}_n\}$, a quantity we denote as $c_i = \sum_{j=1}^{n} \mathbf{W}_j \boldsymbol{e}_i$ where $\{\boldsymbol{e}_i\}$ are the standard basis vectors.

Before we proceed further, note that for any matrices $\mathbf{M}$ and $\mathbf{D}$ with non-negative entries and $\mathbf{D}$ diagonal,

$$\text{Tr}(\mathbf{D}\mathbf{M}) = \sum_i D_{ii} M_{ii} \leq \max_i D_{ii} \cdot \text{Tr}(\mathbf{M}) \, .$$

We can therefore conclude that

$$\text{Tr}(\mathbf{A}^\dagger (\mathbf{A}^\dagger)^{\mathsf{H}}) \leq \max_i \frac{1}{c_i} \cdot \text{Tr}(\mathbf{B}^{-2}) = \frac{1}{\min_i c_i} \cdot \sum_{i=1}^{m} \frac{1}{|\tilde{\boldsymbol{h}}|_i^2} \, .$$

Thus we see that the expression for the covariance is decreased by a factor of at least the square root of minimum coverage. This corresponds to the lower bound on the SNR:

$$SNR \geq \sqrt{\min_i c_i} \cdot \frac{\gamma \bar{s}}{f \sigma_\eta} \, .$$

Where f is defined as in (3.8).

Blur Kernel Optimization

In this section we discuss the reformulation of the optimization problem in (3.10) as a smooth objective with convex constraints. Recall that the optimization problem has the form

$$\min_{\boldsymbol{h}} \quad \sqrt{\frac{1}{m} \sum_{i=0}^{m} \frac{\max_i |\tilde{\boldsymbol{h}}|_i^2}{|\tilde{\boldsymbol{h}}|_i^2}}$$
$$s.t. \quad 0 \leq h_i \leq 1 \, \forall \, i, \quad \sum_i h_i = \gamma \, .$$

First, note that by definition $\tilde{\boldsymbol{h}} = \mathbf{F}\boldsymbol{h}$ where $\mathbf{F}$ represents the discrete Fourier transform (DFT) matrix. Then, we know that $\max_i |\tilde{\boldsymbol{h}}|_i^2$ is the DC component of the signal, which is

equal to γ and therefore fixed for any feasible $\boldsymbol{h}$. Therefore, the blur kernel which maximizes (3.10) is the same as the one that maximizes

$$\min_{\boldsymbol{h}} \ \sum_{i=0}^{m} \frac{1}{(\mathbf{F}_i^{\mathsf{H}} \boldsymbol{h})^2} \ \ s.t. \ \ 0 \le h_i \le 1, \ \ \sum_i h_i = \gamma \, ,$$

where $\mathbf{F}_i$ represents columns of the DFT matrix.

It is possible to use projected gradient methods because the objective function is smooth nearly everywhere and the constraints are convex. At each iteration, there is a gradient step followed by a projection step. The gradient step is defined as

$$\tilde{\boldsymbol{h}}^{k+1} = \boldsymbol{h}^k + \alpha^k \sum_{i=0}^{m} \frac{2}{(\mathbf{F}_i^{\mathsf{H}} \boldsymbol{h}^k)^3} \cdot \mathbf{F}_i \, ,$$

for potentially changing step size α^k. The projection step is defined as

$$\boldsymbol{h}^{k+1} = \Pi_{\mathcal{S}}(\tilde{\boldsymbol{h}}_{k+1}) \, ,$$

where $\mathcal{S}$ is the intersection of the box constraint $\{0 \le h_i \le 1\}$ and the simplex constraint $\{\sum_i h_i = \gamma\}$. Efficient methods for this projection exist [199].

Fundamental DNF Limits

There are fundamental limits on how SNR can be improved by coded illumination. We examine a fundamental lower bound on the DNF to demonstrate this.

Recall that

$$f^2 = \max_i |\tilde{\boldsymbol{h}}|_i^2 \cdot \frac{1}{m} \sum_{i=1}^{m} \frac{1}{|\tilde{\boldsymbol{h}}|_i^2} \, .$$

Then, note that $\frac{1}{m} \sum_{i=1}^{m} \frac{1}{|\tilde{\boldsymbol{h}}|_i^2}$ is the reciprocal of the harmonic mean of $\{|\tilde{\boldsymbol{h}}|_1^2, ..., |\tilde{\boldsymbol{h}}|_m^2\}$. Since the harmonic mean is always less than the arithmetic mean, we have that

$$\frac{1}{m} \sum_{i=1}^{m} \frac{1}{|\tilde{\boldsymbol{h}}|_i^2} \ge \frac{1}{\frac{1}{m} \sum_{i=1}^{m} |\tilde{\boldsymbol{h}}|_i^2} \, .$$

Next, we apply Parseval's and have $\frac{1}{m} \sum_{i=1}^{m} |\tilde{\boldsymbol{h}}|_i^2 = \sum_{i=1}^{m} h_i^2$. Additionally, $\max_i |\tilde{\boldsymbol{h}}|_i$ is the DC component of the signal, which is specified by the constraint $\sum_{i=1}^{m} h_i = \gamma$. As a result,

$$f^2 \ge \frac{\gamma^2}{\sum_{i=1}^{m} h_i^2} \, .$$

Finally, we see that

$$\max_{\boldsymbol{h} \in [0,1]^m} \sum_{i=1}^{m} h_i^2 \ : \ \sum_{i=1}^{m} h_i = \gamma$$

is achieved for binary h and has the maximum value γ. Therefore,

$$f^2 \geq \frac{\gamma^2}{\sum_{i=1}^{m} h_i^2} \geq \frac{\gamma^2}{\gamma} = \gamma \, .$$

That is, the DNF grows at a rate of at least $\sqrt{\gamma}$. As a result, the best achievable SNR (using (3.9)) is

$$SNR \leq \frac{\sqrt{\gamma} \bar{s}_0}{\sqrt{\gamma \bar{s}_0 + \sigma_r^2}} = \sqrt{\bar{s}_0} \sqrt{\frac{\gamma \bar{s}_0}{\gamma \bar{s}_0 + \sigma_r^2}} \, ,$$

This upper bound on SNR increases with γ. In Methods Section 3.2, we discuss an exact closed form for $f(\gamma)$ that yields an expression for optimal multiplexing.

However, if σ_r is much smaller than the total captured signal, i.e. $\sigma_r \ll \gamma \bar{s}_0$, the SNR will not increase with γ, and in fact its maximum value,

$$SNR \leq \sqrt{\bar{s}_0}$$

is achieved by strobed illumination (i.e. $\gamma = 1$). In other words, when signal is large compared with readout noise, strobed will be optimal, regardless of the illumination optimization method.

Derivation of Illumination Throughput

Stop-and-Stare

In the stop-and-stare acquisition strategy, the sample is illuminated for the full dwell time (t_{sns}), which is set by motion stage parameters such as maximum velocity, acceleration, and the necessary stage settle time (v_{stage}, a_{stage}, and t_{settle} respectively), as well as camera readout ($t_{readout}$). These parameters are related to frame rate r_{frame} by the following relationship:

$$t_{sns} = \frac{1}{r_{frame}} - \max(t_{readout}, 2t_{accel} + t_{move})$$

Note that this equation assumes perfect hardware synchronization and instantaneous acceleration ($\frac{\partial a}{\partial t} = \infty$). The variables t_{accel} and t_{move} are defined as:

$$t_{accel} = \frac{v_{stage}}{a_{stage}}$$

$$t_{move} = \frac{d_{frame} - 0.5 * a_{stage} * t_{accel}^2}{v_{stage}}$$

Here the expression $d_{frame} = FOV * (1 - O)$ is the distance between frames, which is determined by the field-of-view of a single frame (FOV) and inter-frame overlap fraction O.

Combining terms, we arrive at an expression for t_{sns}:

$$t_{sns} = \frac{1}{r_{frame}} - \max(t_{readout}, 2\frac{v_{stage}}{a_{stage}} - \frac{d_{frame} - 0.5 * a_{stage} * t_{accel}^2}{v_{stage}})$$

When camera readout time $t_{readout}$ is short, t_{sns} can be simplified to:

$$t_{sns} = \frac{1}{r_{frame}} - 2\frac{v_{stage}}{a_{stage}} - \frac{d_{frame} - 0.5 * \frac{v_{stage}^2}{a_{stage}}}{v_{stage}}$$

Strobed Illumination

The maximum pulse duration for strobed illumination is related to the time required to move a distance of one effective pixel size $\frac{\Delta}{M}$ at a velocity v_{stage}:

$$t_{strobe} = \frac{\frac{\Delta}{M}}{v_{stage}}$$

The stage velocity v_{stage} may be bounded by the motion stage hardware (v_{max}) or by the field of view of the microscope (FOV):

$$v_{stage} = \min(v_{max}, r_{frame}FOV)$$

Coded Illumination

The calculation of t_{coded} for coded illumination is synonymous to the strobed illumination case, weighted by the multiplexing coefficient used to generate the illumination sequence (γ), and using the v_{stage} calculation from the strobed subsection:

$$t_{coded} = \frac{\frac{\gamma\Delta}{M}}{v_{stage}}$$

System Parameters

Parameter	Value
Maximum Motion Stage Velocity	$40\frac{m}{s}$
Motion Stage Acceleration	$400\frac{m}{s^2}$
Motion Stage Settle Time	$0.1s$
Objective Mag / NA	$10\times/0.25NA$
Frame Overlap	20%
Illumination Power	600
Camera Readout Time	26ms
Camera Readout Noise	$3.7\ e^-$
Camera Quantum Efficiency	60%
Camera Pixel Size	$6.5\mu m$
Fluorophore Quantum Yield	79% lux
Illumination repetition Rate	$250kHz$

Bibliography

[1] A. Van Helden, S. Dupré, and R. van Gent, *The origins of the telescope*, vol. 12 (Amsterdam University Press, 2010).

[2] "Milestones in Light Microscopy," (2009). URL `http://www.nature.com/milestones/milelight/pdf/milelight_all.pdf`.

[3] R. Hooke, *Micrographia: or, Some physiological descriptions of minute bodies made by magnifying glasses* (London: J. Martyn and J. Allestry, 1665).

[4] J. Goodman, *Introduction to Fourier optics* (McGraw Hill, 1968).

[5] J. C. Maxwell, *The Scientific Papers of James Clerk Maxwell...*, vol. 2 (University Press, 1890).

[6] H. Gernheim and A. Gernheim, *The history of photography from the camera obscure to the beginning of the modern era.* (McGraw-Hill, 1969).

[7] E. R. Dowski and W. T. Cathey, "Extended depth of field through wave-front coding," Appl. Opt. **34**(11), 1859–1866 (1995).

[8] M. G. Gustafsson, "Surpassing the lateral resolution limit by a factor of two using structured illumination microscopy," Journal of microscopy **198**(2), 82–87 (2000).

[9] M. G. Gustafsson, "Nonlinear structured-illumination microscopy: wide-field fluorescence imaging with theoretically unlimited resolution," Proceedings of the National Academy of Sciences **102**(37), 13,081–13,086 (2005).

[10] E. Betzig, G. H. Patterson, R. Sougrat, O. W. Lindwasser, S. Olenych, J. S. Bonifacino, M. W. Davidson, J. Lippincott-Schwartz, and H. F. Hess, "Imaging intracellular fluorescent proteins at nanometer resolution," Science **313**(5793), 1642–1645 (2006).

[11] M. J. Rust, M. Bates, and X. Zhuang, "Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM)," Nature Methods **3**, 793 (2006). URL `https://doi.org/10.1038/nmeth929`.

[12] J. M. Rodenburg and H. M. Faulkner, "A phase retrieval algorithm for shifting illumination," Applied Physics Letters **85**(20), 4795–4797 (2004).

[13] G. Zheng, R. Horstmeyer, and C. Yang, "Wide-field, high-resolution Fourier Pty-chographic microscopy," Nature Photonics **7**(9), 739–745 (2013). URL `http://www.nature.com/nphoton/journal/v7/n9/abs/nphoton.2013.187.html`.

[14] D. A. Agard, "Optical sectioning microscopy: cellular architecture in three dimensions," Annual review of biophysics and bioengineering **13**(1), 191–219 (1984).

[15] M. Broxton, L. Grosenick, S. Yang, N. Cohen, A. Andalman, K. Deisseroth, and M. Levoy, "Wave optics theory and 3-D deconvolution for the light field microscope," Optics express **21**(21), 25,418–25,439 (2013).

[16] R. Ng, M. Levoy, M. Bredif, G. Duval, M. Horowitz, and P. Hanrahan, "Light Field Photography with a Hand-held Plenoptic Camera," Tech. Rep. CTSR 2005-02, Stanford (2005).

[17] S. R. P. Pavani and R. Piestun, "Three dimensional tracking of fluorescent micropar-ticles using a photon-limited double-helix response system," Optics express **16**(26), 22,048–22,057 (2008).

[18] E. Wolf, "Three-dimensional structure determination of semi-transparent objects from holographic data," Optics Communications **1**(4), 153–156 (1969).

[19] K. Kim, Z. Yaqoob, K. Lee, J. W. Kang, Y. Choi, P. Hosseini, P. T. So, and Y. Park, "Diffraction optical tomography using a quantitative phase imaging unit," Optics letters **39**(24), 6935–6938 (2014).

[20] M. H. Maleki, A. J. Devaney, and A. Schatzberg, "Tomographic reconstruction from optical scattered intensities," JOSA A **9**(8), 1356–1363 (1992).

[21] G. Popescu, T. Ikeda, R. R. Dasari, and M. S. Feld, "Diffraction phase microscopy for quantifying cell structure and dynamics," Optics Letters **31**(6), 775–777 (2006).

[22] B. Bhaduri, H. Pham, M. Mir, and G. Popescu, "Diffraction phase microscopy with white light," Optics Letters **37**(6), 1094–1096 (2012).

[23] S. Witte, A. Plauşka, M. C. Ridder, L. van Berge, H. D. Mansvelder, and M. L. Groot, "Short-coherence off-axis holographic phase microscopy of live cell dynamics," Biomed. Opt. Express **3**(9), 2184–2189 (2012). URL `http://www.osapublishing.org/boe/abstract.cfm?URI=boe-3-9-2184`.

[24] P. S.A., "QUANTITATIVE PHASE MICROSCOPY," (2009). URL `http://www.phasicscorp.com`.

[25] P. Bon, B. Wattellier, S. Monneret, H. Giovanini, and G. Maire, "Method and system for structural analysis of an object by measuring the wave front thereof," (2012). US Patent App. 13/500,385, URL `https://www.google.com/patents/US20120274945`.

[26] B. Allman, K. Nugent, and C. Porter, "An optical system for producing differently focused images," (2005). WO Patent App. PCT/AU2005/000,204, URL `https://www.google.com/patents/WO2005085936A1?cl=en`.

[27] L. Waller, S. S. Kou, C. J. R. Sheppard, and G. Barbastathis, "Phase from chromatic aberrations," Optics Express **18**(22), 22,817–22,825 (2010). URL `http://www.opticsexpress.org/abstract.cfm?URI=oe-18-22-22817`.

[28] E. H. Moore, "On the reciprocal of the general algebraic matrix," Bull. Am. Math. Soc. **26**, 394–395 (1920).

[29] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," SIAM journal on imaging sciences **2**(1), 183–202 (2009).

[30] J. W. Cooley and J. W. Tukey, "An algorithm for the machine calculation of complex Fourier series," Mathematics of computation **19**(90), 297–301 (1965).

[31] A. N. Tikhonov, "On the stability of inverse problems," in *Dokl. Akad. Nauk SSSR*, vol. 39, pp. 195–198 (1943).

[32] E. J. Candes, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted L1 minimization," Journal of Fourier analysis and applications **14**(5-6), 877–905 (2008).

[33] H. L. Taylor, S. C. Banks, and J. F. McCoy, "Deconvolution with the L1 norm," Geophysics **44**(1), 39–52 (1979).

[34] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, *et al.*, "Distributed optimization and statistical learning via the alternating direction method of multipliers," Foundations and Trends® in Machine learning **3**(1), 1–122 (2011).

[35] J. M. Bioucas-Dias and M. A. Figueiredo, "A new TwIST: Two-step iterative shrinkage/thresholding algorithms for image restoration," IEEE Transactions on Image processing **16**(12), 2992–3004 (2007).

[36] S. Ravishankar and Y. Bresler, "Learning sparsifying transforms," IEEE Transactions on Signal Processing **61**(5), 1072–1086 (2013).

[37] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," Physica D: nonlinear phenomena **60**(1-4), 259–268 (1992).

[38] B. Wahlberg, S. Boyd, M. Annergren, and Y. Wang, "An ADMM algorithm for a class of total variation regularized estimation problems," IFAC Proceedings Volumes **45**(16), 83–88 (2012).

[39] U. Kamilov, E. Bostan, and M. Unser, "Wavelet shrinkage with consistent cycle spinning generalizes total variation denoising," IEEE Signal Processing Letters **19**(4), 187–190 (2012).

[40] O. Cossairt, M. Gupta, and S. K. Nayar, "When does computational imaging improve performance?" IEEE transactions on image processing **22**(2), 447–458 (2013).

[41] A. Agrawal and R. Raskar, "Optimal single image capture for motion deblurring," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pp. 2560–2567 (IEEE, 2009).

[42] H. Chernoff, "Locally optimal designs for estimating parameters," The Annals of Mathematical Statistics pp. 586–602 (1953).

[43] F. Zernike, "Phase contrast, a new method for the microscopic observation of transparent objects," Physica **9**, 686–698 (1942).

[44] Smith, "Microscopic interferometry," Research **8**, 385–395 (1955).

[45] G. Zheng, C. Kolner, and C. Yang, "Microscopy refocusing and dark-field imaging by using a simple LED array," Optics Letters **36**(20), 3987–3989 (2011).

[46] D. F. Albeanu, E. Soucy, T. F. Sato, M. Meister, and V. N. Murthy, "LED arrays as cost effective and efficient light sources for widefield microscopy," PLoS One **3**(5), e2146 (2008).

[47] Z. Liu, L. Tian, S. Liu, and L. Waller, "Real-time brightfield, darkfield, and phase contrast imaging in a light-emitting diode array microscope," Journal of Biomedical Optics **19**(10), 106,002 (2014).

[48] L. Tian, J. Wang, and L. Waller, "3D differential phase-contrast microscopy with computational illumination using an LED array," Optics Letters **39**(5), 1326–1329 (2014). URL http://ol.osa.org/abstract.cfm?URI=ol-39-5-1326.

[49] L. Tian and L. Waller, "Quantitative differential phase contrast imaging in an LED array microscope," Optics express **23**(9), 11,394–11,403 (2015).

[50] G. Popescu, Y. Park, W. Choi, R. R. Dasari, M. S. Feld, and K. Badizadegan, "Imaging red blood cell dynamics by quantitative phase microscopy," Blood Cells, Molecules, and Diseases **41**(1), 10–16 (2008).

[51] G. Popescu, Y. Park, N. Lue, C. Best-Popescu, L. Deflores, R. R. Dasari, M. S. Feld, and K. Badizadegan, "Optical imaging of cell mass and growth dynamics," American Journal of Physiology-Cell Physiology **295**(2), C538–C544 (2008).

[52] B. C. Platt and R. Shack, "History and principles of Shack-Hartmann wavefront sensing," Journal of refractive surgery **17**(5), S573–S577 (2001).

[53] F. Zernike, "How I discovered phase contrast," Science **121**(3141), 345 (1955).

[54] W. Yun, M. Feser, and B. Hornberger, "System and method for quantitative reconstruction of Zernike phase-contrast images," (2010). US Patent 7,787,588.

[55] X. Cui, C. Yang, and G. Tearney, "Quantitative differential interference contrast (DIC) microscopy and photography based on wavefront sensors," (2011). US Patent 8,039,776, URL http://www.google.com/patents/US8039776.

[56] Z. Wang, L. Millet, M. Mir, H. Ding, S. Unarunotai, J. Rogers, M. U. Gillette, and G. Popescu, "Spatial light interference microscopy (SLIM)," Optics Express **19**(2), 1016–1026 (2011).

[57] L. Waller, L. Tian, and G. Barbastathis, "Transport of intensity phase-amplitude imaging with higher order intensity derivatives," Optics express **18**(12), 12,552–12,561 (2010).

[58] J. Petruccelli, L. Tian, and G. Barbastathis, "The transport of intensity equation and partially coherent fields," in *Biomedical Optics*, p. BSu3A.69 (Optical Society of America, 2012). URL http://www.opticsinfobase.org/abstract.cfm?URI=BIOMED-2012-BSu3A.69.

[59] Z. Jingshan, R. A. Claus, J. Dauwels, L. Tian, and L. Waller, "Transport of Intensity phase imaging by intensity spectrum fitting of exponentially spaced defocus planes," Optics Express **22**(9), 10,661–10,674 (2014).

[60] Z. Jingshan, L. Tian, P. Varma, and L. Waller, "Nonlinear Optimization Algorithm for Partially Coherent Phase Retrieval and Source Recovery," IEEE Trans. Comput. Imaging (2016).

[61] Y. Awatsuji, M. Sasada, and T. Kubota, "Parallel quasi-phase-shifting digital holography," Appl. Phys. Lett. **85**(6), 1069–1071 (2004).

[62] D. Gross, "Recovering low-rank matrices from few coefficients in any basis," IEEE Trans. Inf. Theory **57**(3), 1548–1566 (2011).

[63] D. Hamilton and C. Sheppard, "Differential phase contrast in scanning optical microscopy," Journal of Microscopy **133**(1), 27–39 (1984).

[64] S. B. Mehta and C. J. Sheppard, "Quantitative phase-gradient imaging at high resolution with asymmetric illumination-based differential phase contrast," Optics Letters **34**(13), 1924–1926 (2009).

[65] L. Tian, J. Wang, and L. Waller, "3D differential phase-contrast microscopy with computational illumination using an LED array," Optics letters **39**(5), 1326–1329 (2014).

[66] D. Lee, S. Ryu, U. Kim, D. Jung, and C. Joo, "Color-coded LED microscopy for multi-contrast and quantitative phase-gradient imaging," Biomedical optics express **6**(12), 4912–4922 (2015).

[67] M. K. Johnson and E. H. Adelson, "Retrographic sensing for the measurement of surface texture and shape," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pp. 1070–1077 (IEEE, 2009).

[68] A. N. R. Claus, P. Naulleau and L. Waller, "Quantitative phase retrieval with arbitrary pupil and illumination," Opt. Express **23**, 26,672–26,682 (2015).

[69] S. Osher, M. Burger, D. Goldfarb, J. Xu, and W. Yin, "An iterative regularization method for total variation-based image restoration," Multiscale Modeling & Simulation **4**(2), 460–489 (2005).

[70] M. Chen, Z. F. Phillips, and L. Waller, "Quantitative differential phase contrast (DPC) microscopy with computational aberration correction," Optics express **26**(25), 32,888–32,899 (2018).

[71] Z. F. Phillips and M. Chen, "Technical Report: Benchmark Technologies Quantitative Phase Target," Benchmark Technologies (2018).

[72] C. Ma, Z. Liu, L. Tian, Q. Dai, and L. Waller, "Motion deblurring with temporally coded illumination in an LED array microscope," Opt. Lett. **40**(10), 2281–2284 (2015). URL http://ol.osa.org/abstract.cfm?URI=ol-40-10-2281.

[73] Z. F. Phillips, M. V. D'Ambrosio, L. Tian, J. J. Rulison, H. S. Patel, N. Sadras, A. V. Gande, N. A. Switz, D. A. Fletcher, and L. Waller, "Multi-contrast imaging and digital refocusing on a mobile microscope with a domed led array," PloS one **10**(5), e0124,938 (2015).

[74] M. Kellman, M. Chen, Z. F. Phillips, M. Lustig, and L. Waller, "Motion-resolved quantitative phase imaging," Biomedical optics express **9**(11), 5456–5466 (2018).

[75] E. D. Barone-Nugent, A. Barty, and K. A. Nugent, "Quantitative phase-amplitude microscopy I: optical microscopy," Journal of Microscopy **206**(3), 194–203 (2002).

[76] M. H. Hayes, *Statistical Digital Signal Processing and Modeling* (Wiley, 1996).

[77] R. Horstmeyer, R. Heintzmann, G. Popescu, L. Waller, and C. Yang, "Standardizing the resolution claims for coherent microscopy," Nature Photon. (2016).

[78] M. Kellman, E. Bostan, N. Repina, and L. Waller, "Physics-based learned design: Optimized coded-illumination for quantitative phase imaging," IEEE Transactions on Computational Imaging (2019).

[79] Z. Phillips, M. CHEN, and L. Waller, "Quantitative Differential Phase Contrast Imaging with Pupil Recovery," in *Optics in the Life Sciences Congress*, p. JTu5A.2 (Optical Society of America, 2017).

[80] J. R. Silvester, "Determinants of block matrices," The Mathematical Gazette **84**(501), 460–467 (2000).

[81] Z. Phillips, M. Chen, and L. Waller, "Single-shot quantitative phase microscopy with color-multiplexed differential phase contrast (cDPC)," PloS one **12**(2), e0171,228 (2017).

[82] Z. E. Perlman, M. D. Slack, Y. Feng, T. J. Mitchison, L. F. Wu, and S. J. Altschuler, "Multidimensional Drug Profiling By Automated Microscopy," Science **306**(5699), 1194–1198 (2004). `http://science.sciencemag.org/content/306/5699/1194.full.pdf`, URL `http://science.sciencemag.org/content/306/5699/1194`.

[83] P. Brodin and T. Christophe, "High-content screening in infectious diseases," Current opinion in chemical biology **15**(4), 534–539 (2011).

[84] M. Bickle, "The beautiful cell: high-content screening in drug discovery," Analytical and bioanalytical chemistry **398**(1), 219–226 (2010).

[85] U. Liebel, V. Starkuviene, H. Erfle, J. C. Simpson, A. Poustka, S. Wiemann, and R. Pepperkok, "A microscope-based screening platform for large-scale functional protein analysis in intact cells," FEBS Letters **554**(3), 394–398 (2003).

[86] W.-K. Huh, J. V. Falvo, L. C. Gerke, A. S. Carroll, R. W. Howson, J. S. Weissman, and E. K. O'shea, "Global analysis of protein localization in budding yeast," Nature **425**(6959), 686 (2003).

[87] J. Peiffer, F. Majewski, H. Fischbach, J. Bierich, and B. Volk, "Alcohol embryo-and fetopathy: Neuropathology of 3 children and 3 fetuses," Journal of the Neurological Sciences **41**(2), 125–137 (1979).

[88] M. Remmelinck, M. B. S. Lopes, N. Nagy, S. Rorive, K. Rombaut, C. Decaestecker, R. Kiss, and I. Salmon, "How could static telepathology improve diagnosis in neuropathology?" Analytical Cellular Pathology **21**(3-4), 177–182 (2000).

[89] M. Alegro, P. Theofilas, A. Nguy, P. A. Castruita, W. Seeley, H. Heinsen, D. M. Ushizima, and L. T. Grinberg, "Automating cell detection and classification in human brain fluorescent microscopy images using dictionary learning and sparse coding," Journal of neuroscience methods **282**, 20–33 (2017).

[90] L. Tian, X. Li, K. Ramchandran, and L. Waller, "Multiplexed coded illumination for Fourier Ptychography with an LED array microscope," Biomed. Opt. Express **5**(7), 2376–2389 (2014). URL `http://www.opticsinfobase.org/boe/abstract.cfm?URI=boe-5-7-2376`.

[91] C. Z. Microscopy, "Zeiss Axio Scan.Z1," (2019). URL `https://www.zeiss.com/microscopy/us/products/imaging-systems/axio-scan-z1.html`.

[92] A. E. Carpenter, T. R. Jones, M. R. Lamprecht, C. Clarke, I. H. Kang, O. Friman, D. A. Guertin, J. H. Chang, R. A. Lindquist, J. Moffat, P. Golland, and D. M. Sabatini, "CellProfiler: image analysis software for identifying and quantifying cell phenotypes," Genome Biology **7**(10), R100 (2006). URL `https://doi.org/10.1186/gb-2006-7-10-r100`.

[93] A. W. Lohmann, R. G. Dorsch, D. Mendlovic, Z. Zalevsky, and C. Ferreira, "Space–bandwidth product of optical signals and systems," J. Opt. Soc. Am. A **13**(3), 470–473 (1996).

[94] L. Tian, Z. Liu, L.-H. Yeh, M. Chen, J. Zhong, and L. Waller, "Computational illumination for high-speed in vitro Fourier ptychographic microscopy," Optica **2**(10), 904–911 (2015).

[95] L. T. Grinberg, R. E. de Lucena Ferretti, J. M. Farfel, R. Leite, C. A. Pasqualucci, S. Rosemberg, R. Nitrini, P. H. N. Saldiva, W. Jacob Filho, and Brazilian Aging Brain Study Group, "Brain bank of the Brazilian aging brain study group—a milestone reached and more than 1,600 collected brains," Cell and Tissue Banking **8**(2), 151–162 (2007). URL `https://doi.org/10.1007/s10561-006-9022-z`.

[96] L. Rayleigh, "XV. On the theory of optical images, with special reference to the microscope," The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science **42**(255), 167–195 (1896).

[97] R. Raskar, A. Agrawal, and J. Tumblin, "Coded exposure photography: motion deblurring using fluttered shutter," ACM Transactions on Graphics (TOG) **25**(3), 795–804 (2006).

[98] Y. E. Nesterov, "A method for solving the convex programming problem with convergence rate $O(1/k^2)$," Dokl. Akad. Nauk SSSR **269**, 543–547 (1983). URL `https://ci.nii.ac.jp/naid/10029946121/en/`.

[99] K. Mitra, O. S. Cossairt, and A. Veeraraghavan, "A Framework for Analysis of Computational Imaging Systems: Role of Signal Prior, Sensor Noise and Multiplexing," IEEE Transactions on Pattern Analysis and Machine Intelligence **36**(10), 1909–1921 (2014).

[100] K. Hagiwara and K. Kuno, "Regularization learning and early stopping in linear networks," in *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*, vol. 4, pp. 511–516 (IEEE, 2000).

[101] P. Yalamanchili, U. Arshad, Z. Mohammed, P. Garigipati, P. Entschev, B. Kloppenborg, J. Malcolm, and J. Melonakos, "ArrayFire - A high performance software library for parallel computing with an easy-to-use API," (2015).

[102] *Bounds for the condition numbers of spatially-variant convolution matrices in image restoration problems* (OSA, Toronto, 2011).

[103] Website:, "https://github.com/zfphil/illuminate," Github (2019).

[104] A. Edelstein, N. Amodaj, K. Hoover, R. Vale, and N. Stuurman, "Computer Control of Microscopes Using µManager," Current Protocols in Molecular Biology **92**(1), 14.20.1–14.20.17 (2010).

[105] T. Kluyver, B. Ragan-Kelley, F. Pérez, B. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. Hamrick, J. Grout, S. Corlay, P. Ivanov, D. Avila, S. Abdalla, and C. Willing, "Jupyter Notebooks – a publishing format for reproducible computational workflows," in *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, F. Loizides and B. Schmidt, eds., pp. 87 – 90 (IOS Press, 2016).

[106] "Nikon Perfect Focus," `https://www.microscopyu.com/applications/live-cell-imaging/nikon-perfect-focus-system`. Accessed: 2019-01-23.

[107] "Zeiss Definite Focus," `https://www.zeiss.com/microscopy/us/products/light-microscopes/axio-observer-for-biology/definite-focus.html`. Accessed: 2019-01-23.

[108] C. S. Murphy and M. W. Davidson, "Light Source Power Levels," `http://zeiss-campus.magnet.fsu.edu/articles/lightsources/powertable.html` (2019). Accessed: 2019-02-13.

[109] A. C. Croce and G. Bottiroli, "Autofluorescence spectroscopy and imaging: a tool for biomedical research and diagnosis," European journal of histochemistry : EJH **58**(4), 2461; 2461–2461 (2014). URL `https://www.ncbi.nlm.nih.gov/pubmed/25578980`.

[110] J. Lippincott-Schwartz, N. Altan-Bonnet, and G. H. Patterson, "Photobleaching and photoactivation: following protein dynamics in living cells." Nature cell biology pp. S7–14 (2003).

[111] M. Y. Berezin and S. Achilefu, "Fluorescence Lifetime Measurements and Biological Imaging," Chemical Reviews **110**(5), 2641–2684 (2010).

[112] A. Köhler, "Ein neues Beleuchtungsverfahren für mikrophotographische Zwecke," Zeitschrift für wissenschaftliche Mikroskopie und für mikroskopische Technik **10**(4), 433–440 (1893).

[113] W. T. Anderson, "Xenon compact arc lamps," JOSA **41**(6), 385–388 (1951).

[114] T. A. Edison, "Electric lamp," (1880). US Patent 223,898.

[115] A. L. Schawlow and C. H. Townes, "Infrared and optical masers," Physical Review **112**(6), 1940 (1958).

[116] V. Poher, H. Zhang, G. Kennedy, C. Griffin, S. Oddos, E. Gu, D. Elson, J. Girkin, P. French, M. Dawson, *et al.*, "Optical sectioning microscopes with no moving parts using a micro-stripe array light emitting diode," Optics Express **15**(18), 11,196–11,206 (2007).

[117] D. Dominguez, L. Molina, D. B. Desai, T. O'Loughlin, A. A. Bernussi, and L. G. de Peralta, "Hemispherical digital optical condensers with no lenses, mirrors, or moving parts," Optics Express **22**(6), 6948–6957 (2014). URL http://www.opticsexpress.org/abstract.cfm?URI=oe-22-6-6948.

[118] K. Guo, S. Dong, P. Nanda, and G. Zheng, "Optimization of sampling pattern and the design of Fourier ptychographic illuminator," Opt. Express **23**(5), 6171–6180 (2015).

[119] D. Brady, M. Gehm, R. Stack, D. Marks, D. Kittle, D. Golish, E. Vera, and S. Feller, "Multiscale gigapixel photography," Nature **486**(7403), 386–389 (2012).

[120] D. L. Marks, P. Llull, Z. Phillips, J. Anderson, S. Feller, E. Vera, H. Son, S.-H. Youn, J. Kim, M. Gehm, *et al.*, "Characterization of the AWARE 10 two-gigapixel wide-field-of-view visible imager," Applied optics **53**(13), C54–C63 (2014).

[121] P. Llull, L. Bange, Z. Phillips, K. Davis, D. L. Marks, and D. J. Brady, "Characterization of the AWARE 40 wide-field-of-view visible imager," Optica **2**(12), 1086–1089 (2015).

[122] H. Zhu, O. Yaglidere, T.-W. Su, D. Tseng, and A. Ozcan, "Cost-effective and compact wide-field fluorescent imaging on a cell-phone," Lab on a Chip **11**, 315–322 (2011).

[123] N. A. Switz, M. V. D'Ambrosio, and D. A. Fletcher, "Low-Cost Mobile Phone Microscopy with a Reversed Mobile Phone Camera Lens," PLOS ONE **9**(5), e95,330 (2014).

[124] Z. J. Smith, K. Chu, A. R. Espenson, M. Rahimzadeh, A. Gryshuk, M. Molinaro, D. M. Dwyre, S. Lane, D. Matthews, and S. Wachsmann-Hogiu, "Cell-phone-based platform for biomedical device development and education applications," PLOS ONE **6**(3), e17,150 (2011).

[125] R. N. Maamari, J. D. Keenan, D. A. Fletcher, and T. P. Margolis, "A mobile phone-based retinal camera for portable wide field imaging," British Journal of Ophthalmology **98**(4), 438–441 (2013).

[126] A. Ozcan, "Mobile phones democratize and cultivate next-generation imaging, diagnostics and measurement tools," Lab on a Chip **14**, 3187–3194 (2014). URL `http://dx.doi.org/10.1039/C4LC00010B`.

[127] S. Vashist, O. Mudanyali, E. Schneider, R. Zengerle, and A. Ozcan, "Cellphone-based devices for bioanalytical sciences," Analytical and Bioanalytical Chemistry **406**(14), 3263–3277 (2014). URL `http://dx.doi.org/10.1007/s00216-013-7473-1`.

[128] R. A. Steenblik and P. H. Steenblik, "Lenses and uses, including microscopes," (2005). US Patent 6,847,480.

[129] J. S. Cybulski, J. Clements, and M. Prakash, "Foldscope: origami-based paper microscope," PloS one **9**(6), e98,781 (2014).

[130] S. A. Boppart and R. Richards-Kortum, "Point-of-care and point-of-procedure optical imaging technologies for primary care and global health," Science Translational Medicine **6**(253), 253rv2–253rv2 (2014).

[131] A. Greenbaum, Y. Zhang, A. Feizi, P.-L. Chung, W. Luo, S. R. Kandukuri, and A. Ozcan, "Wide-field computational imaging of pathology slides using lens-free on-chip microscopy," Science Translational Medicine **6**(267), 267ra175 (2014). `http://stm.sciencemag.org/content/6/267/267ra175.full.pdf`, URL `http://stm.sciencemag.org/content/6/267/267ra175.abstract`.

[132] O. Mudanyali, D. Tseng, C. Oh, S. O. Isikman, I. Sencan, W. Bishara, C. Oztoprak, S. Seo, B. Khademhosseini, and A. Ozcan, "Compact, light-weight and cost-effective microscope based on lensless incoherent holography for telemedicine applications," Lab on a Chip **10**(11), 1417–1428 (2010).

[133] D. Tseng, O. Mudanyali, C. Oztoprak, S. O. Isikman, I. Sencan, O. Yaglidere, and A. Ozcan, "Lensfree microscopy on a cellphone," Lab on a Chip **10**(14), 1787–1792 (2010).

[134] D. N. Breslauer, R. N. Maamari, N. A. Switz, W. A. Lam, and D. A. Fletcher, "Mobile phone based clinical microscopy for global health applications," PloS one **4**(7), e6320 (2009).

[135] A. Skandarajah, C. D. Reber, N. A. Switz, and D. A. Fletcher, "Quantitative imaging with a mobile phone microscope," PloS one **9**(5), e96,906 (2014).

[136] B. Kachar, "Asymmetric illumination contrast: a method of image formation for video light microscopy," Science **227**(4688), 766–768 (1985).

[137] H.-U. Dodt, M. Eder, A. Frick, and W. Zieglgänsberger, "Precisely Localized LTD in the Neocortex Revealed by Infrared-Guided Laser Stimulation," Science **286**(5437), 110–113 (1999). `http://www.sciencemag.org/content/286/5437/110.full.pdf`.

[138] T. N. Ford, K. K. Chu, and J. Mertz, "Phase-gradient microscopy in thick tissue with oblique back-illumination," Nature Methods **9**(12), 1195–1197 (2012).

[139] L. Tian and L. Waller, "3D intensity and phase imaging from light field measurements in an LED array microscope," Optica **2**(2), 104–111 (2015).

[140] A. C. Kak and M. Slaney, *Principle of Computerized Tomographic Imaging* (Society for Industrial and Applied Mathematics, 2001).

[141] M. Guizar-Sicairos and J. R. Fienup, "Phase retrieval with transverse translation diversity: a nonlinear optimization approach," Optics Express **16**(10), 7264–7278 (2008).

[142] A. M. Maiden and J. M. Rodenburg, "An improved ptychographical phase retrieval algorithm for diffractive imaging," Ultramicroscopy **109**(10), 1256–1262 (2009).

[143] A. Tripathi, I. McNulty, and O. G. Shpyrko, "Ptychographic overlap constraint errors and the limits of their numerical recovery using conjugate gradient descent methods," Optics Express **22**(2), 1452–1466 (2014).

[144] A. M. Maiden, M. J. Humphry, and J. M. Rodenburg, "Ptychographic transmission microscopy in three dimensions using a multi-slice approach," Journal of the Optical Society of America A: Optics, Image Science, and Vision **29**(8), 1606–1614 (2012). URL `http://josaa.osa.org/abstract.cfm?URI=josaa-29-8-1606`.

[145] Z. Jingshan, L. Tian, J. Dauwels, and L. Waller, "Partially coherent phase imaging with simultaneous source recovery," Biomedical optics express **6**(1), 257–265 (2015).

[146] J. Zhong, L. Tian, P. Varma, and L. Waller, "Nonlinear optimization algorithm for partially coherent phase retrieval and source recovery," IEEE Transactions on Computational Imaging **2**(3), 310–322 (2016).

[147] Z. Bian, S. Dong, and G. Zheng, "Adaptive system correction for robust Fourier ptychographic imaging," Optics Express **21**(26), 32,400–32,410 (2013). URL `http://www.opticsexpress.org/abstract.cfm?URI=oe-21-26-32400`.

[148] L.-H. Yeh, J. Dong, J. Zhong, L. Tian, M. Chen, G. Tang, M. Soltanolkotabi, and L. Waller, "Experimental robustness of Fourier ptychography phase retrieval algorithms," Opt. Express **23**(26), 33,214–33,240 (2015). URL `http://www.opticsexpress.org/abstract.cfm?URI=oe-23-26-33214`.

[149] L. Bian, G. Zheng, K. Guo, J. Suo, C. Yang, F. Chen, and Q. Dai, "Motion-corrected Fourier ptychography," Biomed. Opt. Express **7**(11), 4543–4553 (2016). URL `http://www.osapublishing.org/boe/abstract.cfm?URI=boe-7-11-4543`.

[150] G. Satat, M. Tancik, O. Gupta, B. Heshmat, and R. Raskar, "Object classification through scattering media with deep learning on time resolved measurement," Optics express **25**(15), 17,466–17,479 (2017).

[151] G. Zheng, X. Ou, R. Horstmeyer, and C. Yang, "Characterization of spatially varying aberrations for wide field-of-view microscopy," Optics express **21**(13), 15,131–15,143 (2013).

[152] R. Horstmeyer, X. Ou, J. Chung, G. Zheng, and C. Yang, "Overlapped Fourier coding for optical aberration removal," Opt. Express **22**(20), 24,062–24,080 (2014). URL `http://www.opticsexpress.org/abstract.cfm?URI=oe-22-20-24062`.

[153] X. Ou, G. Zheng, and C. Yang, "Embedded pupil function recovery for Fourier ptychographic microscopy," Optics Express **22**(5), 4960–4972 (2014). URL `http://www.opticsexpress.org/abstract.cfm?URI=oe-22-5-4960`.

[154] J. Chung, J. Kim, X. Ou, R. Horstmeyer, and C. Yang, "Wide field-of-view fluorescence image deconvolution with aberration-estimation from Fourier ptychography," Biomed. Opt. Express **7**(2), 352–368 (2016).

[155] X. Ou, R. Horstmeyer, G. Zheng, and C. Yang, "High numerical aperture Fourier ptychography: principle, implementation and characterization," Optics express **23**(3), 3472–3491 (2015).

[156] T. J. Holmes, "Blind deconvolution of quantum-limited incoherent imagery: maximum-likelihood approach," J. Opt. Soc. Am. A **9**(7), 1052–1061 (1992).

[157] P. Sarder and A. Nehorai, "Deconvolution methods for 3-D fluorescence microscopy images," IEEE Signal Processing Magazine **23**(3), 32–45 (2006).

[158] G. Ayers and J. C. Dainty, "Iterative blind deconvolution method and its applications," Optics letters **13**(7), 547–549 (1988).

[159] A. J. Bell and T. J. Sejnowski, "An information-maximization approach to blind separation and blind deconvolution," Neural computation **7**(6), 1129–1159 (1995).

[160] T. F. Chan and C.-K. Wong, "Total variation blind deconvolution," IEEE transactions on Image Processing **7**(3), 370–375 (1998).

[161] A. Levin, "Blind motion deblurring using image statistics," in *Advances in Neural Information Processing Systems*, pp. 841–848 (2006).

[162] von F. Zernike, "Beugungstheorie des schneidenver-fahrens und seiner verbesserten form, der phasenkontrastmethode," Physica **1**(7), 689 – 704 (1934).

[163] M. Chen, L. Tian, and L. Waller, "3D differential phase contrast microscopy," Biomedical optics express **7**(10), 3940–3950 (2016).

[164] Z. Liu, L. Tian, S. Liu, and L. Waller, "Real-time brightfield, darkfield, and phase contrast imaging in a light-emitting diode array microscope," Journal of Biomedical Optics **19**(10), 5–19 (2014). URL `http://dx.doi.org/10.1117/1.JBO.19.10.106002`.

[165] J. A. Rodrigo and T. Alieva, "Rapid quantitative phase imaging for partially coherent light microscopy," Optics express **22**(11), 13,472–13,483 (2014).

[166] H. Lu, J. Chung, X. Ou, and C. Yang, "Quantitative phase imaging and complex field reconstruction by pupil modulation differential phase contrast," Optics express **24**(22), 25,345–25,361 (2016).

[167] D. C. Liu and J. Nocedal, "On the limited memory BFGS method for large scale optimization," Mathematical programming **45**(1-3), 503–528 (1989).

[168] J. Chung, H. Lu, X. Ou, H. Zhou, and C. Yang, "Wide-field Fourier ptychographic microscopy using laser illumination source," Biomedical optics express **7**(11), 4787–4802 (2016).

[169] M. Chen, L. Tian, and L. Waller, "3D differential phase contrast microscopy," Biomedical Optics Express **7**, 3940 (2016).

[170] X. Ou, R. Horstmeyer, C. Yang, and G. Zheng, "Quantitative phase imaging via Fourier ptychographic microscopy," Optics Letters **38**(22), 4845–4848 (2013).

[171] S. Dong, K. Guo, P. Nanda, R. Shiradkar, and G. Zheng, "FPscope: a field-portable high-resolution microscope using a cellphone lens," Biomed. Opt. Express **5**(10), 3305–3310 (2014). URL `http://www.opticsinfobase.org/boe/abstract.cfm?URI=boe-5-10-3305`.

[172] P. Thibault, M. Dierolf, O. Bunk, A. Menzel, and F. Pfeiffer, "Probe retrieval in ptychographic coherent diffractive imaging," Ultramicroscopy **109**(4), 338–343 (2009).

[173] J. Sun, Q. Chen, Y. Zhang, and C. Zuo, "Efficient positional misalignment correction method for Fourier ptychographic microscopy," Biomed. Opt. Express **7**(4), 1336–1350 (2016).

[174] J. Liu, Y. Li, W. Wang, H. Zhang, Y. Wang, J. Tan, and C. Liu, "Stable and robust frequency domain position compensation strategy for Fourier ptychographic microscopy," Opt. Express **25**(23), 28,053–28,067 (2017).

[175] A. Maiden, M. Humphry, M. Sarahan, B. Kraus, and J. Rodenburg, "An annealing algorithm to correct positioning errors in ptychography," Ultramicroscopy **120**, 64–72 (2012).

[176] F. Zhang, I. Peterson, J. Vila-Comamala, A. Diaz, F. Berenguer, R. Bean, B. Chen, A. Menzel, I. K. Robinson, and J. M. Rodenburg, "Translation position determination in ptychographic coherent diffraction imaging," Optics express **21**(11), 13,592–13,606 (2013).

[177] J. Dou, Z. Gao, J. Ma, C. Yuan, Z. Yang, and L. Wang, "Iterative autofocusing strategy for axial distance error correction in ptychography," Optics and Lasers in Engineering **98**(Supplement C), 56 – 61 (2017). URL `http://www.sciencedirect.com/science/article/pii/S0143816617301409`.

[178] R. Eckert, L. Tian, and L. Waller, "Algorithmic Self-calibration of Illumination Angles in Fourier Ptychographic Microscopy," in *Imaging and Applied Optics 2016*, p. CT2D.3 (Optical Society of America, 2016).

[179] G. Satat, B. Heshmat, D. Raviv, and R. Raskar, "All Photons Imaging Through Volumetric Scattering," Scientific Reports **6**, 33,946 (2016).

[180] A. Pan, Y. Zhang, T. Zhao, Z. Wang, D. Dan, M. Lei, and B. Yao, "System calibration method for Fourier ptychographic microscopy," Journal of Biomedical Optics **22**(9), 11 – 22 (2017). URL `http://dx.doi.org/10.1117/1.JBO.22.9.096005`.

[181] T. M. Turpin, L. H. Gesell, J. Lapides, and C. H. Price, "Theory of the synthetic aperture microscope," in *Advanced Imaging Technologies and Commercial Applications*, vol. 2566, pp. 230–241 (International Society for Optics and Photonics, 1995).

[182] J. Di, J. Zhao, H. Jiang, P. Zhang, Q. Fan, and W. Sun, "High resolution digital holographic microscopy with a wide field of view based on a synthetic aperture technique and use of linear CCD scanning," Appl. Opt. **47**(30), 5654–5659 (2008).

[183] R. Horstmeyer, J. Chung, X. Ou, G. Zheng, and C. Yang, "Diffraction tomography with Fourier ptychography," Optica **3**(8), 827–835 (2016).

[184] J. Sun, Q. Chen, Y. Zhang, and C. Zuo, "Sampling criteria for Fourier ptychographic microscopy in object space and frequency space," Opt. Express **24**(14), 15,765–15,781 (2016).

[185] C. Kuang, Y. Ma, R. Zhou, J. Lee, G. Barbastathis, R. R. Dasari, Z. Yaqoob, and P. T. C. So, "Digital micromirror device-based laser-illumination Fourier ptychographic microscopy," Opt. Express **23**(21), 26,999–27,010 (2015). URL `http://www.opticsexpress.org/abstract.cfm?URI=oe-23-21-26999`.

[186] S. Sen, I. Ahmed, B. Aljubran, A. A. Bernussi, and L. G. de Peralta, "Fourier ptychographic microscopy using an infrared-emitting hemispherical digital condenser," Appl. Opt. **55**(23), 6421–6427 (2016).

[187] S. Hell, G. Reiner, C. Cremer, and E. H. Stelzer, "Aberrations in confocal fluorescence microscopy induced by mismatches in refractive index," Journal of microscopy **169**(3), 391–405 (1993).

[188] S. Kang, P. Kang, S. Jeong, Y. Kwon, T. D. Yang, J. H. Hong, M. Kim, K.-D. Song, J. H. Park, J. H. Lee, M. J. Kim, K. H. Kim, and W. Choi, "High-resolution adaptive optical imaging within thick scattering media using closed-loop accumulation of single scattering," Nature Communications **8**(1), 2157 (2017). URL `https://doi.org/10.1038/s41467-017-02117-8`.

[189] A. Shanker, A. Wojdyla, G. Gunjala, J. Dong, M. Benk, A. Neureuther, K. Goldberg, and L. Waller, "Off-axis Aberration Estimation in an EUV Microscope Using Natural Speckle," in *Imaging and Applied Optics 2016*, p. ITh1F.2 (Optical Society of America, 2016).

[190] C. Dammer, P. Leleux, D. Villers, and M. Dosire, "Use of the Hough transform to determine the center of digitized X-ray diffraction patterns," Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms **132**(1), 214 – 220 (1997).

[191] J. Cauchie, V. Fiolet, and D. Villers, "Optimization of an Hough transform algorithm for the search of a center," Pattern Recognition **41**(2), 567 – 574 (2008).

[192] H. K. Yuen, J. Princen, J. Illingworth, and J. Kittler, "A comparative study of Hough transform methods for circle finding," in *Proc. 5th Alvey Vision Conf., Reading (31 Aug*, pp. 169–174 (1989).

[193] E. Davies, *Machine Vision: Theory, Algorithms and Practicalities*, 3rd ed. (Morgan Kauffmann Publishers, 2004).

[194] M. Born and E. Wolf, *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*, 7th ed. (Cambridge University Press, 1999).

[195] J. Sun, C. Zuo, L. Zhang, and Q. Chen, "Resolution-enhanced Fourier ptychographic microscopy based on high-numerical-aperture illuminations," in *Scientific reports* (2017).

[196] T. Falk, D. Mai, R. Bensch, Ö. Çiçek, A. Abdulkadir, Y. Marrakchi, A. Böhm, J. Deubner, Z. Jäckel, K. Seiwald, A. Dovzhenko, O. Tietz, C. Dal Bosco, S. Walsh, D. Saltukoglu, T. L. Tay, M. Prinz, K. Palme, M. Simons, I. Diester, T. Brox, and O. Ronneberger, "U-Net: deep learning for cell counting, detection, and morphometry," Nature Methods **16**(1), 67–70 (2019). URL `https://doi.org/10.1038/s41592-018-0261-2`.

[197] Y. Xue and N. Ray, "Cell Detection in Microscopy Images with Deep Convolutional Neural Network and Compressed Sensing," arXiv preprint arXiv:1708.03307 (2017).

[198] Z. Phillips, R. Eckert, and L. Waller, "Quasi-dome: A self-calibrated high-na led illuminator for fourier ptychography," in *Imaging Systems and Applications*, pp. IW4E–5 (Optical Society of America, 2017).

[199] M. D. Gupta, S. Kumar, and J. Xiao, "L1 projections with box constraints," arXiv preprint arXiv:1010.0141 (2010).