University of California
Santa Barbara

Computational Imaging Methods for Improving Resolution in Biological Microscopy

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Electrical & Computer Engineering

by

Kevin G. Chan

Committee in charge:

Professor Michael Liebling, Chair
Professor Hua Lee
Professor B. S. Manjunath
Professor Pradeep Sen

March 2017

The Dissertation of Kevin G. Chan is approved.

Professor Hua Lee

Professor B. S. Manjunath

Professor Pradeep Sen

Professor Michael Liebling, Committee Chair

March 2017

Computational Imaging Methods for Improving Resolution in Biological Microscopy

by

Kevin G. Chan

For Lauren

Acknowledgements

First and foremost, I would like to express my sincere gratitude to my advisor, Prof. Michael Liebling, for his guidance and support throughout my entire graduate career. I appreciate the opportunity he's given me to work in his lab, as well as all the insights he's imparted to me over the years. Without him, this thesis would not have been possible. I also would like to thank him and his family, Maeva, Alice, and Mina, for the warmth, kindness, and generosity they showed me when I felt like a stranger in a strange land.

In addition, I would like to thank Prof. Hua Lee, Prof. Pradeep Sen, Prof. John Shynk, and Prof. B. S. Manjunath, for being on my Ph.D. committee. Their feedback and their expertise are much appreciated.

I would also like to thank Nicholas Högasten, Austin Richards, Stephanie Lin, and everyone else at FLIR Systems for the valuable internship opportunity.

I am also grateful to all the people at UCSB who have supported me and my research over the years. Specifically, I thank Val de Veyra, Student Affairs Manager of the ECE Department, for all her support and kindness throughout my graduate career. I also thank Ronni Wynn, facility manager for the Animal Resource Center at UCSB, for all her assistance with zebrafish care. I also thank Dr. Sebastian Streichan at the Kavli Institute for Theoretical Physics for our collaborations and the use of his custom single plane illumination microscope.

I also thank all the previous members of the Systems Bioimaging Lab – Nikhil, Michael, Jungho, Sandeep, Chieh, Jeff, and John – for being wonderful colleagues and friends, both in and out of the lab.

Finally, I thank my dad, whose guidance has made me who I am today, and Lauren, whose unwavering support keeps me motivated for the future.

Kevin G. Chan

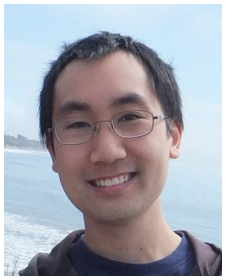

CONTACT INFORMATION

Email: kevin@kchan.me
Phone: (510) 557-3243
Web: http://www.kchan.me
LinkedIn: linkedin.com/in/kevingchan

EDUCATION

University of California, Santa Barbara, Santa Barbara, CA
Ph.D., Electrical and Computer Engineering, 2017

- Adviser: Dr. Michael Liebling
- Group: Systems Bioimaging Laboratory
- Area of Study: Signal & Image Processing
- Thesis: Computational imaging methods for improving resolution in biological microscopy

University of California, Santa Barbara, Santa Barbara, CA
M.S., Electrical and Computer Engineering, 2013

Harvey Mudd College, Claremont, CA
B.S., Engineering, 2011

WORK EXPERIENCE

Idiap Research Institute, Martigny, Switzerland **2015 - 2016**
Computational Imaging Intern

- Designed and assembled a computational imaging system with active illumination for temporal superresolution.
- Investigated the performance of different active illumination codes using computational simulations.
- Developed a video reconstruction algorithm capable of achieving temporal superresolution by a factor of up to 6×.

FLIR Systems, Goleta, CA **2015**
Video, Signal Processing, and Algorithms Intern

- Developed image processing simulations for infrared camera systems.
- Implemented single image superresolution for low-cost, low-resolution infrared camera sensors.
- Implemented a spatially-variant, point spread function-aware algorithm for infrared image deblurring.

RESEARCH EXPERIENCE

Systems Bioimaging Laboratory, UC Santa Barbara **2012-present**
Graduate Student Researcher

- I am currently working on biological image processing algorithms, including temporal superresolution, deconvolution, tomographic reconstruction, and blood flow video analysis.

- I am developing computational imaging methods that combine novel hardware with image processing algorithms for cardiac fluorescence microscopy. I apply these methods to imaging experiments with in vivo microscopy of live transgenic zebrafish.
- I develop bioimage processing tools for ImageJ and Imaris with Java and Matlab.

Computer Science Department, Harvey Mudd College **2010**
Undergraduate Researcher

- I helped develop LogiSketch, an educational software system written in C# for sketching and simulating digital logic circuits on a tablet PC.
- I implemented a decision tree algorithm to increase recognition speed and a context-based refinement algorithm to increase recognition accuracy.
- I implemented on-line learning with a Bayes Classifer to improve recognition robustness and adapt to the user's individual drawing style.

Laboratory of Neuroimaging, UCLA **2009**
Undergraduate Researcher

- I collaborated with a multidisciplinary team to analyze registration of tetrahedral-based volumetric meshes of the brain.
- I implemented various metrics in Matlab for localizing and quantifying changes during registration.
- I applied my analysis methods to MRI data from a study of Alzheimer's Disease from the Alzheimer's Disease Neuroimaging Initiative (ADNI).

Teaching Experience

UC Santa Barbara, Santa Barbara, CA
Teaching Assistant **2011-2015**

- ECE 2A, 2B, 2C: Circuits, Devices, and Systems
- ECE 15A: Fundamentals of Logic Design
- ECE 178: Digital Image and Video Processing
- ECE 278B: Principles of Biological Microscopy

Research Mentorship Program **2012, 2013**

- I mentored 1-2 high school students for 6 weeks in the summer as they conducted a research project in the Systems Bioimaging Lab.

Condor Techs **2013**

- I mentored 4 students from Oxnard College for 2 weeks in the summer as they conducted a short research project in the Systems Bioimaging Lab.

Skills

- Matlab, Java, C#, C/C++, Python
- Microsoft Office, LaTeX, SVN
- ImageJ, Imaris, Adobe Photoshop, Adobe Illustrator

Publications

K. G. Chan, S. Calinon, and M. Liebling, "Temporal superresolution imaging of repeating processes using a single camera and active illumination," *IEEE Transactions on Computational Imaging*, (submitted).

K. G. Chan, S. J. Streichan, L. A. Trinh, and M. Liebling, "Simultaneous temporal superresolution and denoising for cardiac fluorescence microscopy," *IEEE Transactions on Computational Imaging*, vol. 2, no. 3, pp. 348–358, 2016.

K. G. Chan and M. Liebling, "A point-spread-function-aware filtered back-projection algorithm for focal-plane-scanning optical projection tomography," in *IEEE International Symposium on Biomedical Imaging*, 2016.

N. Chacko, **K. G. Chan**, and M. Liebling, "Intensity-based point-spread-function-aware registration for multi-view applications in optical microscopy," in *IEEE International Symposium on Biomedical Imaging*, 2015.

K. G. Chan and M. Liebling, "Estimation of divergence-free 3D cardiac blood flow in a zebrafish larva using multi-view microscopy," in *IEEE International Symposium on Biomedical Imaging*, 2015.

K. Chan, L. Trinh, and M. Liebling, "A temporal superresolution method applied to low-light cardiac fluorescence microscopy," in *Proceedings of the IEEE Asilomar Conference on Signals, Systems and Computers*, 2013.

Abstract

Computational Imaging Methods for Improving Resolution in Biological Microscopy

by

Kevin G. Chan

Optical microscopy is an essential tool for biological research, as it allows for non-invasive imaging of small animals. However, optical microscopy has its limits. Due to the low light level, fluorescence microscopy prohibits high speed imaging, making it difficult to study fast dynamic biological processes. In addition, optical blur due to the diffraction of light results in limited spatial resolution, particularly when using objective lenses with low numerical apertures. In this thesis, we propose computational imaging methods to overcome these limitations using a combination of novel image acquisition procedures and reconstruction algorithms.

The first part of this thesis deals with improving temporal resolution in fluorescence microscopy to image rapid, repeating processes. We take advantage of multiple acquisitions, each taken with different time delays or temporally modulated illumination patterns, to recover high frequency information that is lost with traditional imaging. We demonstrate our method to image the beating heart in live embryonic zebrafish with reduced motion blur and high resolution in time.

The second part of this thesis deals with reducing spatial blur in optical projection tomography, a form of optical microscopy that uses multiple 2D projections to reconstruct a 3D image of an object. We propose a method to reduce the optical distortion (as characterized by the system's optical point spread function) that can be implemented with a scanning acquisition approach combined with a modified filtered backprojection algorithm for reconstruction. We demonstrate our method to image blood vessels in

larval zebrafish with high spatial resolution and reduced out-of-focus blur.

The final part of this thesis deals with the dimensional limitation of 2D sensors for measuring 3D motion in microscopy. We propose a method to combine two-dimensional motion estimates from multiple views to recover out-of-plane velocity and reconstruct a divergence-free, three-dimensional velocity field. We demonstrate our method to measure, for the first time, dynamic blood flow in 3D inside the beating heart of a live zebrafish using optical microscopy.

This thesis provides new tools that integrate custom image acquisition procedures and image reconstruction algorithms to overcome the resolution limitations – temporal, spatial, and out-of-plane velocity resolution – in optical microscopy. The methods presented in this thesis, in particular the single camera, active illumination method for temporal superresolution in fluorescence microscopy, will be directly applicable to a broad range of biological studies and will open up new perspectives for imaging small organisms in 3D (and time) with high spatio-temporal resolution.

Contents

Chapter 1

Introduction

With the development of digital cameras, optical microscopy has become an indispensable tool for quantitative measurements in biological research. While there have been many recent advances in microscope architecture and optical design, there is still a need for better imaging solutions as scientists continually seek to perform measurements on smaller scales with higher resolution.

Recently, computational imaging has gained popularity as a new way of thinking about image processing. Rather than considering image processing as a separate step, independent of image acquisition, in computational imaging, image acquisition systems and image reconstruction algorithms are jointly designed. In computational microscopy, we apply the computational imaging mindset to microscopy system design. By considering image acquisition and image processing interdependently, we can design new integrated solutions for biomicroscopy that surpass the limits of traditional microscopes.

1.1 Motivation

By nature, biology is a field of science that is heavily driven by observation. In order to study and understand a complex biological system, it is first necessary to observe and measure it in its natural (or close to natural) setting. For example, to study the developing heart in small animals, it is necessary to observe the heart as non-invasively as possible and to quantitatively measure it with sufficiently high resolution to distinguish the features of interest. Optical microscopy is a widely used tool to perform such observations, for the heart as well as for many other biological systems. However, every optical imaging system has a limited resolving power, which prohibits measurement of features at a scale beyond the system's resolution limit.

In this thesis, we adhere to a strict definition of resolution limit: the smallest distance between which two separate features are still distinguishable after imaging. We consider a system with a resolution limit smaller than our features of interest to have (sufficiently) high resolution, and vice versa. Typically, this refers to spatial resolution, as any optical imaging system will produce some degree of optical blur due to the diffraction of light [1]. However, we take a more general definition of the separation distance between features, as features can be separated not only in space, but also in time, as well as in the feature dimension itself.

In this thesis, we consider three different types of resolution in biological microscopy imaging systems. First, temporal resolution refers to the minimum separation in time between two events such that they are still distinguishable. Temporal resolution is limited by the frame rate and exposure time of the camera or sensor, and it imposes restrictions on the speed of dynamic processes that can be observed and recorded. Second, spatial resolution refers to the minimum Cartesian distance separating two point sources such that they are still distinguishable, and it is limited by the diffraction of light and the

numerical aperture of the imaging system. Spatial resolution imposes restrictions on the smallest biological structures that can be imaged and measured. Finally, out-of-plane velocity resolution refers to the loss of 3D motion information when capturing images with a 2D sensor. Projecting 3D velocity down to two dimensions results in many 3D vectors becoming indistinguishable, as they are all projected onto the same 2D vector. Out-of-plane velocity resolution imposes limitations on our ability to study 3D biological motion, since any out-of-plane components of velocity are lost.

1.2 Contribution and Thesis Organization

In this thesis, we propose methods to improve resolution – temporal, spatial, and out-of-plane velocity resolution – in various forms of optical microscopy. These methods are based on the inversion of a forward imaging model, which we have flexibility in designing with computational imaging. We summarize our contributions below.

- **Improving Temporal Resolution in Cardiac Fluorescence Microscopy Using Active Illumination (Chapter 2)**

  We propose two computational imaging approaches for increasing temporal resolution beyond an imaging system's limit when imaging repeating processes. Our first approach uses constant illumination, is compatible with any conventional microscope without requiring hardware modifications, and is able to improve temporal resolution by a factor of approximately 1.6× in physical experiments (2.8× in noise-free simulations). Our second approach uses active, coded illumination, and is able to improve resolution by a factor of 2.5× in physical experiments (6.3× in noise-free simulations).

  Previous methods for temporal superresolution either require multiple cameras

[2, 3, 4] or a specially modified camera [5, 6, 7, 8]. In contrast, we use a single un-modified camera, and we instead use illumination modulation to reveal high frequency temporal information. Additionally, we propose minimizing an energy function with ℓ_1 norms for both the data fidelity term and the regularization term. While this is computationally more expensive than minimizing an ℓ_2 norm for a least-squares solution (as is commonly done), it produces a reconstruction that is less sensitive to outlier data and is qualitatively sharper. We demonstrate both our methods in practice for live imaging of the developing zebrafish heart under low-light fluorescence microscopy.

- **Improving Spatial Resolution in Optical Projection Tomography Using Point-Spread-Function-Aware Filtered Backprojection (Chapter 3)**

  We propose a direct inversion formula for optical projection tomography (OPT) that accounts for optical blur due to the diffraction of light. This approach improves spatial resolution in the tomographic reconstruction and reduces out-of-focus blur. We utilize a modified image acquisition procedure in which we scan the focal plane during acquisition to achieve a greater depth-of-field. We show that, when following such an acquisition procedure, our reconstruction approach is a regularized inverse of the forward model, and it can be implemented efficiently with a modified filtered backprojection algorithm. Previous methods for OPT reconstruction assumed straight-ray geometry [9, 10], which is a valid assumption for x-ray computed tomography, but is incorrect for OPT due to the diffraction of light. While several methods have acknowledged this, they have focused on ad-hoc solutions to mitigate the effects of light diffraction [11, 12]. In contrast, we propose an analytically-derived inversion formula based on a forward imaging model that incorporates the system's point spread function (PSF).

- **Improving Out-of-Plane Velocity Resolution in Cardiac Flow Velocimetry Using Multi-View Imaging (Chapter 4)**

  We propose a computational algorithm to reconstruct a three-dimensional velocity field from multiple two-dimensional velocity projections acquired from different views. This method allows us to measure, for the first time, three-dimensional blood flow through the live zebrafish heart using non-invasive optical microscopy. Due to the high velocities of blood cells, previous methods for measuring blood flow in microscopy only focused on flow in two dimensions [13, 14]. While methods have been proposed to recover three-dimensional flow with MRI [15, 16] and ultrasound imaging [17], our method is the first to measure three-dimensional blood flow with optical microscopy.

Chapter 2

Improving Temporal Resolution in Cardiac Fluorescence Microscopy Using Active Illumination

Abstract[1]

Fast, dynamic processes pose a challenge to conventional imaging, particularly in low-light settings such as fluorescence microscopy where the camera frame rate is limited due to insufficient light in the scene. In such a setting, videos are often acquired with long exposure times, resulting in poor temporal resolution, motion blur, or, possibly, temporal aliasing. Here, we propose two computational imaging approaches for temporal superresolution, one using constant illumination and one using active, coded illumination. These methods recover a high temporal resolution video from multiple low-resolution acquisitions acquired with a single camera. Using our approach, we observed, in practice, a resolution improvement factor of 1.6 with our constant illumination approach and a factor of 2.5 with our coded illumination approach. We also applied our methods to in

[1]This chapter is based on [18], [19],and [20].

vivo fluorescence imaging of the beating heart in embryonic zebrafish, where we observed a noticeable reduction of motion blur.

2.1 Introduction

When imaging rapid, dynamic processes, a long exposure time results in motion blur. Furthermore, a long exposure time also leads to a low sampling rate, which can result in temporal aliasing. For many applications, it is not possible to increase the sampling rate, either because of limitations of the camera technology, or because there is insufficient light in the process. Such extreme imaging situations are common in biological microscopy, where live samples are observed with a high ratio of velocity to required resolution. For example, a biological structure moving at a speed of 1000 $\mu\mathrm{m}/\mathrm{s}$ (such as cardiac muscle cells observed directly within the developing embryo) requires imaging beyond 1000 frames per second when the target resolution is 1 μm. This is due to the motion blur caused by the object's movement, which, at an exposure of $1/1000$ of a second, already amounts to a distance $1\ \mu\mathrm{m} = 1000\ \mu\mathrm{m}/\mathrm{s} \times 1/1000\ \mathrm{s}$ [21]. To study sub-cellular structures moving at similar speeds, even higher frame rates may be required, with the limitation that only light emitted by fluorescent molecules in the sample itself (and not reflected from an external source) is available. Given these challenges, we are interested in temporal superresolution methods which can reconstruct a high effective frame rate video from multiple low frame rate, low temporal resolution videos. We assume that each low temporal resolution video is acquired with a sufficient photon count per pixel to avoid the signal being overwhelmed by measurement noise. Additionally, we assume that the fluorescence emission response is linear, so that doubling the excitation intensity results in doubling the fluorescence emission. This is a valid assumption for single-photon imaging, provided the excitation intensity is within a range to avoid fluorophore saturation. We propose two such temporal

superresolution methods for cardiac fluorescence microscopy: one that is compatible with any standard microscope without the need for additional hardware [19], and one using active illumination with a custom, programmable light source [20]. The contributions of our methods are as follows:

1. We propose two temporal superresolution methods based on ℓ_1 minimizations: one using constant illumination that is compatible with any standard microscopy setup and requires no hardware modifications, and one that uses active (coded) illumination that provides greater resolution improvement.

2. We provide a FIJI software plugin to demonstrate our constant illumination temporal superresolution method.

3. For our active illumination method, we propose a closed-form expression to generate a well-conditioned set of active illumination codes for any reconstruction factor K.

4. We include temporal registration as part of our superresolution methods to correct for unsynchronized acquisitions.

5. We report simulation results under common experimental conditions showing that it is possible to achieve a resolution improvement factor of 2.8× and 6.3×, respectively, for our our constant illumination and active illumination methods.

6. In physical experiments, due to noise and registration error, we observed a practical resolution improvement factor of up to 1.6× for our constant illumination method and 2.5× for our active illumination method.

This chapter is organized as follows. In Section 2.2, we discuss several methods related to ours. In Section 2.4, we present the acquisition and superresolution method. In Section 2.5, we investigate the resolution improvement this technique offers, and we

demonstrate its performance in computational simulations and physical experiments. In Section 2.6, we summarize the advantages of our methods and discuss some details to consider when using them.

2.2 Related methods

Table 2.1: A summary of methods for imaging rapid, dynamic events

	Our Constant Method [19]	Our Active Method [20]	Shechtman *et al.* [2]	Agrawal *et al.* [3]	Pournaghi & Wu [4]	Bub *et al.* [24]	Reddy *et al.* [5]	Veeraraghavan *et al.* [6]	Holloway *et al.* [7]	Koller *et al.* [8]	Raskar *et al.* [27]	Ma *et al.* [28]
Uses a single camera	✓	✓				✓	✓	✓	✓	✓	✓	✓
Does not require special modified cameras	✓	✓	✓									✓
Can handle complex scene motion	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		
Does not rely heavily on prior knowledge/sparsity of signal	✓	✓	✓	✓		✓						
Uses temporally modulated illumination		✓						✓				✓

Methods for reconstructing a high frame rate video from a low frame rate video can be categorized into three groups: temporal interpolation, temporal rebinning (interleaving), and temporal superresolution. While methods have been proposed for sophisticated interpolation [22] and rebinning [23], these methods do not truly improve temporal resolution. For example, if one were to capture a video of a stationary light source whose intensity oscillated with a period equal to the camera's integration time, each frame in the resulting video would be identical, and these methods would be unable to recover the oscillation of the light intensity. Temporal superresolution methods, on the other hand, seek to improve resolution by using additional information about the imaging system to reconstruct a temporally superresolved video. In this section, we briefly describe

some of the existing approaches for temporal superresolution, and we summarize these for comparison in Table 2.1.

2.2.1 Multi-camera temporal superresolution

In [2], Shechtman *et al.* proposed a staggered exposure approach using multiple synchronized cameras to achieve temporal superresolution. Agrawal *et al.* extended this method by using different binary temporal sampling patterns for each camera, making the inverse superresolution problem well-posed [3]. More recently, Pournaghi and Wu further extended this multi-camera approach for temporal superresolution by using different binary sampling patterns for each pixel (or column of pixels) and exploiting spatio-temporal sparsity [4]. While these methods successfully reconstruct videos with improved temporal resolution, they require multiple synchronized cameras. As a result, the cost of such systems increases with the number of cameras used, which may end up being prohibitive for the sensitive cameras required in fluorescence microscopy. Our active illumination approach expands on the binary sampling patterns proposed by Agrawal *et al.* but instead uses a single camera system.

2.2.2 Single camera temporal superresolution

Several other works have proposed methods for temporal superresolution using a single, modified camera. For example, Bub *et al.* proposed a method for temporal superresolution by offsetting pixel exposure times during a single frame capture [24]. However, with this method, increasing temporal resolution required sacrificing spatial resolution. In [**?**], Reddy *et al.* proposed a compressed sensing camera architecture with programmable, multiplexed pixels that exploited the spatio-temporal redundancy in videos. Unfortunately, both of these methods require a modified camera system with per-

pixel exposure control, which is not commonly available in commercial camera systems.

Quite a few other single-camera, compressive sensing imaging systems have been proposed for temporal superresolution. In [6], Veeraraghavan *et al.* proposed a coded strobing method with a reconstruction algorithm assuming Fourier domain sparsity to image high-speed, periodic events. While our active illumination approach also utilizes coded sampling for periodic signals, we do not require Fourier-domain sparsity during the reconstruction step. This allows us to perform our reconstruction with only a few (e.g. $Q = 4$, where Q will be formally defined later in the text) acquired cycles. In addition, rather than using pseudo-random binary codes which can be time-consuming to optimize, we propose a closed-form expression for a well-conditioned set of binary modulation codes. Other compressive sensing systems include the flutter shutter video camera, which Holloway *et al.* combined with a total variation prior and a learned-dictionary video prior to synthesize high effective frame rate videos [7]. Also, Koller *et al.* developed a modified camera with a binary mask in the optical path and demonstrated a prototype capable of reconstructing 10 frames of video from a single coded image [8]. In contrast to these compressive sensing approaches, we do not require our underlying signal to be sparse, nor do we rely as heavily on our underlying signal conforming to a prior model. Instead, we assume that we can acquire additional observations of our signal so that the problem is no longer underdetermined.

2.2.3 Active illumination stroboscopy

Active illumination has been previously used for stroboscopy to freeze the motion of objects with extremely short, precisely-timed flashes of light [25]. Such a method has also been applied to cardiac microscopy [26]. In contrast, we use active *coded* illumination, in combination with a computational image reconstruction algorithm, to improve temporal

resolution for repeating processes captured with a single camera and ungated acquisitions. This allows for a lower peak illumination intensity compared to conventional strobing.

2.2.4 Coded motion deblurring

Motion deblurring methods are a class of methods related to temporal superresolution. While they do not improve temporal resolution in the strictest sense, they seek to reduce one of the symptoms of poor temporal resolution—spatial motion blur. In [27], Raskar *et al.* proposed coded motion deblurring (flutter shutter), where the camera exposure is modulated by a binary code during a single frame. This changes the point spread function of the motion blur and makes the spatial deconvolution problem well-posed. Ma *et al.* extended on this work by temporally encoding the illumination rather than the exposure and using a non-binary code [28]. By using a coded exposure rather than a constant exposure, these methods improved the invertability of the point spread function and made the associated deconvolution well-posed. However, these works are only capable of *spatially* deblurring uniform motion blur along a straight line (with a known point spread function), while our methods focus on temporal resolution improvement and do not impose any restrictions on the motion in the scene.

2.2.5 Translational spatial superresolution

While it is not a method for increasing temporal resolution, translational spatial superresolution has similarities with our constant illumination approach for temporal superresolution. In the spatial superresolution methods of, for example, Sroubek et al. [29], Farsiu et al. [30], and Ben-Ezra et al. [31], several low spatial resolution images with translational shifts were combined into a high spatial resolution image. Here, we instead reconstruct a *temporal* superresolution sequence from multiple low temporal resolution

image sequences.

2.2.6 Spatial structured illumination

While it is also not a method for increasing temporal resolution, structured illumination microscopy has some methodological parallels with our active illumination approach for temporal superresolution. Structured illumination is an approach for spatial superresolution in optical microscopy proposed by Gustafsson, in which spatially modulated light is used to illuminate an object. In a diffraction-limited microscopy system, spatially modulating the illumination allows access to high frequency information that can be used to reconstruct a superresolution image overcoming the limited spatial bandwidth of optical microscopy by a factor of up to 2 [32] or more [33]. Configurations with more general spatial light patterns have also been proposed [34], as well as configurations that capitalize on the non-linear relationship between excitation and fluorescence emission [35]. Alternative reconstruction techniques, which include estimation of hyper-parameters, have also been proposed [36].

In our approach, we use temporally structured illumination (i.e. active illumination) for superresolution in time rather than in the spatial dimensions. This imposes a different set of challenges, such as sub-resolution temporal registration of modulated signals, which we solve as part of our superresolution algorithm. In addition, the time-limited light integration of the shutter aperture differs from the frequency-limited setting of a lens.

2.3 Problem Formation

For a single pixel at location $\mathbf{x} = (x, y, z)$, we assume the image intensity is temporally repeating, and that a single cycle can be described by a discrete time sequence, $I_{\text{ref}}[\mathbf{x}, n], n = 0, \ldots, N-1$, where N is the number of samples covering the duration of

one cycle, and where the sample spacing T_{H} is sufficiently small to prevent temporal aliasing of the fundamental frequency. We wish to measure samples of this high temporal resolution sequence, which we place into a vector of length N (one vector for every pixel position $\mathbf{x} = (x, y, z)$),

$$\mathbf{f}_{\mathbf{x}} = \left(I_{\mathrm{ref}}[\mathbf{x}, 0], \ldots, I_{\mathrm{ref}}[\mathbf{x}, N-1]\right)^{\top}. \tag{2.1}$$

Since in low-light, live fluorescence microscopy, a long exposure time is needed to capture enough photons, this high temporal resolution sequence is not directly accessible. Instead, we acquire Q low resolution sequences $I_{\mathrm{acq}}^{(q)}$ of length M ($q = 1, \ldots, Q$), each covering the duration of a full period but with a larger sample spacing T_{L} ($T_{\mathrm{L}} > T_{\mathrm{H}}$), from which we will reconstruct the higher sampling rate sequence $\mathbf{f}_{\mathbf{x}}$. We assume that the ratio between sampling steps, $\frac{T_{\mathrm{L}}}{T_{\mathrm{H}}} = \frac{N}{M} = K$, is an integer. The q-th acquisition at pixel position $\mathbf{x} = (x, y, z)$,

$$\mathbf{g}_{\mathbf{x}}^{(q)} = \left(I_{\mathrm{acq}}^{(q)}[\mathbf{x}, 0], \ldots, I_{\mathrm{acq}}^{(q)}[\mathbf{x}, M-1]\right)^{\top}, \tag{2.2}$$

is an M-length sequence, where we model $I_{\mathrm{acq}}^{(q)}[\mathbf{x}, m]$ as a discrete convolution of the high temporal resolution sequence I_{ref} with the system's temporal integration filter along the time index followed by a decimation operation. Specifically,

$$I_{\mathrm{acq}}^{(q)}[\mathbf{x}, m] = \left(I_{\mathrm{ref}}[\mathbf{x}, \cdot] \circledast h_{\mathrm{int}}^{(q)} \circledast s_q\right)_{\downarrow K}[\mathbf{x}, m] + v_q[\mathbf{x}, m], \tag{2.3}$$

where $v_q[\mathbf{x}, m]$ is an additive noise term, K is the decimation factor, $h_{\mathrm{int}}^{(q)}$ represents the system's temporal light integration filter, and s_q applies a circular sub-frame shift such that

$$s_q[n] = \alpha_q \delta[n + \xi_q] + (1 - \alpha_q)\, \delta[n + \xi_q + 1], \tag{2.4}$$

to account for the fact that the camera starts capturing the first frame at an arbitrary time $\sigma_q \in \mathbb{R}$ after the beginning of the periodic sequence. In (2.4), $\delta[n]$ is the Kronecker delta function, and the circular sub-frame time shift σ_q is separated into an integer shift $\xi_q = \lfloor \sigma_q \mod N \rfloor$ and a fractional shift $\alpha_q = (\sigma_k \mod N) - \xi_q$, where mod denotes the modulo operation.

2.4 Proposed Methods

We propose two temporal superresolution methods, one using constant illumination and one using active illumination. While our two methods share a common model of the problem, they differ in their model of the system's temporal light integration filter, h_{int}. As a result, they have different models for image acquisition, as described in Sections 2.4.1.1 and 2.4.2.2 and illustrated in Figure 2.1. They also differ in the way the Q low resolution sequences are acquired. In our constant illumination approach, we assume that all low resolution sequences are acquired as part of a single long video, so that there is no delay between the end of one acquisition and the start of the next. On the other hand, for our coded illumination approach, we assume that each low resolution sequence is acquired separately with an unknown phase shift between them. Finally, our two methods use different approaches to temporal registration and superresolution reconstruction, as will be described in the following sections.

2.4.1 Constant Illumination Imaging System

2.4.1.1 Image Acquisition Model

Under constant illumination, we capture a long video sequence, $I_{\text{acq}}[\mathbf{x}, n]$ for $n = 1, \ldots, MQ$, containing $I_{\text{acq}}^{(q)}$ for $q = 1, \ldots, Q$. These are Q repetitions of our underlying

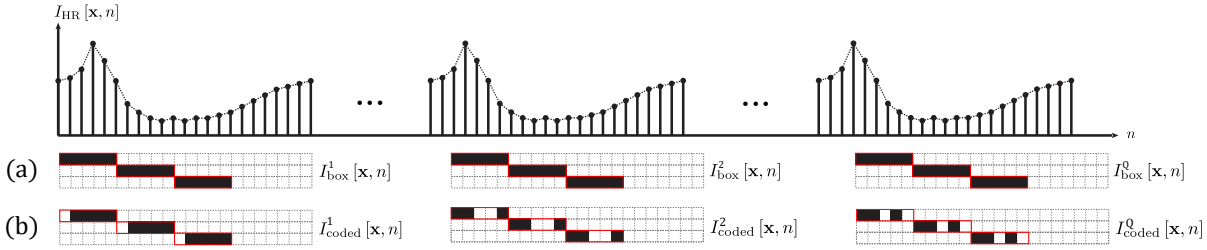

Figure 2.1: We assume that we can acquire Q cycles of a temporally repeating process. (a) Under constant illumination, each acquired frame can be considered as a summation of K high resolution frames. (b) With coded illumination, each high resolution frame is multiplied by a binary weight prior to summation. Using a different illumination code for each acquired cycle allows us to better reconstruct the high resolution signal $I_{\mathrm{HR}}[\mathbf{x}, n]$. Since we do not synchronize our acquisitions with the underlying signal, each acquired cycle may begin at a different phase of the repeating cycle.

repeating process of interest. We assume that at a given coordinate $\mathbf{x}$, the underlying signal, $I[\mathbf{x}, n]$, at any time point $n \in \mathbb{Z}$ (not only $0 \leq n < N$), is a temporally repeating signal such that $I[\mathbf{x}, n]$ is either equal (up to some additive noise) to a matching time point in I_{ref}, or can be interpolated from I_{ref}, provided that I_{ref} is chosen to be a complete cycle. More specifically, using linear interpolation, our assumption translates to:

$$\begin{aligned} I[\mathbf{x}, n] = &\left(1 - \left(w[n] - \lfloor w[n] \rfloor\right)\right) I_{\mathrm{ref}}[\mathbf{x}, \lfloor w[n] \rfloor] \\ &+ \left(w[n] - \lfloor w[n] \rfloor\right) I_{\mathrm{ref}}[\mathbf{x}, \lceil w[n] \rceil] \\ &+ v_s[\mathbf{x}, n], \end{aligned} \tag{2.5}$$

where $w[n]$ is the real-valued "index" in I_{ref} that corresponds to $I[\mathbf{x}, n]$ ($0 \leq w[n] < N$), and $v_s[\mathbf{x}, n]$ is an additive noise term.

We then model the temporal light integration filter as a box filter,

$$h_{\mathrm{int}}[n] = \begin{cases} \frac{1}{N_{\mathrm{int}}} & 0 \leq n < N_{\mathrm{int}} \\ 0 & \text{otherwise} \end{cases}, \tag{2.6}$$

that accounts for the camera's integration time T_{int} and integration window width $N_{\mathrm{int}} =$

round $\left(\frac{T_{\text{int}}}{T_{\text{L}}}K\right)$.

Using discrete matrix-vector notation similar to that used in [29], we can express our image acquisition forward model in Equation (2.3) at pixel location $\mathbf{x} = (x, y, z)$ as

$$\mathbf{g}_{\mathbf{x}} = \mathbf{D}_K \mathbf{H} \mathbf{S} \mathbf{f}_{\mathbf{x}} + \mathbf{v}_{\mathbf{x}}, \tag{2.7}$$

where $\mathbf{g}_{\mathbf{x}} = \left(\mathbf{g}_{\mathbf{x}}^{(1)}, \ldots, \mathbf{g}_{\mathbf{x}}^{(Q)}\right)^{\top}$ is a vector containing the observed samples, $\mathbf{g}_{\mathbf{x}}^{(q)}$ and $\mathbf{f}_{\mathbf{x}}$ are described in Equations (2.2) and (2.1) respectively, $\top$ denotes transposition, $\mathbf{D}_K$ is an $MQ \times NQ$ matrix that downsamples by K, $\mathbf{H}$ is an $NQ \times NQ$ circulant matrix with shifted versions of the filter $h_{\text{int}}[n]$ as its rows, $\mathbf{S}$ is an $NQ \times N$ matrix with elements

$$\mathbf{S}_{ij} = \begin{cases} 1 - (w[i] - \lfloor w[i] \rfloor) & j = \lfloor w[i] \rfloor \\ w[i] - \lfloor w[i] \rfloor & j = \lceil w[i] \rceil \\ 0 & \text{otherwise} \end{cases} \tag{2.8}$$

that replicates and interpolates the high-resolution reference sequence $\mathbf{f}_{\mathbf{x}}$ according to sub-frame positions $w[n]$ (unknown in practice), and $\mathbf{v}_{\mathbf{x}}$ is an additive noise vector that accounts for any measurement noise or intrinsic variability in the underlying process.

2.4.1.2 Sub-resolution Temporal Registration

The first task to invert Equation (2.7) is estimating the best matching sub-frame indices, $w[n]$, for each sample $n = 0, \ldots, MQ - 1$. To do this, we first define a new set of sub-frame indices in the low-resolution signal,

$$w_{\ell}[n] = \frac{1}{K} w[n], \tag{2.9}$$

which we split into a large, whole-sample integer shift $\bar{w}_\ell[n]$ and a smaller sub-sample shift Δ_n,

$$w_\ell[n] = \bar{w}_\ell[n] + \Delta_n, \tag{2.10}$$

where $|\Delta_n| < 1$. We then use a two-step process to separately estimate $\bar{w}_\ell[n]$ and Δ_n that is similar to the approach used in [37], with the exception that we explicitly estimate the sub-sample shift.

We first identify a low-resolution reference cycle $\hat{I}_{\text{ref}}[\mathbf{x}, m], m = 0, \ldots, M-1$ from the long sequence $I_{\text{acq}}[\mathbf{x}, n]$. For simplicity, we consider the first $M = \frac{N}{K}$ time points in $I_{\text{acq}}[\mathbf{x}, n]$ to be our reference cycle $\hat{I}_{\text{ref}}[\mathbf{x}, m] = I^{(1)}_{\text{acq}}[\mathbf{x}, m]$. In practice, the reference cycle can be user-defined by its first and last frame, determined automatically by taking the M first frames (with M an estimate of the cardiac frequency obtained from the peak frequency component in the Fourier transform of the entire signal), or a combination of both (user-adjustment following automatic period estimation). While the sub-frame accuracy of the reference sequence is not necessary for successful super-resolution, the assumption in Equation (2.5), which states that each time point in the long sequence can be obtained by interpolating two frames from the reference cycle, should not be violated.

To determine the whole-sample shift $\bar{w}_\ell[n]$, we find the best match $\hat{I}_{\text{ref}}[\mathbf{x}, \bar{w}_\ell[n]]$ to $I_{\text{acq}}[\mathbf{x}, n]$. Specifically, we use a dynamic programming synchronization algorithm [38] to find $\tilde{w}[n] \in \mathbb{Z}$ for $n = 0, \ldots, MQ-1$, such that

$$\mathcal{Q} = \sum_{n=0}^{N-1} \sum_{\mathbf{x}} \left| I_{\text{acq}}[\mathbf{x}, n] - \hat{I}_{\text{ref}}[\mathbf{x}, (\tilde{w}[n] \bmod M)] \right| \tag{2.11}$$

is minimized with respect to $\tilde{w}[n]$, under the constraint that

$$w_{\min} \leq \tilde{w}[n] - \tilde{w}[n-1] \leq w_{\max}, \tag{2.12}$$

where mod denotes the modulo operator, $M = \frac{N}{K}$, and $w_{\min}$ and $w_{\max}$ are positive integers ($0 < w_{\min} \leq w_{\max}$) that allow for nonuniform temporal warping, restrict $\tilde{w}[n]$ to be monotonically increasing, and limit excessive dilation. In practice, $w_{\min}$ and $w_{\max}$ can be used to restrict the search space, and thereby reduce the computational burden, when bounds on the temporal variability can be estimated. The whole-sample shift is then given by

$$\bar{w}_\ell[n] = \tilde{w}[n] \bmod M. \tag{2.13}$$

To determine the sub-sample shift Δ_n, we approximate the reference signal at spatial location $\mathbf{x}$ and sub-integer index $w_\ell[n] = \bar{w}_\ell[n] + \Delta_n$ with a first order Taylor series and a finite difference approximation to the derivative,

$$\begin{aligned} I_{\text{acq}}[\mathbf{x}, n] \approx \hat{I}_{\text{ref}}[\mathbf{x}, \bar{w}_\ell[n]] + \\ \Delta_n \left(\hat{I}_{\text{ref}}[\mathbf{x}, \bar{w}_\ell[n+1]] - \hat{I}_{\text{ref}}[\mathbf{x}, \bar{w}_\ell[n]] \right). \end{aligned} \tag{2.14}$$

Using a least-squares approach, we take all spatial locations $\mathbf{x}$ into account to find

$$\begin{aligned} \Delta_n = \arg\min_s \sum_{\mathbf{x}} ((I_{\text{acq}}[\mathbf{x}, n] - \hat{I}_{\text{ref}}[\mathbf{x}, \bar{w}_\ell[n]] \Big) \\ - \left(\hat{I}_{\text{ref}}[\mathbf{x}, \bar{w}_\ell[n+1]] - \hat{I}_{\text{ref}}[\mathbf{x}, \bar{w}_\ell[n]] \right) s \Big)^2. \end{aligned} \tag{2.15}$$

This is a reverse interpolation problem using a piecewise linear approximation to the underlying function $\hat{I}_{\text{ref}}[\mathbf{x}, n]$. Instead of using linear interpolation to find the value of the function at a given location, we wish to find the location at which the function takes on a given value, assuming that the function is piecewise linear. In practice, we use a higher order approximation (such as the cubic approximation described in Appendix 2.A.1) rather than a linear approximation to the underlying function.

Once we obtain $\bar{w}_\ell[n]$ and Δ_n, we can determine the best matching sub-integer refer-

ence index

$$w[n] = K\left(\bar{w}_\ell[n] + \Delta_n\right) \tag{2.16}$$

and populate the matrix $\mathbf{S}$ in Equation (2.8).

2.4.1.3 Superresolution Reconstruction

The task of superresolution reconstruction is an inverse problem to recover the unknown $\mathbf{f_x}$ from the measured $\mathbf{g_x}$. We solve this independently at each spatial location $\mathbf{x} = (x, y, z)$ by minimizing an ℓ_1 cost function,

$$\hat{\mathbf{f}}_\mathbf{x}(\lambda) = \arg\min_{\mathbf{f_x}} \|\mathbf{D}_K\mathbf{H}\mathbf{S}\mathbf{f_x} - \mathbf{g_x}\|_1 + \lambda\|\mathbf{\Gamma}\mathbf{f_x}\|_1, \tag{2.17}$$

where λ is a regularization weighting constant, and $\mathbf{\Gamma}$ is a Tikhonov regularization second-order difference operator that is chosen to favor temporally smooth solutions,

$$\mathbf{\Gamma} = \begin{bmatrix} 1 & -1 & 0 & \dots & 0 & 0 \\ -1 & 2 & -1 & \dots & 0 & 0 \\ & & & \ddots & & \\ 0 & 0 & 0 & \dots & -1 & 1 \end{bmatrix}, \tag{2.18}$$

and all other matrices are given by the forward imaging model in Equation (2.7). This is a strictly one-dimensional temporal reconstruction problem, and we solve this independently at each pixel without any spatial constraints.

The minimization in Equation (2.17) is equivalent to the minimization problem:

$$\begin{aligned}\hat{\mathbf{f}}_{\mathbf{x}} = \underset{\mathbf{f}_{\mathbf{x}},\tilde{\mathbf{y}}}{\arg\min} \sum_i \tilde{\mathbf{y}}_i \;\; \text{s.t.} \\ -\tilde{\mathbf{y}} \leq \left[\begin{array}{c} \mathbf{D}_K\mathbf{H}\mathbf{S} \\ \hline \\ \lambda\mathbf{\Gamma} \end{array}\right] \mathbf{f}_{\mathbf{x}} - \left[\begin{array}{c} \mathbf{g}_{\mathbf{x}} \\ \hline \\ \mathbf{0} \end{array}\right] \leq \tilde{\mathbf{y}},\end{aligned} \tag{2.19}$$

where $\tilde{\mathbf{y}}$ is a helper variable vector that bounds the data fidelity error and the regularization error. Since the double-sided inequality in Equation (2.19) can be rewritten as a single-sided inequality,

$$\begin{bmatrix} \mathbf{D}_K\mathbf{H}\mathbf{S} \\ -\mathbf{D}_K\mathbf{H}\mathbf{S} \\ \lambda\mathbf{\Gamma} \\ -\lambda\mathbf{\Gamma} \end{bmatrix} \mathbf{f}_{\mathbf{x}} - \tilde{\mathbf{y}} \leq \begin{bmatrix} \mathbf{g}_{\mathbf{x}} \\ -\mathbf{g}_{\mathbf{x}} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \tag{2.20}$$

we can pose the minimization in Equation (2.17) as a linear programming problem subject to an inequality constraint,

$$\begin{aligned}\hat{\mathbf{y}} = &\min \mathbf{c}^\top \mathbf{y} \\ &\text{s.t. } \mathbf{A}\mathbf{y} \leq \mathbf{b} \\ &\text{and } \mathbf{y} \geq \mathbf{0},\end{aligned} \tag{2.21}$$

with the following matrix definitions:

$$
\begin{aligned}
\mathbf{c}^{\top} &= \left[\begin{array}{c|c} \mathbf{0}_{1\times N} & \mathbf{1}_{1\times(MQ+N)} \end{array}\right], \\
\mathbf{y} &= \left[\begin{array}{c} (\mathbf{f_x})_{N\times 1} \\ \hline \\ \tilde{\mathbf{y}}_{(MQ+N)\times 1} \end{array}\right], \\
\mathbf{A} &= \begin{bmatrix} \mathbf{D}_K\mathbf{H}\mathbf{S} & -\mathbf{I}_{MQ\times MQ} & \mathbf{0}_{MQ\times N} \\ -\mathbf{D}_K\mathbf{H}\mathbf{S} & -\mathbf{I}_{MQ\times MQ} & \mathbf{0}_{MQ\times N} \\ \lambda\mathbf{\Gamma} & \mathbf{0}_{N\times MQ} & -\mathbf{I}_{N\times N} \\ -\lambda\mathbf{\Gamma} & \mathbf{0}_{N\times MQ} & -\mathbf{I}_{N\times N} \end{bmatrix}, \\
\mathbf{b} &= \begin{bmatrix} \mathbf{g_x} \\ -\mathbf{g_x} \\ \mathbf{0}_{N\times 1} \\ \mathbf{0}_{N\times 1} \end{bmatrix}.
\end{aligned}
\tag{2.22}
$$

We find the solution to this linear programming problem by the simplex method [39]. To reconstruct the full temporal superresolution video, we apply the optimization at each pixel location independently. We then produce the complete temporal superresolution sequence,

$$\hat{I}_{\mathrm{SR}}[\mathbf{x}, n] = \hat{I}_{\mathrm{SR}}[x, y, z, n] = \hat{\mathbf{f}}_{\mathbf{x}}(\lambda), \tag{2.23}$$

by assembling the solutions at all pixel locations $\mathbf{x} = (x, y, z)$. Since each pixel location can be treated independently, the computation time can be significantly reduced through parallelization on multi-core processors or multi-node clusters.

2.4.1.4 Optimal Regularization Weighting

If, for each value of λ, the energy of each of the two terms in the minimization (2.17) is computed separately and a locus is recorded in a plot with the two energies in the x and y axes, respectively, one can empirically observe that the resulting curve has the approximate shape of the letter L. The *optimal* choice for λ is the corner of this L-curve [40], as it provides a good balance between the two terms of Equation (2.17). We find this corner by finding the value of $\lambda \in [\lambda_{\min}, \lambda_{\max}]$ that minimizes the two-group linear regression error,

$$\begin{aligned} \lambda_{\text{opt}} = \arg\min_{\lambda^\star} & \sum_{\lambda=\lambda_{\min}}^{\lambda^\star} \Big(L(\lambda) - \hat{L}_1(\lambda) \Big)^2 \\ & + \sum_{\lambda=\lambda^\star}^{\lambda_{\max}} \Big(L(\lambda) - \hat{L}_2(\lambda) \Big)^2, \end{aligned} \tag{2.24}$$

where

$$L(\lambda) = \Big(\|\mathbf{D}_K \mathbf{H} \mathbf{S} \hat{\mathbf{f}}_{\mathbf{x}}(\lambda) - \mathbf{g}_{\mathbf{x}}\|_1, \lambda \|\boldsymbol{\Gamma} \hat{\mathbf{f}}_{\mathbf{x}}(\lambda)\|_1 \Big) \tag{2.25}$$

is a point on the L-curve, and $\hat{L}_1$, $\hat{L}_2$ are least-squares linear regressions to $L(\lambda)$ in the ranges $[\lambda_{\min}, \lambda^\star]$ and $[\lambda^\star, \lambda_{\max}]$, respectively. The motivation behind this approach is that the L-curve's corner location is the point that best separates the L-curve into two straight-line regions.

2.4.2 Active Illumination Imaging System

2.4.2.1 Coded Exposure Acquisition

Our active illumination method has similarities with the coded exposure approach by Agrawal *et al.* [3], in which a camera's shutter is rapidly opened and closed during the integration duration of the sensor. To begin, we briefly summarize the coded exposure

approach in more general terms so as to better relate it to our method and to highlight similarities and differences. At each pixel location $\mathbf{x} = (x, y)$, the process of interest can be modeled as a discrete, N-length, high resolution signal, $I_{\text{ref}}[\mathbf{x}, n], n = 0, \ldots, N-1$. As illustrated in Figure 2.1, when capturing video with a constant exposure camera, each acquired frame can be considered as a summation of K high resolution frames,

$$I_{\text{acq}}[\mathbf{x}, m] = \sum_{k=0}^{K-1} I_{\text{ref}}[\mathbf{x}, n+k], \tag{2.26}$$

where $n = mK$ and $m = 0, \ldots, M-1$ for $M = {}^{N}/_{K}$. In contrast, when capturing video with a coded exposure, due to the open/closed nature of camera exposure, each high resolution frame is multiplied by a binary weight $c[k]$ before summation,

$$I_{\text{acq}}[\mathbf{x}, m] = \sum_{k=0}^{K-1} c[k]\, I_{\text{ref}}[\mathbf{x}, n+k]. \tag{2.27}$$

By capturing the same sequence with Q different, carefully chosen binary codes (different sets of $c_q[k]$ for $q = 1, \ldots, Q$), the system

$$\mathbf{b}_m = \mathbf{A}\mathbf{f}_m, \tag{2.28}$$

where $\mathbf{b}_m$ is a $Q \times 1$ vector containing the m-th samples of the coded sequences $I^q_{\text{acq}}[\mathbf{x}, m]$,

$$\mathbf{b}_m = \begin{bmatrix} I^1_{\text{acq}}[\mathbf{x}, m] \\ \vdots \\ I^Q_{\text{acq}}[\mathbf{x}, m] \end{bmatrix}, \tag{2.29}$$

and where $\mathbf{A}$ is a $Q \times K$ sampling matrix formed by the binary codes $c_q[k]$,

$$\mathbf{A} = \begin{bmatrix} c_1[0] & \dots & c_1[K-1] \\ \vdots & & \vdots \\ c_Q[0] & \dots & c_Q[K-1] \end{bmatrix}, \tag{2.30}$$

is well-posed and can be solved for the $K \times 1$ vector $\mathbf{f}_m$ containing samples of the high resolution temporal sequence,

$$\mathbf{f}_m = \begin{bmatrix} I_{\text{ref}}[\mathbf{x}, mK] \\ I_{\text{ref}}[\mathbf{x}, mK+1] \\ \vdots \\ I_{\text{ref}}[\mathbf{x}, mK+K-1] \end{bmatrix}. \tag{2.31}$$

For $K = Q = 4$ and when $c[k]$ are limited to be binary, Agrawal *et al.* [3] found the optimal coded sampling matrix (with the lowest condition number) for codes with at least 50% light efficiency to be

$$\mathbf{A} = \begin{bmatrix} 1 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix}. \tag{2.32}$$

2.4.2.2 Image Acquisition Model

Rather than coding the camera exposure with a shutter, we instead propose to temporally code the illumination. Temporally coding the illumination has the advantage

that the scene is not illuminated during the "closed exposure" code intervals (i.e. when $c_q[k] = 0$) and thus is not subject to unnecessary illumination. In settings such as fluorescence imaging, where each fluorescent molecule can undergo only a limited number of excitation-emission cycles before becoming inactive, this ensures that no emitted photons are wasted. However, while exposure coding allowed for parallelization of different codes across multiple cameras, only a single code can be applied at a time with illumination coding. As a result, the codes must be applied sequentially to a repeating process.

Following the notation from our forward model in Equation (2.3), the temporal integration filter $h_{\text{int}}^{(q)}$ now takes the form of our binary illumination codes,

$$h_{\text{int}}^{(q)}[n] = \begin{cases} c_q[n] & 0 \leq n < K \\ 0 & \text{otherwise} \end{cases}. \tag{2.33}$$

As with our constant illumination method, we assume that the underlying signal $I[\mathbf{x}, n]$ is temporally repeating, and that we are able to capture Q repetitions of this signal, $I_{\text{acq}}^{(q)}$ for $q = 1, \ldots, Q$. However, unlike in Section 2.4.1.1, where we assumed that all repetitions are captured as part of one long video sequence, here we assume that each repetition is captured separately with a different arbitrary sub-resolution shift $\tilde{s}_q$ applied to I_{ref},

$$I_{\text{acq}}^{q}[\mathbf{x}, n] = \sum_{k=0}^{K-1} c_q[k]\, I_{\text{ref}}[\mathbf{x}, m + k + \tilde{s}_q]. \tag{2.34}$$

In practice, the temporal shift between different coded acquisitions may not correspond to an integer $\tilde{s}_q$ number of high-resolution frames, but instead may be a fractional shift,

$$\tilde{s}_q = K s_q, \tag{2.35}$$

where s_q is the corresponding low-resolution, real-valued shift. In such a case, we assume

that I^q_{acq} can be generated by using linearly interpolated samples of I_{ref}, just as in Equation (2.3). We further assume that sequences can be coarsely pre-aligned to within one temporal frame so that $|s_q| < 1$.

Assuming periodic boundary conditions on I_{ref} and following the matrix-vector notation in [19], we can rewrite our forward imaging model at pixel location $\mathbf{x} = (x, y, z)$ as

$$\mathbf{g_x} = \mathbf{CSf_x}, \tag{2.36}$$

where $\mathbf{g_x}$ contains the observed, coded samples,

$$\mathbf{g_x} = \left(\mathbf{g}^{(1)}_{\mathbf{x}}, \ldots, \mathbf{g}^{(Q)}_{\mathbf{x}}\right)^\top, \tag{2.37}$$

for

$$\mathbf{g}^{(q)}_{\mathbf{x}} = \left(I^{(q)}_{\text{acq}}\left[\mathbf{x}, 0\right], \ldots, I^{(q)}_{\text{acq}}\left[\mathbf{x}, M-1\right]\right)^\top, \tag{2.38}$$

$\mathbf{f_x}$ is the superresolution sequence for which we wish to solve,

$$\mathbf{f_x} = \left(I_{\text{ref}}\left[\mathbf{x}, 0\right], \ldots, I_{\text{ref}}\left[\mathbf{x}, N-1\right]\right)^\top, \tag{2.39}$$

$\mathbf{C}$ is an $MQ \times NQ$ matrix with shifted and zero-padded versions of the codes c_q as its rows, and $\mathbf{S}$ is an $NQ \times N$ matrix that replicates and interpolates the superresolution sequence and is of the form

$$\mathbf{S} = \begin{bmatrix} \mathbf{S}_1 \\ \vdots \\ \mathbf{S}_Q \end{bmatrix} \tag{2.40}$$

where each sub-matrix $\mathbf{S}_q$ $(q = 1, \ldots, Q)$ is a square matrix with elements

$$\mathbf{S}_q(i,j) = \begin{cases} 1 - |s_q| & j = i \\ |s_q| & j = i+1 \text{ if } s_q > 0, \\ & j = i-1 \text{ if } s_q < 0 \\ 0 & \text{otherwise.} \end{cases} \tag{2.41}$$

2.4.2.3 Superresolution Reconstruction

To accurately reconstruct the superresolution sequence, $\mathbf{f_x}$, we must first estimate the unknown sub-resolution shifts s_q for $q = 1, \ldots, Q$. In [19], we solved this as a temporal registration problem for image sequences under constant illumination. However, with the addition of coded illumination, sequences acquired under different illumination codes can no longer be directly registered. Instead, with coded illumination, we consider estimating the sub-resolution shifts s_q along with the superresolution reconstruction itself as an optimization problem.

We first assume that the acquired sequences all begin at approximately the same phase of the repeating process, so that the sub-resolution shifts are in the range $-1 < s_q < 1$. We then identify the optimal sub-resolution shifts by minimizing the objective function,

$$\hat{\mathbf{s}} = \underset{-\mathbf{1}<\mathbf{s}<\mathbf{1}}{\arg\min} \|\mathbf{CS}\hat{\mathbf{f}}_{\mathbf{x}}(\mathbf{s}) - \mathbf{g_x}\|_1 + \alpha\|\mathbf{\Gamma}\hat{\mathbf{f}}_x(\mathbf{s})\|_1, \tag{2.42}$$

where, for a particular set of shifts $\mathbf{s} = (s_1, \ldots, s_Q)^\top$,

$$\hat{\mathbf{f}}_{\mathbf{x}}(\mathbf{s}) = \underset{\mathbf{f_x}}{\arg\min} \|\mathbf{CSf_x} - \mathbf{g_x}\|_1 + \beta\|\mathbf{\Gamma f_x}\|_1, \tag{2.43}$$

and where $\mathbf{1}$ is a $Q \times 1$ vector of all ones, α and β are regularization weighting constants,

Γ is a regularization matrix of the form

$$\mathbf{\Gamma} = \begin{bmatrix} 1 & -1 & 0 & \dots & 0 & 0 \\ -1 & 2 & -1 & \dots & 0 & 0 \\ & & & \ddots & & \\ 0 & 0 & 0 & \dots & -1 & 1 \end{bmatrix}, \tag{2.44}$$

and $\mathbf{S}$ is a linear interpolation matrix defined by the shifts s_q (as described in Equations (2.40) and (2.41)). Once we determine the set of optimal shifts $\hat{\mathbf{s}} = (\hat{s}_1, \dots, \hat{s}_Q)^\top$, we reconstruct the superresolution sequence,

$$I_{\text{ref}}[\mathbf{x}, n] = \hat{\mathbf{f}}_{\mathbf{x}}(\hat{\mathbf{s}}), \tag{2.45}$$

independently at each spatial location $\mathbf{x} = (x, y, z)$ by minimizing the ℓ_1 cost function in Equation (2.43) with the optimal shifts $\hat{\mathbf{s}}$ defining the matrix $\mathbf{S}$.

2.4.2.4 Code Selection

For a given reconstruction factor K and number of acquisitions Q, we propose a set of binary illumination codes for $q = 1, \dots, Q$,

$$c_q[k] = \begin{cases} \mathcal{R}\left(\frac{k-1}{2K}\right) & q = 1, \\ \mathcal{R}\left(\frac{k}{K}\left\lfloor\frac{q}{2}\right\rfloor + \psi\right) & q > 1, \end{cases} \tag{2.46}$$

where $\lfloor\cdot\rfloor$ denotes the floor function, where

$$\psi = \begin{cases} 0 & \text{for } q \text{ even}, \\ \frac{1}{4} & \text{for } q \text{ odd}, \end{cases} \tag{2.47}$$

and where $\mathcal{R}(x)$ is a periodic function with period 1,

$$\mathcal{R}(x) = \begin{cases} 1 & \text{if } 0 \leq x < \frac{1}{2}, \\ 0 & \text{if } \frac{1}{2} \leq x < 1. \end{cases} \tag{2.48}$$

These illumination codes apply a multi-frequency binary temporal modulation to the underlying signal. In the case of $K = Q = 4$, these codes are equivalent to the optimal binary codes proposed in [3]. More importantly, this closed-form expression allows us to generate a set of well-conditioned codes for any choice of K or Q. Having a closed-form code-generating function is valuable, because searching for the optimal code from the space of all possible codes, even when restricted to only binary codes as in [3], becomes prohibitively time-consuming when $K > 4$ or $Q > 4$.

2.5 Experiments

To validate our methods, we conducted both computational simulations and in vivo experiments. During the in vivo experiments, we imaged spatially static (but with a temporally varying intensity) as well as moving samples. We applied our methods to 2D+time microscopy image sequences acquired using an epi-fluorescence microscopy setup (illumination and light collection occur through the same microscope objective). To demonstrate our constant illumination method with 4D imaging, we used a multiview selective plane illumination microscope [41] to image live zebrafish embryos. For both methods, we reconstructed a full resolution, temporal superresolution video by solving the minimization problem in Equation (2.17) and Equation (2.43) using the `cplexlp` solver from the IBM CPLEX optimization package [42]. In our experiments with constant illumination, we assumed loose bounds on the signal's temporal warping with $w_{\min} = 1$ and $w_{\max} = \frac{M}{2}$. We obtained all our ℓ_1 reconstructions through independent temporal processing on a

pixel-by-pixel basis without any spatial post-processing. For each ℓ_1 reconstruction, unless otherwise mentioned, we determined the value of λ by performing the reconstruction on a small, manually-selected region of interest with several different values and selecting the optimal value according to Equation (2.24). All experiments and procedures involving zebrafish embryos were performed following standard techniques [43] under a protocol approved by the Institutional Animal Care and Use Committee at the University of California, Santa Barbara.

2.5.1 Quantifying Resolution Improvement

2.5.1.1 Constant Illumination Temporal Superresolution

To quantify the constant illumination method's temporal resolution improvement experimentally, we imaged a static fluorescent sample (drawn on a glass coverslip with a Sharpie Accent Highlighter, Sanford L.P., Oak Brook, III.) illuminated by a time-varying light source. The illumination (and also the emitted fluorescence) followed a repeating temporal chirp signal, where each cycle is a sinusoid with its frequency increasing linearly from 0 Hz to 30 Hz over two seconds. We imaged the sample with a Leica DMI6000B inverted microscope and an HCX PL S-APO 20×/0.50 air objective. We acquired fourteen low temporal resolution cycles, each with a known uniform temporal shift, at 30 frames per second using a Hamamatsu ImageEM C9100-13 EM-CCD camera. EM-CCD cameras have a much higher sensitivity and gain than regular CCD cameras, and therefore are well adapted for fluorescence microscopy despite the relatively low frame rate at full resolution and full field (512×512 pixels per frame). At 30 frames per second, the second half of the sequence (corresponding to illumination frequencies between 15 Hz to 30 Hz) is aliased because the camera sampling rate is too slow to accurately capture the rapid flickering of the sample. Naive interpolation and resampling cannot recover these aliased

high frequencies. However, by using 14 low resolution cycles (with known uniform shifts) to reconstruct a sequence with a $K = 4\times$ temporal magnification factor, our temporal superresolution method recovers oscillations up to 24 Hz (Figure 2.2, Supplementary movie 1). This is equivalent to a 1.6-fold increase in bandwidth, which is consistent with the theoretically-derived practical superresolution limit in [44].

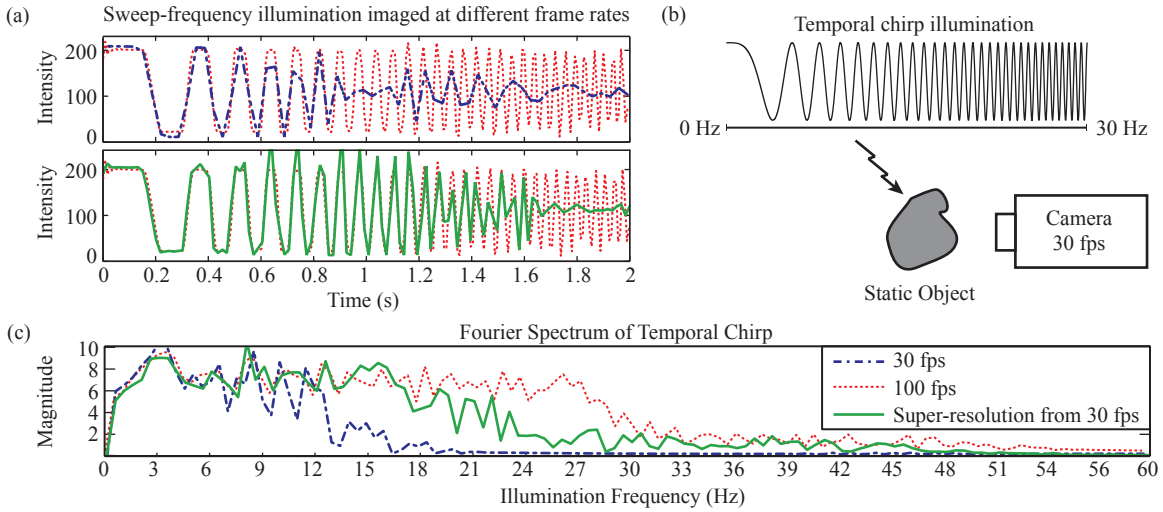

Figure 2.2: We measure a static fluorescent sample's mean intensity over time while under a linear sweep-frequency (chirp) illumination from 0 Hz to 30 Hz, as illustrated in (b). Sampling a linear chirp signal below the Nyquist rate results in temporal aliasing, as shown in (a: top). A non-aliased acquisition taken at 100 fps is shown for comparison. Fourteen sequences acquired at 30 fps are used for temporal superresolution reconstruction. The superresolution sequence, shown in (a: bottom) reveals rapid, high frequency changes in intensity that are normally lost when imaging at 30 fps. A comparison of the signals' frequency spectrum in (c) reveals a superresolution improvement in bandwidth over the 30 fps sequence by a factor of $1.6\times$.

We conducted further simulations to explore the relationship between the number of cycles used for reconstruction and the superresolution performance. In these simulations, we used a one-dimensional temporal chirp signal, with its frequency increasing linearly from 0 Hz to 150 Hz over one second, as the original high-resolution signal. The signal was low-pass filtered and sampled at 150 Hz so that frequencies above 75 Hz were lost. For reconstruction, we first assumed the shifts $w[n]$ to be known (no registration error, $\sigma = 0$)

and uniformly distributed over one frame interval. We measured resolution improvement by comparing the bandwidth of the reconstructed signal (with a 2× reconstruction factor, $K = 2$) to the bandwidth of the observed, low-resolution signal, and we repeated the simulation using up to 60 cycles for the reconstruction. As expected, we found that the resolution improvement increases significantly as more cycles are included, but eventually, including additional cycles provides negligible improvement and is limited to a resolution improvement factor of approximately 1.9× when $K = 2$ (Figure 2.3 and $\sigma = 0$ curve, see further description in Section 2.5.2).

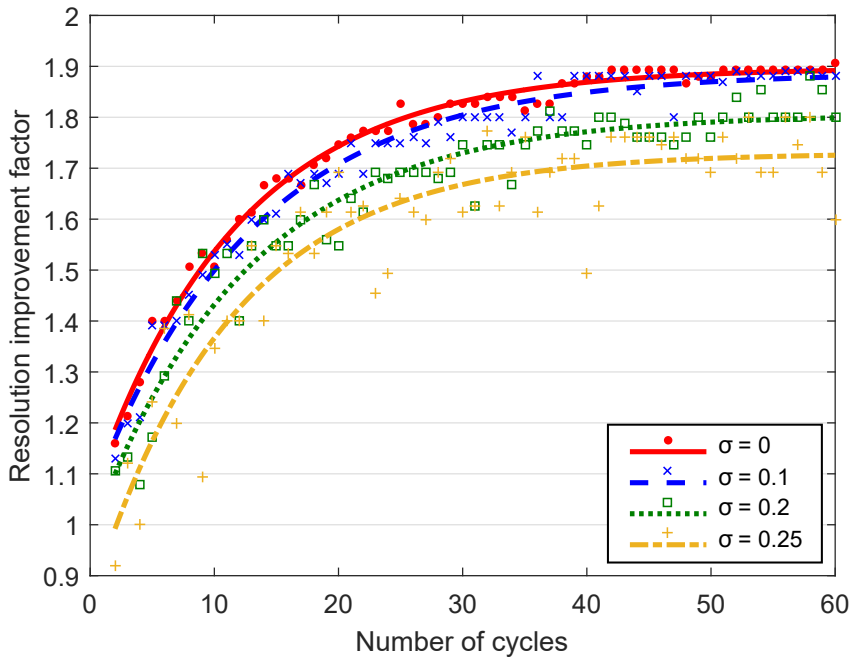

Figure 2.3: Our simulations reconstructing a temporal chirp signal (with a 2× reconstruction factor, $K = 2$) show that, relative to the number of cycles used for the reconstruction, the resolution improvement follows an exponential rise to a maximum limit of approximately 1.9× when the registration is perfect ($\sigma = 0$ frames). In the presence of zero-mean, normally distributed registration error with a standard deviation of σ, the resolution improvement also follows an exponential rise, but to a lower maximum limit.

2.5.1.2 Active Illumination Temporal Superresolution

To quantify our coded illumination method's resolution improvement, we again performed computational simulations using a one-dimensional temporal chirp signal as our high-resolution ground-truth signal, I_{ref}. Following Equation (2.34), we applied our proposed illumination codes c_q to generate $Q = 4$ low-resolution coded sequences $I_{\text{acq}}^{(q)}$. We then applied our superresolution method to reconstruct the high-resolution chirp signal I_{ref} with a reconstruction factor of $K = 4$, $K = 8$, and $K = 12$. We repeated this both when all coded acquisitions are synchronized ($s_q = 0$) and when there is a sub-resolution temporal shift, $|s_q| > 0$, between each acquisition that we recover during the reconstruction process. For the latter case, s_q was drawn from a uniform random distribution between $[-\sigma, \sigma]$ with $\sigma = 1/2$ (corresponding to a shift of up to ± 0.5 low-resolution frames). For each selection of K, Q, and σ, we computed the median condition number and mean resolution gain over 20 trials (Table 2.2). For comparison, we also performed this simulation with random binary illumination codes and with our constant illumination method. For the case of $K = Q = 4$, since our proposed codes are identical to the optimal codes, we also compared our reconstructions with the least-squares reconstruction proposed in [3]. Since the least-squares method in [3] assumes that all acquisitions are synchronized, it does not recover the random shifts s_q, and as a result does not perform as well as our ℓ_1 method.

Additionally, since one of the advantages of our approach is that we can easily generate a set of codes for any choice of the reconstruction factor K and the number of codes/acquisitions Q, we repeated this simulation with $Q = 8$ and $Q = 12$ (with $K = 4$, $K = 8$, and $K = 12$ for each). Selecting a value of $Q > K$ made the problem overdetermined and improved both the condition number of the sampling matrix and the resolution improvement factor. For these values of K and Q, it is computationally intractable to

Table 2.2: Resolution improvement for different code types with different reconstruction factors K and number of acquisitions Q.

		$Q = 4$				$Q = 8$		$Q = 12$	
		$\sigma = 0$		$\sigma = 1/2$		$\sigma = 1/2$		$\sigma = 1/2$	
		Condition Number	Resolution Gain	Condition Number	Resolution Gain	Condition Number	Resolution Gain	Condition Number	Resolution Gain
$K = 4$	Constant Code	10^{16}	2.0	10^{18}	1.7	10^{16}	1.9	10^{15}	2.1
	Random Code	6.8	3.1	54	1.7	6.6	2.5	4.9	3.0
	Proposed Code c_q	2.3	3.7	9.2	2.7	5.3	3.5	4.8	3.7
	Optimal Code[1] [3]	2.3	3.7	2.3	1.7	-	-	-	-
$K = 8$	Constant Code	10^{16}	3.6	10^{16}	1.8	10^{47}	2.1	10^{15}	2.6
	Random Code	3.7	1.9	4.8	1.6	10^2	2.3	15	2.9
	Proposed Code c_q	2.5	4.2	6.4	2.8	10^2	4.7	14	6.3
$K = 12$	Constant Code	10^{14}	3.7	10^{14}	1.9	10^{15}	2.8	10^{18}	2.8
	Random Code	3.8	1.5	4.7	1.8	17	2.7	10^2	2.8
	Proposed Code c_q	2.7	4.2	5.8	2.6	22	4.5	10^2	5.1

compute the optimal code, and even performing a suboptimal greedy search can be time consuming. With our proposed codes c_q, we can generate a set of codes for any choice of K or Q that remains well-conditioned when $\sigma = 0$. When $\sigma > 0$, our sampling matrix $\mathbf{A} = \mathbf{CS}$ is no longer guaranteed to be well-conditioned (due to the random shifts s_q), but the condition number of $\mathbf{A}$ can be improved by choosing $Q > K$.

From these results (tabulated in Table 2.2), we make several observations. First, modulating the illumination with a binary code (even a completely random binary code) produces a sampling matrix that is many orders-of-magnitude better conditioned than that produced by constant illumination, even when acquisitions are unsynchronized ($\sigma > 0$). Second, on average, completely random binary codes do not provide as much resolution improvement as our proposed binary codes. Due to the random nature of these codes, they can produce sampling matrices that are not well-conditioned or are rank-deficient. Third, for a given reconstruction factor K, using additional codes (for additional acquisitions) such that $Q > K$ makes the problem overdetermined, improves the condition number of the sampling matrix, and improves the resolution gain in the reconstruction.

Fourth, for a given number of coded acquisitions Q, increasing the reconstruction factor K beyond Q makes the problem more underdetermined, and, on average, does not improve resolution. Finally, our proposed binary codes consistently outperform both constant illumination as well as binary codes chosen at random, and our proposed method is able to improve resolution by 6.3× when $K = 8$ and $Q = 12$. For comparison, our constant illumination method is only able to improve resolution by 2.8× when $K = Q = 12$. For the case of $Q = 4$, $K = 12$, and $\sigma = 0.5$ (similar conditions to those used in Section 2.5.1.1), we observed a resolution gain of 1.9× with our constant illumination approach and 2.7× with our proposed coded illumination approach. Beyond $K = 8$, the resolution improvement of our method approaches a limit, as registration error has a greater effect on the reconstruction. For example, a registration error of 0.15 low-resolution frames (observed in Figure 2.5), corresponds to only about half of a high-resolution frame when $K = 4$, but corresponds to more than a full high-resolution frame when $K = 8$.

2.5.2 Effect of Registration Error

2.5.2.1 Constant Illumination Temporal Superresolution

The performance of the sub-frame temporal registration step is critical to the superresolution reconstruction quality. Since sub-resolution registration accuracy is difficult to measure in practice (due to a lack of a ground truth in experimental data), we conducted a simulation on synthetic data to quantify the registration error. We generated a reference signal, $\hat{I}_{\text{ref}}[n] = \cos(2\pi f(n\Delta T))$, and a warped test signal, $\hat{I}[n] = \cos(2\pi f(n\Delta T + \varepsilon[n]))$, where $\varepsilon[n]$ is a random shift drawn from a uniform distribution between $-\Delta T/2$ and $\Delta T/2$, f is the signal frequency, and ΔT is the sample spacing. For consistency with our in vivo imaging experiments in which we image the zebrafish heart (beating at approximately 3 beats per second) at 30 frames per second, we chose $f = 3$ Hz

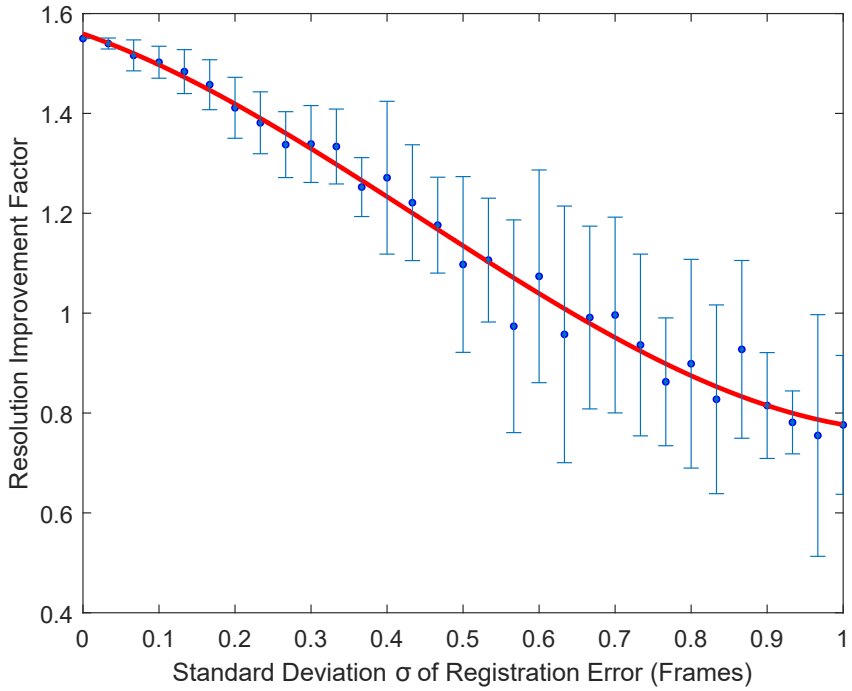

Figure 2.4: We reconstruct a temporal chirp with a 2× reconstruction factor from 10 cycles under different levels of registration error σ. As the standard deviation of the registration error increases, the resolution improvement in the reconstructed signal decreases. Beyond a standard deviation $\sigma = 0.5$ frames, the reconstructed resolution is worse than the original resolution (resolution improvement factor $< 1\times$).

and $\Delta T = 1/30$ s. We registered the test signal $\hat{I}[n]$ to the reference signal $\hat{I}_{\text{ref}}[n]$ and compared the registered time indices to the true time indices to determine the registration error. In this simulation, for a sample size of 1000 time points, we observed that 80% of registered samples are within ± 0.02 frames of the true temporal shift.

We next performed a simulation to characterize the effect of registration error on the resolution improvement in the reconstruction. We generated 10 cycles of a repeating temporal chirp signal with a known sample index $w[n]$ at each sample. We then low-pass filtered and downsampled the signal by a factor of 2, and performed the reconstruction using incorrect sample indices, $w[n] + \varepsilon[n]$, where $\varepsilon[n]$ is a random error drawn from a zero-mean normal distribution with a standard deviation σ. We repeated this with various values of σ, and for each value of σ, we repeated this simulation 8 times. As σ

increases, the resolution improvement decreases, and past $\sigma = 0.5$ frames, the resolution of the reconstruction is worse than the original resolution due to the registration step matching dissimilar samples together (Figure 2.4). We then repeated the simulation in which we varied the number of cycles used for reconstruction, this time imposing a random registration error to each sample, drawn from a normal distribution with standard deviation $\sigma = 0.1$, 0.2, and 0.25. As we increase the number of cycles, the resolution improvement still increases, but the maximum resolution improvement limit decreases as σ increases (Figure 2.3). We can use this result to estimate the registration error in the experimental data used in Figure 2.2. Since we used 14 cycles in our reconstruction in Figure 2.2, and we observed a resolution improvement of approximately 1.6×, we estimate our registration error to have a standard deviation of approximately 0.1 frames.

2.5.2.2 Active Illumination Temporal Superresolution

In our active illumination temporal superresolution method, since we do not assume that our acquisitions are gated with the process to be imaged (though knowing the gating values would allow us to bypass or initialize the registration), each observed cycle of the repeating process may be acquired with a different sub-resolution shift s_q. We assume that the acquisitions can be coarsely registered (e.g. by cross-correlation or, possibly, in some applications, gating) as a pre-processing step, such that $|s_q| < 1$. As a part of our reconstruction algorithm, we estimate s_q through a constrained minimization in Equation (2.42). We evaluated the accuracy of this shift optimization by simulating acquisitions with different shifts and comparing the recovered shift $\hat{s}_q$ to the true shift s_q. In this simulation, we used a step function as the high-resolution signal, and we chose $K = 4$, $Q = 8$, $\alpha = 1$, and $\beta = 0.1$. The shifts s_q were chosen at random from a uniform distribution between $[-1, 1]$. We used Matlab's `fmincon` with the interior-point algorithm [45] to perform the shift optimization, and we recorded the registration

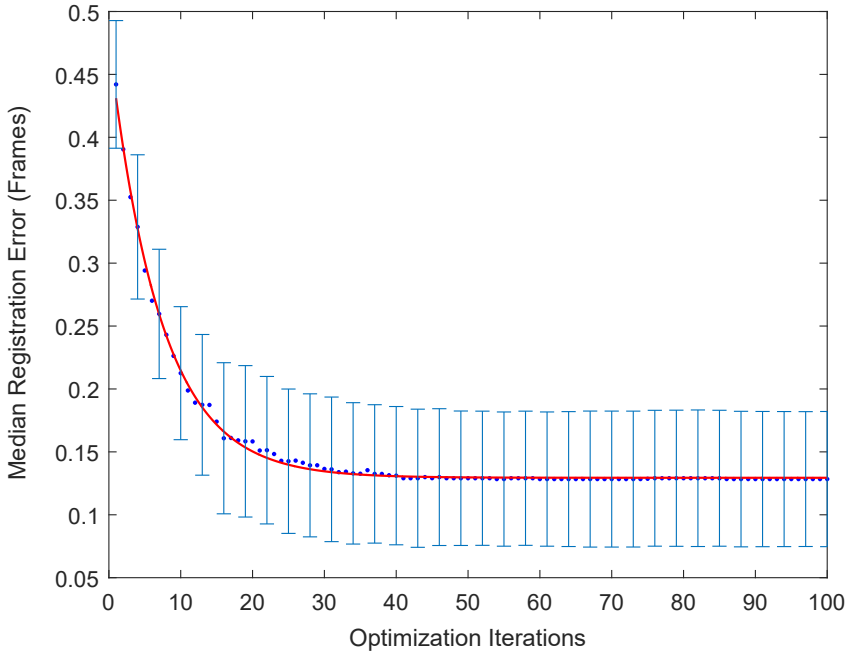

Figure 2.5: Using Matlab's `fmincon` with the interior-point algorithm to minimize Equation (2.42), we observed that, on average, the median absolute registration error converges within 30 iterations to an error of between 0.1 to 0.2 low-resolution frames. We averaged this over 100 trials, each with a random set of shifts s_q drawn from a uniform distribution between $[-1, 1]$.

error at each iteration of the optimization. We repeated this 100 times, each with a different randomly selected set of shifts $\mathbf{s}$. On average, for this choice of parameters, the registration error converged in approximately 30 iterations to an error of between 0.1 to 0.2 frames (Figure 2.5).

To observe how such a registration error affects the reconstruction itself, we performed computational simulations using a one-dimensional temporal chirp signal as our high-resolution ground-truth signal, I_{ref}. Following Equation (2.34), we applied our proposed illumination codes c_q to generate $Q = 4$ low-resolution coded sequences $I_{\text{acq}}^{(q)}$. In this simulation, s_q was split into a known shift ϕ_q and an unknown shift δ_q ($s_q = \phi_q + \delta_q$). The known shift ϕ_q was randomly drawn from a uniform distribution between $[-1/2, 1/2]$, and the unknown shift δ_q was randomly drawn from a uniform distribution between

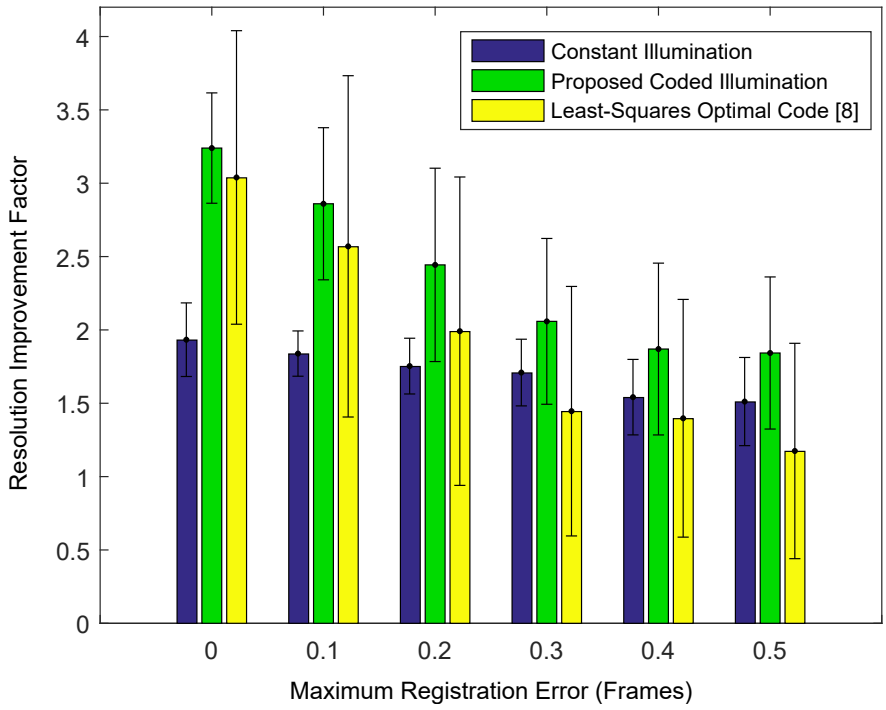

Figure 2.6: Using our superresolution method in Equation (2.43), we simulated reconstructing a temporal chirp signal (with $K = Q = 4$) subject to different amounts of registration error in the low-resolution acquisitions. For each level of registration error, we repeated this simulation 100 times, each time with a different set of random error. For each set of registration errors, we simulated the acquisition and reconstruction under constant illumination [19], our proposed binary codes c_q, as well as the least-squares coded sampling method in [3]. Note that, for $K = Q = 4$, our proposed codes are identical to the optimal codes from [3]. However, our method uses a regularized ℓ_1 based reconstruction while [3] uses a least-squares reconstruction.

$[-\sigma, \sigma]$. We then applied Equation (2.43) to reconstruct the high resolution chirp signal with a reconstruction factor of $K = 4$, and we calculated the resolution improvement of the reconstruction (i.e. the ratio of the recovered bandwidth to the bandwidth of the constant illumination acquisition). We measured the bandwidth of the chirp signal as the number of accurate Fourier coefficients whose magnitude is within $\tau_m = 50\%$ of the ground truth's magnitude, and whose phase is within $\tau_p = 1\%$ of the ground truth's phase. Since we did not apply the shift optimization in Equation (2.42) to estimate s_q, we performed the reconstruction with $\hat{s}_q = \phi_q$ for $q = 1, \ldots, Q$. We repeated this for

$\sigma = 0$, 0.1, 0.2, 0.3, 0.4, and 0.5, and for each value of σ, we repeated the simulation 100 times, each time with a different set of randomly drawn shifts ϕ_q and δ_q. For comparison, we also repeated this with the constant illumination code [19] as well as the least-squares method proposed by Agrawal *et al.* [3]. While our proposed codes are identical to the optimal codes in [3] for $K = Q = 4$, we utilize a regularized ℓ_1 reconstruction instead of a least-squares reconstruction, which performs more consistently and provides greater resolution improvement on average. As expected, a larger registration error results in a worse reconstruction and a lower resolution improvement factor (Figure 2.6). Also, a registration error of between 0.1 to 0.2 frames, which we observed in Figure 2.5, corresponds to approximately a 2.5× resolution improvement for our proposed coded illumination.

2.5.3 Comparing Reconstruction Quality

2.5.3.1 Constant vs. Active Illumination

To demonstrate our method in practice, we assembled a coded illumination imaging system consisting of a Seoul Semiconductor high power LED [46], controlled by an Arduino microcontroller [47] and synchronized to a Thorlabs DCC3240C camera [48]. We programmed the Arduino to control the LED illumination intensity according to our set of codes c_q (with $K = 8$ and $Q = 24$), such that the code length K overlapped with the camera's exposure time, which we set to 16.6 milliseconds for a frame rate of 60 frames per second. Using this imaging setup, we acquired videos of a moving target (playing cards) under both constant illumination and coded illumination. To ensure that the target followed the same motion path in all videos, and to eliminate the need for a preliminary spatial registration step, we used a Rethink Robotics Baxter robot [49] to hold the target and move them in a consistent manner following an arc across the camera's field of view. For both our constant illumination and coded illumination superresolution methods, we

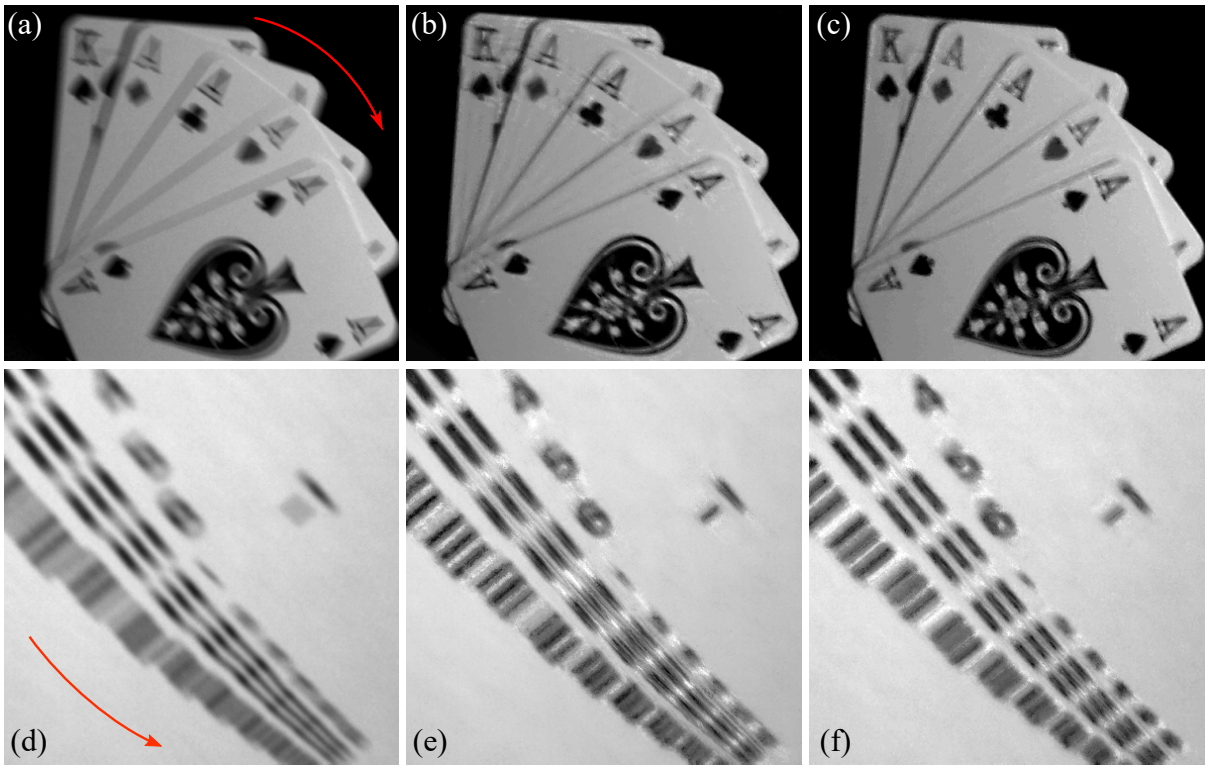

Figure 2.7: We imaged a hand of cards (a,b,c) and a resolution test target (d,e,f) moving in an arc across the camera's field of view. (a,d) An image from the original sequence, captured at 60 frames per second with an exposure time of 16.6 ms, is blurred due to the motion of the cards/target. Arrows denote the direction of motion. (b,e) We used 24 cycles of similar motions captured under constant illumination to reconstruct a temporal superresolution sequence with a reconstruction factor of $K = 8$. While the reconstruction reduces the motion blur, it also results in "ringing" or "ghosting" artifacts around the cards' letters and suits in (b) and around the smaller Ronchi rulings in (e). (c,f) We used 24 cycles captured under our proposed coded illumination with $K = 8$ and $Q = 24$ to reconstruct a temporal superresolution sequence. The coded reconstruction significantly reduces the motion blur without as severe "ringing" artifacts as present in the constant illumination reconstruction.

used 24 videos to reconstruct an 8× temporal superresolution sequence. For our constant illumination superresolution method, we performed the reconstruction with $\lambda = 0.3$ (Figure 2.7(b)). For our active illumination superresolution method, we performed the reconstruction with regularization parameters $\alpha = 0.2$ and $\beta = 0.2$ (Figure 2.7(c)). Both reconstructions visibly reduce the motion blur compared to the original, low-resolution acquisition (Figure 2.7(a)). However, because of its lower condition number and increased numerical stability, our proposed coded illumination reconstruction contains noticeably fewer "ringing" artifacts compared to the constant illumination reconstruction.

We also repeated this experiment with a modified USAF 1951 resolution test target in which the groups of bars were laid out along an arc matching our motion path (Figure 2.7(d,e,f)). In the low-resolution acquisition, the smallest resolvable set of bars (group -3, element 6) corresponded to a resolution limit of 0.22 line pairs per millimeter. In our coded reconstruction, the smallest resolvable set of bars (group -1, element 2) corresponded to a resolution limit of 0.56 line pairs per millimeter, or a 2.5× resolution improvement factor over the low-resolution acquisition.

2.5.3.2 ℓ_1 vs. ℓ_2

An important aspect in our method is the ℓ_1 norm in the minimization function, Equation (2.17) and (2.43). Minimizing an ℓ_1 norm is more computationally expensive than minimizing an ℓ_2 norm (for a least-squares solution). The latter is much easier to implement and faster to compute, but the solution is more affected by outlier data. Outlier data can be caused by the shot noise typical in low-light fluorescence microscopy, irregularity in the repeating signal (such as arrhythmia), or registration error. We compared our proposed ℓ_1 constant illumination reconstruction with the corresponding least-squares reconstruction (i.e. the solution to Equation (2.17), except with ℓ_2 norms replacing both of the ℓ_1 norms) and with the result of temporal interpolation to evaluate the recon-

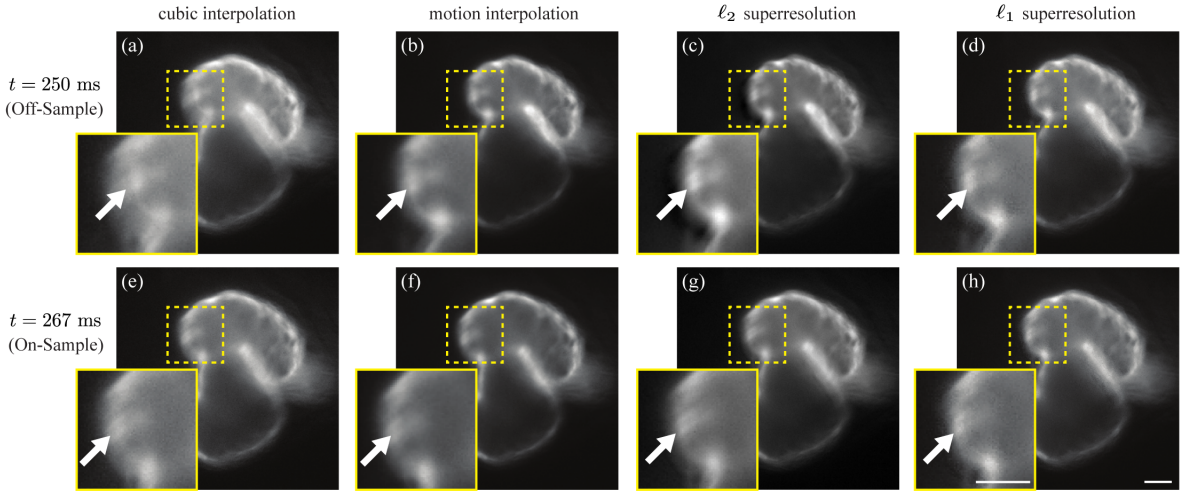

Figure 2.8: The cardiac cycle of a live fluorescent *Gt(tpm4-citrine)*^{ct31a} zebrafish embryo (3 dpf) is imaged at 30 fps with a 20× objective. (a-d) An off-sample frame halfway in between two original video frames (between $t = 233$ms and $t = 267$ms) and (e-h) its neighboring on-sample frame (at $t = 267$ms) are reconstructed by (a,e) cubic interpolation, (b,f) motion interpolation using Adobe After Effects [50], (c,g) minimizing an ℓ_2 version of our constant illumination cost function with $\lambda = 0.1$, and (d,h) minimizing our proposed ℓ_1 cost function (Equation (2.17)) with $\lambda = 0.15$. These choices for λ produce similar background noise levels in both the ℓ_1 and ℓ_2 reconstructions. Arrows denote a bright region in the heart wall that is blurred due to the motion of the heart. A comparison of this region in each reconstruction shows that our proposed ℓ_1 method is best at reducing this motion blur, both for off-sample and on-sample frames. Scalebars are 50 μm.

structed image quality for both on-sample and off-sample time points. Specifically, we used both ℓ_1 and ℓ_2 norms to reconstruct an image sequence of the cardiac cycle for a three day-old (three days post-fertilization) *Gt(tpm4-citrine)*^{ct31a} zebrafish embryo [51] from a low temporal resolution sequence acquired at 30 frames per second [18]. We chose λ for the ℓ_1 reconstruction according to Equation (2.24), and we chose λ for the least-squares reconstruction to provide a similar background noise level as that of the ℓ_1 reconstruction. We compared this to both cubic interpolation and motion interpolation (using Adobe After Effects [50]).

Results show that our proposed method is best able to reduce motion blur (Figure 2.8). In fact, results show that the least-squares approach tends to over-smooth the solution, which can worsen the problem of motion blur. Additionally, while interpolation methods can be used to create higher rate sequences, they do not actually improve the effective temporal resolution, and they cannot reduce temporal aliasing or motion blur. Naive interpolation (using linear or cubic interpolation) can introduce additional spatial blur as a result of blending neighboring frames. Motion interpolation can avoid these frame-blending artifacts when calculating intermediate frames, but does not improve the resolution of frames that fall on integer samples. Our method is able to reconstruct a true higher rate sequence without additional spatial blurring in intermediate frames and with reduced motion blur in both on-sample and off-sample frames (Figure 2.8).

We also repeated this comparison for our coded illumination approach with a simulation in which we used a video of the beating heart in a live Tg(*cmlc2*:EGFP) fluorescent zebrafish embryo captured at 60 frames per second (under conditions similar to those in [19]) as the high-resolution ground truth signal I_{ref}. We simulated coded acquisition with $Q = 11$ codes of length $K = 8$ (for a low-resolution frame rate of 7.5 fps), and we applied each code c_q to a different cardiac cycle. By using different cardiac cycles for each coded acquisition, we preserve any noise or temporal variations of the cardiac cycle in our

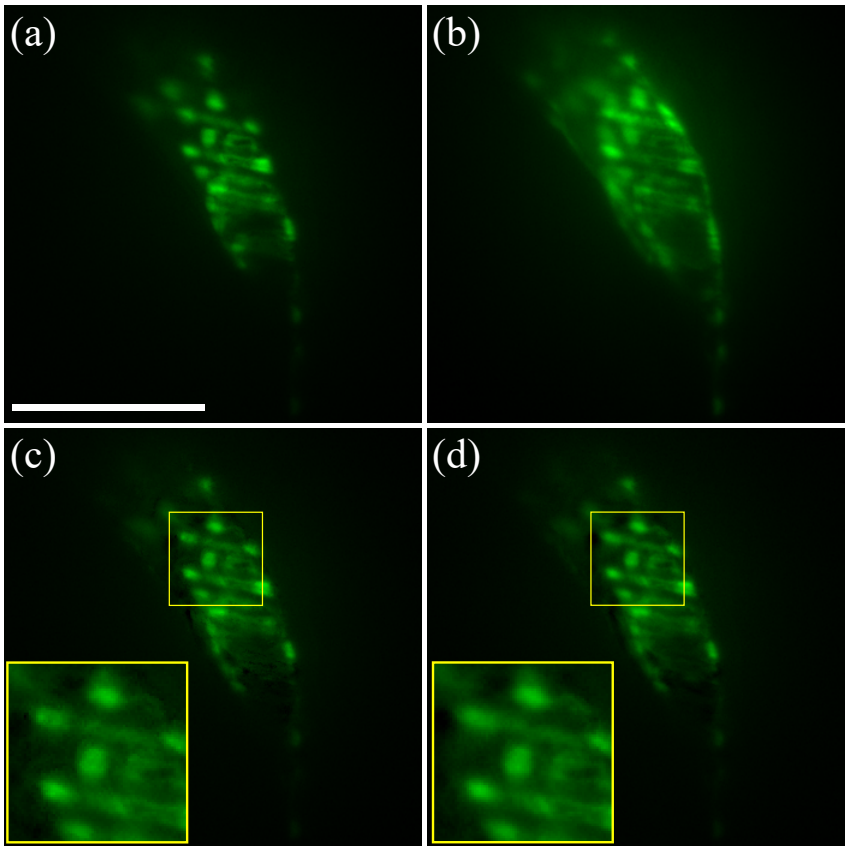

Figure 2.9: We simulated coded illumination for fluorescence microscopy with a video of the beating heart in a zebrafish embryo captured at 60 fps. We generated $Q = 11$ coded acquisitions at 7.5 fps with each code applied to a different cardiac cycle. From these coded acquisitions, we reconstructed a single high resolution cardiac cycle (with $K = 8$) using our proposed method, as well as an ℓ_2 version of our method. (a) A single frame from the original video at 60 fps. (b) A simulated coded image for $q = 11$. (c) A single frame from our ℓ_1-based reconstruction corresponding to the cardiac phase shown in (a). (d) The corresponding frame from the least-squares version of our algorithm. While both the ℓ_1 and ℓ_2-based approaches are able to reconstruct a high-resolution cardiac cycle, the ℓ_1-based reconstruction produces a spatially and temporally sharper result. Scalebar is 100 μm.

low-resolution coded sequences. Using a single frame from each coded cardiac cycle, we used our proposed algorithm to reconstruct $K = 8$ high-resolution frames (Figure 2.9). We also performed the reconstruction using a least-squares version of our algorithm in which the ℓ_1 norms in Equations (2.42) and (2.43) were replaced with ℓ_2 norms. In both cases, we performed the reconstructions with $\alpha = \beta = 0.01$. In both reconstructions, individual cells, which were blurred together in the low-resolution image, are distinctly visible. However, the least-squares reconstruction produces a slightly more spatially and temporally smooth result, which is counter-productive for improving temporal resolution and reducing motion blur.

2.5.4 Application to Fluorescence Microscopy

2.5.4.1 Constant Illumination Temporal Superresolution

We applied our constant illumination temporal superresolution method to reconstruct a 3D time series of a beating heart in a live Tg(*cmlc2:eGFP*) zebrafish embryo (at 2.5 days post-fertilization) expressing green fluorescent protein in the heart [52]. We acquired a dataset consisting of 125 z-slices (with 2 μm spacing between adjacent slices) at 60 frames per second using a custom-built multiview selective plane illumination microscope with two illumination arms and two detection arms [41]. Using this setup, we illuminated the sample through both illumination arms with a 3-μm thick light sheet produced by rapidly scanning a Cobolt MLD 488 nm laser beam through the sample. We acquired images from a single view with a Nikon CFI Apo LWD 25×/1.1 water dipping objective and a Hamamatsu ORCA-Flash 4.0 V2 camera. To reduce file size and speed up processing time, we downsampled the images in the x and y directions to 180 pixels × 180 pixels. Prior to our temporal superresolution reconstruction, we synchronized the z-slices using the method in [38]. We used nine heartbeat cycles to reconstruct a temporal

superresolution sequence with twice the original sampling rate and with regularization parameter $\lambda = 0.5$. In addition to temporal superresolution, our method also reduces noise, as shown in Figure 2.10.

To quantify the noise reduction of our temporal superresolution method, we calculated the peak signal-to-noise ratio (SNR), $\text{PSNR} = 10\log_{10}\left[\max\left(I^2\right)/\sigma^2\right]$, and the contrast-to-noise ratio (CNR), $\text{CNR} = 10\log_{10}\left[\left(\mu_1 - \mu_2\right)/\sigma\right]$, where I is the entire 3D + time image sequence, σ is the standard deviation of a background region of the image sequence, and μ_1 and μ_2 are, respectively, the mean pixel intensities in appropriately chosen heart and background regions of interest. As shown in Table 2.3, our method performs better than the wavelet-based PURE-LET method (which specifically takes into account Poisson-type noise, yet does not take advantage of temporal redundancy) [54] and our previous multicycle method [53] at denoising the experimental cardiac fluorescence image sequences.

Table 2.3: Denoising comparisons for data in Figure 2.10.

	SNR [dB]	CNR [dB]
Raw Synchronized	38.17	12.59
Multicycle Median Denoising [53]	43.64	15.17
PURE-LET Denoising [54]	43.86	15.80
Temporal Superresolution	**44.10**	**15.93**

2.5.4.2 Active Illumination Temporal Superresolution

We applied our active illumination temporal superresolution method to reconstruct a 2D time series of a beating heart in a live Tg(*cmlc2:eGFP*) zebrafish embryo (at 39 hours post-fertilization) expressing green fluorescent protein in the heart [52]. We acquired $Q = 24$ low-resolution coded image sequences at 10 frames per second using a Leica DMI6000B inverted microscope with a HCX PL S-APO 10x/0.30 air objective

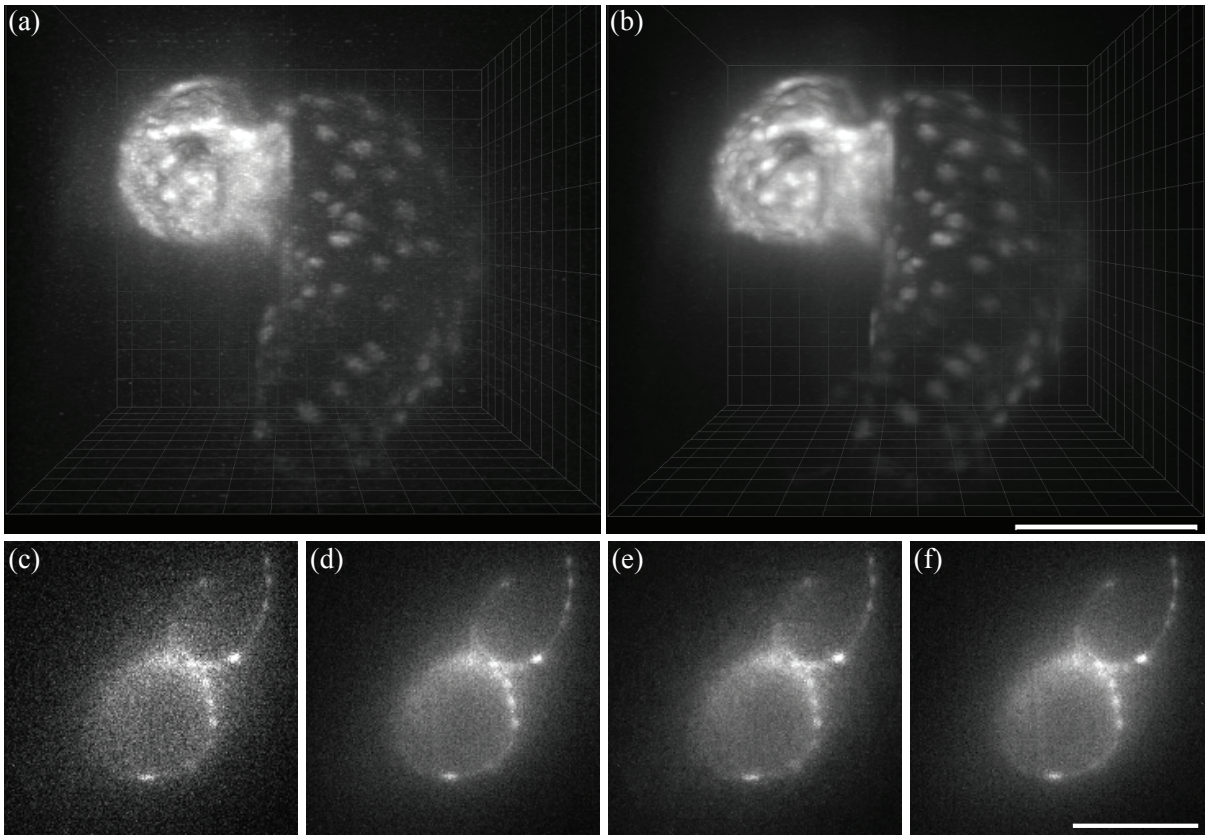

Figure 2.10: We imaged the heart of a live, 2.5 dpf, Tg(*cmlc2:eGFP*) zebrafish embryo in fluorescence at 60 frames per second. (a) 125 z-slices are synchronized to reconstruct a 3D volume. Due to the low illumination intensity and the short integration time used during acquisition, the resulting image is severely corrupted by Poisson–type noise. (b) Our temporal superresolution reconstruction is able to simultaneously temporally superresolve the image sequence and remove much of the noise. We used nine cardiac cycles with a regularization parameter $\lambda = 0.5$ to reconstruct a single denoised heart beat with an effective sampling rate of 120 frames per second. See supplementary movie 5 for the full video showing the cardiac cycle. (c) An image with low signal-to-noise ratio from a single timepoint and z-slice is shown from the original image sequence. (d) The image is denoised using the multicycle denoising method in [53]. (e) The image is denoised using PURE-LET denoising [54]. (f) The corresponding denoised image is shown from our temporal superresolution reconstruction with $\lambda = 0.5$. Quantitative denoising results are tabulated in Table **??**. Scale bar is 100 μm.

and custom-built light source (consisting of a blue high power LED controlled by an Arduino microcontroller [47]) synchronized to a Hamamatsu ImageEM C9100-13 EM-CCD camera. From these 24 coded image sequences of the cardiac cycle, we applied our algorithm in Equation (2.43) with $\beta = 0.21$ to reconstruct an $8\times$ superresolution sequence. We also repeated this with our constant illumination temporal superresolution method, using 24 low-resolution image sequences acquired at 10 frames per second under constant illumination to reconstruct an $8\times$ superresolution sequence following Equation (2.17) with $\lambda = 0.22$. In both cases, β and λ were determined using the L-curve approach in Equation (2.24). In both reconstructions, the motion blur of fast-moving cardiac muscle cells is reduced, but to a greater extent in the coded illumination reconstruction (Figure 2.11).

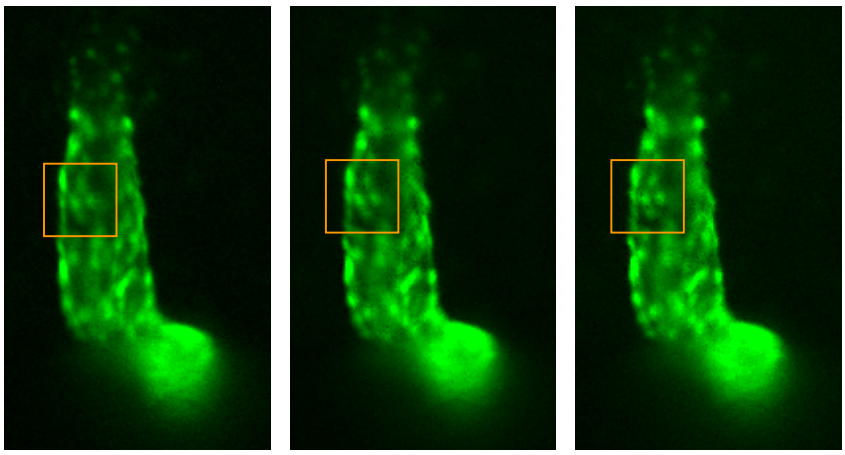

Figure 2.11: A comparison of the superresolution reconstructions under constant illumination (b) and coded illumination (c) with $K = 8$ and $Q = 24$. A corresponding low resolution image under constant illumination (a) is also shown. Motion blur is reduced more in our coded reconstruction than in the reconstruction under constant illumination.

2.6 Conclusion

In summary, we have presented two methods for temporal superresolution of repeating processes. Our first method, designed for constant illumination with any conventional microscope, improves both temporal resolution and signal-to-noise ratio without any additional hardware or hardware modifications. We achieve denoising by combining multiple cycles of the signal through the ℓ_1 minimization in Equation (2.17). Additionally, this minimization also uses these multiple cycles to simultaneously achieve temporal superresolution (by a factor of 1.6 in experimental data), which is not possible with standard denoising algorithms. While our method is limited to repeating processes such as the cardiac cycle, we do not require the process to be strictly periodic, as our method handles nonuniform temporal warping in the repeating process. Our software demonstrating this superresolution algorithm, in the form of a Fiji plugin, is available online for download at [55].

Our second method for temporal superresolution combines a coded illumination approach for video acquisition with an ℓ_1 reconstruction algorithm. In addition, we propose a generalizable set of binary modulation codes for any reconstruction factor K and number of acquisitions Q. With this method, we do not assume that acquisitions are precisely synchronized (e.g. via gating). However, when gating acquisitions such that they are precisely synchronized (so that $s_q = 0$), our proposed set of modulation codes remains well-conditioned even for relatively large values of K and Q (e.g. $K = Q = 12$). Precisely gating acquisitions (which may be possible for some applications) would eliminate the need for our temporal registration step and could also improve the resolution improvement factor in the reconstruction (depending on the gating accuracy).

One of the advantages of our coded illumination approach is its light efficiency. Unlike in coded exposure approaches [3, 4], in which the object is still illuminated during closed-

exposure intervals, we do not illuminate our object unnecessarily. Moreover, compared to temporal superresolution under constant illumination [19], for the same number of acquisitions Q, we illuminate the object with approximately only half the total energy. This is especially advantageous in fluorescence biomicroscopy, where biological samples are often sensitive to harsh (with high peak intensity) illumination and where fluorescent molecules can only be excited a limited number of times before becoming inactive. By using modulated illumination, we can extend the lifetime of the samples' fluorescent molecules, allowing for longer-duration and less-damaging imaging sessions.

We anticipate our methods to be particularly useful for low-light fluorescence imaging of repeating biological processes, such as the beating heart in developing organisms. In such settings, a long exposure time is often necessary to acquire a sufficient number of photons for an acceptable signal-to-noise ratio. Unfortunately, when imaging dynamic processes (e.g. the heart), a long exposure time also results in motion blur due to the motion of the sample. With our approach, we can reconstruct a video with high temporal resolution and reduced motion blur, even in a low-light setting when the individual low-resolution sequences are acquired with a long exposure time.

2.A Appendix

2.A.1 Reverse Cubic Interpolation for Sub-Sample Shift Estimation

Rather than using a piecewise linear approximation to solve for the sub-sample shift Δ_n in Section 2.4.1.2, in practice it is preferable to use a higher order approximation, such as a monotone cubic approximation [56], in which

$$\begin{aligned} I_{\text{acq}}\left[\mathbf{x}, n\right] \approx & \hat{I}_{\text{ref}}\left[\mathbf{x}, \bar{w}\left[n\right]\right] H_1(\Delta_n) + \\ & \hat{I}_{\text{ref}}\left[\mathbf{x}, \bar{w}\left[n+1\right]\right] H_2(\Delta_n) + \\ & D\left[\mathbf{x}, \bar{w}\left[n\right]\right] H_3(\Delta_n) + \\ & D\left[\mathbf{x}, \bar{w}\left[n+1\right]\right] H_4(\Delta_n), \end{aligned} \tag{2.49}$$

where $D\left[\mathbf{x}, \bar{w}[n]\right]$ is the temporal derivative of $\hat{I}_{\text{ref}}$ at $\bar{w}[n]$ (which we calculate using the method in [56]), and $H_k(\Delta_n), k = 1, \ldots, 4$ are the cubic Hermite basis functions evaluated at Δ_n. Assuming that $0 < \Delta_n < 1$,

$$H_1(\Delta_n) = 2\Delta_n^3 - 3\Delta_n^2 + 1, \tag{2.50}$$

$$H_2(\Delta_n) = -2\Delta_n^3 + 3\Delta_n^2, \tag{2.51}$$

$$H_3(\Delta_n) = \Delta_n^3 - 2\Delta_n^2 + \Delta_n, \tag{2.52}$$

$$H_4(\Delta_n) = \Delta_n^3 - \Delta_n^2. \tag{2.53}$$

Given the image sequence $I_{\text{acq}}\left[\mathbf{x}, n\right]$ and the whole-sample, integer shift $\bar{w}[n]$, we

calculate Δ_n with the following minimization:

$$\Delta_n = \arg\min_t \mathcal{C}(t), \tag{2.54}$$

where

$$\begin{aligned}\mathcal{C}(t) = \sum_{\mathbf{x}} \Big(& \hat{I}_{\text{ref}}\left[\mathbf{x}, \bar{w}\left[n\right]\right] H_1(t) \\ & + \hat{I}_{\text{ref}}\left[\mathbf{x}, \bar{w}\left[n+1\right]\right] H_2(t) \\ & + D\left[\mathbf{x}, \bar{w}\left[n\right]\right] H_3(t) \\ & + D\left[\mathbf{x}, \bar{w}\left[n+1\right]\right] H_4(t) \\ & - I_{\text{acq}}\left[\mathbf{x}, n\right]\Big)^2,\end{aligned} \tag{2.55}$$

which is a minimization of a scalar polynomial function and can be solved by finding the roots of $\frac{\mathrm{d}}{\mathrm{d}t}\mathcal{C}(t)$.

Chapter 3

Improving Spatial Resolution in Optical Projection Tomography Using Point-Spread-Function-Aware Filtered Backprojection

Abstract[2]

Traditional optical projection tomography requires the use of low numerical aperture (NA) objectives in order to achieve a large depth of field and approximate parallel projection geometry. However, low NA objectives suffer from poor resolution, resulting in blur in the reconstructed image. In this chapter, we present a modified filtered backprojection method suitable for focal-plane-scanning optical projection tomography (FPS-OPT), where each projection is obtained by scanning through focal planes during collection. We show that FPS-OPT has an exact inversion formula akin to a filtered backprojection, but that incorporates the system's point-spread-function to recover a deblurred 3D volume.

[2]This chapter is based on [57].

With simulations, we demonstrate that FPS-OPT permits the use of high numerical aperture objectives that lead to more accurate images. We further illustrate the technique on experimentally acquired data from a fluorescently-labeled zebrafish larva, which shows that our approach reduces out-of-focus blur.

3.1 Introduction

Optical projection tomography (OPT) is a 3D microscopy technique that has been used to image small transparent animals (up to a few millimeters) such as cleared mouse embryos [58] and zebrafish embryos [59]. Similar to computed tomography (CT), optical projection tomography uses multiple projections through a sample to reconstruct a 3D volume. However, unlike CT, which uses x-rays that travel in approximately straight lines, OPT uses visible light and microscope optics that accept light rays over a range of angles. In order to apply traditional tomographic reconstruction techniques to OPT, researchers typically use low numerical aperture (NA) objectives with OPT to reduce the acceptance angle of the system and achieve approximately straight-line projections. However, low NA objectives have worse lateral resolution than high NA objectives, thus limiting the lateral resolution of OPT systems.

The problem with high NA objectives is that they have a shallow depth of field, making them unsuitable for traditional OPT imaging of thick samples. To adapt OPT for use with high NA objectives, Miao et al. [60] proposed scanning the focal plane through the entire sample to create pseudoprojections with improved resolution compared to traditional (low NA) OPT projections. Additionally, for high-NA OPT imaging with a selective plane illumination microscope, Bassi et al. used high-pass filtering and weighted averaging of multiple focal slices to create an extended depth of field projection [61]. These pseudoprojections and extended depth of field projections are then used

to reconstruct a 3D image using filtered backprojection (FBP) [9].

While the image formation process through convolution with an extended PSF (rather than line integrals) has been acknowledged in OPT and methods have been proposed to mitigate its effect [11, 12, 62], they are ad-hoc methods that focus on fixed focal plane OPT, for which an analytically-derived inversion formula is lacking. Other methods include a well-defined forward model, but require computationally intensive iterative algorithms for reconstruction [63]. In this chapter, we instead consider the acquisition procedure of Miao et al. [60], which we will refer to as *focal-plane-scanning OPT* (FPS-OPT), and we derive an analytic inversion formula that fully incorporates the point-spread-function.

This chapter is organized as follows. In Section 3.2, we present the image formation process in optical projection tomography. In Section 3.3, we present our proposed image acquisition and reconstruction approach. In Section 3.4.1, we characterize our method with a 2D phantom and point-spread-functions of various NAs. In Section 3.4.2, we illustrate our method on data acquired with FPS-OPT of blood vessels in a fluorescent zebrafish embryo.

3.2 Problem Formation

Assuming minimal optical attenuation and scattering, the ideal image formation process in optical projection tomography can be described by the Radon transform of the underlying object. For a 2D object $f(x, y)$, the Radon transform at a particular angle θ is the set of line integrals at a distance s from the origin (Fig. 3.1 (a)),

$$Rf(s, \theta) = \iint_{\mathbb{R}^2} f(x, y)\delta(x\cos\theta + y\sin\theta - s)\,\mathrm{d}x\,\mathrm{d}y, \tag{3.1}$$

where $\delta(\cdot)$ is the Dirac delta function, and $R(s,\theta)$ is the projection of $f(x,y)$ along the direction θ. However, this model is only accurate when (i) the entire sample is contained in the system's depth of field, and (ii) the system's optical point spread function (PSF) is infinitely narrow in the lateral (x-y) direction. In OPT, this model is inaccurate, particularly the second condition. The image formation model in OPT is more accurately described by that of widefield imaging, in which a focal plane is sampled from the result of the convolution between the object and a rotated version of the system's PSF (Fig. 3.1 (b)),

$$\tilde{I}(s,\theta) = (f * T_\theta\{h\})(s\cos\theta, s\sin\theta), \tag{3.2}$$

where $T_\theta\{\cdot\}$ is a transformation operator that rotates a function by an angle θ about the origin, h is the point-spread-function of the system, and s is the distance from the origin (which we assume to lie on the focal plane). In such a case, the standard filtered backprojection (FBP) algorithm can no longer be used to accurately recover $f(x,y)$.

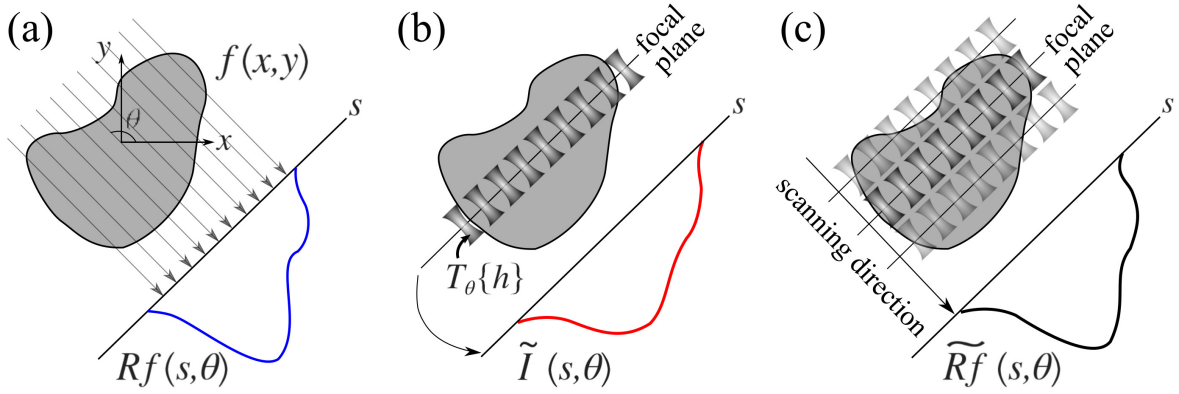

Figure 3.1: (a) In traditional tomography, acquisitions are assumed to be straight line projections through the entire object. (b) In OPT (for the 2D case), acquisitions are better described as the result of the 2D convolution between the object and the PSF sampled on a single focal plane. (c) In FPS-OPT, each acquisition is obtained by summing up images obtained while scanning through all focal planes.

3.3 Proposed Method

3.3.1 Increasing the Depth of Field

To address the first problem when the sample is not completely contained in the system's depth of field, we follow [60] to create a pseudoprojection by scanning the focal plane through the entire sample (FPS-OPT). This can be done either by linearly sweeping the focal plane through the sample during a single camera integration period, or by acquiring a focal stack and retrospectively integrating along the projection direction by digitally averaging the slices. By scanning the focal plane during imaging, we manually perform a Radon transform as in Fig. 3.1 (c),

$$\widetilde{Rf}(s,\theta) = \iint_{\mathbb{R}^2} (f * T_\theta\{h\})(x,y) \cdot \delta(x\cos\theta + y\sin\theta - s)\,\mathrm{d}x\,\mathrm{d}y. \tag{3.3}$$

However, the imaging model still differs from the ideal model in Fig. 3.1 (a), because for each angle θ, the Radon transform is taken of a differently blurred underlying image due to the rotated point-spread-function $T_\theta\{h\}$.

3.3.2 Deblurring with PSF-Aware Filtered Backprojection

Applying standard filtered backprojection to Equation (3.3) will not allow us to recover $f(x,y)$. Rather, it will reconstruct a blurred version of $f(x,y)$ due to the effect of the optical point-spread-function h. However, recall that according to the Fourier slice theorem, the M-dimensional Fourier transform of a projection (X-ray transform) of an N-dimensional function onto M dimensions is equivalent to an M-dimensional slice of that function's N-dimensional Fourier transform [64],

$$\mathcal{F}^M \circ \mathcal{P}_\theta^{N\to M} = \mathcal{S}_\theta^{N\to M} \circ \mathcal{F}^N, \tag{3.4}$$

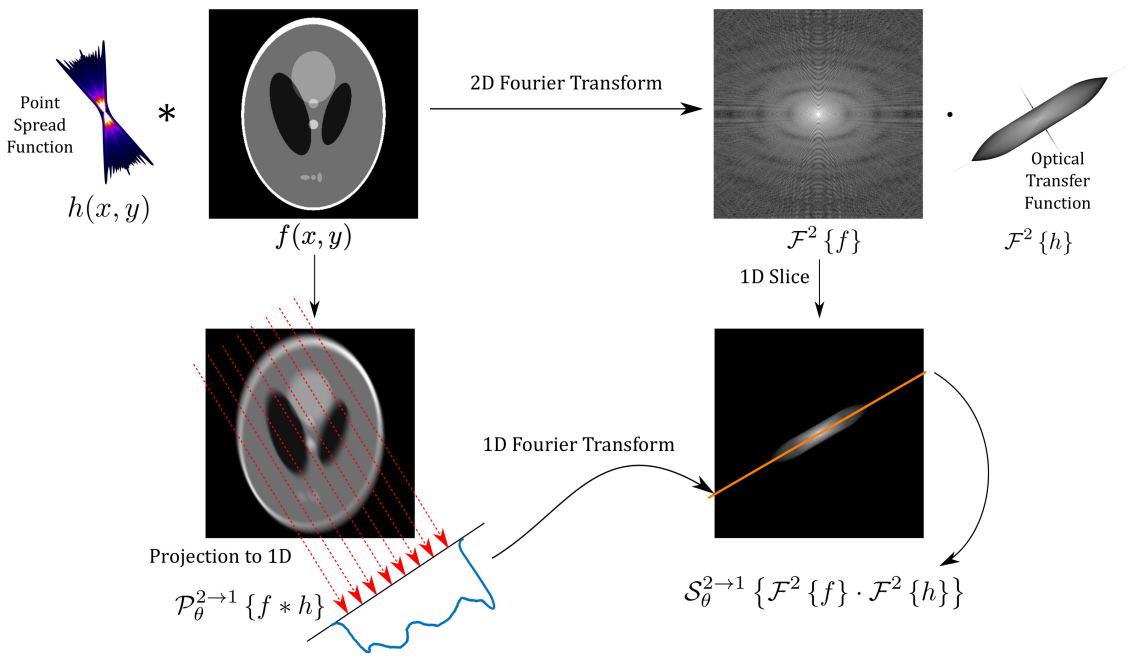

Figure 3.2: The Fourier transform of an object's projection is equivalent to a slice from the underlying object's Fourier transform. The inclusion of the system's point spread function introduces an extra convolution in space, or equivalently, an extra multiplication in the Fourier domain.

where $\mathcal{F}^M$ denotes an M-dimensional Fourier transform, $\mathcal{P}_\theta^{N\to M}$ denotes a projection (X-ray transform) at an angle θ from N dimensions to M dimensions, $\mathcal{S}_\theta^{N\to M}$ denotes an M-dimensional slice normal to the direction of θ from N dimensions, and $\mathcal{F}^N$ denotes an N-dimensional Fourier transform. For consistency with Figure 3.1, the remainder of this section considers the case where $M = 1$ and $N = 2$. However, we note that this also generalizes to higher dimensions, specifically to the case where $M = 2$ and $N = 3$ for 3D OPT.

As the Radon transform in Equation (3.3) is a projection operator $\mathcal{P}_\theta^{2\to1}$, by the Fourier-slice theorem, its Fourier transform is equivalent to a slice from the Fourier

transform of the PSF-blurred, underlying image (as illustrated in Figure 3.2),

$$\mathcal{F}^1\left\{\widetilde{Rf}(s,\theta)\right\} = \mathcal{S}_\theta^{2\to1}\left\{\mathcal{F}^2\left\{f * T_\theta\left\{h\right\}(x,y)\right\}\right\}. \tag{3.5}$$

Since the convolution between two functions is equivalent to the product of their Fourier transforms, and since the point-wise product can be performed after the slicing operator,

$$\mathcal{F}^1\left\{\widetilde{Rf}(s,\theta)\right\} = \mathcal{S}_\theta^{2\to1}\left\{\mathcal{F}^2\left\{f(x,y)\right\}\right\} \cdot \mathcal{S}_\theta^{2\to1}\left\{\mathcal{F}^2\left\{T_\theta\left\{h\right\}(x,y)\right\}\right\}. \tag{3.6}$$

Applying the Fourier-slice theorem once more to the PSF-term in the product allows us to separate the Fourier-slice of the desired, underlying function:

$$\mathcal{S}_\theta^{2\to1}\left\{\mathcal{F}^2\left\{f(x,y)\right\}\right\} = \frac{\mathcal{F}^1\left\{\widetilde{Rf}(s,\theta)\right\} \cdot \mathcal{F}^{1*}\left\{\mathcal{P}_\theta^{2\to1}\left\{T_\theta\left\{h\right\}(x,y)\right\}\right\}}{|\mathcal{F}^1\left\{\mathcal{P}_\theta^{2\to1}\left\{T_\theta\left\{h\right\}(x,y)\right\}\right\}|^2}. \tag{3.7}$$

At this point, we recall that the filtered back projection (FBP) algorithm [9] can be used to invert the Radon transform by expressing the reconstructed image as:

$$f(x,y) \approx \sum_{i=1}^{K} Q_{\theta_i}(x\cos\theta_i + y\sin\theta_i), \tag{3.8}$$

where $Q_{\theta_i}(s)$ is the projection at an angle θ after Fourier domain filtering with a ramp filter $W = |\omega|$,

$$\mathcal{F}^1\left\{Q_\theta(s)\right\} = \mathcal{S}_\theta^{2\to1}\left\{\mathcal{F}^2\left\{f(x,y)\right\}\right\} \cdot W. \tag{3.9}$$

Substituting (3.7) in the above, each filtered backprojection can be expressed in terms of the FPS-OPT-measured blurred projection $\widetilde{Rf}(s,\theta)$. Specifically, we use

$$\mathcal{F}^1\left\{Q_\theta(s)\right\} = \mathcal{F}^1\left\{\widetilde{Rf}(s,\theta)\right\} \cdot \tilde{W}, \tag{3.10}$$

where $\tilde{W}$ is the product of the ramp filter W and a (regularized) version of the expression in Equation (3.7),

$$\tilde{W} = \frac{W \cdot \mathcal{F}^{1*}\left\{\mathcal{P}_{\theta}^{2\to 1}\left\{T_{\theta}\left\{h\right\}(x,y)\right\}\right\}}{|\mathcal{F}^{1}\left\{\mathcal{P}_{\theta}^{2\to 1}\left\{T_{\theta}\left\{h\right\}(x,y)\right\}\right\}|^2 + \lambda|\mathcal{F}^{1}\left\{r(s)\right\}|^2}, \tag{3.11}$$

where λ is the regularization weighting parameter and r is a high-pass regularization filter. This regularization term in the denominator stabilizes the inverse filter and prevents it from growing large when the optical transfer function is close to zero. With this result, we can now apply the FBP in (3.8), to reconstruct the pre-blurred image $f(x, y)$ from the blurred projections $\widetilde{Rf}(s, \theta)$.

3.4 Experiments

3.4.1 Simulations with Shepp-Logan Phantom

To evaluate our method, we simulated an optical projection tomography system in Matlab using the spline-based Radon transform [65] and a Shepp-Logan phantom as a test image. In this simulation, we used 180 uniformly spaced projections over a total of 180° degrees (for an angular spacing of 1°). For each projection angle, we convolved the Shepp-Logan phantom with a rotated point-spread-function prior to computing the Radon transform. We obtained 2D point-spread-functions by extracting the central y-z plane from 3D PSFs generated with the Born & Wolf model [1, 66]. We compared image reconstruction quality in terms of peak signal-to-noise ratio, $\text{PSNR} = 10\log_{10}\left(\max(\tilde{f})^2/\text{MSE}\right)$ where $\tilde{f}$ is the reconstructed image, and MSE is the mean squared error between $\tilde{f}$ and the original ground truth image.

In our simulations, we compared three different reconstruction algorithms: traditional filtered backprojection (FBP), simultaneous algebraic reconstruction technique (SART)

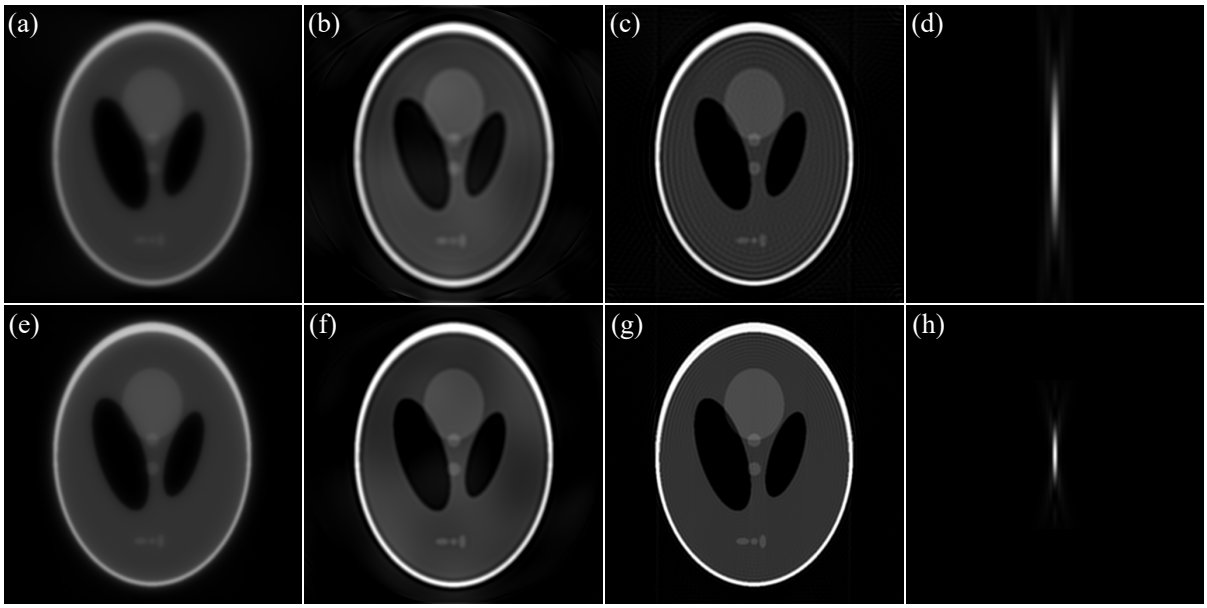

Figure 3.3: We compare traditional filtered backprojection (a,e) with SART (b,f) and with our PSF-aware filtered backprojection (c,g) using point spread functions generated with numerical apertures of 0.3 (a,b,c) and 0.5 (e,f,g). The corresponding PSF for NAs of 0.3 and 0.5 are shown in (d) and (h), respectively.

[10, 9], and our proposed PSF-aware filtered backprojection. In our SART implementation, we incorporated the blur of the point spread function into the forward projection model rather than using a simple straight-ray projection model. We computed all SART reconstructions using 10 iterations. For our PSF-aware filtered backprojection, we performed each simulation with multiple values of λ and selected the λ that produced the reconstruction with the highest PSNR.

3.4.1.1 Constant Magnification

We first considered the scenario where we varied the numerical aperture of the OPT system without changing the sampling resolution (magnification factor). We chose the sampling resolution to be 100 nm per pixel in both directions. It is well known that, for the same magnification factor, a higher numerical aperture system allows for better resolution than a lower numerical aperture system. This is due to the fact that the

lateral width of a point-spread-function, as defined by its full width at half maximum (FWHM), is inversely proportional to the system's NA. Additionally, the axial width of a point-spread-function is inversely proportional to the square of the system's NA. We compared our PSF-aware filtered backprojection to traditional FBP and SART under this scenario for numerical apertures of 0.3 and 0.5 (Figure 3.3). In each case, our method outperformed standard FBP and SART in terms of PSNR (Table 3.1 (a)). Note that as the PSF's lateral width increases with lower NA, the inverse filter also becomes more numerically unstable and requires a greater regularization weighting λ. Additionally, the reconstructed image contains more visible ringing artifacts at a lower NA (Figure 3.3(b)).

3.4.1.2 Variable Magnification

We also considered the scenario where we varied the NA while also varying the sampling resolution to keep the PSF width the same (isolating the effect of PSF shape). Depending on the NA, the sample spacing was adjusted such that the lateral width (FWHM) of the PSF covered the same number of samples. This is similar to imaging with different magnification objectives, since a higher magnification microscope objective will typically also have a higher NA, yet camera sampling would ideally be adjusted to the optical resolution (PSF width). We compared our PSF-aware FBP to traditional FBP under this scenario for numerical apertures of 0.1, 0.7, and 1.3. In each case, the sampling resolution was adjusted to be 1/4 of the lateral width of the PSF. Under these conditions, higher NAs perform worse with traditional FBP since the PSF becomes more hourglass-shaped than low NA PSFs. When instead using our PSF-aware FBP, we were able to reconstruct images with higher PSNRs than both traditional FBP and SART for all three numerical apertures (Table 3.1 (b)).

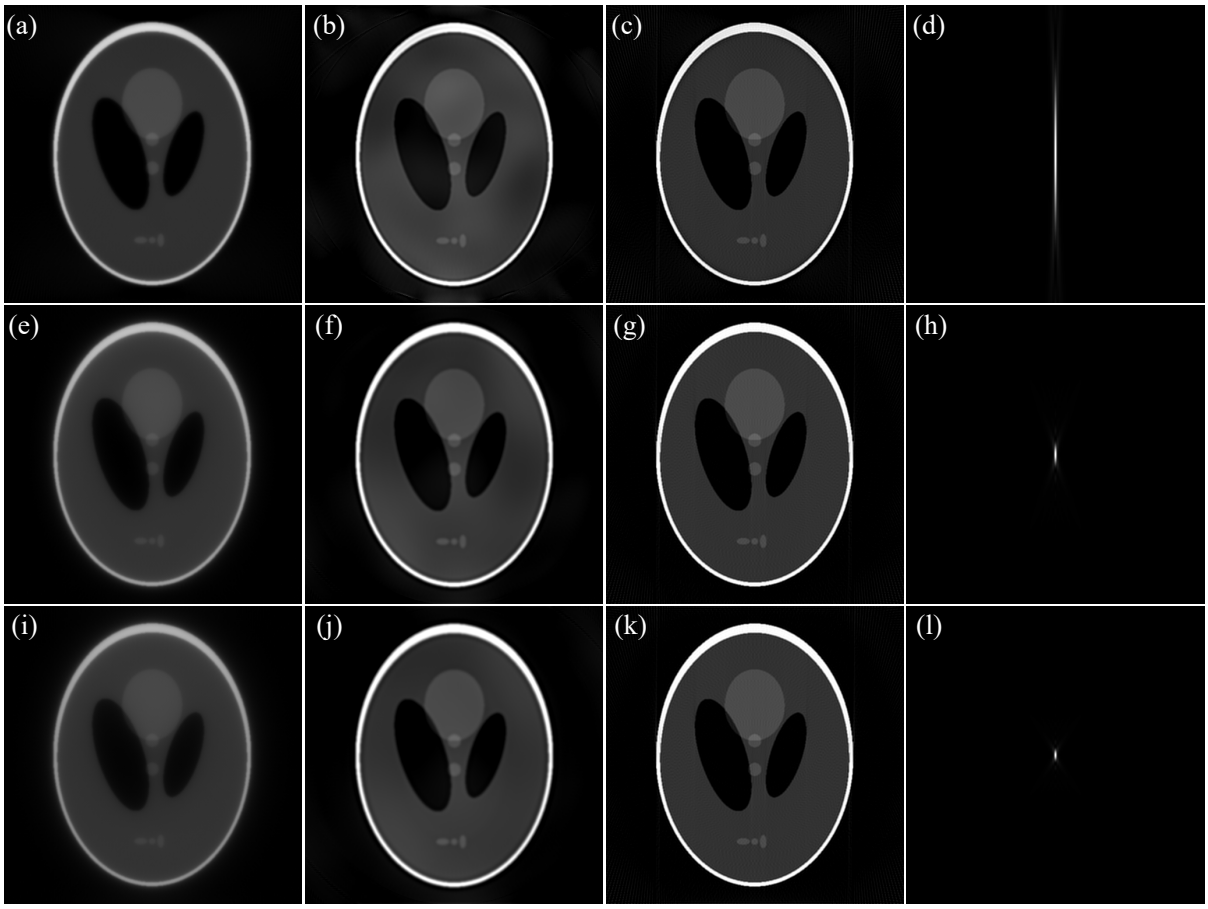

Figure 3.4: We compare traditional filtered backprojection (a,e,i) with SART (b,f,j) and with our PSF-aware filtered backprojection (c,g,k) using point spread functions generated with numerical apertures of 0.1 (a,b,c), 0.7 (e,f,g), and 1.3 (h,i,j). The corresponding PSF for NAs of 0.1, 0.7 and 1.3 are shown in (d), (h), and (l), respectively. For each point spread function, the sampling resolution was adjusted to be $1/4$ of the lateral width of the PSF.

Table 3.1: PSNR comparisons between traditional FBP, SART, and our PSF-aware FBP with different numerical apertures.

		PSNR (dB)		
	NA	FBP	SART	PSF-Aware FBP
(a) Constant Magnification	0.3	18.85	21.05	**23.46**
	0.5	20.12	22.84	**25.23**
(b) Variable Magnification	0.1	22.20	23.23	**25.44**
	0.7	20.56	23.92	**26.08**
	1.3	19.78	23.91	**26.03**

3.4.1.3 Number of Projections

An important factor in tomographic imaging is the number of acquired projections and the angular spacing between them. Using fewer projections allows for faster imaging due to the reduced acquisition (and reconstruction) time, but often results in "star-like" artifacts, especially with filtered backprojection. To evaluate the effect of the number of projections on our reconstruction algorithm, we performed a simulation with the same 1.3 NA settings used in Section 3.4.1.2 (where the sampling resolution was adjusted to be 1/4 of the PSF width). We acquired sets of 30, 60, 90, and 180 uniformly projections over 180°, for angular spacings of 6°, 3°, 2°, and 1°, respectively. For each set of projections, we compared the reconstructions using standard filtered backprojection (FBP), SART, and our PSF-aware filtered backprojection. When there are few projections separated by a large angular spacing (60 projections separated by 3°, and 30 projections separated by 6°), our PSF-aware FBP reconstruction contains severe "star" artifacts (Figure 3.4.1.3 (i,j)). In these cases, iterative reconstruction with SART provides better PSNR than both traditional and our PSF-aware filtered backprojection. However, when the angular spacing between projections is sufficiently small ($\leq 2°$), our PSF-aware FBP outperforms SART (Table 3.2).

Table 3.2: PSNR comparisons between traditional FBP, SART, and our PSF-aware FBP reconstructions using different numbers of acquired projections.

	PSNR (dB)		
# Projections	FBP	SART	PSF-Aware FBP
30	18.96	**19.67**	15.60
60	19.65	**22.10**	20.70
90	19.76	23.13	**23.83**
180	19.78	23.91	**26.03**

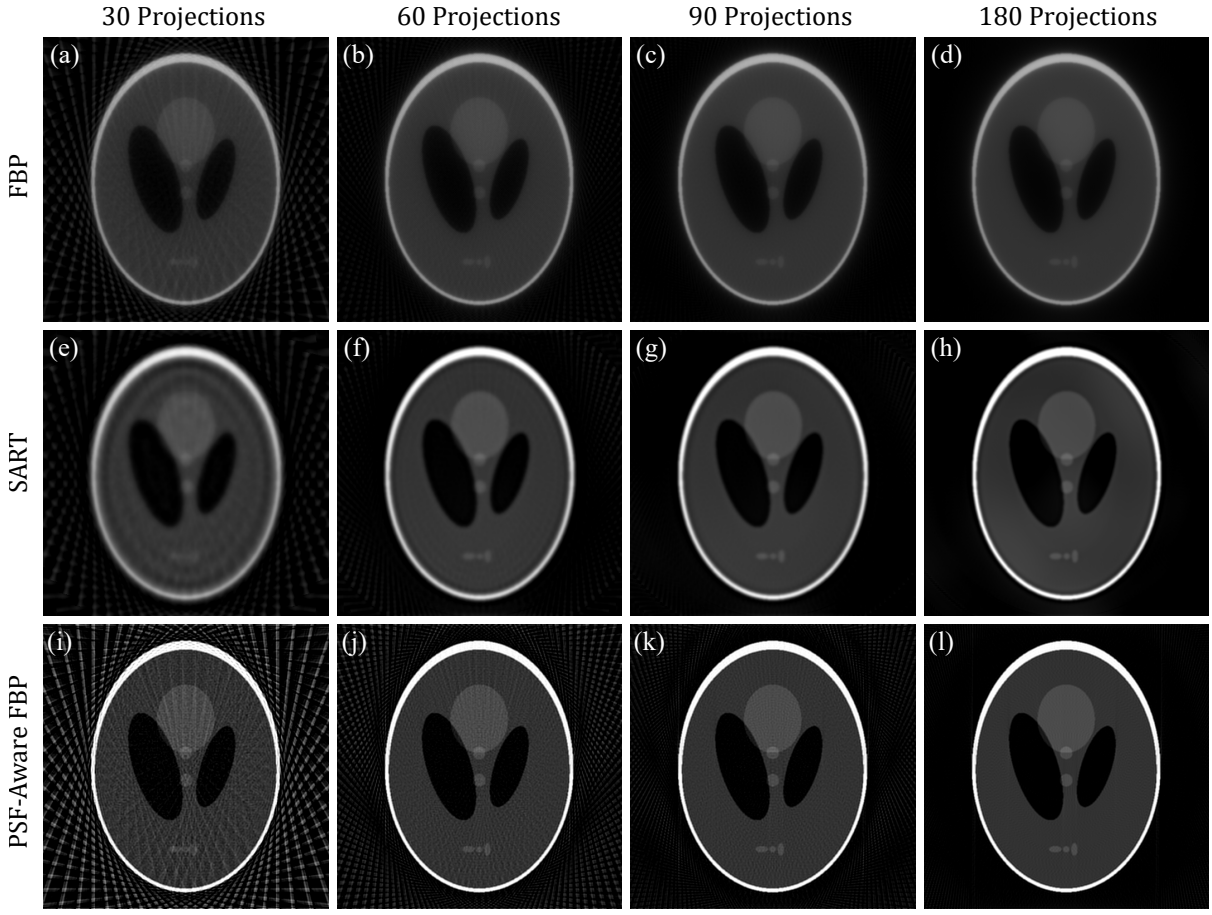

Figure 3.5: We compare traditional FBP (a,b,c,d) with SART (e,f,g,h) and with our PSF-aware FBP (i,j,k,l) to reconstruct an image from 30 projections (a,e,i), 60 projections (b,f,j), 90 projections (c,g,k), and 180 projections (d,h,l). Quantitative results are tabulated in Table 3.2.

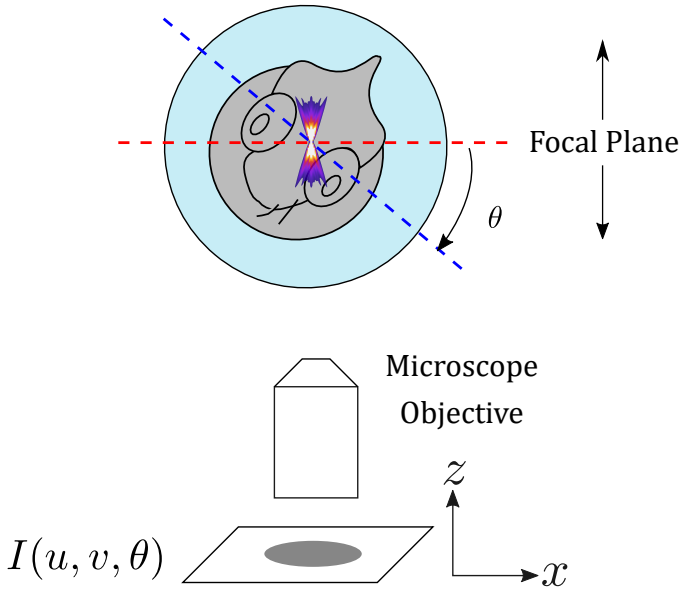

Figure 3.6: Diagram of the OPT rotational acquisition procedure for zebrafish. For each projection angle θ, the focal plane is scanned through the entire zebrafish to create a full projection, even with a shallow depth-of-field objective.

3.4.2 Application to Blood Vessel Imaging

To demonstrate our approach in practice, we imaged the tail of a 64 hpf (hours post-fertilization) Tg(*fli1a*:eGFP) zebrafish that expresses green fluorescence in its blood vessels. We acquired 1600 projections over 360° (with 8 focal slices per projection, over a total depth of 1000 μm) using a custom-built rotational stage consisting of a stepper motor connected to a fluorinated ethylene propylene (FEP) tube that has a refractive index close to that of water. We mounted the larval zebrafish inside the FEP tube in a 1% low melting-point agar solution. We placed this rotational stage on a Leica DMI6000B inverted widefield microscope and imaged in fluorescence with a 10×/0.3 dry objective and a Hamamatsu C9100-13 EM-CCD camera. Figure 3.6 shows a sketch of the image acquisition setup and rotation orientation.

At each angle, we acquired a projection of the zebrafish fluorescence emission using

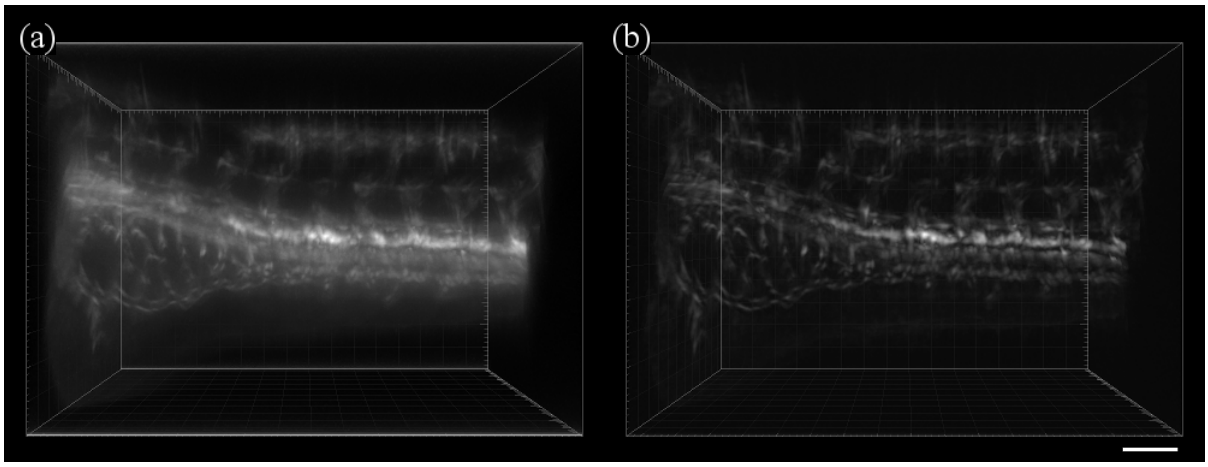

Figure 3.7: We compare (a) standard FBP with (b) our PSF-aware FBP for 3D fluorescence OPT imaging of a Tg(*fli1a*:eGFP) zebrafish tail. Our proposed PSF-aware FBP algorithm reconstructs an image with less out-of-focus blur. Scalebar is 100 μm.

FPS-OPT, and we computed the projection's center of mass to determine the position of its rotational axis [67]. From these, we reconstructed a 3D volume using both traditional FBP and our PSF-aware FBP (Fig. 3.7). Our PSF-aware FBP reconstruction contains significantly less out-of-focus blur compared to the traditional FBP reconstruction.

3.5 Conclusion

In conclusion, we have derived a direct inversion algorithm for focal plane scanning optical projection tomography (FPS-OPT) that can be implemented using a modified filtered backprojection algorithm that incorporates knowledge of the system's point-spread-function. Through simulations with a 2D phantom, we showed that our modified filtered backprojection offers a noticeable improvement in PSNR compared to traditional FBP and SART reconstructions. Additionally, our proposed reconstruction method is computationally inexpensive compared to SART. For our simulations with a 512×512 phantom in Section 3.4.1, our proposed FBP algorithm took approximately 1.8 seconds (the same time as traditional FBP) to reconstruct an image from 180 projections, whereas SART

(with 10 iterations) took, on average, approximately 50 minutes. Finally, we applied our method to image the tail of a larval zebrafish, where we demonstrate that our PSF-aware FBP algorithm is able to reconstruct an image with significantly reduced out-of-focus blur compared to a traditional FBP reconstruction.

Chapter 4

Improving Out-of-Plane Velocity Resolution in Cardiac Flow Velocimetry Using Multi-View Imaging

Abstract[3]

Conventional fluid flow estimation methods for in vivo optical microscopy are limited to two-dimensions and are only able to estimate the components of flow parallel to the imaging plane. This limits the study of flow in more intricate biological structures, such as the embryonic zebrafish heart, where flow is three-dimensional. To measure three-dimensional blood flow, we propose an algorithm to reconstruct a 3D, divergence-free flow map from multiple 2D flow estimates computed from image stacks captured from different views. This allows us to estimate the out-of-plane velocity component that is normally lost with single-view imaging. This chapter describes our 3D flow reconstruction

[3]This chapter is based on [68]

algorithm, evaluates its performance on a simulated velocity field, and demonstrates its application to in vivo cardiac imaging within a live zebrafish larva.

4.1 Introduction

Blood flow in the embryonic heart plays a critical role to ensure normal development, and perturbations to the normal flow can lead to severe heart defects [13]. Measuring these flows in 3D has remained a challenge due to the limited acquisition speed of conventional 3D imaging modalities and the rapid motion of blood cells in the heart. New microscope designs using electrically tunable lenses have begun to address this issue and have demonstrated acquision rates of up to 30 volumes per second [69]. However, such designs require a tradeoff between temporal and axial sampling, so increasing the number of volumes per second is only possible by taking volumes with fewer z-slices. In optical microscopy, blood flow estimation is typically performed on 2D video sequences acquired at very high frame rates (400-1000 frames per second for zebrafish) [13, 14]. However, this 2D approach is only able to measure the components of velocity parallel to the acquisition plane, and any out-of-plane motion is lost, as illustrated in Fig. 4.1(a-c).

A number of methods have recently been proposed to recover three-dimensional flow with MRI [15, 16] and ultrasound imaging [17]. Unfortunately, these imaging modalities often lack the spatio-temporal resolution necessary for imaging blood flow in small organisms and developing embryos. In comparison, optical microscopy offers high resolution and allows use of fluorescent probes to image specific biological structures and functions of interest. Several 3D particle image velocimetry (PIV) methods have also been proposed, including holographic PIV [70] and defocusing PIV [71]. However, these methods have yet to be demonstrated with in vivo microscopy to measure 3D blood flow.

In this chapter, we present a method to reconstruct 3D, divergence-free flow fields

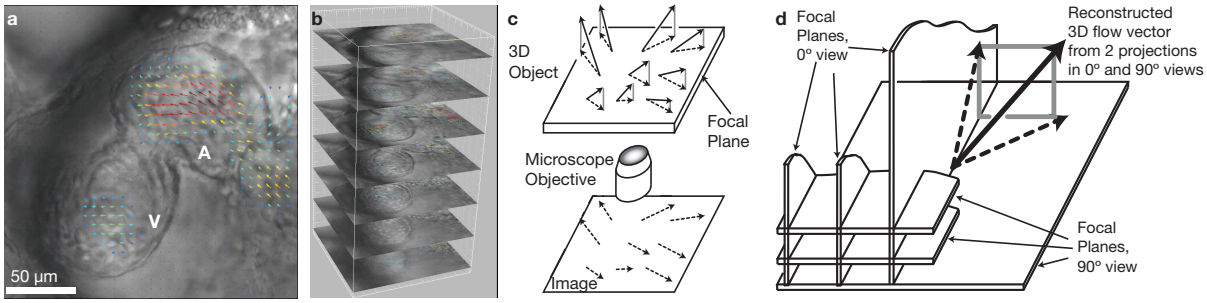

Figure 4.1: (a) Still frame of a high-speed video of the beating heart in a larval zebrafish (A: atrium, V: ventricle). Blood flow was estimated using an optical flow technique. (b) 2D flow can be estimated at multiple depths in the sample by adjusting the focal plane. (c) Flow estimation algorithms can only recover the in-plane component (projection, dashed arrows) of the true velocity vector (black arrows). The out-of-plane component (gray line connecting the in-plane and true velocity vectors) is inaccessible. (d) The out-of-plane component can be measured by imaging the object from a different view (e.g. 90°).

from multiple 2D projections acquired from different rotated views (e.g. 0° and 90° as shown in Fig. 4.1(d)). At a high level, our method is similar to the multi-view, divergence-free method used by Liu et al. to measure 3D motion of muscle tissue using MRI [72]. Our method also has similarities with other multi-view flow reconstruction methods [73, 74] that formulate the reconstruction problem as a constrained and regularized inverse problem. Unlike these methods, however, we do not use explicit constraints or separate regularization terms, but rather we directly reconstruct flow fields using radial basis functions which guarantee our reconstructed flow to be divergence-free, a common assumption for flow estimation. This also allows us to represent the 3D flow field with relatively few coefficients, making the method computationally tractable.

This chapter is organized as follows. In Section 4.2, we describe the image formation model and the problem with single-view imaging for measuring 3D flow. In Section 4.3 we present the multi-view 3D flow reconstruction method as a quadratic minimization problem. In Section 4.4.1, we evaluate our method using a simulated flow field. In Section 4.4.2, we discuss the experimental acquisition procedure and demonstrate our approach

with in vivo microscopy to produce volumetric maps of 3D blood flow in the beating heart of a developing zebrafish larva.

4.2 Problem Formation

Given a three-dimensional vector field $\mathbf{v}(\mathbf{x}, t)$, at every position $\mathbf{x} \in \mathbb{R}^3$, we consider an orthogonal projection of this vector field onto a plane:

$$\mathbf{v}_{\mathbf{n}_k}(\mathbf{x}, t) = \mathbf{P}_{\mathbf{n}_k} \{\mathbf{v}(\mathbf{x}, t)\} \tag{4.1}$$

$$= \mathbf{v}(\mathbf{x}, t) - \langle \mathbf{v}(\mathbf{x}, t), \mathbf{n}_k \rangle \, \mathbf{n}_k, \tag{4.2}$$

where $\mathbf{P}_{\mathbf{n}_k}\{\cdot\}$ is an operator that projects a vector onto the plane with unit normal vector $\mathbf{n}_k$, and $\langle \cdot, \cdot \rangle$ is an inner product between two vectors. This situation reflects conventional imaging and estimation of 2D flow, where out-of-plane flow velocities (i.e. $\langle \mathbf{v}(\mathbf{x}, t), \mathbf{n}_k \rangle \, \mathbf{n}_k$) are lost.

In optical microscopy, when imaging thin, flat samples (that can be approximated as two-dimensional), such as cells under a coverslip, the velocity vector field can be measured by acquiring a 2D + time image sequence and using, for example, an optical flow [75] or other motion estimation algorithm. However, when imaging thick samples where motion is three-dimensional, traditional 2D imaging restricts our out-of-plane velocity resolution and can only provide information about motion in the focal plane.

4.3 Proposed Method

To recover a three-dimensional velocity field $\mathbf{v}(\mathbf{x}, t)$, we consider acquiring K two-dimensional projections $\mathbf{v}_{\mathbf{n}_k}(\mathbf{x}, t)$, for $k = 1, \ldots, K$, each on a different rotated plane.

From these K projections, we recover the 3D vector field $\mathbf{v}(\mathbf{x}, t)$ with the following minimization:

$$\hat{\mathbf{v}}(\mathbf{x}, t) = \underset{\tilde{\mathbf{v}}(\mathbf{x},t)}{\arg\min} \sum_{k=1}^{K} \left(\mathbf{P}_{\mathbf{n}_k} \left\{\tilde{\mathbf{v}}(\mathbf{x}, t)\right\} - \mathbf{v}_{\mathbf{n}_k}(\mathbf{x}, t)\right)^2. \tag{4.3}$$

This equation ensures data consistency, i.e. the least-squares error is minimized when the projected estimate matches the measured field. However, we also wish to enforce fluid incompressibility in our flow reconstruction. Using a divergence-free interpolation method based on radial basis functions [76], we require that $\hat{\mathbf{v}}(\mathbf{x}, t)$ satisfy

$$\hat{\mathbf{v}}\left(\mathbf{x}, t\right) = \sum_{j=1}^{M} \mathbf{\Phi}\left(\mathbf{x} - \mathbf{m}_j\right) \mathbf{c}_j\left(t\right), \tag{4.4}$$

where $\mathbf{c}_j$ $(j = 1, \ldots, M)$ are vectorial radial basis coefficients, $\mathbf{m}_j$ are their corresponding node locations, and $\mathbf{\Phi}$ is a matrix-valued radial basis function given by

$$\mathbf{\Phi}\left(\mathbf{r}\right) = \left[\left(1 - \frac{\|\mathbf{r}\|^2}{2\alpha^2}\right)\mathbf{I} + \frac{1}{2\alpha^2}\mathbf{r}\mathbf{r}^\top\right] e^{-\frac{\|\mathbf{r}\|^2}{2\alpha^2}}, \tag{4.5}$$

where $\mathbf{I}$ is the identity matrix, and α is a real, positive-valued parameter that controls the smoothness of the vector field.

Combining Equations (4.3) and (4.4), we obtain the following modified minimization:

$$\hat{\mathbf{c}}_j\left(t\right) = \underset{\mathbf{c}_j(t)}{\arg\min} \sum_{k=1}^{K} \left(\mathbf{P}_{\mathbf{n}_k} \left\{\sum_{j=1}^{M} \mathbf{\Phi}(\mathbf{x} - \mathbf{m}_j)\mathbf{c}_j(t)\right\} - \mathbf{v}_{\mathbf{n}_k}(\mathbf{x}, t)\right)^2, \tag{4.6}$$

which can be solved using the conjugate gradient method. This minimization equation solves for the vectorial radial basis coefficients that produce the 3D vector field best matching our observed vector projections. After solving for the divergence-free coefficients $\hat{\mathbf{c}}_j(t)$, the 3D vector field can be interpolated using Equation (4.4).

4.4 Experiments

4.4.1 Simulation

We evaluated our method using the following divergence-free vector field:

$$\mathbf{v}(\mathbf{x}) = \begin{bmatrix} \gamma_1 yz \\ \gamma_1 xz \\ \gamma_2 \cos\left(\gamma_3 \left(x+y\right)\right) \end{bmatrix}, \tag{4.7}$$

sampled on a 10×10×10 grid from -1 mm to 1 mm in each direction (Fig. 4.2(a)), and where γ_1, γ_2, γ_3 are constants such that $\mathbf{v}(\mathbf{x})$ is in units of mm/s. Specifically, $\gamma_1 = 1\ \text{mm}^{-1} \cdot \text{s}^{-1}$, $\gamma_2 = 1\ \text{mm} \cdot \text{s}^{-1}$, and $\gamma_3 = 1\ \text{mm}^{-1}$. We computed two $10 \times 10 \times 10$ focal stacks from two different views of this vector field: one where vectors were projected onto slices parallel to the xy-plane (Fig. 4.2(b)) and one where vectors were projected onto slices with normal vector $\mathbf{n} = \left[-\sin\left(45°\right), 0, \cos\left(45°\right)\right]^\top$, corresponding to a 45° rotation of the xy-plane about the y-axis (Fig. 4.2(c)). In each view, we added zero-mean white Gaussian noise to both the magnitude and phase of the projected vectors. The standard deviation of the magnitude and phase noise are, respectively, $\sigma_{\text{mag}} = 0.05$ mm/s and $\sigma_\phi = 45°$. From these two views, we were able to recover the 3D vector field (shown in Fig. 4.2(d)) with a total mean squared error (for all three vector components) of 0.13 mm/s. All vector fields were visualized using ParaView 4.2.0 [77].

Additionally, we explored how the angle between different views affected the reconstruction accuracy. Using again the vector field given in Eq. 4.7 and taking two views separated by an angle θ, we applied our method to reconstruct a 3D vector field. We repeated this with different values of θ and found mean squared errors of 0.17 mm/s, 0.14 mm/s, and 0.06 mm/s for $\theta = 15°$, 30°, 90°, respectively. Unsurprisingly, as θ ap-

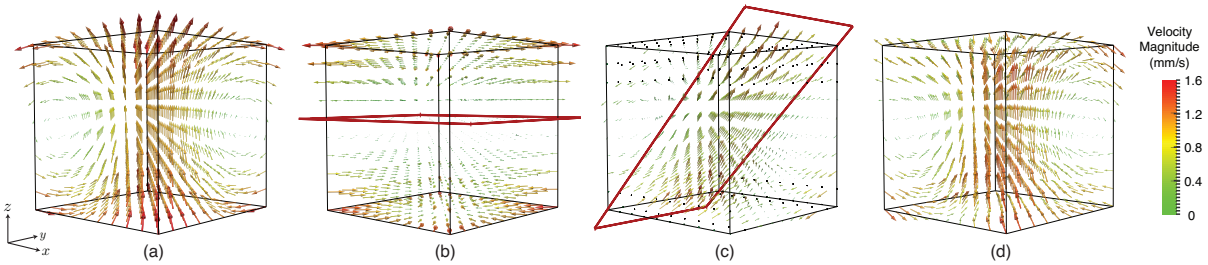

Figure 4.2: (a) We simulated a $10 \times 10 \times 10$ 3D divergence-free velocity field. (b,c) We compute $10 \times 10 \times 10$ focal stacks from two views of this velocity field, rotated by 45°. In each view, the 3D vectors on each focal slice are projected onto that 2D focal plane. In each stack, the center slice is outlined in red. (d) Our method is able to combine these two views to reconstruct the original 3D velocity field with a mean squared error of 0.13 mm/s.

proaches 90°, the mean squared error decreases. This suggests that, in a 3D flow imaging experiment, it is ideal to acquire data from orthogonal views. When this is not possible (due to a lack of optical access), it is advisable to take views with the largest separation angle possible.

4.4.2 In-Vivo Application

To demonstrate our method with in vivo microscopy, we immersed a 60 hpf (hours post fertilization) zebrafish larva in a 1.2% low melting point agarose, 0.016% tricaine (MS-222) solution and placed it inside a tube made of fluorinated ethylene propylene (FEP) which has a refractive index close to that of water. We then placed this tube on the stage of a Leica DMI6000B inverted microscope equipped with an HCX PL S-APO 20×/0.50 dry objective. The tube was oriented perpendicular to the optical axis, and its axis of rotation was aligned parallel to the y-axis of the focal plane. Using a stepper motor, we rotated the tube to image the zebrafish heart from three different views, each rotated by an additional 18°. At each view, we acquired movies with 512×512 pixels at 500 frames per second, covering 3 heartbeats, at 10 different z-slices with 15 μm between each slice. After acquisition, we computationally synchronized the movies and extracted

a channel containing only blood cells using the algorithm described in [14]. We then estimated flow velocity vectors (parallel to the imaging plane) at each z-slice using the Lucas-Kanade optical flow algorithm [75], as implemented in FlowJ [78]. Fig. 4.3 shows an example of 2D velocities estimated from three different views (for visibility, only a single plane is shown for each view).

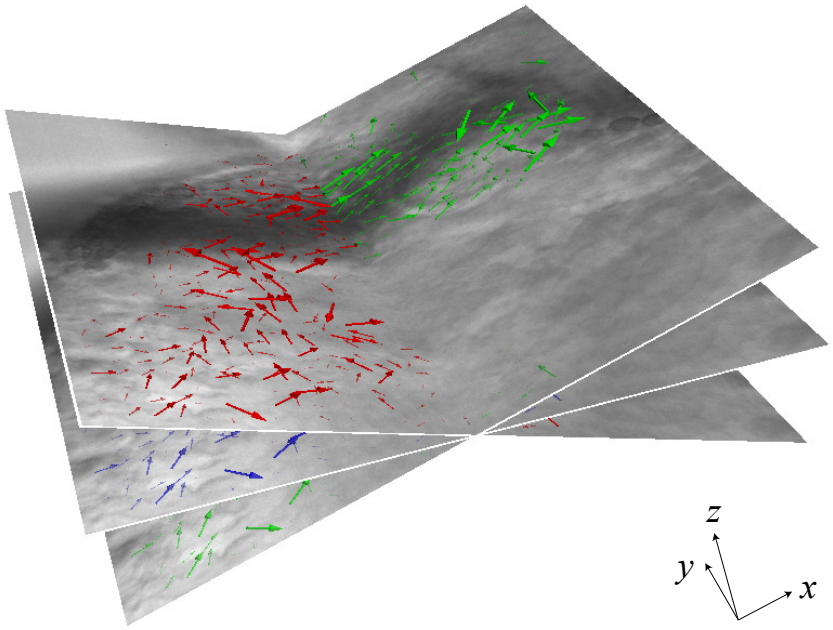

Figure 4.3: We acquired focal stacks of the zebrafish heart from three views: $-18°$, $0°$, and $18°$ relative to the xy-plane and rotated about the y-axis. At each view, 2D optical flow was used to estimate velocity vectors in each plane. For visibility, a single slice is shown for each view.

After estimating all 2D velocity fields, we estimated the 3D velocity field using the algorithm described in Section 4.2. We obtained a set of $8\times8\times4$ uniformly spaced vector coefficients from the minimization in Eq. 4.6 and used it to interpolate the 3D velocity field onto a $128\times128\times64$ grid. Fig. 4.4 shows the reconstructed blood flow velocity field during the atrial contraction phase of the heart. For a 60 hpf zebrafish embryo, we observed a maximum blood flow velocity of approximately 4 mm/s as blood is pumped from the atrium to the ventricle. This value falls between the maximum AV velocities

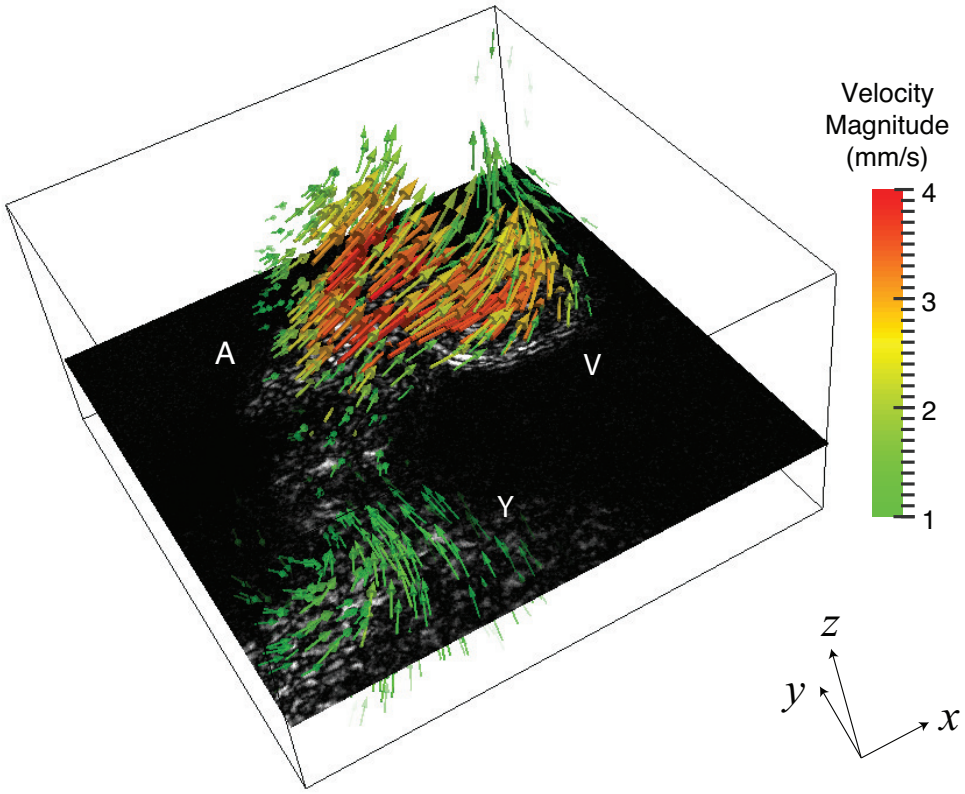

Figure 4.4: We combine 2D flow estimates from three different views to recover a divergence-free 3D velocity map of blood flow through the heart of a zebrafish larva (A: atrium, V: ventricle, Y: yolk).

of 0.9 mm/s and 5 mm/s for a 37 hpf and 4.5 dpf (days post fertilization) embryo, respectively, measured by Hove et al. [13]. We repeated this procedure at all timepoints to characterize three-dimensional blood flow over an entire cardiac cycle.

4.5 Conclusion

Since blood flow is inherently three-dimensional, current 2D imaging methods, which cannot measure out-of-plane motion, are inadequate to fully characterize complex flow trajectories. Multi-view imaging allows one to recover the out-of-plane component of motion and measure 3D flow. In this chapter, we demonstrate a new method for combining multi-view 2D flow estimates to recover a divergence-free, 3D flow field. Since our method starts from 2D vector fields, we rely on accurate 2D flow separation and motion estimation algorithms as a pre-processing step. Additionally, since the normal vectors $\mathbf{n}_k$ are assumed to be known, any rotational imperfections could also result in errors in the recovered 3D flow. This may be resolved with careful calibration, better volumetric image registration, or perhaps by including the projection angles (the normal vectors) into the optimization framework itself. We foresee our method to be applicable to not only transmitted light microscopy, but also to other modalities such as fluorescence microscopy (with fluorescently labeled blood cells).

Chapter 5

Conclusion

5.1 Summary of Contributions

In conclusion, we have presented a set of computational imaging methods for improving resolution – temporal, spatial, and out-of-plane velocity resolution – in biological microscopy. These methods use modified data acquisition procedures and solve an inverse problem to reconstruct images that surpass the limits of conventional microscopes. We summarize the novelties of our methods below.

To image the beating heart at high speeds in low-light fluorescence microscopy, we proposed an active illumination image acquisition procedure and an ℓ_1 reconstruction algorithm that utilizes multiple differently illuminated images to achieve sub-frame temporal resolution. Since an active illumination light source may not be readily available to many microscopy users, we also proposed a translational temporal superresolution method using constant illumination that is compatible with any conventional microscope and does not require any specialized hardware. Other methods for temporal superresolution either use multiple cameras [2, 3] or specially modified cameras [5, 7, 8] and are ill-suited to microscopy. In contrast, our temporal superresolution methods in Chapter 2

use only a single camera with the requirement that we image repeating processes. To develop our temporal superresolution methods for repeating processes, we proposed sub-frame temporal registration procedures to align multiple cycles of a repeating process with sub-sample accuracy. In our constant illumination method, temporal registration is performed as a separate pre-processing step. In our active illumination method, temporal registration is performed jointly as part of the superresolution optimization. Also, while [3] and [6] use temporal exposure modulation (which has similarities with our illumination modulation), they use a computationally expensive exhaustive search to determine their camera exposure codes, whereas we propose a closed-form expression to generate a set of well-conditioned modulation codes for any desired reconstruction factor.

To image millimeter-scale three-dimensional fluorescent samples with high spatial resolution and without out-of-focus blur, we proposed a point-spread-function-aware filtered backprojection algorithm combined with focal-plane scanning acquisition for optical projection tomography (OPT). While focal plane scanning had been previously proposed for OPT [60], previous methods for focal plane scanning OPT (FPS-OPT) have ignored the optical point spread function (PSF) during tomographic reconstruction. Several previous methods for traditional (single-plane) OPT reconstruction have acknowledged the presence of the PSF during image formation. However, these methods use ad-hoc filtering to reduce PSF blur, and they lack an exact inversion formula based on a forward imaging model [11, 12, 62]. Other methods include a well-defined forward model, but require computationally intensive iterative algorithms for reconstruction [63]. In Chapter 3, we showed that focal plane scanning results in a modified forward imaging model that has a direct, computationally inexpensive inverse. With our method, we can reconstruct a "deconvolved" three dimensional image without need for 3D deconvolution algorithms, as the optical PSF is accounted for within the filtered backprojection algorithm itself.

To measure three-dimensional blood flow velocity in live, in-vivo microscopy, we pro-

posed a divergence-free 3D vector field reconstruction algorithm that combines velocity estimates from two (or more) views. Previous flow-estimation methods for optical microscopy have been limited to two-dimensional flow [13, 14]. While methods have been proposed to measure 3D flow with MRI [15, 16] and ultrasound imaging [17], these imaging modalities often lack the spatio-temporal resolution necessary for imaging blood flow in small organisms and developing embryos. Our method has similarities with other multi-view flow reconstruction methods [73, 74] that formulate the reconstruction problem as a constrained and regularized inverse problem. However, our approach is unique in its use of radial basis functions to model a divergence-free vector field. This assumption reduces the computational complexity of the inverse problem, since it allows us to represent the 3D flow field with relatively few coefficients. Our method allowed us to measure, for the first time, 3D blood flow at high spatio-temporal resolution in live larval zebrafish using optical microscopy.

5.2 Future Outlook

We have presented a series of tools to overcome the limitations of specific imaging modalities within optical microscopy. Future work may include extending these tools to other imaging applications. For example, optical astronomical imaging may benefit from our constant illumination temporal superresolution method in Chapter 2, since it is also characterized by low photon counts, long exposure times, and periodic motions. Future work may also include combining these tools. Specifically one could use our temporal superresolution method in Chapter 2 in combination with our OPT reconstruction method in Chapter 3, to achieve both high spatial and high temporal resolution. Another possible direction of future work is to apply our cardiac flow velocimetry method in Chapter 4 to conduct more in-depth, quantitative analyses of how blood flow changes

over development (and even across indivduals). In addition, one could apply our 3D flow reconstruction algorithm to measure blood flow dynamics in abnormal development (e.g. with a heart defect), possibly providing new insights into cardiovascular disease.

Bibliography

[1] M. Born and E. Wolf, *Principles of Optics*, ch. 8, pp. 436–445. Cambridge University Press, 7 ed., 2003.

[2] E. Shechtman, Y. Caspi, and M. Irani, *Space-time super-resolution*, *IEEE Trans. Patt. An. Mach. Intel.* **27** (April, 2005) 531–545.

[3] A. Agrawal, M. Gupta, A. Veeraraghavan, and S. G. Narasimhan, *Optimal coded sampling for temporal super-resolution*, in *2010 IEEE Conf. Comput. Vis. Patt. Recog. (CVPR)*, pp. 599–606, June, 2010.

[4] R. Pournaghi and X. Wu, *Coded acquisition of high frame rate video*, *IEEE Trans. on Image Proc.* **23** (Dec, 2014) 5670–5682.

[5] D. Reddy, A. Veeraraghavan, and R. Chellappa, *P2C2: Programmable pixel compressive camera for high speed imaging*, in *2011 IEEE Conf. Comput. Vis. Patt. Recog. (CVPR)*, pp. 329–336, June, 2011.

[6] A. Veeraraghavan, D. Reddy, and R. Raskar, *Coded strobing photography: Compressive sensing of high speed periodic videos*, *IEEE Trans. Patt. An. Mach. Intel.* **33** (2011), no. 4 671–686.

[7] J. Holloway, A. C. Sankaranarayanan, A. Veeraraghavan, and S. Tambe, *Flutter shutter video camera for compressive sensing of videos*, in *IEEE Int. Conf. Comp. Photo. (ICCP)*, pp. 1–9, April, 2012.

[8] R. Koller, L. Schmid, N. Matsuda, T. Niederberger, L. Spinoulas, O. Cossairt, G. Schuster, and A. K. Katsaggelos, *High spatio-temporal resolution video with compressed sensing*, *Opt. Express* **23** (Jun, 2015) 15992–16007.

[9] A. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. IEEE, 1988.

[10] A. Andersen and A. Kak, *Simultaneous algebraic reconstruction technique (SART): A superior implementation of the ART algorithm*, *Ultrasonic Imaging* **6** (1984), no. 1 81 – 94.

[11] J. Walls, J. Sled, J. Sharpe, and R. Henkelman, *Resolution improvement in emission optical projection tomography*, *Phys. Med. Biol.* **52** (2007), no. 10 2775–2790.

[12] L. Chen, J. McGinty, H. Taylor, L. Bugeon, J. Lamb, M. Dallman, and P. French, *Incorporation of an experimentally determined MTF for spatial frequency filtering and deconvolution during optical projection tomography reconstruction*, *Opt. Express* **20** (Mar, 2012) 7323–7337.

[13] J. Hove, R. Koster, A. Forouhar, G. Acevedo-Bolton, S. Fraser, and M. Gharib, *Intracardiac fluid forces are an essential epigenetic factor for embryonic cardiogenesis*, *Nature* **421** (Jan., 2003) 172–177.

[14] S. Bhat, J. Ohn, and M. Liebling, *Motion-based structure separation for label-free high-speed 3-D cardiac microscopy*, *IEEE Trans. Image Process.* **21** (Aug, 2012) 3638–3647.

[15] L. Wigström, L. Sjöqvist, and B. Wranne, *Temporally resolved 3D phase-contrast imaging*, *Magn. Reson. Med.* **36** (1996), no. 5 800–803.

[16] M. Markl, A. Frydrychowicz, S. Kozerke, M. Hope, and O. Wieben, *4D flow MRI*, *J. Magn. Reson. Im.* **36** (2012), no. 5 1015–1036.

[17] J. Jensen and P. Munk, *A new method for estimation of velocity vectors*, *IEEE Trans. Ultrason., Ferroelectr., Freq. Control* **45** (1998).

[18] **K. Chan**, L. Trinh, and M. Liebling, *A temporal superresolution method applied to low-light cardiac fluorescence microscopy*, in *Proceedings of the IEEE Asilomar Conference on Signals, Systems and Computers*, 2013.

[19] **K. G. Chan**, S. J. Streichan, L. A. Trinh, and M. Liebling, *Simultaneous temporal superresolution and denoising for cardiac fluorescence microscopy*, *IEEE Transactions on Computational Imaging* **2** (2016), no. 3 348–358.

[20] **K. G. Chan**, S. Calinon, and M. Liebling, *Temporal superresolution imaging of repeating processes using a single camera and active illumination*, *IEEE Transactions on Computational Imaging* ((submitted)).

[21] J. Vermot, S. E. Fraser, and M. Liebling, *Fast fluorescence microscopy for imaging the dynamics of embryonic development*, *HFSP J.* **2** (June, 2008) 143–155.

[22] A. Bruderlin and L. Williams, *Motion signal processing*, in *Proc. of SIGGRAPH '95*, pp. 97–104, ACM, 1995.

[23] H. Bruder, S. Schaller, B. Ohnesorge, and T. Mertelmeier, *High-temporal-resolution volume heart imaging with multirow computed tomography*, *Proc. SPIE* **3661** (1999) 420–432.

[24] G. Bub, M. Tecza, M. Helmes, P. Lee, and P. Kohl, *Temporal pixel multiplexing for simultaneous high-speed, high-resolution imaging*, *Nature Methods* **7** (March, 2010) 209–211.

[25] H. E. Edgerton, "Stroboscopic." `http://edgerton-digital-collections.org/techniques/stroboscope`. Accessed Mar. 3, 2016.

[26] D. W. Staudt, J. Liu, K. S. Thorn, N. Stuurman, M. Liebling, and D. Y. R. Stainier, *High-resolution imaging of cardiomyocyte behavior reveals two distinct steps in ventricular trabeculation*, *Development* **141** (2014), no. 3 585–593.

[27] R. Raskar, A. Agrawal, and J. Tumblin, *Coded exposure photography: Motion deblurring using fluttered shutter*, *ACM Trans. Graph.* **25** (July, 2006) 795–804.

[28] 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** (May, 2015) 2281–2284.

[29] F. Sroubek, G. Cristobal, and J. Flusser, *A unified approach to superresolution and multichannel blind deconvolution*, *IEEE Trans. Image Proces.* **16** (Sept., 2007) 2322–2332.

[30] S. Farsiu, M. Robinson, M. Elad, and P. Milanfar, *Fast and robust multiframe super resolution*, *IEEE Trans. Image Proces.* **13** (Oct., 2004) 1327–1344.

[31] M. Ben-Ezra, A. Zomet, and S. K. Nayar, *Video super-resolution using controlled subpixel detector shifts*, *IEEE Trans. Patt. An. Mach. Intel.* **27** (June, 2005) 977–987.

[32] M. G. L. Gustafsson, *Surpassing the lateral resolution limit by a factor of two using structured illumination microscopy*, *J. Microsc.* **198** (2000), no. 2 82–87.

[33] M. G. L. Gustafsson, *Nonlinear structured-illumination microscopy: Wide-field fluorescence imaging with theoretically unlimited resolution*, *Proc. Nat. Acad. Sc.* **102** (2005), no. 37 13081–13086.

[34] Q. S. Hanley, J. Verveer, M. J. Gemkow, D. Arndt-Jovin, and T. M. Jovin, *An optical sectioning programmable array microscope implemented with a digital micromirror device*, *Journal of Microscopy* **196** (1999), no. 3 317–331.

[35] R. Heintzmann, *Saturated patterned excitation microscopy with two-dimensional excitation patterns*, *Micron* **34** (2003), no. 6-7 283–291.

[36] F. Orieux, E. Sepulveda, V. Loriette, B. Dubertret, and J. C. Olivo-Marin, *Bayesian estimation for optimized structured illumination microscopy*, *IEEE Trans Image Process* **21** (Feb, 2012) 601–614.

[37] H. Takeda, P. Milanfar, M. Protter, and M. Elad, *Super-resolution without explicit subpixel motion estimation*, *IEEE Trans. Image Proces.* **18** (Sept, 2009) 1958–1975.

[38] M. Liebling, J. Vermot, A. S. Forouhar, M. Gharib, M. E. Dickinson, and S. E. Fraser, *Nonuniform temporal alignment of slice sequences for four-dimensional imaging of cyclically deforming embryonic structures*, in *Proc. 3rd IEEE Int. Symp. Biomed. Imag. (ISBI'06)*, pp. 1156–1159, April,, 2006.

[39] G. B. Dantzig, A. Orden, and P. Wolfe, *The generalized simplex method for minimizing a linear form under linear inequality restraints.*, *Pacific J. Math.* **5** (1955), no. 2 183–195.

[40] P. Hansen, *Analysis of discrete ill-posed problems by means of the L-curve*, *SIAM Review* **34** (1992), no. 4 561–580.

[41] U. Krzic, S. Gunther, T. E. Saunders, S. J. Streichan, and L. Hufnagel, *Multiview light-sheet microscope for rapid in toto imaging*, *Nature Methods* **9** (2012), no. 7 730–733.

[42] ILOG-CPLEX, "High-performance software for mathematical programming and optimization." `http://www.ilog.com/products/cplex`, 2005.

[43] M. Westerfield, *The Zebrafish Book: A Guide for the Laboratory Use of Zebrafish (Brachydanio Rerio), 4th Edition*. Univ. of Oregon Press, Eugene, 2000.

[44] Z. Lin and H. Y. Shum, *Fundamental limits of reconstruction-based superresolution algorithms under local translation*, *IEEE Trans. Patt. An. Mach. Intel.* **26** (Jan., 2004) 83–97.

[45] T. F. Coleman and Y. Li, *An interior trust region approach for nonlinear minimization subject to bounds*, *SIAM Journal on Optimization* **6** (1996), no. 2 418–445.

[46] Seoul Semiconductor. `http://www.seoulsemicon.com/`. Accessed: Sep. 11, 2016.

[47] Arduino. `https://www.arduino.cc/`. Accessed: Sep. 11, 2016.

[48] Thorlabs, Inc. `https://www.thorlabs.com/`. Accessed: Sep . 11, 2016.

[49] Rethink Robotics, "Baxter." `http://www.rethinkrobotics.com/baxter/`. Accessed: Sep 11, 2016.

[50] Adobe Systems Inc., "Adobe After Effects CS5." `http://www.adobe.com/products/aftereffects.html`. Accessed Nov. 30, 2015.

[51] L. A. Trinh, T. Hochgreb, M. Graham, D. Wu, F. Ruf-Zamojski, C. S. Jayasena, A. Saxena, R. Hawk, A. Gonzalez-Serricchio, A. Dixson, E. Chow, C. Gonzales, H. Y. Leung, I. Solomon, M. Bronner-Fraser, S. G. Megason, and S. E. Fraser, *A versatile gene trap to visualize and interrogate the function of the vertebrate proteome*, *Genes & Development* **25** (Nov, 2011) 2306–2320.

[52] C.-J. Huang, C.-T. Tu, C.-D. Hsiao, F.-J. Hsieh, and H.-J. Tsai, *Germ-line transmission of a myocardium-specific gfp transgene reveals critical regulatory elements in the cardiac myosin light chain 2 promoter of zebrafish*, *Developmental Dynamics* **228** (2003), no. 1 30–40.

[53] S. Bhat, I. V. Larina, K. V. Larin, M. E. Dickinson, and M. Liebling, *Multiple-cardiac-cycle noise reduction in dynamic optical coherence tomography of the embryonic heart and vasculature*, *Opt. Lett.* **34** (Dec, 2009) 3704–3706.

[54] F. Luisier, C. Vonesch, T. Blu, and M. Unser, *Fast interscale wavelet denoising of poisson-corrupted images*, *Signal Processing* **90** (February, 2010) 415–427.

[55] "`http://sybil.ece.ucsb.edu/pages/temporalsuperresolution`."

[56] F. N. Fritsch and R. E. Carlson, *Monotone piecewise cubic interpolation*, *SIAM J. Num. Analys.* **17** (1980), no. 2 238–246.

[57] **K. G. Chan** and M. Liebling, *A point-spread-function-aware filtered backprojection algorithm for focal-plane-scanning optical projection tomography*, in *IEEE International Symposium on Biomedical Imaging*, 2016.

[58] J. Sharpe, U. Ahlgren, P. Perry, B. Hill, A. Ross, J. Hecksher-Sørensen, R. Baldock, and D. Davidson, *Optical projection tomography as a tool for 3D microscopy and gene expression studies*, *Science* **296** (2002), no. 5567 541–545.

[59] R. Bryson-Richardson, S. Berger, T. Schilling, T. Hall, N. Cole, A. Gibson, J. Sharpe, and P. Currie, *FishNet: an online database of zebrafish anatomy*, *BMC Biol.* **5** (2007), no. 34.

[60] Q. Miao, J. Hayenga, M. Meyer, T. Neumann, A. Nelson, and E. Seibel, *Resolution improvement in optical projection tomography by the focal scanning method*, *Opt. Lett.* **35** (Oct, 2010) 3363–3365.

[61] A. Bassi, B. Schmid, and J. Huisken, *Optical tomography complements light sheet microscopy for in toto imaging of zebrafish development*, *Development* **142** (2015), no. 5 1016–1020.

[62] J. Guo, Y. Yang, D. Dong, L. Shi, H. Hui, M. Xu, J. Tian, and X. Liu, *A projection selection method to improve image quality in optical projection tomography*, in *Int. Conf. of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 206–209, Aug, 2014.

[63] U. S. Kamilov, I. N. Papadopoulos, M. H. Shoreh, A. Goy, C. Vonesch, M. Unser, and D. Psaltis, *Optical tomographic image reconstruction based on beam propagation and sparse regularization*, *IEEE Transactions on Computational Imaging* **2** (March, 2016) 59–70.

[64] Z. Liang and D. M. Jr., *Partial Radon transforms*, *IEEE Trans. Image Proces.* **6** (Oct, 1997) 1467–1469.

[65] S. Horbelt, M. Liebling, and M. Unser, *Discretization of the Radon transform and of its inverse by spline convolutions*, *IEEE Trans. Med. Imag.* **21** (April, 2002) 363–376.

[66] H. Kirshner, F. Aguet, D. Sage, and M. Unser, *3-D PSF fitting for fluorescence microscopy: Implementation and localization application*, *J. Microsc.* **249** (Jan, 2013) 13–25.

[67] S. Azevedo, D. Schneberk, J. Fitch, and H. Martz, *Calculation of the rotational centers in computed tomography sinograms*, *IEEE Trans. Nucl. Sc.* **37** (Aug, 1990) 1525–1540.

[68] **K. G. Chan** and M. Liebling, *Estimation of divergence-free 3D cardiac blood flow in a zebrafish larva using multi-view microscopy*, in *IEEE International Symposium on Biomedical Imaging*, 2015.

[69] F. Fahrbach, F. Voigt, B. Schmid, F. Helmchen, and J. Huisken, *Rapid 3D light-sheet microscopy with a tunable lens*, *Opt. Express* **21** (Sep, 2013) 21010–21026.

[70] Y. Pu, X. Song, and H. Meng, *Off-axis holographic particle image velocimetry for diagnosing particulate flows*, *Exp. in Fluids* **29** (2000), no. 1 117–128.

[71] F. Pereira, M. Gharib, D. Dabiri, and D. Modarress, *Defocusing digital particle image velocimetry: a 3-component 3-dimensional DPIV measurement technique. Application to bubbly flows*, *Exp. in Fluids* **29** (2000), no. 1 078–084.

[72] X. Liu, K. Abd-Elmoniem, M. Stone, E. Murano, J. Zhuo, R. Gullapalli, and J. Prince, *Incompressible deformation estimation algorithm (IDEA) from tagged MR images*, *IEEE Trans. Med. Imag.* **31** (Feb, 2012) 326–340.

[73] A. Falahatpisheh, G. Pedrizzetti, and A. Kheradvar, *Three-dimensional reconstruction of cardiac flows based on multi-planar velocity fields*, *Exp. in Fluids* **55** (2014), no. 11.

[74] M. Arigovindan, M. Sühling, C. Jansen, P. Hunziker, and M. Unser, *Full motion and flow field recovery from echo doppler data*, *IEEE Trans. Med. Imag.* **26** (Jan, 2007) 31–45.

[75] B. Lucas and T. Kanade, *An iterative image registration technique with an application to stereo vision*, in *Int. Joint Conf. Artificial Intel.*, pp. 674–679, 1981.

[76] O. Skrinjar, A. Bistoquet, J. Oshinski, K. Sundareswaran, D. Frakes, and A. Yoganathan, *A divergence-free vector field model for imaging applications*, in *IEEE Int. Symp. Biomed. Imag.*, pp. 891–894, June, 2009.

[77] A. Squillacote, *The ParaView Guide: A Parallel Visualization Application*. Kitware, Inc., 3rd ed., 2008.

[78] M. Abramoff, W. Niessen, and M. Viergever, *Objective quantification of the motion of soft tissues in the orbit*, *IEEE Trans. Med. Imag.* **19** (Oct, 2000) 986–995.