Holographic Microscopy for Soft Matter and Biophysics

A DISSERTATION PRESENTED
BY
THOMAS G. DIMIDUK
TO
THE DEPARTMENT OF PHYSICS

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN THE SUBJECT OF
PHYSICS

HARVARD UNIVERSITY
CAMBRIDGE, MASSACHUSETTS

NOVEMBER 2016

Thesis advisor: Professor Vinothan N. Manoharan Thomas G. Dimiduk

Holographic Microscopy for Soft Matter and Biophysics

Abstract

I discuss a series of advancements I have made towards making digital holographic microscopy into a useful tool for experimental scientists in soft-matter physics and biophysics. Digital holograms can be recorded with simple hardware at high speed to capture three-dimensional information about the dynamics of aqueous suspensions of colloidal particles, cells, viruses or other microscopic objects.

The challenge of working with holograms is that they map information about the objects non-locally onto an interference pattern. Therefore, post processing is needed to extract the information from a hologram. Traditionally this has been done by reconstructions, effectively shining light back through the hologram to obtain a representation of the recorded objects. More recently Ovryn and Izen (JOSA A 2000) and Lee, Grier and coworkers (Opt. Express 2012) have shown that more precise information can be recovered by physically modelling the light scattering that creates the hologram and solving a constrained inverse-scattering problem to obtain information about the scatterers such as their position or size. This technique gives precise results but requires a scattering model for the objects under observation. It therefore requires significant expertise to set up and implement.

Thesis advisor: Professor Vinothan N. Manoharan Thomas G. Dimiduk

In this dissertation I present several advances that improve upon this state-of-the art. First, I present a simple, inexpensive, portable, battery-powered holographic microscope that is suitable for imaging biological samples inside an incubator. Next, I describe a method using a general scattering model called the discrete dipole approximation to analyze holograms of non-spherical particles. Because this analysis is computationally expensive, I present a new method based on analyzing a random subset of the pixels of a hologram. This method, which significantly speeds up computation, is the basis for a framework based on Bayesian inference that gives a more intuitive and rigorous way of specifying prior information and presenting uncertainty in results, which I present at the end.

The motivating thread through this thesis is building tools to enable new experiments using holography and making it easier for scientists who are not experts in holography and light scattering to use holography as a tool to do the science that interests them. In support of these goals, I have implemented all of the computational techniques and physical models in an open source library, HoloPy, to make it as easy as possible for other scientists to use them.

Contents

Prior Publications

This thesis is primarily based on the following publications:

- Chapter 2: T. G. Dimiduk, A. Wang, J. Fung, S. Razavi, I. Kretzschmar, K. Chaudhary, and V. N. Manoharan, "Using the discrete dipole approximation and holographic microscopy to measure rotational dynamics of non-spherical colloidal particles," Journal of Quantitative Spectroscopy and Radiative Transfer, vol. 146, pp. 499–509, Oct. 2014.
- Chapter 3: T. G. Dimiduk, R. W. Perry, J. Fung, and V. N. Manoharan, "Random-subset fitting of digital holograms for fast three-dimensional particle tracking," Applied Optics, vol. 53, pp. G177–G183, Sept. 2014.
- T. G. Dimiduk and V. N. Manoharan, "Bayesian approach to analyzing holograms of colloidal particles," Optics Express, Submitted for publication.

Other works to which the author has contributed in graduate school include:

- J. Fung, R. W. Perry, T. G. Dimiduk, and V. N. Manoharan, "Imaging Multiple Colloidal Particles by Fitting Electromagnetic Scattering Solutions to Digital Holograms," Journal of Quantitative Spectroscopy and Radiative Transfer, vol. 113, pp. 2482–2489, Dec. 2012.
- R. W. Perry, G. Meng, T. G. Dimiduk, J. Fung, and V. N. Manoharan, "Real-space studies of the structure and dynamics of self-assembled colloidal clusters," Faraday Discussions, 2012.
- A. Carpio, T. G. Dimiduk, M. L. Rapún, and V. Selgas. Noninvasive imaging of 3D micro and nanostructures by topological methods. SIAM Journal on Imaging Sciences. In press.

Acknowledgments

I am thankful for the many people who have helped me with science and with life throughout grad school. Without them I would have accomplished much less, and been much less happy doing it. I will not take the space here to properly thank everyone, but know that if you helped with science or were a good friend, I appreciate it.

I should obviously start with my advisor Vinny Manoharan. While we didn't always communicate well, you taught me to be a much better communicator. You also helped me grow from an enthusiastic young physics graduate into a proper scientist. You gave me the space to learn how to work independently, but were always available when I needed advice (even if I didn't remember to ask for it as often as I should have). You helped me to learn to think, speak, and write more precisely. Thank you for introducing me to distributed version control, unit testing, and the python scientific stack, and giving me the space to grow as a programmer. Your good taste in software design was as valuable in its own way as your expertise in physics. Lastly, thank you for fostering such an amazing social environment in the group and for your excellent taste in people as students and postdocs. I ended up liking all of the people I worked with in lab and becoming close friends with many of them.

Next I should thank my collaborators and friends Rebecca Perry, Anna Wang, Jerome

Fung, and Viva Horowitz. I literally cannot imagine what grad school would have been like without you. It is obvious from our shared publications how much we worked together, but it is not obvious how much I enjoyed working with all of you. Becca: your enthusiasm, curiosity and wonder never ceased to amaze me. My biggest regret when we switched to hot-desking was that I was no longer sitting next to you to have you excitedly show me the latest cool thing you had found or done. You were already an excellent scientist, but it was just fun watching how fast you picked up programming and became a proper hacker. I glad to have been able to help you along that road, though I am sure you would have gotten there without me. Anna: Thank you for finding some useful stuff to do with my DDA work even though it didn't pan out for biological work like I was hoping. Your perspective on life was sometimes foreign to me (and no I don't mean the Australian part), but I am a better person for letting some of it rub off on me. Jerome: Thank you for introducing me to the nuts and bolts of holography and scattering. I could never keep up with you on scattering theory, but I am glad I finally convinced you of the value of functions. Viva, your enthusiasm and curiosity was inspiring, even a little overwhelming sometimes. I wish more of your care and detail mindedness had rubbed off on me. You got to see too much of me at my worst but somehow still hold me in excessively high regard. I don't deserve you.

If I tried to properly thank all of the wonderful people I worked with in the Manoharan lab I would never finish these acknowledgments, but here is the attempt I have the energy left for after way too much writing. Thank you Dave and Ryan for getting me started

with microscopy and holography. Thank you Jesse for all the brainstorming and ideas we shared, I am a smarter person when talking to you, and you help me recognize strengths I didn't realize I had. Sofia, your enthusiasm and infectious curiosity were a ray of sunshine from the right end of the office. Nick, thanks for all the fun you brought to lab, and for making sure I knew I was appreciated. Aaron and Solomon, I am glad to see holography will carry on in your capable hands. Emily, Ellen, Nabila, Toly, Andrew, Victoria, Annie, Reese, Jin-Gyu, Ben, Kosta, Guangnan, I never worked closely with you nor got to know you as well as I would have preferred but you made lab a better place and I am glad our to have spend the time I did with you.

My route through grad school was not without its bumps. I am eternally grateful to all of my friends and family who helped keep me sane for a most of the long slog and then helped me piece myself back together when I fell apart in spite of their best efforts. I would like to call out my friends Cheryl Agaskar-Sorace, Ameya Agaskar, Lusann Yang, Dale Winter, and Emily Russel, and Alan Morse. You all had your own tribulations through grad school, know that you were a big part of me making it through mine.

I should also thank my counselors Carmen Cruz and David Abramson. When I got too far down for even the best of friends and family to help, you somehow patched me back together and got me going again. I would not have finished this degree without your help.

Jeff and Sally, thank you for being such amazing siblings. Jeff: I probably don't need to tell you this, but you in many ways edge out my parents as the oldest and best

influence on how I think, especially about numbers, logic, and solving problems. On top of that, no one could ask for a more fun brother to hang out with. Sally: we were both too busy to talk as often as we should have, but your levelheaded sensibility combined with deep resilience was a source of strength for me. Ellen and Ben, thank you for being such great partners to my siblings and good friends to me.

Lastly, I thank my parents. Words can't do the gratitude I owe you justice, but I know you understand anyway. In addition to being the reason I am in the world at all, you are also the reason I am a scientist and a big part of the reason I am any good at it. I had the fortune of having two parents I could talk to about my research, and who could actually give me useful advice on it. It also ended up being far more important than I would have preferred that you were still willing to take care of me when I needed it even as an adult. Thank you Mom and Dad getting me here, for sharing in the highs, and for helping me through the lows, and for everything.

Co-Authors

The following authors contributed to the following chapters:

Chapter 2 Ekaterina A. Kosheleva, David Kaz, Ryan McGorty, Emily J. Gardel, and Vinothan N. Manoharan

Chapter 3 Anna Wang, Jerome Fung, Sepideh Razavi, Ilona Kretzschmar, Kundan Chaudhary, Jennifer A. Lewis, and Vinothan N. Manoharan

Chapter 4 Rebecca W. Perry, Jerome Fung, and Vinothan N. Manoharan

Chapter 5 Vinothan N. Manoharan

1

Introduction

Humans are naturally visual creatures. To a perhaps surprising degree, what we study is determined by what we can look at. It is not surprising, then, that there have been many Nobel prizes awarded for tools to help us see things better (Betzig, Hell and Moerner for super-resolved flouresence microscopy; Boyle and Smith for the charge-coupled device (CCD); Lauterbur and Mansfield for magnetic resonance imaging (MRI); Ruska,

Binning, and Rohrer for the Electron Microscope and the scanning tunneling microscope; Klug for crystallographic electron microscopy; Gabor for Holography; Cormack and Hounsfield for computer assisted tomography)[1] and several more for related technologies. This thesis was made possible by the work leading to two of those imaging prizes.

Holography, invented by Dennis Gabor in 1948[2], is a way to capture a three-dimensional image of an object from the phase of light. Because light oscillates at frequencies around 10^{14} Hz it is not possible to measure the phase directly; therefore, a hologram captures information about the phase from interference. The interference pattern records the electric field at a specific plane. To recover the field, Gabor proposed what is now known as optical reconstruction, which involves shining light back through the hologram. Light that is diffracted by the hologram forms a three-dimensional real image of the object captured by the hologram.

There are several challenges in holography that have limited it to mostly niche applications. Because holography records information in very fine interference fringes, one has to record holograms at very high resolution. Until recently that meant using specialized photographic media. Besides requiring specialized equipment, these media generally require large light irradiance for exposure, necessitating long exposure times. This requirement imposes strict constraints on the stability of the apparatus, since any movement of order the wavelength of the light over the exposure can blur out the fringes in the hologram. Also, it is difficult to obtain quantitative information from optical re-

constructions, further limiting their use in science.

CCD and CMOS digital sensors solve the first of these problems. Modern digital sensors have the resolution and pixel density to be useful for capturing holograms, enabling a new era of digital holography[3]. CCD and CMOS sensors can be much more light sensitive than the old photographic media, which, combined with easy availability of high powered lasers, allows us to record holograms more quickly. The enhanced speed shortens the time over which motion can blur the fringes during recording and relaxes the constraints on the stability of the apparatus. It has taken holography from a relatively slow way to record an image to one of the fastest ways to record a three-dimensional scene. My colleagues have captured holograms at rates in excess of 1000 frames per second[4].

Unlike many other three-dimensional imaging techniques, such as confocal microscopy, holographic microscopes need not be complex. They require high precision and careful alignment, but only a few optical elements. Therefore, as I show in Chapter 2 it is possible to build simple holographic instruments. Although these offer no particular advantage in quantitative analysis, they show that new ways of making useful measurements with holograms can make the technique more widely applicable.

Recording holograms digitally also makes it possible to do digital processing which addresses the problem of making quantitative measurements with holograms. For example, there is a digital analog of optical reconstruction called numerical reconstruction[3] (Chapter 2). This technique enables the use of holography as a kind of large-depth-of-

field microscopy for imaging swimming microorganisms[5–7], tracking microspheres[8–11], tracking flow[12], imaging cells[13,14] other imaging other biological samples[15–17]. However, for objects of size comparable to the wavelength of light, this method leads to imaging artifacts[18]. These artifacts and other difficulties in working with reconstructions have not prevented experts from doing useful work with holograms, but they do limit the wider application of holography.

The need for higher precision measurements led Ovryn and Izen[19,20] and later Lee, Grier and coworkers[21] to an approach that involves predicting holograms from Lorenz–Mie scattering theory and solving an inverse scattering problem to characterize the particle recorded in the hologram. The approach put forth by Lee and coworkers is to use the Levenberg–Marquardt algorithm[22] to minimize a χ^2 error metric (sum over squared pixel differences) and fit for the parameters of the Lorenz–Mie theory (particle position, radius, and refractive index) that best match the recorded hologram. These parameters provide a high-precision measurement of that particle.

This technique has proved useful for tracking[4,23], characterizing[23–26] and monitoring[27,28] particles and for microrefractometry[29] but it suffers from some limitations. Lorenz–Mie theory is only valid for spherical particles, and in fact for only a single sphere. It is possible to use the theory for multiple spheres by linearly superimposing the scattered electric fields from each. However, this approach neglects any multiple scattering or near-field effects of scatterers in close proximity and thus is only valid for well-separated particles. With my assistance, my colleagues Jerome Fung and Rebecca

Perry overcame this limitation by implementing an exact multi-sphere superposition scattering theory that enabled them to make new measurements of dense clusters of particles[30,31].

As the reader is hopefully well aware, only a small fraction of things in the world are spheres, or even well approximated by spheres. Jokes about spherical cows aside, most physicists are not satisfied with considering only spherical objects. Fortunately, it is possible, though somewhat more difficult, to model light scattering from other shapes. One general technique for modelling light scattering from arbitrarily shaped objects is the discrete dipole approximation (DDA)[32]. In the DDA objects are represented as a discrete three-dimensional array of point dipoles. One then solves for a self-consistent polarization of these dipoles an incident electric field. This solution can be remarkably accurate; the only error introduced is from the discretization. I have adapted the DDA for use in modeling and fitting holograms, and with my colleague Anna Wang, have used this technique to study the diffusion of non-spherical particles, as described in Chapter 3.

Unfortunately, these scattering models can take a significant amount of time to compute. Because of the thousands of evaluations required to find optimal parameters, fitting holograms can take a prohibitively long time; particularly when working with the DDA or with larger clusters of spheres. To help address this problem, I developed a technique for fitting holograms using only a randomly selected fraction of its pixels. Although one must capture holograms at high resolutions to record fine fringes, the

resulting holograms have a large amount of redundant data. Randomly selecting pixels can sample all of the useful information while avoiding many wasteful evaluations. I describe this technique in detail in Chapter 4 and demonstrate its effectiveness and accuracy.

So far I have glossed over a detail that significantly limits these scattering-model-based approaches to analyzing holograms. The space covered by the many-dimensional parameterizations of scattering objects is vast and can contain local minima of the χ^2 error metric. Thus, algorithms such as Levenberg–Marquardt will only converge if given an initial guess sufficiently close to the correct answer. This turns out to be a significant impediment; it is a difficult task to determine such an initial guess from a hologram. The need to manually initialize fits is a major limiting factor in analysis of holographic data, particularly for more complex systems. A skilled expert can find, with heroic effort, satisfactory initial guesses to fit a cluster of six particles[30], but that expertise took years to develop; therefore, manual guessing is hardly a scalable approach to wide application.

There is another problem with the parameterizations used in these fits. Lee and coworkers[21] were forced to introduce a non-physical parameter α to account for some optical effects in the recording of the hologram. Furthermore, one must always fit for all parameters of the scattered object or objects, even if one is interested in, say, only the position or size of a particle. Covariance between the parameters can mean that the parameters not of interest can affect the uncertainty of those that are. Parameters that

are required in the model but are not of interest are often called "nuisance parameters" because they interfere with analysis. These parameters are such a nuisance because, in classical frequentist statistics, there isn't really anything you can do about them.

Fortunately there is a way to deal with both of these problems, and I describe it in Chapter 5. Bayesian statistical techniques provide a more powerful way to include prior information, such as the initial guesses used in fitting holograms, and enable one to integrate over, or marginalize, nuisance parameters. Bayesian priors are a natural way to describe what one knows about a parameter. Combined with a subset based tempering approach I introduce in Section 5.3.1, the method allows us to infer parameter values starting with only information that is easy to obtain. Then, by marginalizing over α and any other parameters not of interest, one can accurately assess the uncertainty of parameters of interest and the covariances between them.

These techniques greatly expand the space of phenomena that can be investigated using powerful scattering-based approaches. Combined with my open source implementation in the software package HoloPy[33], they also greatly reduce the expertise required to use these techniques, making them available to a much wider variety of scientists interested in dynamics of microscopic in soft matter physics and biophysics.

2

A Simple and Inexpensive Holographic Microscope for Biological Imaging

Recent advances in optical microscopy have led to enhanced resolution and sensitivity, but at the cost of simplicity and portability. Although modern microscopes

can optically section a sample, detect single molecules, and even break the diffraction limit, they are still bulky instruments costing tens or hundreds of thousands of dollars. This often puts the modern optical microscope in a time-shared central facility instead of in every lab. Common add-ons such as environmental chambers, micromanipulators, or microfluidic devices further drive up costs and complicate sharing. But for many experiments in biology, materials science, and physics that do not require all the features offered by the latest generation of microscopes, it makes more sense to use simpler, cheaper microscopes that can be added onto other specialized apparatus. Such "commodity" could enable new experiments and even new applications for microscopy. They could also help make advanced imaging available to laboratories in the developing world, to small businesses, and to students.

Coherent imaging methods such as holography may be the most promising route to simple microscopes. Holographic microscopy is a widefield phase imaging technique similar to differential interference contrast and phase contrast microscopy. It captures information about both the phase and amplitude of light scattered from a sample. As a result, a relatively simple optical setup, in conjunction with a modern desktop computer, can create and visualize three-dimensional reconstructions of a sample[3]. Unlike conventional widefield microscopes, the image given by holographic reconstruction is linearly related to optical path lengths. This allows holography to avoid the contrast artifacts introduced by traditional widefield phase techniques, thereby enabling quantitative interpretation of the images. Furthermore, the enhanced depth of field of holographic

microscopes eliminates the need to change focus. Thus, a low cost holographic microscope can compete favorably with standard widefield microscopy for many applications.

Holographic microscopes are especially advantageous for three-dimensional (3D) imaging. Compared to the current standard for 3D microscopy, the confocal microscope, a holographic microscope is much simpler, requiring no moving parts to obtain a 3D image, and potentially much faster, because a 3D reconstruction can be obtained from a single 2D image. Holography therefore offers a low cost, less invasive alternative to confocal microscopy for three-dimensional imaging of samples where chemically specific labeling is not needed.

Here we show how advances in semiconductor technology have made it possible to build an inexpensive, self-contained, compact digital holographic microscope with good resolution and large depth of field. Our microscope uses conventional optical parts, an inexpensive diode laser, a hobbyist microcontroller board, and a consumer grade digital camera. The parts can be purchased on eBay or hobbyist websites for a total cost of less than $1000. Reconstruction is done on a standard desktop computer. The microcontroller synchronizes the laser and camera so that the sample is exposed to light for only 80 ns per hologram, resulting in a total exposure of about 20 nJ of laser light, compared to tens of millijoules for a typical 3D confocal stack[34]. At the same time, the microscope requires no pinholes or spatial filters, making it easy to align and robust to vibration. We show that the instrument can take 3D images of living cells and colloidal particle suspensions with a resolution better than 3 µm and a depth of field of more

than 150 µm.

Early holographic microscopes required precision equipment such as water-cooled lasers, spatial filters, and vibration isolators, but diode lasers and CMOS detectors have made these components unnecessary. In many ways a diode laser is an ideal light source for microscopy; it gives a high-quality, narrow-band, easily collimated beam from a compact, efficient illumination source. Lasers have not seen much use in widefield microscopy because their long coherence length imposes coherent noise and interference fringes. But in holography coherence becomes an advantage, allowing the extraction of phase information. Laser diodes suitable for holographic microscopy are now mass-produced for consumer electronics such as CD writers and laser printers. At the same time, CMOS detectors with pixel density sufficient to capture fine interference patterns with only modest magnification are now available in mass-market cameras, and modern desktop computers can reconstruct the captured megapixel holograms in less than a second. The combination of a pulsed laser diode and sensitive camera allows fast exposures that eliminate the need for vibration isolation. This confluence of technologies enables us to construct a "commodity" microscope with nonetheless impressive capabilities.

Holographic microscopes have been used to study biological systems, but the instruments are more complicated and more expensive than the one we show here. These other approaches include both off-axis and inline configurations. Off-axis microscopes can more easily achieve high resolution, but because they have two optical paths, they are sensitive to vibration and require careful alignment[35–38]. Inline holography gener-

ally involves a simpler optical train. Lensless versions of inline holographic microscopes have been developed for a few applications in biology, particularly analysis of swimming microorganisms[5,6,8–10,15]. These lensless instruments can be quite simple, but they impose restrictions on the geometry of the apparatus, and they require imaging cameras with very high resolution. Therefore many recent inline apparatuses for examining eukaryotic cells or small particles use an objective lens[7,11–13,16]. We choose to use an inline configuration with an objective lens because it allows us maintain good resolution while using a consumer-grade digital camera for imaging.

Other holographic microscopes have been constructed to be simple and low-cost, but ours is distinct in that it achieves these goals without a significant tradeoff in image quality or resolution. Thus, although simple microscopes can be made with lensless design and LED illumination[9], we choose to use a laser diode, which, while more expensive than an LED, has a high coherence length at a cost (∼$40) that is still small compared to that of a digital camera. An instrument recently developed by Seo and coworkers[14] also focused on simplicity, low cost, and biological imaging, but as the authors note, the instrument is engineered for cytometry, not high-resolution imaging.

Our instrument has two other distinguishing features. First, acquisition is self-contained: the instrument is internally powered and internally controlled, and it stores images on-board. This allows it to operate autonomously in a closed or remote space. Second, we use consumer and hobbyist components throughout to keep the cost low and position the microscope to benefit from the rapid improvements in these technologies.

The result is a simple, flexible and inexpensive instrument that can be incorporated into another experiment or apparatus at minimal cost and effort.

2.1 Methods

2.1.1 Design of the holographic microscope

Our instrument uses an in-line configuration[39] comprising only a laser, camera, objective, sample stage, beam telescope and microcontroller, as shown in Figure 2.1. Twin-image effects are minimized because our samples are relatively thin (tens or hundreds of µm) compared to the optical distances (multiple mm).

A battery powered 785 nm ADL 78901 laser diode provides illumination. This diode outputs a single transverse mode and has a measured coherence length in excess of 5 cm. For the images shown here, we used either 2.3 mW of laser power and 80 µs pulses or 0.235 mW and 20 ms pulses. We measured the power with a Thorlabs PM 100 power meter. An aspheric lens collimates the diode output, and a 3.75× beam-expanding telescope provides uniform illumination across the objective aperture. We did not need to incorporate a spatial filter because the diode produces a beam of sufficient quality. This greatly simplifies alignment.

A battery-powered Arduino board (Adafruit Industries) employing an Amtel ATMega 168 microcontroller controls both the laser and the camera. An analog output on the microcontroller drives a switching transistor which controls power to the laser (Figure

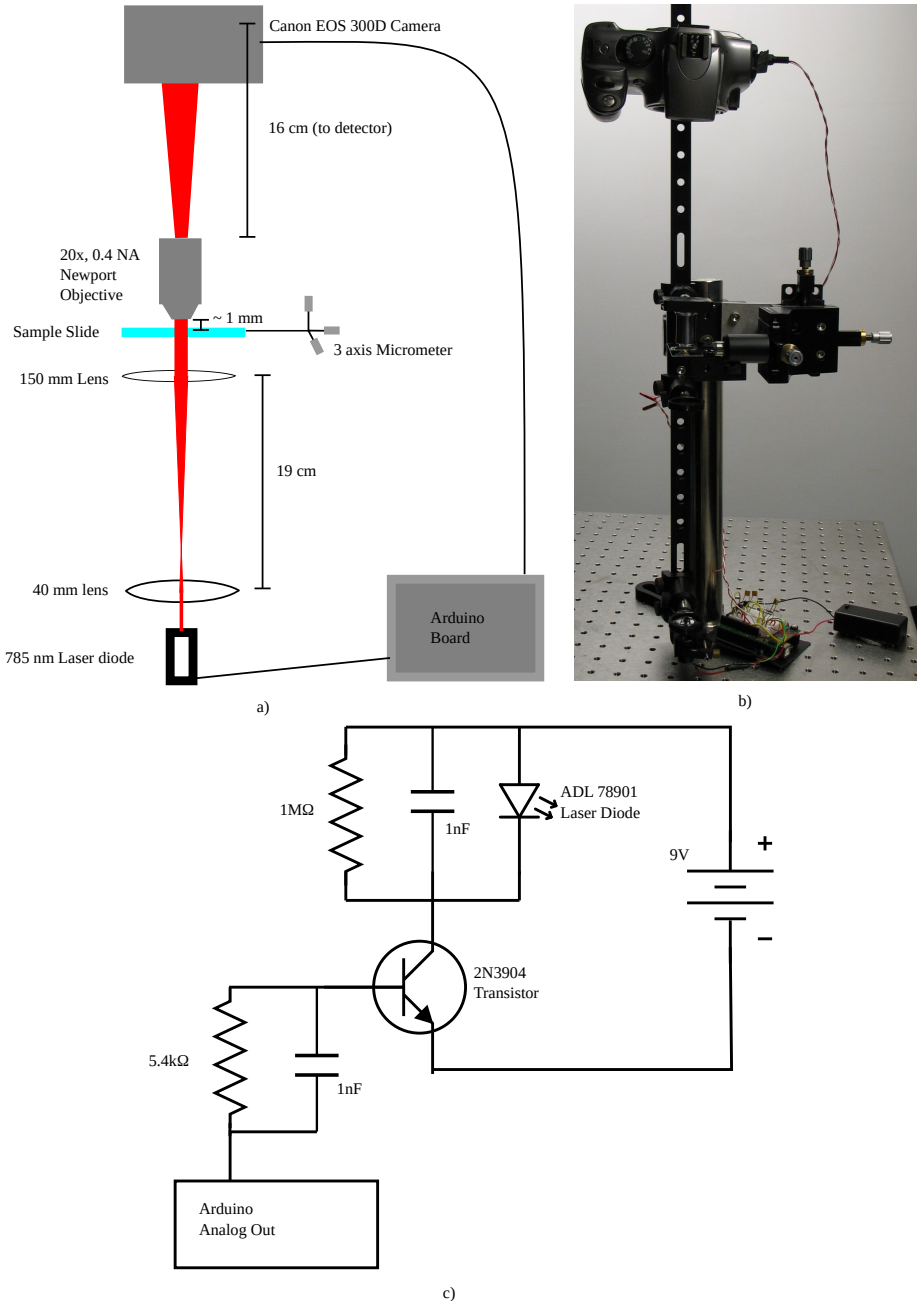

Figure 2.1: a) Microscope diagram, showing optically relevant distances such as that between the objective and the camera sensor. The beam is approximately collimated entering and exiting the beam telescope. b) Photograph of microscope. c) Drive circuitry for the ADL 78901 laser diode. The 1 MΩ resistor and 1 nF capacitor protect the laser diode from current spikes. The 5.5 $\mathrm{k}\Omega$ resistor and 1 nF capacitor smooth the Arduino Pulse Width Modulation (PWM) output into an analog voltage. The 2N3904 Transistor controls current through the diode.

1c). Because the diode is on for only a short time per exposure, no temperature control is needed. The microcontroller also triggers the camera shutter using the camera's external shutter release port. Using a microcontroller for these functions allows us to program arbitrary sequences with varying illumination and frame timing. The programs can be uploaded to the microcontroller from a desktop computer using a USB port.

Samples are mounted to a 3-axis micrometer stage for translation and focusing, and a 20×, 0.4 NA Newport objective magnifies the hologram. We take care to keep all optics free from dust and fingerprints since light scattered off these blemishes adds additional noise fringes. Because we use only collimated or weakly diverging beams, there are wide tolerances on most of the optical distances.

We record holograms with a Canon EOS 6.3 megapixel 300D digital SLR camera. This camera has a 3088×2056 pixel (22.7 mm×15.1 mm) CMOS sensor with 7.35 µm square pixels. Like most consumer cameras, it has an IR blocking filter. We removed this filter using the procedure given in [21] to obtain good sensitivity with our 785 nm light source. To minimize interference fringes from lens reflections, we image light from the objective directly onto the CMOS sensor. We store the images as RAW files on an 8 GB Compact Flash II 133X card (Kingston Technology). Although the camera can take a burst of up to four images in a little over a second, it requires about 15 seconds to save each image to the memory card. Newer consumer cameras store images significantly faster than our five-year old model.

The camera, microcontroller, and laser are all battery powered so that the instrument

can acquire data without any outside connection. Ate the end of an experiment, we transfer the holograms to a desktop computer for reconstruction.

2.1.2 Reconstruction

We use the convolution method to reconstruct holograms[3]. The reconstructed image is

$$\Gamma(\xi, \eta, d) = \int_{-\infty}^{\infty} h(x, y) g(\xi, \eta, x, y, d) dx\, dy \tag{2.1}$$

where h is the observed hologram and g is the impulse response at a distance d. This equation contains two sets of planar coordinates: x and y describe the position at the sensor, and ξ and η describe the position in the reconstructed image. The impulse response g is

$$g(\xi, \eta, x, y, d) = \frac{\frac{i}{\lambda} \exp\left(i \frac{2\pi}{\lambda} \sqrt{d^2 + (x+\xi)^2 + (y+\eta)^2}\right)}{\sqrt{d^2 + (x+\xi)^2 + (y+\eta)^2}} \tag{2.2}$$

where λ is the wavelength in medium. The reconstruction, Γ, is a complex two dimensional array whose elements represent both phase and amplitude of the reconstructed field. The amplitude can be extracted as $|\Gamma|$, and the phase as

$$\phi(\xi, \eta, d) = \arg(\Gamma). \tag{2.3}$$

Phase extraction is useful because many objects, including biological cells, are weakly

absorbing but have spatially varying indices of refraction. Although these samples have little effect on the amplitude of transmitted light, they introduce a noticeable phase shift. Phase images therefore frequently give better contrast for biological materials. When reconstructing a phase image we examine reconstructions closely spaced over half a wavelength to determine which phase gives the best contrast for the features of interest. We then reconstruct a volume by stepping integer multiples of the wavelength from that distance, maintaining the same relative phase and therefore the same contrast.

As discussed in a paper by Sheng and coworkers[11], to a good approximation the effect of the lens is simply to rescale the lateral dimensions of the hologram by the objective magnification. We account for this effect by reducing the dimensions of the numerical pixels used in reconstruction by the same factor. Thus we reconstruct as if we had a detector with 0.34 µm pixels that lies at the focus of the objective. The reconstruction distance d is then measured from where the effective detector lies.

We extract RAW images from the camera to 24-bit, 3-channel TIFF files using the `dcraw` software package[40]. We use only the red channel in analysis because it has the highest sensitivity to IR illumination. In order to obtain a square image suitable for fast Fourier transforms, we extract a 1024×1024 or 2048×2048 pixel subimage centered on features of interest. We perform reconstructions on a desktop computer with a 3 GHz Intel Core 2 Duo CPU using custom software written in Python with the SciPy numerical library[41].

2.1.3 Sample preparation

We used a USAF 1951 resolution test target (Edmund Optics) to measure the resolution of the microscope. The target consists of line groups of decreasing size; the size of each group is given by a table from the manufacturer's website. We determine resolution by examining the reconstructed image and finding the smallest element where the lines are distinct.

We show images from two types of samples: microsphere suspensions and living cells. The microsphere samples consist of 2 µm Invitrogen sulfate latex particles in Knox Gelatine. We prepared gelatin with one gram of Gelatine powder per 40 mL of tap water, and added ∼100 µL of 8% w/v particles to ∼100 mL of gelatin solution. We made a thick sample (several hundred micrometers) using silicon vacuum grease (Dow Corning) as a spacer between a standard microscope slide and a number one cover slip. We flowed hot gelatin-particle solution into the gap by capillary action and then sealed the cover slip to the slide using five minute epoxy. The gelatin then cools and solidifies, preventing the particles from sedimenting or diffusing.

The living cell samples are wet mounts of cheek cells (Squamous Epithelium) obtained from swabbing the inside of the mouth. The cells from the swab were placed on a microscope slide, diluted in tap water, and covered with a number one cover slip. Because we did not use an isotonic buffer, we imaged the cells within thirty minutes of mounting, before they were lysed by the hypotonic medium.

2.2 Results and Discussion

Using the USAF resolution grid shown in Figure 2.2, we measured the resolution of our reconstructed images to be about 2.8 µm. Resolution in a holographic microscope is limited both by the diffraction limit of the microscope objective used and by the ability of the camera to detect and resolve closely spaced interference fringes. The theoretical resolution for our 0.4 NA objective is about 1 µm resolution with 785 nm light. Therefore the resolution appears to be limited by the number and quality of the interference fringes detected by the camera.

To investigate the effect of the camera resolution, we performed two reconstructions on the same hologram, one with a 1024×1024 pixel subimage (shown in Figures 2.2a and b) and another with a 2048×2048 subimage (shown in Figures 2.2c and d). In reconstructions based on the smaller subimage, group 6 element 3 is barely visible, giving a resolution of 6.2 µm. When the larger field is reconstructed, group 7 element 4 is visible, giving a resolution of 2.76 µm. This indicates that the camera can detect the high spatial frequency information outside the smaller subimage. Wavefront curvature does not appear to significantly affect the reconstructed image quality; the grid remains in focus throughout the field of view in the larger reconstruction.

The remainder of the information encoding spatial frequencies down to 1 µm is contained in fringes that either lie outside the reconstructed area or are too small, too blurry, or too faint to reconstruct accurately. This could be due to spherical aberration

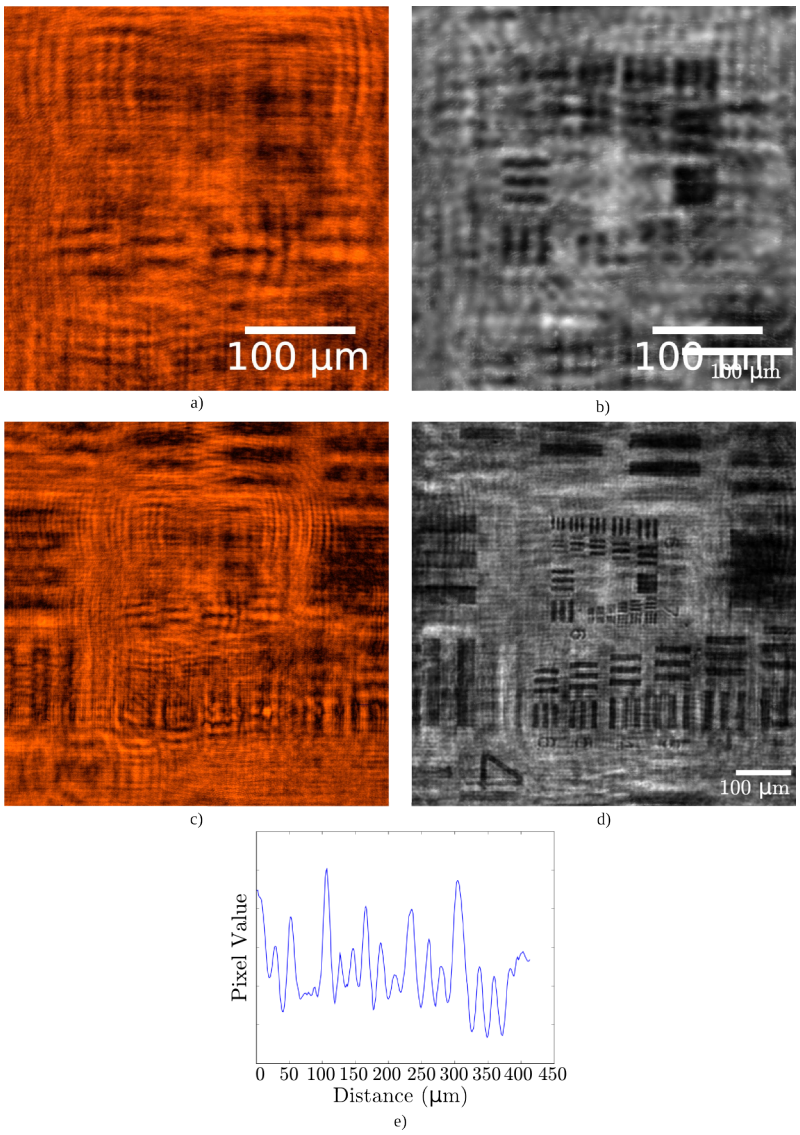

Figure 2.2: Reconstruction of a USAF 1951 Resolution target to establish resolution of the instrument. a) 1024×1024 section of the hologram. b) Reconstruction of (a) at a distance of 3.92569 mm. c) 2048×2048 section of the same hologram. d) Reconstruction of (c), also at 3.92569 mm. e) Intensity profile along a linear cross-section through group 7, showing that the lines in the first four elements are resolvable. This hologram was taken with 2.3 mW of laser power and ∼80 $\mu\mathrm{s}$ exposure. Author's Note: In this figure and throughout this thesis I have placed scale bars on reconstructions but not on holograms. Holograms are not images in the traditional sense and scale bars give a false impression that one can make straightforward estimates of object size by inspection.

from the imaging optics, to nonuniform illumination or to the limited size and resolution of the sensor. Given the sharpness of the reconstructed images and our observation that the resolution is not significantly improved by higher numerical aperture objectives, we believe that the resolution is currently limited by the pixel density of the camera. A recent professional-grade or a medium-format digital SLR camera, both of which are available commercially, may well be able to improve the instrument resolution at the same magnification.

To demonstrate the large depth of field made possible by holography, we image an ~ 200 µm thick sample containing 2 µm sulfate latex spheres embedded in gelatin. Figure 2.3a shows a hologram of this sample, and Figures 2.3b and c show reconstructions at two different planes 160 µm apart. The white arrows indicate particles that are in focus at the two different planes; the instrument is able to capture both of them in the same hologram. Fig 3d shows a volume reconstruction of this sample. The particles appear elongated due to an inherent effect of inline holography, as discussed by Sheng and coworkers[11]. This elongation increases for holograms of objects further from the focus of the objective. In our reconstructions the particles are elongated by a factor of about 20 at a distance of 3-4 mm from the focus of the objective, down to only a factor of 2 when near the focus. This variable elongation complicates quantitative axial measurements, but accurate transverse slices can still be obtained at any desired axial position.

Finally, we demonstrate that the instrument can image live cells. Figure 2.4 shows a

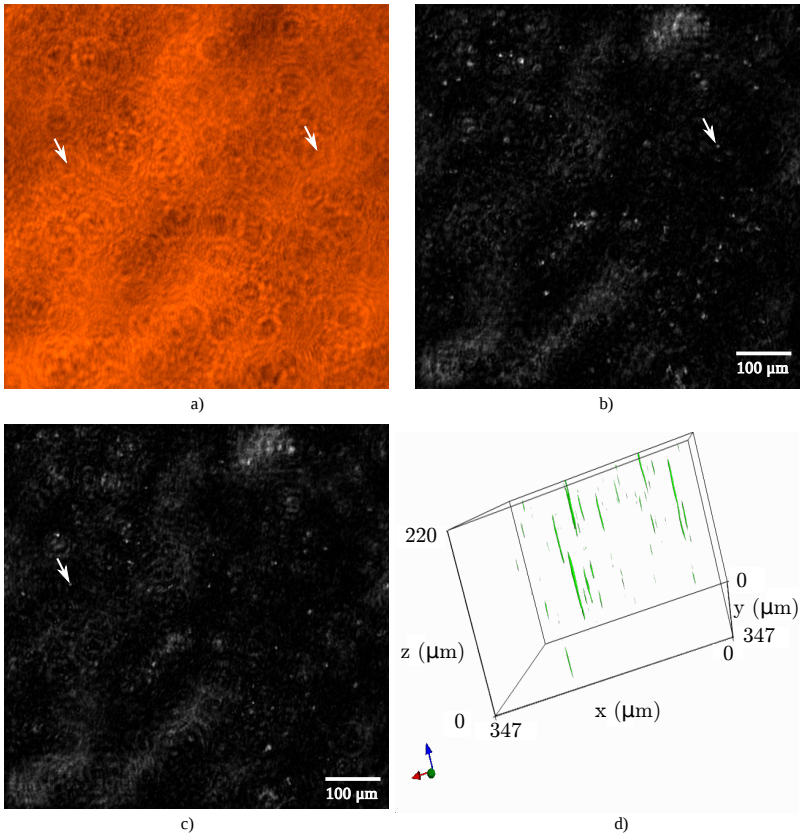

Figure 2.3: Sulfate latex beads in gelatin slab of thickness $\sim 200\ \mu\mathrm{m}$. a) Hologram. b) and c) Reconstruction at two planes separated by 160 $\mu\mathrm{m}$ (0.14 mm and 0.3 mm) in which different particles are visible. White arrows show two different particles separated axially by 160 $\mu\mathrm{m}$ in the reconstructed image. d) Isosurface volume reconstruction of this sample. The particles do not appear spherical because of elongation induced by inline holography. This hologram was taken with 0.235 mW of laser power and 20 ms exposure time.

hologram and reconstructions of squamous epithelium cells. We show both amplitude and phase images reconstructed from a 2048×2048 section of a hologram which contains 8 of these cells. Single cells can be visualized by zooming in on the reconstructed image. The nucleus is visible in both the amplitude and phase images, but the cell boundaries are clearer in the phase image. This is because the cell membrane and cytoplasm scatter weakly but are thick enough to affect the optical path length. Reconstruction at a different depth (Figure 2.4f) shows a plane below the nucleus, demonstrating the ability to optically section. Some subcellular features are more visible in the amplitude image. Because both amplitude and phase images are available from the same hologram, one can choose the image that gives better contrast for the subcellular features of interest.

2.3 Conclusions

We have demonstrated a simple, inexpensive, self-contained holographic microscope capable of imaging unstained live cells with a depth of field greater than one hundred microns, a resolution better than 3 microns, and a total exposure of 200nJ per 3D reconstruction. We obtain sufficient contrast and resolution to identify single cells and subcellular features. For many applications, such a simple holographic microscope can replace a much larger and more expensive instrument, while at the same time offering unique advantages such as customized time sequences of exposures and quantitative phase contrast. Our results suggest that the current consumer-grade digital cameras may be able to bring the resolution of the microscope close to the diffraction limit.

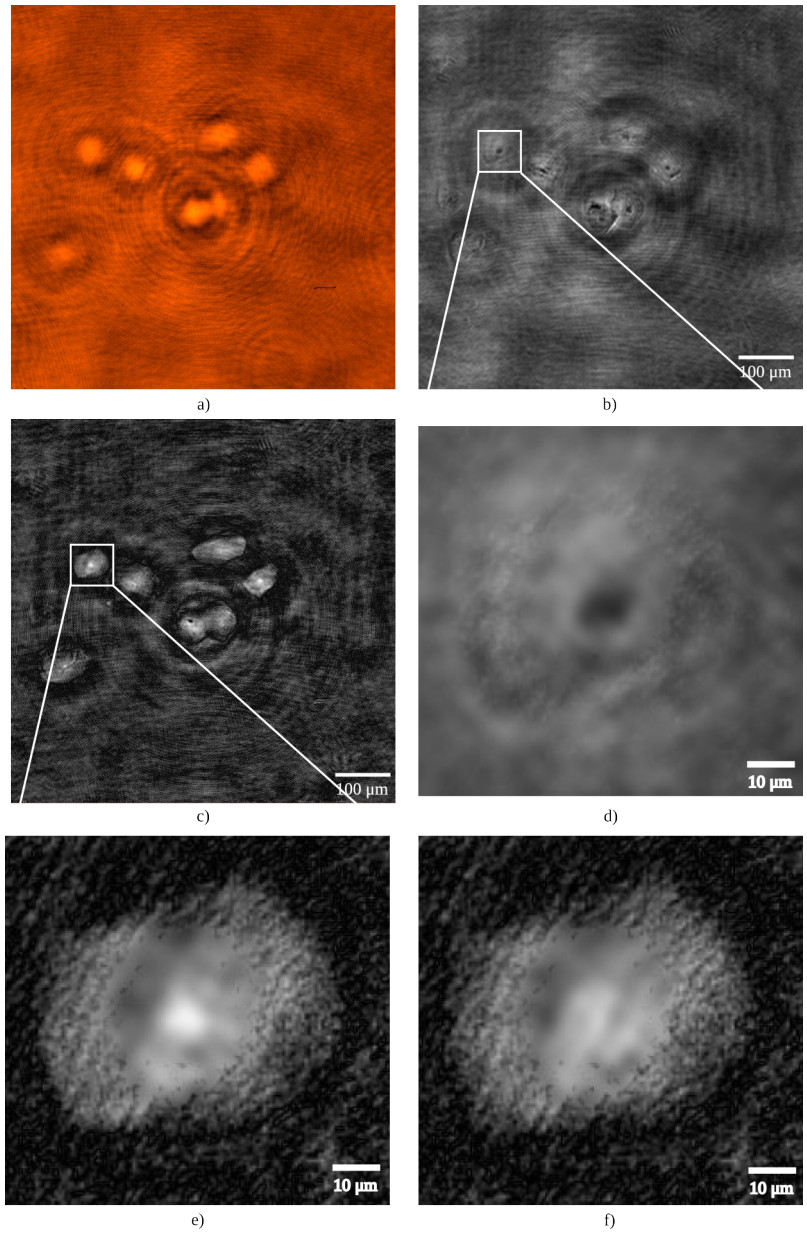

Figure 2.4: Squamous epithelium cells wet mounted and imaged approximately $f/2$ past the normal working distance of the objective. a) 2048×2048 region of hologram. b) Amplitude image reconstructed from (a) at 3.998 mm. c) Phase image reconstructed at 3.998 mm. d) Close up of boxed region from (b), showing nucleus (dark) and outer edge of cell. e) Close up of boxed region from (c), showing nucleus and edge of cell in phase contrast. f) Close up of boxed region from (c) at 4.055 mm, showing a region below the nucleus. This hologram was taken with 2.3 mW of laser power and ~ 80 µs exposure.

But the possibilities for such a microscope go beyond simply replacing widefield or confocal microscopes. Portable and inexpensive microscopes can enable entirely new kinds of experiments. For example, holographic microscopy may be especially well suited for long duration imaging of living cell growth and organization. Ordinarily such experiments are done by placing an environmental chamber around a bulky microscope. A holographic microscope could be built inside a cell culture chamber, where it can follow movement in three dimensions without refocusing while at the same time subjecting the cells to minimal light flux. It could function autonomously over many days, a period that is generally infeasible with an expensive microscope and environmental chamber. The large depth of field and the ability to optically section samples may enable new studies of three-dimensional cultures, an important new direction in understanding tissue and tumor growth[42]. Furthermore the short exposure time may allow the microscope to capture three-dimensional dynamics on much smaller time scales[21,43] are possible in confocal microscopy.

Holographic microscopy has not yet caught on in biology. We aim to lower the entry barrier for researchers without expertise in optics so that new experiments and new scientific discoveries become possible. To this end, the microscope we show here is only a first step; much remains to be done. In particular, the resolution needs to be improved to image smaller organelles and subcellular structures, the frame rate needs to increase, and reconstruction must happen in real-time. We believe that improvements in these areas will accompany the rapid evolution of consumer electronics.

Acknowledgments

We thank Jose A. Dominguez-Caballero, Nick Loomis, George Barbastathis, Jeff Lichtman, and Stan Leibler for helpful discussions. This research was funded by the National Science Foundation under CAREER award number CBET-0747625. This material is based upon work supported under National Science Foundation Graduate Research Fellowships (TGD and EJG). DK and RM were supported by Department of Defense NDSEG Fellowships.

3

Measuring non-spherical particle dynamics with the DDA and holography

Measurements of the dynamics of colloidal particles are key to understanding the mechanisms of colloidal aggregation[44] and self-assembly[45]. Microscopic measurements

of the diffusion of individual particles can furthermore be used to infer interactions between particles[46] as well as the local rheological properties of the medium in which the particles are suspended[47]. Most microscopic measurements are constrained to two dimensions because of the limited depth of field of wide-field microscopy[48]. Confocal microscopy can be used to capture three-dimensional (3D) dynamics of spherical and non-spherical particles[49,50], but the time required to scan the beam through a 3D sample limits these measurements to large particles or particles in a viscous fluid.

Holographic microscopy is an alternative technique that can in principle capture 3D colloidal dynamics with acquisition times orders of magnitude smaller than those of confocal microscopy. In a holographic microscope, light from a coherent source scatters from the sample and interferes with a reference wave—which can simply be the transmitted, undiffracted beam. The interference pattern, or hologram, contains phase information about the scattered wave and can be used to determine the 3D position of the particle. Because holograms can be captured as fast as a camera allows and then processed offline, the technique can be used to probe millisecond or even microsecond dynamics. The 3D information is typically recovered through optical or numerical reconstruction[51]. Cheong and co-workers[52] used this method to track the 3D translation and rotation of a high-aspect-ratio (1:25) copper oxide nanoparticle through numerical reconstruction and a skeletonization procedure. However, for wavelength-scale particles the precision of reconstructions is limited owing to distortions in the reconstructed volume[18,53].

We demonstrate a precise way to measure the dynamics of wavelength-scale, non-spherical colloidal particles from holographic measurements. Our approach, which uses an inverse-scattering analysis rather than reconstruction, follows the work of Lee and coworkers[21], who showed that fitting a Lorenz–Mie scattering solution to a measured hologram yields nanometer-scale-precision measurements of the positions of spherical particles. Researchers from our group, including Fung and coworkers[54,55] and Perry and coworkers[30] later showed that the translational, rotational, vibrational, and nonequilibrium dynamics of several interacting wavelength-scale spheres could be measured with high precision by fitting Mackowski and Mishchenko's T-matrix solution[56] to holograms of multiple colloidal particles. All of these measurements were limited to spheres, or collections of spheres, because the fitting technique requires a solution for the scattered field, and exact solutions are known for very few particle morphologies.

One can, however, perform approximate numerical scattering calculations for a wide variety of particles using the discrete dipole approximation (DDA) of Purcell and Pennypacker[32]. We show that numerical scattering calculations using the DDA can be fit to holograms of non-spherical particles, allowing us to track 3D translational and rotational dynamics at high precision, even for particles with subtle asymmetries. We use these measurements to determine translational and rotational diffusion coefficients and show that the measured values agree well with theoretical calculations, validating the technique. Although DDA calculations are orders of magnitude slower than calculations of exact scattering solutions, parallelization of the scattering calculations and the fitting

algorithm can significantly reduce the analysis time.

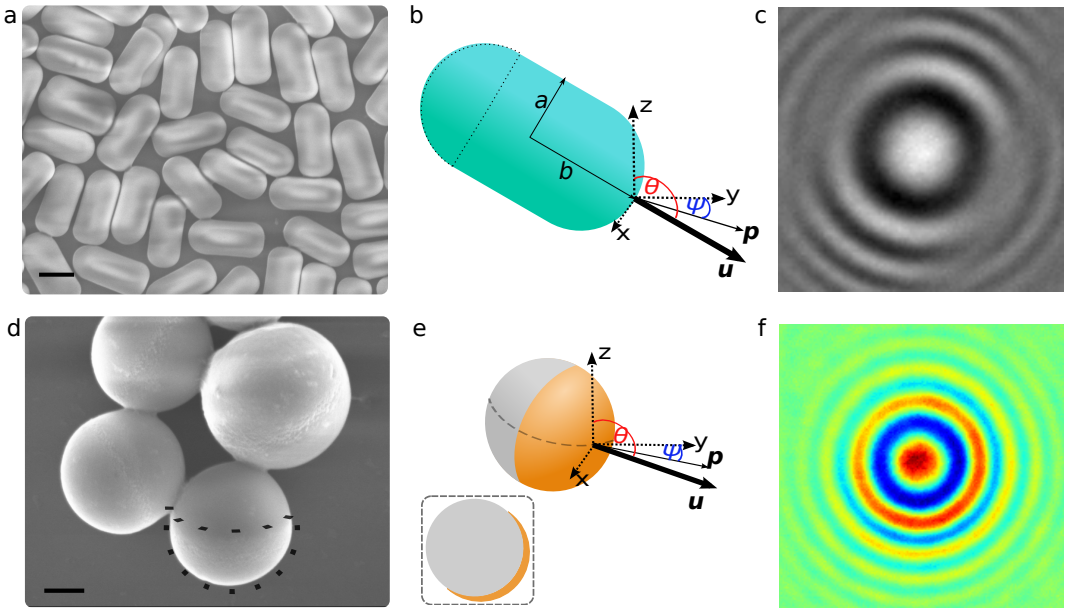

Figure 3.1: The non-spherical particles used in our experiments. Top row, silica rods; bottom row, Janus particles consisting of a polystyrene sphere coated with a 50-nm-thick layer of TiO_2. The orientation of both types of particles is defined by the orientational unit vector $\mathbf{u}$ and its projection onto the $x-y$ plane, $\mathbf{p}$. The z-axis coincides with the direction of the incident plane wave illumination. (a) Scanning electron micrograph of the rods. Scale bar is 1 μm. (b) The rod is modeled as a spherocylinder with semi-minor-axis length a and semi-major-axis length b. (c) Measured hologram of a single rod. (d) Scanning electron micrograph of the Janus particles. The dotted line outlines the TiO_2 cap. Scale bar is 300 nm. (e) The Janus particle is modeled as a plain polystyrene sphere with a hemispherical cap of TiO_2 (orange). A cross section perpendicular to the equator is also shown. (f) Measured hologram of a Janus particle, shown in false color to highlight asymmetry.

3.1 Experimental Methods

We use two types of anisotropic particles in this study: silica rods and polystyrene/TiO_2 Janus particles suspended in water (Figure 3.1). The silica rods are synthesized using a modified one-pot method recently reported by Kuijk and coworkers[57,58]. In the initial

growth step, we make silica rods with a length (L) of 1.45 ± 0.06 µm and diameter (D) of 0.29 ± 0.02 µm, which are then used as seeds in subsequent growth steps[59] to produce the final silica rods. The size of the rods, as determined by scanning electron microscopy (SEM), is 1.0±0.2 µm (minor axis) by 2.0±0.2 µm (major axis) (Figure 3.1a). We fabricate the Janus particles by depositing 50 nm of TiO_2 onto sulfate-terminated polystyrene (PS) particles (Invitrogen), as described in reference[60]. The polystyrene particle diameter, as determined by SEM, is 900 ± 100 nm (neglecting anomalously large particles). The TiO_2 cap covers approximately one hemisphere of the particle, as shown in Figure 3.1d.

For holographic imaging, we suspend the particles in deionized water at approximately 10^{-5} volume fraction and place them in sample cells consisting of two No. 1 glass coverslips (VWR) separated and sealed by vacuum grease (Dow Corning). To validate the DDA method, we also image a 0.95-µm-radius polystyrene sphere (Invitrogen) diffusing in a 54% v/v glycerol solution.

We capture holograms with a digital holographic microscope built on a Nikon (TE2000-E) inverted microscope, as shown in Figure 3.2 and described in detail in papers by Kaz and coworkers[4] and Fung and coworkers[54]. When imaging the Janus particles we use a 100× oil-immersion objective (Nikon) with a high numerical aperture (1.40) that allows us to maximize detail in the holograms. When imaging the silica rods, which sediment rapidly owing to their density, we use a 60× water-immersion objective (Nikon), which has a lower numerical aperture (1.20) but a larger working distance that allows us to

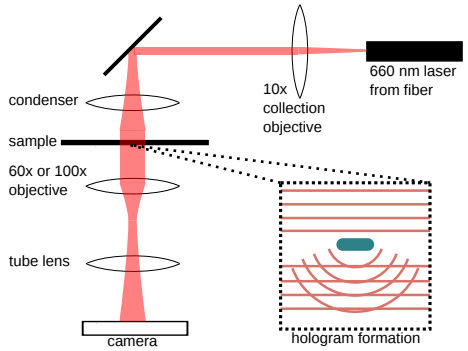

Figure 3.2: Schematic of our digital in-line holographic microscope. A series of lenses, including the condenser, shapes light from a 660 nm laser diode (Opnext HL6545MG) into a plane wave to illuminate the sample. An objective collects the transmitted light and light scattered from the sample, and a high-speed camera (Photon Focus) captures the hologram formed by the interference between scattered and transmitted waves.

capture longer trajectories. To obtain true bulk-diffusion measurements, we retain data only from particles that remain at least 20 µm away from the sample cell boundaries throughout the trajectory.

In a typical experiment we use 50 mW of laser power and a 0.015 ms exposure time per camera frame for the 60× lens or 0.05 ms for the 100× lens. These exposure times are short enough to minimize blurring due to Brownian motion. We capture holograms with a Photon Focus MVD-1024E-160 camera at 100 frames per second, store them in RAM using a frame grabber (EPIX PIXCI E4), and then save to disk for further processing. For each trajectory we also record a background image from the same region of the sample cell before or after the particle is present to account for scattering and illumination defects in the optical train.

3.2 Fitting holograms using the DDA

In contrast to holograms of spheres, holograms of non-spherical particles do not have azimuthal symmetry about the central maximum (see Figure 3.1c, f). The azimuthal asymmetry encodes information about the particle shape and orientation, while the radial spacing of the interference fringes encodes the particle position. To quantitatively extract this information, we fit a scattering model to the holograms.

In our fitting procedure, we compute holograms from a model, compare the computed holograms to experimentally recorded ones, and iterate until the model matches the data. To compute the hologram, we model the physical process of hologram formation: the interference of scattered and reference fields. In an inline hologram, the reference field is approximately equal to the incident field $\mathbf{E}_i$, as long as the density of scatterers is low. In our experiments, we work at low concentrations of particles, so that we can assume $\mathbf{E}_i$ is a constant plane wave. The observed intensity is therefore

$$I_{\text{holo}} = |\mathbf{E}_i + \mathbf{E}_s|^2, \tag{3.1}$$

where $\mathbf{E}_s$ is the scattered field, which we must compute to simulate the hologram. Previous work[4,21,30,54] used numerical implementations of the Lorenz–Mie solution or a multi-sphere superposition solution[56] to compute $\mathbf{E}_s$. However, these exact solutions are limited to spheres or collections of spheres. For the non-spherical particles used here,

we must compute $\mathbf{E}_s$ approximately. We do this using the discrete dipole approximation, as implemented by the open-source scattering code ADDA[61].

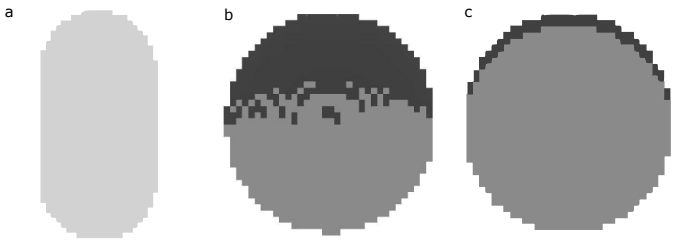

Figure 3.3: Voxelations of the particles. White represents vacuum, light gray is silica, medium gray is polystyrene and dark gray is TiO_2. Each dipole is represented by a square. (a) Side view of a rod with a dipole size of 44 nm, a semi-minor-axis length of 500 nm and a semi-major-axis length of 1 µm. (b) Side view of a Janus particle with a dipole size of 24 nm, inner particle radius of 439 nm and a maximum cap thickness of 50 nm. (c) Cross section of the same Janus particle.

In the DDA, a scatterer is represented as an array of point dipoles. Therefore the first step of our data analysis is to discretize the particle, dividing it into an array of voxels (volumetric pixels). To ensure validity of the DDA, we discretize the scatterer using at least 10 dipoles per wavelength in the scattering medium, as recommended by the ADDA documentation[61]. The voxelations of both particles are shown in Figure 3.3. For the Janus particle, the dipole size is comparable to the thickness of the TiO_2 layer; but because the shape and size of the cap varies from particle to particle, a smaller dipole size would not necessarily increase accuracy.

To compute a hologram, we calculate the scattering angles for each pixel in the detector, then invoke ADDA to compute the scattering matrix of the voxelated particle at each scattering angle. We obtain the electric field $\mathbf{E}_s(i,j)$ at each pixel (i,j) on the detector from the scattering matrices, angles, and distances. Finally, we numerically

interfere the computed $\mathbf{E}_s(i,j)$ with $\mathbf{E}_i(i,j)$, a plane wave, to obtain the hologram. All of these steps are implemented in our open-source hologram processing code, HoloPy[33].

Once we have created a DDA scattering model for a particle, we use nonlinear minimization to fit the model to measured, background-divided, normalized holograms, thereby obtaining measurements of the positions, refractive index, and geometrical parameters of individual particles. Formally, we use the Levenberg–Marquardt algorithm to minimize the objective function

$$f(\{p\}) = \sum_{i,j} \left| I_{\text{measured}}(i,j) - I_{\text{computed}}\left(i,j;\{p\}\right) \right|^2 = \chi^2 \tag{3.2}$$

where

$$I_{\text{measured}} = \frac{I_{\text{data}}}{I_{\text{background}}} \frac{\overline{I}_{\text{background}}}{\overline{I}_{\text{data}}},$$

and $\{p\}$ is the set or a subset of the parameters used by the DDA model to calculate the scattering from the particle. The minimization typically requires several iterations to converge. We then use the fitted position and orientations obtained from a time-series of holograms to compute the diffusion coefficients.

Because the two types of particles have different geometries, we must use different sets of parameters to model them. We model the rod, which has rounded ends as shown in Figure 3.1a, as a dielectric spherocylinder with refractive index n, hemisphere radius a (equal to the length of the semi-minor-axis), and semi-major-axis length b. The full

set of parameters is

$$\{p_{\text{rod}}\} = \{\mathbf{r}, n, a, b, \theta, \psi, \alpha\}, \tag{3.3}$$

where $\mathbf{r}$ is the center of mass position, θ and ψ are Euler angles describing particle orientation (see Figure 3.1b), and α is a normalization constant[21]. We model the Janus particle (Figure 3.1d) as a dielectric sphere with refractive index n_{PS} and radius a_{Janus} capped by a hemispherical dielectric shell with index n_{TiO_2} and thickness t. Here the set of parameters is

$$\{p_{\text{Janus}}\} = \{\mathbf{r}, n_{\text{PS}}, n_{\text{TiO}_2}, a_{\text{Janus}}, t, \theta, \psi, \alpha\}, \tag{3.4}$$

where $\mathbf{r}$ is the center of the sphere, θ and ψ are the Euler angles shown in Figure 3.1e, and α is again a normalization constant.

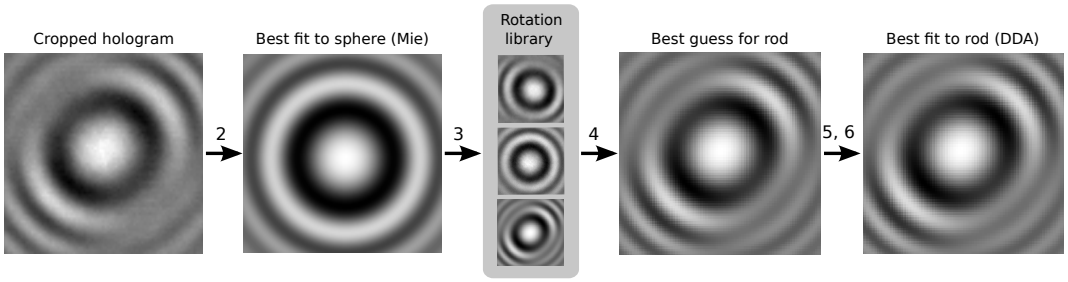

Figure 3.4: Procedure for fitting DDA models to experimental holograms of non-spherical particles to extract position and orientation. The numbers refer to the steps in the text.

To assess the validity of the fits, we examine the coefficient of determination, R^2, which measures how much of the variation from the mean value of one hologram is

captured in another:

$$R^2 = 1 - \frac{\sum_{i,j} |I_1(i,j) - I_2(i,j)|^2}{\sum_{i,j} \left|I_1(i,j) - \bar{I}_1\right|^2}, \tag{3.5}$$

where $\bar{I}_1 = 1$ is the mean value of the normalized hologram. R^2 is a useful quantity because unlike χ^2, which is sensitive to the amplitude of the hologram fringes, the R^2 values for good fits do not vary markedly across experimental systems with different scattering cross sections[31]. We use R^2 to assess how well DDA reproduces the results of exact Lorenz–Mie calculations for holograms of spheres and to assess how well the holograms from our model match data from experiments.

Because DDA calculations take at least an order of magnitude longer than Lorenz–Mie calculations for scatterers of the same size, we aim to fit all the frames in parallel. This goal is complicated by convergence requirements. In practice, we have found that the minimization does not converge unless the initial guess for $\{p\}$ is close to the global minimum of the objective function. Therefore our usual procedure[54,55] is to fit the frames sequentially: we use reconstruction to manually obtain an initial guess for the first frame of the trajectory and then use the best-fit result for each frame as the initial guess for the next. Here such a process would take too long. A DDA calculation takes on the order of 30 seconds on a single CPU; since about 200 such calculations are required for each minimization to converge, it would take 20 weeks to sequentially fit a 2000-frame trajectory. We therefore use the procedure described below and illustrated in Figure 3.4 to obtain the initial guesses for each frame and fit the frames in parallel:

1. *Select the most asymmetric region:* The particle orientation is primarily encoded in the rotational asymmetry of a hologram. We maximize sensitivity to the particle orientation by cropping the hologram to leave only the most asymmetric region, the first fringe or two around the central maximum. We use a Hough transform-based algorithm[62] to locate the center of the hologram. For our experiments the first fringe is located within a 140×140 pixel area.

2. *Find approximate position with a Lorenz–Mie fit:* We obtain an estimate of the particle's position by approximating it as a sphere with parameters $\{p\} = \{\mathbf{r}_{\text{Mie}}, n, a_{\text{Mie}}, \alpha\}$ and fitting the Lorenz–Mie solution to the hologram. This fit is very fast (order of 10 seconds per frame) and although χ^2 is large, $\mathbf{r}_{\text{Mie}}$ provides a good initial guess for the position of the particle's center.

3. *Make rotation library:* We use the DDA to calculate holograms of the non-spherical particle over a range of Euler angles θ, ψ, assuming a center position $\mathbf{r}_{\text{Mie}}$.

4. *Use library to guess angles:* We compare this library to the experimental hologram and determine the Euler angles θ, ψ that yield the highest R^2 value.

5. *Initial DDA fit:* We fit the DDA model to the measured hologram using the guess for $\mathbf{r}_{\text{Mie}}$ from step 2, θ, ψ from step 4, bulk values for the index or indices, and an estimate of the geometrical parameters a and b or a_{Janus} and t from SEM images. We allow all parameters to vary during the fit. In this step, we fit the first 500 frames of a trajectory, discard the fits that fail to converge, then calculate the average best-fit refractive indices and radii.

6. *Final DDA fit:* Because the particle's physical properties should not change across a trajectory, we fix the refractive indices and radii to the average values from step 5 and fit the cropped holograms for an entire trajectory to obtain the position and orientation as a function of time.

This procedure removes the serial dependency between the frames in a time series, allowing us to send each frame in the initial (step 5) and final (step 6) DDA fits to an individual CPU on a computing cluster. However, removing the serial dependency

potentially introduces artifacts in the trajectories of the particles, since the results from each frame are not used to constrain the results of the next. We must therefore detect and correct spurious fit results.

There are two kinds of errors. The first arises because the rod particles are symmetric, having two degenerate orientational unit vectors pointing in opposite directions. In a trajectory constructed from a series of parallel fits, the fitted orientational unit vector may therefore flip almost 180° between frames. Using the probability density function for rotational displacements described in reference[55] and an estimated rotational diffusion coefficient $D_{\text{est}} = 0.3\ \text{s}^{-1}$, we estimate that the probability of a 90° or greater change in angle between frames is zero, to within machine precision. Therefore we can correct artificial flips by examining pairs of frames in the trajectory and flipping the vector in the second frame by 180° if the angle changes by more than 90° between frames.

The second type of spurious fit result arises because the Levenberg–Marquardt algorithm is sensitive to noise. Depending on the initial guess, a fit can converge to a secondary minimum in which the orientational unit vector $\mathbf{u}$ points in a direction mirrored about the x-y plane from that of the previous frame. A small number of these bad fits systematically increases the apparent diffusion coefficient.

We therefore reject contributions to the mean-square displacement where the probability of the measured angular displacement is less than 0.0001. Again, we calculate the probability of observing a given displacement using the probability density function for

rotational displacements described in reference[55] and an estimated rotational diffusion coefficient $D_{\text{est}} = 0.3\ \text{s}^{-1}$. Using $D_{\text{est}} = 0.2\ \text{s}^{-1}$ or $D_{\text{est}} = 0.4\ \text{s}^{-1}$ changes the coefficients extracted from the data by less than 7%, an effect which we include in the error on the measured diffusion coefficients. This procedure eliminates poor fit results while minimizing the bias of results toward our estimated diffusion coefficient.

3.3 Results and Discussion

We demonstrate the validity and usefulness of the DDA as a tool for holography and particle tracking through a series of experimental and computational tests. First we compare the results of DDA and Lorenz–Mie calculations for scattering from spherical particles, which the Lorenz–Mie solution models exactly. Next we fit the DDA model to holograms of a rod and a Janus particle, and we obtain translational and rotational diffusion coefficients from the mean-square displacements measured across time-series of holograms. We compare the measured diffusion coefficients to predicted values to demonstrate that we accurately and precisely capture the particle motion.

3.3.1 Validation of DDA calculations and fits

We first examine how the approximations involved in our DDA implementation affect the quality of the fits and calculated holograms. To do this, we compare holograms of spheres calculated with our DDA model to those calculated from the Lorenz–Mie solution. We are interested in two effects: first, how aliasing introduced by the voxelation affects the

precision of fitting holograms and, second, how the absence of a near-field calculation in ADDA affects the accuracy of the holograms. By "near-field" we mean the part of the scattered field that does not follow an asymptotic $1/r$ decay. ADDA calculates only the asymptotic portion of the scattered field. In previous work[31,54] we found that it was necessary to include the full radial dependence of $\mathbf{E}_s$ to accurately fit holograms of particles close to the focal plane.

We test the validity of the scattering model, including our voxelation scheme, by fitting the DDA model and the Lorenz–Mie solution to 893 measured holograms of a 0.95-µm-radius polystyrene sphere (Invitrogen) freely diffusing in a 54% v/v glycerol solution. We fix both the refractive index and radius of the particle and allow only the center position of the sphere to vary. Throughout the trajectory, the particle remains 15 µm to 18 µm from the focal plane.

We find that the best-fit x- and y-coordinates obtained by fitting the DDA model differ negligibly (0.3 ± 0.1 nm) from those obtained from Lorenz–Mie, while the best-fit z-coordinates differ by 194 ± 1 nm, a significant offset. This difference in apparent axial position might arise from voxelation errors or the asymptotic approximation for the radial dependence of the scattered field. The Lorenz–Mie model includes non-asymptotic corrections for the scattered field, while the DDA model does not.

The offset in the z-coordinate does not affect our dynamical measurements, which depend only on the displacement of the particle, not its absolute position. We find that the total displacements $\|\mathbf{r}(t) - \mathbf{r}(0)\|$ for the DDA and Lorenz–Mie models, as

calculated from the best-fit coordinates for the entire trajectory, agree to within 0.5 nm (Figure 3.5a) for most of the frames analyzed. The maximum deviation in displacement is 2 nm, which is within the reported accuracy of holographic measurements based on inverse-scattering analysis[21]. Thus the systematic error in the absolute particle position appears to cancel when calculating the displacement, and the resulting displacement measurements are nearly as precise as those obtained by fitting the data to the exact Lorenz–Mie solution.

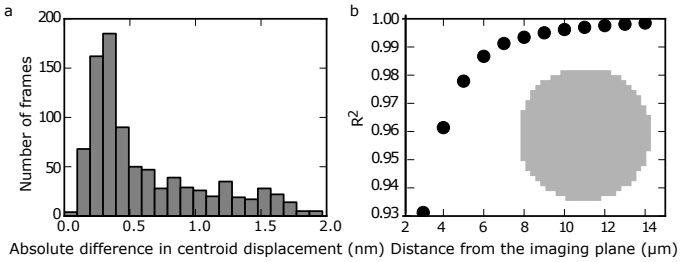

Figure 3.5: Comparison of the DDA model and Lorenz–Mie solution for fitting holograms of a single sphere. (a) Histogram of differences between displacements of a 950-nm-radius diffusing sphere measured by fitting the DDA and Lorenz–Mie models to measured holograms. (b) R^2 for pairs of holograms calculated from the exact Lorenz–Mie solution and the DDA model of a 450-nm-radius sphere at various distances from the imaging plane. Inset: voxelation of the 450-nm-radius sphere used in the DDA model. The voxel size is 34 nm. The particle refractive index is 1.585.

To further examine the impact of the asymptotic approximation for the scattered field, we compare holograms computed from the DDA model and Lorenz–Mie solutions for a 450-nm-radius sphere as a function of distance from the focal plane. The agreement between the two models, as measured by the R^2 value, increases with distance from the imaging plane, as shown in Figure 3.5b. These data suggest that near-field effects start to vary significantly with distance when the particle is 5 μm or less from the focal plane.

To ensure the validity of our DDA fits, we take all measurements with the particle at least 5 µm from the imaging plane.

As another check on the validity of the technique, we examine the average of fitted values for the optical properties and sizes of our non-spherical particles. We find $n = 1.495 \pm 0.012$ for the rod, $n_{\mathrm{PS}} = 1.581 \pm 0.046$ for the polystyrene particle and $n_{\mathrm{TiO_2}} = 2.74 \pm 0.23$ for the shell of the Janus particle, all of which are close to bulk values for the materials. The best-fit dimensions of the rod are $2a = 1.002 \pm 0.037$ µm and $2b = 2.158 \pm 0.143$ µm, in agreement with the measurements from SEM. The fitted dimensions of a Janus particle are $a_{\mathrm{Janus}} = 443 \pm 20$ nm, in excellent agreement with the SEM images, and $t = 47 \pm 8$ nm, in good agreement with the expected value based on the TiO_2 deposition process.

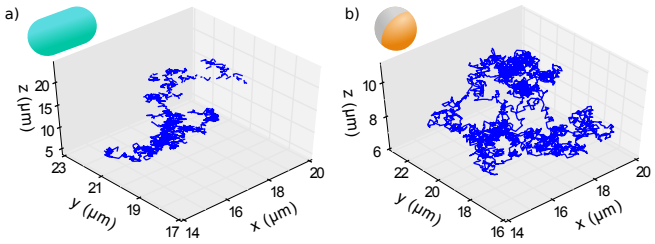

Figure 3.6: Trajectories of particles as determined from DDA fits. (a) Rod trajectory captured at 100 frames per second for 20 seconds. (b) Janus particle trajectory captured at 100 frames per second for 30 seconds.

Qualitatively, the trajectories of the particles we obtain by fitting the DDA model to time-series of holograms are consistent with stochastic Brownian motion (Figure 3.6). Furthermore, the rod appears to sediment at 0.9 µm/s whereas the Janus particle diffuses about a steady height. This behavior agrees qualitatively with our expectations for

the two systems: the rod should sediment more rapidly because the density difference between the rod and solvent is about three times larger than that of the Janus particle, and the rod has about three times the volume of the Janus particle.

Finally, to quantitatively investigate the validity of the DDA fits, we determine the rotational and translational diffusion coefficients from the data and compare them to theoretically predicted values, as described in the following sections.

3.3.2 Rods

When viewed in the laboratory frame, the motion of a rod is complicated because the viscous drag along any laboratory axis depends on the orientation of the particle. But in the frame of the particle, translational diffusion can be broken down into two components, one parallel to the long axis of the particle ($\parallel$) and the other along either of the two principal axes perpendicular to the long axis ($\perp$). Due to the optical symmetry of the particle, we cannot observe rotations about the major axis. We can therefore measure only one rotational diffusion coefficient, corresponding to rotation about either of the two degenerate principal axes that are coincident with the minor axis.

To measure these diffusion coefficients, we use the best-fit positions obtained from the holograms to calculate the translational and angular mean-square displacements in the particle frame as a function of lag time. We average these displacements over time and calculate the uncertainties associated with these averages using a block decorrelation method[63]. Finally, we fit the mean-square displacements using the calculated uncer-

tainties as weights. The uncertainties are represented as error bars in the figures. In the following sections we report the formulas used to calculate and fit the displacements, as well as our measured and predicted values for the diffusion coefficients.

Theoretical predictions

For clarity, we denote all predicted values of diffusion coefficients with a prime symbol ($'$). The diffusion coefficients of a spherocylinder can be expressed as power series in the aspect ratio $\omega = b/a$, as discussed in the appendix of a paper by Marchenko and coworkers[64], which is adapted from work by Yoshizaki and Yamakawa[65] and Norisuye and coworkers[66]:

$$D_t' \approx \frac{k_B T}{6\pi\eta b}\big(\ln\omega + 0.3863 + 0.6863/\omega - 0.0625/\omega^2 - 0.01042/\omega^3 - 0.000651/\omega^4 + 0.0005859/\omega^5\big) \quad (3.6)$$

where $D_t = (2D_\perp + D_\parallel)/3$. We compare our results to the predicted value for D_t, which is a linear combination of the measured diffusion coefficients. From the same reference, we obtain the rotational diffusion coefficient:

$$D_r' \approx \frac{3k_B T}{8\pi\eta b^3}\left(\ln\omega + 2\ln 2 - \frac{11}{6} + \frac{\ln 2}{\ln(1+\omega)}\left[\frac{1}{3} - 2\ln 2 + \frac{11}{6} - |\mathbf{a}|\right] + \mathbf{a}\cdot\boldsymbol{\Omega}\right) \quad (3.7)$$

where

$$\mathbf{a} = [13.04468, -62.6084, 174.0921, -218.8365, 140.26992, -33.27076],$$

and

$$\mathbf{\Omega} = [\omega^{-1/4}, \omega^{-2/4}, \omega^{-3/4}, \omega^{-4/4}, \omega^{-5/4}, \omega^{-6/4}].$$

Because the rods are slightly asymmetric—flatter on one end of the rod than the other (see Figure 3.1a)—we also compare our results for the translational and rotational diffusion coefficients to theoretical predictions for a cylinder with flat ends[67]:

$$D'_{\perp} \approx \frac{k_B T}{8\pi\eta b}\left(\ln\omega + 0.839 + 0.185/\omega + 0.233/\omega^2\right) \tag{3.8}$$

$$D'_{\parallel} \approx \frac{k_B T}{4\pi\eta b}\left(\ln\omega - 0.207 + 0.980/\omega - 0.133/\omega^2\right) \tag{3.9}$$

$$D'_{r} \approx \frac{3k_B T}{8\pi\eta b^3}\left(\ln\omega - 0.662 + 0.917/\omega - 0.050/\omega^2\right). \tag{3.10}$$

We use the laboratory temperature and uncertainty, $21 \pm 2°\mathrm{C}$, the corresponding viscosity for water, and the dimensions of the particle from the hologram fits to calculate the predicted diffusion coefficients and their uncertainties. As the values of rod's fitted dimensions are not normally distributed, we do not include them when propagating the errors.

Translational diffusion

To extract translational diffusion coefficients from the data, we first calculate the mean-square displacements parallel and perpendicular to the major axis. For short lag times τ we calculate

$$\Delta r^2{}_{\parallel}(\tau) = \left\langle ((\mathbf{r}(t+\tau) - \mathbf{r}(t)) \cdot \mathbf{u}(t))^2 \right\rangle = 2D_{\parallel}\tau + 2\epsilon_{\parallel}^2 \tag{3.11}$$

$$\Delta r^2{}_{\perp}(\tau) = \left\langle |(\mathbf{r}(t+\tau) - \mathbf{r}(t)) \times \mathbf{u}(t)|^2 \right\rangle = 4D_{\perp}\tau + 4\epsilon_{\perp}^2 \tag{3.12}$$

where the angle brackets denote a time average over all contributing pairs from a single trajectory. We then determine the diffusion coefficients $D_{\parallel}$ and $D_{\perp}$ and measurement errors $\epsilon_{\parallel}$ and $\epsilon_{\perp}$ by fitting a linear model to the measured mean-square displacements.

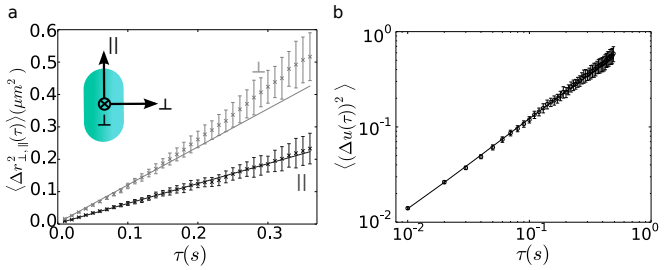

Figure 3.7: Measured mean-square displacements for the silica rods. (a) Mean-square displacements parallel and perpendicular to the major axis. The solid lines are fits to $2D_{\parallel}\tau + 2\epsilon_{\parallel}^2$ and $4D_{\perp}\tau + 4\epsilon_{\perp}^2$. (b) Mean-square displacement of the orientational unit vector $\mathbf{u}$. The solid line is a fit to the function $2 - 2\exp(-2D_r\tau) + 2\epsilon_r^2$.

Our measured results agree with those predicted by the spherocylinder model. We obtain $D_{\parallel} = 0.307 \pm 0.028$ µm^2/s, $D_{\perp} = 0.293 \pm 0.025$ µm^2/s and $D_t = 0.298 \pm 0.026$

$\mu m^2/s$ (see Figure 3.7 and Table 3.1). These values are larger than the predictions for a cylinder with flat ends ($D'_{\parallel} = 0.301 \pm 0.013$ $\mu m^2/s$ and $D'_{\perp} = 0.265 \pm 0.012$ $\mu m^2/s$) but agree quantitatively with the spherocylinder model ($D'_t = 0.297 \pm 0.016$ $\mu m^2/s$), suggesting that the rods are better modeled as spherocylinders.

The fit, which is constrained by the small uncertainties at shorter lag times, falls outside the error range for the perpendicular direction at larger lag times (Figure 3.7). However, the fit for the parallel component is well within the error bars throughout the entire range of lag times. The apparent augmented motion in the perpendicular direction may be due to radiation pressure from the incident beam or coupling between sedimentation and diffusion[68].

Quantity	Experiment	Expected	Source
a (nm)	501 ± 19	500 ± 50	SEM
b (nm)	1079 ± 72	1000 ± 100	SEM
n	1.495 ± 0.012	≤ 1.54	bulk value
D_t ($\times 10^{-13}$ $m^2 s^{-1}$)	0.298 ± 0.026	0.297 ± 0.016	spherocylinder model
D_r (s^{-1})	0.311 ± 0.034	0.311 ± 0.017	spherocylinder model
$D_{\parallel}/D_{\perp}$	1.04 ± 0.05	1.13	cylinder model

Table 3.1: Measured and expected values of parameters for silica rods. Theoretical predictions for the spherocylinder model are based on relations in reference[64], and those for the cylinder model on relations in reference[67].

The measured ratio $D_{\parallel}/D_{\perp} = 1.05 \pm 0.04$ depends only on the aspect ratio of the particle. Our value is smaller than that predicted by the cylinder model ($D'_{\parallel}/D'_{\perp} = 1.13$). The small discrepancy in the measured and predicted values of $D_{\parallel}/D_{\perp}$ is not surprising, given how much the measured $D_{\parallel}$ and $D_{\perp}$ values differ individually from those predicted

by the cylinder model. We are not able to calculate $D_\parallel/D_\perp$ for the spherocylinder model because we do not have explicit expressions for the two individual components. In general, however, we expect $D_\parallel/D_\perp$ to be of order unity for particles that have an aspect ratio of 1:2. Therefore our measured value is physically reasonable.

From the linear fit to the mean-square displacement, we find that the precision to which we track the rod's center of mass is 35 nm or better ($\epsilon_\parallel = 28.3$ nm and $\epsilon_\perp = 35.4$ nm), which is around the same size as one voxel (approximately 35 nm). Though about an order of magnitude lower than the precision to which single spheres can be tracked, this precision compares favorably with that of tracking more complex scatterers, such as spheres in clusters[30,54].

Rotational diffusion

To determine the rotational diffusion coefficient, we first calculate the mean-square displacement of the orientational unit vector $\mathbf{u}$:

$$\Delta\mathbf{u}^2(\tau) = \left\langle (\mathbf{u}(t+\tau) - \mathbf{u}(t))^2 \right\rangle = 2\left(1 - e^{-2D_r\tau}\right) + 2\epsilon_r^2. \tag{3.13}$$

We then fit $\Delta\mathbf{u}^2(\tau)$ to find the diffusion coefficient D_r and its angular uncertainty ϵ_r.

We find $D_r = 0.311 \pm 0.034$ rad^2/s (Figure 3.7), which is larger than that predicted by the cylinder model ($D_r' = 0.205 \pm 0.008$ rad^2/s) but in quantitative agreement with the spherocylinder model ($D_r' = 0.311 \pm 0.017$ rad^2/s). The error obtained from the fits

is $\epsilon_r = 0.0265$ rad, yielding an angular tracking precision of 1.5°. Like the translation results, these results suggest that the rods are modeled well, both hydrodynamically and optically, as spherocylinders. Furthermore the tracking precision is better than that reported for nanorods using reconstructions (approximately 3°) and dimers using multisphere superposition solutions (3.4°)[52,54].

3.3.3 Janus particles

We find that it is more difficult to fit our scattering model to the Janus particle holograms, as they are only weakly asymmetric. In particular, we find that the best-fit polar angle θ is often spurious. There are two local minima in the objective function, corresponding to two polar angles reflected about the x-y plane, and two different z-coordinates: if we calculate a hologram of a Janus particle with $\theta = 0$ and $z = 8.0$ µm, and we fit to it using an initial guess of $\theta \geq \pi/2$, the fit converges to $\theta = \pi$ and $z = 7.811$ µm. We find an $R^2 = 0.9997$ and a per-pixel χ^2 of 3.8×10^{-6}, showing that the best-fit and original holograms are essentially identical. These calculations suggest that errors in the best-fit polar angle are correlated with those of the z-coordinate. Noise in the hologram could therefore cause the fitting algorithm to converge to either local minimum. Indeed, we estimate our noise floor for the 8-bit images to be at least an order of magnitude larger at $(1/255)^2 = 1.5 \times 10^{-5}$ [55]. Increasing the asymmetry of the hologram, for example by increasing the aspect ratio or using a metal-coated Janus particle as in Anthony and coworkers' 2006 paper[48], should eliminate this problem. But

for the results shown below, we determine dynamical data only from the best-fit x- and y-coordinates and the azimuthal angle ψ.

Theoretical predictions

We use the Stokes-Einstein and Stokes-Einstein-Debye relations[69,70] to model the translational and rotational diffusion of the Janus particles, which are approximately spherical:

$$D_t' = k_B T / 6\pi\eta a \tag{3.14}$$

$$D_r' = k_B T / 8\pi\eta a^3. \tag{3.15}$$

Because the hydrodynamic radii of polymer particles—like the PS core in our Janus particle—tend to be larger than the radii measured optically or with electron microscopy due to charged or hairy surfaces[71,72], we do not directly compare the measured diffusion coefficients to values predicted from the theory. Instead, we calculate the effective radius a_{eff} from both the measured translation and rotational diffusion coefficients using Equations (3.14) and (3.15). We then compare these values to the best-fit diameter from the holograms and to each other.

Translational diffusion

To determine the translational diffusion coefficients for the Janus particle, we ignore the translational motion in the direction parallel to the imaging axis (z-axis) because

the polar angle θ has an uncertainty that affects the best-fit z-position, as discussed previously. We treat the Janus particle as a sphere and calculate the translational diffusion coefficient from the mean-square displacement projected onto the $x - y$ plane:

$$\left\langle \Delta x^2(\tau) + \Delta y^2(\tau) \right\rangle = 4D_{\perp,\text{Janus}}\tau + 4\epsilon^2 \tag{3.16}$$

where ϵ is the tracking precision.

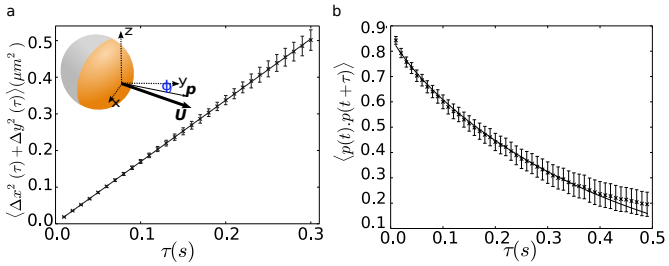

Figure 3.8: (a) Measured mean-square displacement of the Janus particle in the axes perpendicular to the imaging axis. The solid line is a fit to $4D\tau + 4\epsilon^2$. (b) Measured autocorrelation of the particle's projected orientational unit vector $\mathbf{p}$. The solid line is a fit to an multi-exponential decay.

By fitting the mean-square displacement to a linear model, we obtain $D_{\perp,\text{Janus}} = 0.419 \pm 0.003\ \mu\text{m}^2/\text{s}$ (see Figure 3.8). We then use Equation 3.14 at $21 \pm 2°\text{C}$ to obtain an effective particle radius $a_{\text{eff}} = 524 \pm 15$ nm, which is larger than the optical radius from holography (443 ± 20 nm for bare polystyrene, 490 ± 20 nm with the TiO_2 layer). As discussed above, the larger hydrodynamic radius is expected for polymer particles[71,72]. These results are summarized in Table 3.2.

We obtain a tracking precision of $\epsilon = 21.8$ nm from the fits for the directions perpendicular to the imaging axis, which is smaller than the size of one voxel (approximately 35

Quantity	Experiment	Expected	Source
a_{Janus} (nm)	443 ± 20	450 ± 50	SEM
t (nm)	47 ± 8	50	deposition parameters
n_{TiO_2}	2.74 ± 0.23	2.3–2.8	bulk value
n_{PS}	1.581 ± 0.046	≤ 1.585	bulk value
Translational a_{eff} (nm)	524 ± 15	–	
Rotational a_{eff} (nm)	523 ± 6	–	

Table 3.2: Measured and expected values of parameters for Janus particles.

nm) and again on par with the tracking precision of complex scatterers such as spheres in clusters[30,54].

Rotational diffusion

Although the best-fit polar angle θ is unreliable, we can extract the rotational diffusion coefficients from the data by considering only the azimuthal angle ψ. To do this, we derive an expression for the autocorrelation of the projection of the orientational unit vector onto the $x-y$ plane $\mathbf{p}(t) = \cos\psi(t)\,\hat{\mathbf{x}} + \sin\psi(t)\,\hat{\mathbf{y}}$:

$$\langle \mathbf{p}(t) \cdot \mathbf{p}(t+\tau) \rangle = \frac{1}{4} \sum_{\ell=1}^{\infty} \frac{2\ell+1}{\ell(\ell+1)} \left(S_\ell^1\right)^2 \exp\left[-\ell(\ell+1) D_r \tau\right]. \tag{3.17}$$

where $S_\ell^1 \equiv \int_{-1}^{1} P_\ell^1(x)\,dx$ and $P_\ell^1(x)$ are associated Legendre polynomials of order 1 (see appendix in Section 3.5 for further details). This expression is a multiexponential that depends on D_r and can be evaluated numerically.

To determine the rotational diffusion coefficient we fit Equation 3.17 to the data and obtain $D_{r,\text{Janus}} = 1.15 \pm 0.05$ s^{-1} (see Figure 3.8). From Equation 3.15 and a

temperature of $21 \pm 2°\text{C}$, we calculate an effective particle radius of $a_{\text{eff}} = 523 \pm 6$ nm, in excellent agreement with the $a_{\text{eff}} = 524 \pm 15$ nm obtained from translational motion. These results are summarized in Table 3.2. The quantitative agreement between the hydrodynamic radii calculated from the translational and rotational diffusion coefficients demonstrates the internal consistency of our fitting method, showing that the technique can effectively track the rotation and translation of Janus particles in water, despite the small asymmetry of the particles and holograms.

Although we do not use the z-coordinate or the polar angle θ to obtain these results, they would nonetheless be difficult to obtain with a traditional optical microscope. The particles diffuse rapidly, moving more than 5 µm in z throughout the trajectory, which might take them too far out of focus during a 2D microscopic measurement. Furthermore, the fit to the scattering model allows a quantitative determination of the azimuthal angle with no calibration required. Tracking this angle could be useful in other studies, for example to determine how quickly a particle orients in response to an external field or another object.

3.4 Conclusions and future work

We have shown a new technique to measure the 3D translational and rotational dynamics of colloidal particles. Our results show that holographic microscopy can capture the 3D position and orientation of non-spherical colloidal particles, and that these variables can be tracked precisely by fitting scattering models based on the discrete dipole ap-

proximation to the measured holograms. The technique has high precision and temporal resolution, as evidenced by the measured rotational diffusion coefficient for the Janus particle, which is an order of magnitude larger than previously measured 3D diffusion coefficients of a freely-diffusing particle.

We can address many of the limitations of our approach by improving our DDA implementation and fitting procedure. We intend to improve the accuracy of the measurements by antialiasing the voxelation and incorporating near-field corrections to the scattering calculations, as described in reference[73]. Spurious fits might be suppressed with an additional fitting pass that enforces a physically plausible trajectory, perhaps assisted by a Kalman filter[74].

Because the DDA is applicable to scatterers with arbitrary size, shape, and refractive index profiles, our technique could be used to measure the dynamics of a wide variety of particles. This makes it suitable for a number of different applications, including microrheology, measurements of interactions between non-spherical particles, and fundamental studies of colloidal self-assembly and bacterial motion.

Acknowledgments

We thank Jennifer A. Lewis for helpful discussions and feedback on the manuscript this chapter draws from.

This work was funded by the Harvard MRSEC through NSF grant no. DMR-0820484. Scattering calculations and hologram fitting were performed on the Odyssey cluster,

managed by the Harvard FAS Sciences Division Research Computing Group.

SR and IK acknowledge support from the National Science Foundation through awards CBET-1067501 and CBET-1264550. AW thanks The University of Sydney Travelling Scholarship for support. TD acknowledges support from a National Science Foundation Graduate Research Fellowship. KC was supported by the AFOSR MURI award FA9550-12-1-0471.

3.5 Projection Correlation Function $\langle \mathbf{p}(t) \cdot \mathbf{p}(t+\tau) \rangle$

In this Appendix, we derive Eq. 3.17, with which we measure the rotational diffusion coefficient D_r of a Janus particle in Sec. 3.3.3.

The isotropic rotational diffusion of a particle can be quantified by studying the trajectory on the unit sphere of a unit vector $\mathbf{u}$ fixed to the particle. Computing the autocorrelation $\langle \mathbf{u}(t) \cdot \mathbf{u}(t+\tau) \rangle$, where τ is a lag time, from experimental data allows the measurement of the rotational diffusion coefficient D_r. This requires tracking the entire 3D orientation of the particle. Here we consider how D_r can be measured when only a two-dimensional azimuthal projection of $\mathbf{u}$ is observed, as is the case for the Janus particles.

In our experiments, we observe the normalized projection of $\mathbf{u}$ onto the laboratory $x-y$ plane, which is perpendicular to the optical axis. In spherical polar coordinates, where the tip of $\mathbf{u}(t)$ has coordinates $\theta(t)$ and $\psi(t)$ on the unit sphere, this projection is given by $\mathbf{p}(t) = \cos\psi(t)\,\hat{\mathbf{x}} + \sin\psi(t)\,\hat{\mathbf{y}}$. From the data, we can then compute the

autocorrelation of $\mathbf{p}(t)$:

$$\langle \mathbf{p}(t\ +\ \tau)\ \cdot\ \mathbf{p}(t)\rangle \quad = \quad \langle \cos(\psi(t\ +\ \tau))\cos(\psi(t))\ +\ \sin(\psi(t\ +\ \tau))\sin(\psi(t))\rangle. \quad (3.18)$$

To simplify the notation, we will use primes to denote angles at time $t+\tau$; the unprimed angles θ and ψ are at time t. Thus,

$$\langle \mathbf{p}(t+\tau)\cdot\mathbf{p}(t)\rangle = \langle \cos\psi'\cos\psi + \sin\psi'\sin\psi\rangle. \quad (3.19)$$

We show that D_r can be determined from experimental measurements of $\langle \mathbf{p}(t+\tau)\cdot\mathbf{p}(t)\rangle$ by calculating this autocorrelation for a particle undergoing isotropic rotational diffusion characterized by D_r. We neglect translation-rotation coupling and therefore ignore the translational diffusion of the particle. Let $f(\theta,\psi;t)\,d\Omega$ be the probability of finding $\mathbf{u}$ in the solid angle $d\Omega$ near (θ,ψ) at time t. The probability density f is governed by a rotational Fick's law[75]:

$$\frac{\partial f}{\partial t} = D_r\left(\frac{1}{\sin\theta}\frac{\partial}{\partial\theta}\left(\sin\theta\frac{\partial f}{\partial\theta}\right) + \frac{1}{\sin^2\theta}\frac{\partial^2 f}{\partial\psi^2}\right). \quad (3.20)$$

The operator on the right is the Laplacian on the unit sphere. Computing $\langle \mathbf{p}(t+\tau)\cdot\mathbf{p}(t)\rangle$ requires knowing the transition probability density $K(\theta,\psi,\theta',\psi';\tau)$ for $\mathbf{u}$ to move from (θ,ψ) to (θ',ψ') after a lag time τ. If we assume that the distribution of initial orientations (θ,ψ) is uniform, such that $f = 1/(4\pi)$, then using Eq. 3.19 the

autocorrelation of $\mathbf{p}$ will be given by[75]

$$\langle \mathbf{p}(t) \cdot \mathbf{p}(t + \tau) \rangle = \iint (\cos\psi \cos\psi' + \sin\psi \sin\psi') \frac{K(\theta, \psi, \theta', \psi'; \tau)}{4\pi} \, d\Omega \, d\Omega'. \quad (3.21)$$

The transition probability K is given by the probability density $f(\theta', \psi'; \tau)$, governed by Eq. 3.20, with the following initial condition:

$$f(\theta', \psi'; 0) = \frac{\delta(\theta' - \theta)\delta(\psi' - \psi)}{\sin\theta'} \quad (3.22)$$

where δ denotes the Dirac delta function. Separation of variables leads to the following solution for K:

$$\begin{aligned} K(\theta, \psi, \theta', \psi'; \tau) = &\sum_{\ell=0}^{\infty} C_{\ell 0} P_\ell(\cos\theta') \exp\left[-\ell(\ell+1) D_r \tau\right] \\ &+ \sum_{\ell=1}^{\infty} \sum_{m=1}^{\ell} \sum_{p=1}^{2} C_{\ell m}^{(p)} Y_{\ell m}^{(p)}(\theta', \psi') \exp\left[-\ell(\ell+1) D_r \tau\right]. \end{aligned} \quad (3.23)$$

Here, $P_\ell(\cos\theta')$ is a Legendre polynomial, and the $Y_{\ell m}^{(p)}$ are real spherical harmonics[76]:

$$Y_{\ell m}^{(p)}(\theta', \psi') \equiv \begin{cases} P_\ell^m(\cos\theta') \cos m\psi' & \text{if } p = 1 \\ P_\ell^m(\cos\theta') \sin m\psi' & \text{if } p = 2. \end{cases} \quad (3.24)$$

The initial condition in Eq. 3.22 results in the expansion coefficients being

$$C_{\ell 0} = \frac{2\ell+1}{4\pi} P_\ell(\cos\theta) \tag{3.25}$$

for the azimuthally symmetric ($m = 0$) terms and

$$C_{\ell m}^{(p)} = \frac{2\ell+1}{2\pi} \frac{(\ell-m)!}{(\ell+m)!} Y_{\ell m}^{(p)}(\theta,\psi) \tag{3.26}$$

for the remaining terms.

Consider the first term on the right side of Eq. 3.21 which contains $\cos\psi\cos\psi'$. By orthogonality of $\cos\psi$ and $\cos\psi'$, only terms in K with $m = 1$ and $p = 1$ contribute to the integral. Integration over ψ and ψ' contributes two factors of π, and so we obtain

$$\begin{aligned}\iint \cos\psi\cos\psi' \frac{K(\theta,\psi,\theta',\psi';\tau)}{4\pi}\, d\Omega\, d\Omega' = \\ \sum_{\ell=1}^{\infty} \left(\frac{2\ell+1}{\ell(\ell+1)} \frac{1}{8\pi^2} \exp\left[-\ell(\ell+1)D_r\tau\right] \pi^2 \right. \\ \left. \times \int P_\ell^1(\cos\theta)\sin\theta\, d\theta \int P_\ell^1(\cos\theta')\sin\theta'\, d\theta' \right).\end{aligned} \tag{3.27}$$

Integrating the remaining term of Eq. 3.21 over ψ and ψ' gives exactly the same result. Defining

$$S_\ell^1 \equiv \int_{-1}^{1} P_\ell^1(x)\, dx, \tag{3.28}$$

we obtain Eq. 3.17 in the manuscript:

$$\langle \mathbf{p}(t) \cdot \mathbf{p}(t+\tau) \rangle = \frac{1}{4} \sum_{\ell=1}^{\infty} \frac{2\ell+1}{\ell(\ell+1)} \left(S_\ell^1\right)^2 \exp\left[-\ell(\ell+1) D_r \tau\right]. \tag{3.29}$$

Note that $S_\ell^1 = 0$ for even ℓ due to the parity of $P_\ell^1(x)$. To evaluate our result numerically, we use DiDonato's recursion relation for S_ℓ^1 [77]:

$$S_{\ell+2}^1 = \frac{\ell(\ell+2)}{(\ell+1)(\ell+3)} S_\ell^1. \tag{3.30}$$

Unlike the results obtained by prior workers on this problem, our solution can be easily computed and used to measure D_r from experimental data. Saragosti *et al.* obtain a series expression equivalent to Eq. 3.29, but their solution contains complicated angular integrals that are left unevaluated [78]. They therefore determine D_r from a $\ell = 1$ approximation of Eq. 3.29. While the $\ell = 1$ term indeed dominates when $D_r\tau$ is large, the autocorrelation can be measured most precisely near $\tau = 0$, at which the number of independent angular displacements is largest. In this regime, the single-exponential approximation of Saragosti *et al.* fails, and our full solution is necessary.

We verified our result, Eq. 3.29, by computing $\langle \mathbf{p}(t) \cdot \mathbf{p}(t+\tau) \rangle$ for simulated rotational trajectories of particles undergoing rotational diffusion. We simulated rotational diffusion using the algorithm of Beard and Schlick [79]. Figure 3.9 shows the autocorrelations computed from simulated trajectories with two different D_r along with fits to Eq. 3.29.

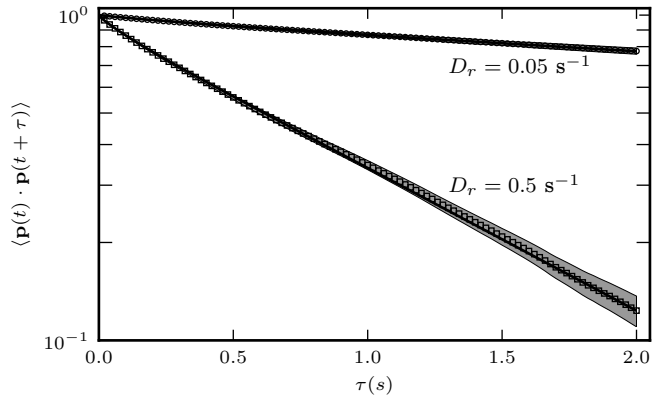

Figure 3.9: Autocorrelation $\langle \mathbf{p}(t) \cdot \mathbf{p}(t+\tau) \rangle$ for simulated particle undergoing rotational diffusion with $D_r = 0.05\ \mathrm{s}^{-1}$ (open circles) and $D_r = 0.5\ \mathrm{s}^{-1}$ (open squares). Solid lines are best fits to Eq. 3.29. Shaded gray regions denote error bars on the autocorrelations, calculated using a block decorrelation technique[63].

We find excellent agreement between the simulated autocorrelations and best fits to Eq. 3.29; the values of D_r determined from the best fits agree with the simulation input values to 0.5% or better.

4

Random-Subset Fitting of Digital Holograms for Fast 3D Particle Tracking

Digital holographic microscopy (DHM) is a powerful tool for measuring three-dimensional dynamics of microscale objects suspended in a fluid. A holographic micro-

scope records a two-dimensional (2D) image, or hologram, resulting from interference between a reference light field and the light scattered from a sample. The 2D hologram encodes phase information that can be used to infer the 3D structure of the sample. Because holograms can be captured with a fast camera, DHM enables measurements of 3D fluctuations and motion on short timescales, making it particularly useful for measuring dynamics of swimming organisms[6,7,16,80], living cells[13], and colloidal particles[81,82].

The principal challenge in DHM is recovering the 3D information from the 2D hologram. The most common method is numerical reconstruction, the digital descendant of Gabor's optical reconstruction[2]. Numerical reconstruction techniques use scalar diffraction theory to back-propagate light from a measured hologram, yielding a 3D image of the sample[3,83]. However, the 3D image can contain artifacts if the sizes of the objects are on the order of the wavelength of light[18,84], as is the case in colloidal suspensions.

An inverse-problem approach can yield more precise and accurate measurements. Ovryn and Izen[19] and Lee and coworkers[21] measured positions and optical properties of spherical colloidal particles by fitting a Lorenz–Mie scattering solution to holograms, taking the particle position, radius, and refractive index as fitting parameters. Fitting holograms using scattering solutions takes advantage of known information about the particle, such as its shape, and avoids artifacts that lead to systematic errors[23]. This inverse approach has been used for high-precision measurements in fluid dynamics[62] and colloid and interface science[4]. Recent work has extended the fitting technique to clusters of spheres[30,31] and non-spherical particles[85].

However, the fitting technique is computationally intensive. For example, Kaz, McGorty and coworkers[4] needed about 100 CPU-hours to fit a single particle trajectory consisting of 14,000 frames. For a cluster of particles, the computational burden is even larger: sequentially fitting 14,000 frames of a six-particle cluster[30] would take several CPU-months.

How can we reduce the time to fit? One strategy is to parallelize. Although holograms can be processed more quickly using the parallel processing inherent to a GPU[62], it is still difficult to process a set of such holograms in parallel. This is because the nonlinear optimization methods used to fit holograms require initial guesses for the particle position, radius, and index of refraction for each frame. The initial guesses can be obtained from the results of the previous frame, but only if the frames are fit sequentially. For clusters containing many particles, automated guesses are difficult because multiple scattering and near-field effects make the hologram sensitive to small changes in particle positions.

A second strategy is to reduce the size of the hologram. Fitting a scattering solution to a hologram is a highly over-determined problem: typically of order 10 parameters are extracted from a hologram containing of order 10000 pixels, each of which represents an intensity measurement. The premise of our approach is that a far smaller number of measurements can sufficiently constrain a fit. This premise was explored by Seifi and coworkers, who used staged resampling[86] and pixel masks[87] to reduce the number of pixels in a hologram and thereby reduce the number of computations required to

evaluate the scattering solution. Here we pursue an alternative approach that requires no prior knowledge about the spatial-frequency spectrum of a hologram.

We show that randomly sampling a subset of pixels in the hologram reduces the fitting time by an order of magnitude or more, with little effect on accuracy or convergence. The random-subset approach requires only a single sampling step, is simple to implement, works with any scattering model, and can be combined with GPU computing or resampling methods to reduce computing time even further. The method also allows us to analyze a time series of holograms in parallel. To do this, we first fit a small random subset of each hologram in a time series, then use the results as initial guesses to fit the full holograms in parallel.

4.1 Background

We first outline the method of fitting scattering solutions to in-line digital holograms, following Lee and coworkers[21]. A digital hologram $\mathbf{H}$ captures phase information through interference of a reference field (E_{ref}) with the field scattered from an object (E_{scat}):

$$\begin{aligned} \mathbf{H} &= |E_{\mathrm{ref}} + E_{\mathrm{scat}}|^2 \\ &= |E_{\mathrm{ref}}|^2 + E_{\mathrm{ref}}^* E_{\mathrm{scat}} + E_{\mathrm{ref}} E_{\mathrm{scat}}^* + |E_{\mathrm{scat}}|^2 \end{aligned} \tag{4.1}$$

Before fitting, we record the background intensity, which we assume to to be equal to the reference wave intensity $|E_{\mathrm{ref}}|^2$. We then calculate a normalized hologram $\mathbf{h}$:

$$\mathbf{h} = \frac{\mathbf{H}}{|E_{\mathrm{ref}}|^2} = 1 + \frac{E_{\mathrm{scat}}}{E_{\mathrm{ref}}} + \frac{E_{\mathrm{scat}}^*}{E_{\mathrm{ref}}^*} + \left|\frac{E_{\mathrm{scat}}}{E_{\mathrm{ref}}}\right|^2 \tag{4.2}$$

We extract information about a physical system from $\mathbf{h}$ through an inverse-problem approach. We fit a scattering model to $\mathbf{h}$ using a nonlinear optimization algorithm; that is, we minimize the objective function

$$g\left([x],[y],[z],[n],[r],\alpha \in \Re\right) = |\mathbf{h}_{\mathrm{model}}([x],[y],[z],[n],[r],\alpha) - \mathbf{h}_{\mathrm{rec}}|^2 \tag{4.3}$$

where the bracketed variables indicate vectors of parameters (one per particle), $\mathbf{h}_{\mathrm{rec}}$ is the recorded hologram, $\mathbf{h}_{\mathrm{model}}$ is the calculated hologram given parameter values, and α is an additional fitting parameter needed for convergence[21]. Depending on the shape and structure of the particles, other parameters may be needed. For example, homogeneous nonspherical particles are characterized by their orientations as well as positions. In what follows we ignore these additional parameters with no loss of generality.

We compute $\mathbf{h}_{\mathrm{model}}$ from a light scattering solution. For a single particle we use the Lorenz–Mie solution[21]. For clusters of multiple particles, we use an adaptation[31] of Mackowski's multisphere superposition solution[88], which takes into account all near- and far-field coupling between the scattered fields. If the numerical minimization of the

objective function converges, the result of this fitting procedure is a measurement of the parameters $[x]$, $[y]$, $[z]$, $[n]$, and $[r]$.

For larger clusters, we generally use fixed values for the index and radius of the particles, as determined from manufacturer specifications or other experiments. The objective function g therefore simplifies to:

$$g\left([x],[y],[z],\alpha \in \Re\right) = \left|\mathbf{h}_{\text{model}}([x],[y],[z],\alpha) - \mathbf{h}_{\text{rec}}\right|^2 \tag{4.4}$$

We use local optimization based on the Levenberg–Marquardt algorithm[22] to minimize g. Owing to the large parameter space and narrow minimum, local optimization algorithms fare better than global ones. However, local approaches require an initial guess for the set of parameters and converge only if this guess is close to the minimum of the objective function. Obtaining a sufficiently good guess is difficult. For clusters of spheres we must manually generate an initial guess for the particle positions, using numerical reconstruction as a guide. However, when studying a time series, we can use the results from fitting one frame as the initial guess for the next, provided that the capture rate is fast relative to particle motion. This approach minimizes manual intervention but precludes processing the frames in parallel.

Fitting is computationally expensive because it requires repeated evaluation of $\mathbf{h}_{\text{model}}$, which, like $\mathbf{h}_{\text{rec}}$, is an array of intensity values. The objective function g in Equation 4.4

is therefore a sum over a large number of pixels:

$$g\left([x],[y],[z],\alpha\right) = \sum_{i,j} \left(h_{\mathrm{model}}(i,j;[x],[y],[z],\alpha) - h_{\mathrm{rec}}(i,j)\right)^2 \tag{4.5}$$

where the subscripts i, j index over pixels in the recorded hologram. To evaluate g, we must calculate h_{model} at each pixel (i, j) on the detector. This corresponds to calculating the scattered field at a set of locations in space. Calculating the scattered field at each pixel is computationally intensive; in both the Lorenz–Mie and multisphere superposition solutions, the field at each point is given by a series expansion in special functions that themselves require many operations to calculate. On a single core of a 2.67 GHz Intel Xeon X5650 CPU, it takes approximately one second to compute a 200×200 pixel hologram of a single 1-µm-diameter particle lying 15 µm from the focal plane. Fitting a single hologram of such a particle takes 15 to 30 seconds. It takes much longer to fit the multisphere superposition solution to a hologram of a cluster: about 3 seconds to compute a hologram of a six-particle cluster of 1.3-µm spheres, and over 500 seconds to fit a single hologram.

4.2 Approach

Our approach reduces the fitting time by ignoring most of the pixels of a hologram, thereby reducing the number of evaluations required to calculate the objective function g. Consider the hologram shown in Figure 4.1a. It contains a large number of pixels, but

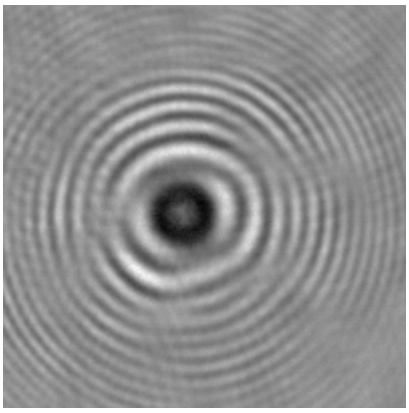

(a) hologram of a cluster of six 1.3-μm-diameter spheres (**h**)

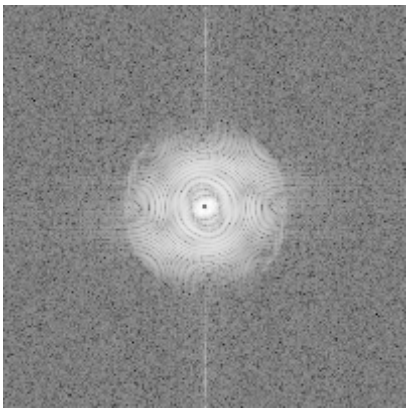

(b) Fourier transform of (a), log amplitude

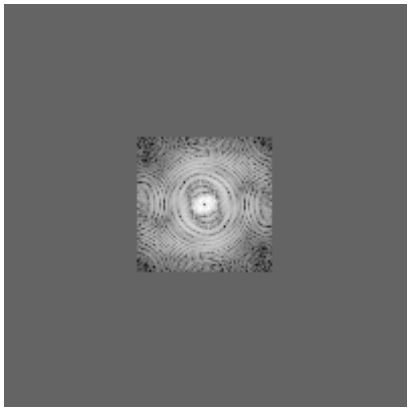

(c) Fourier transform of (a) after downsampling, log amplitude

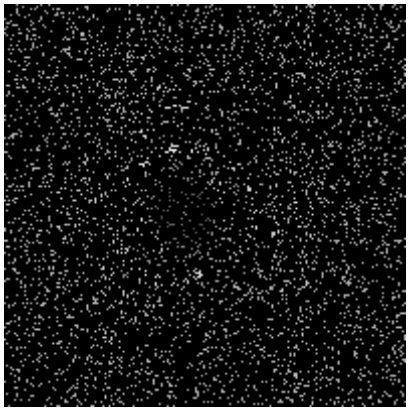

(d) a random selection of ten percent of the pixels of (a)

Figure 4.1: Several representations of the information in a hologram of six 1.3-µm-diameter polystyrene spheres in an aqueous suspension. All of these representations are 200×200 pixels, except (c), which is 67×67 pixels.

only a fraction of these are needed to determine the particle position and properties. For example, much of the information is at low spatial frequencies, as shown in Figure 4.1b. Downsampling techniques like those of Seifi and coworkers[86] take advantage of the fact that, in many holograms, the largest spatial frequencies are the ones most affected by noise and can be discarded with little effect on the fitted parameters. But even after downsampling (Figure 4.1c), a hologram retains symmetries and large patches of nearly uniform value, suggesting that it can be compressed further.

Instead of downsampling, we simply choose a random subset of the pixels in the hologram (Figure 4.1d), and fit a scattering solution to these. This method requires little information about the sample or the optical train, yet is surprisingly effective. The inspiration for this technique comes from compressed sensing, where it has been shown that random sampling is often nearly as efficient as optimal sampling[89–91].

We examined the effectiveness of this technique on two experimentally recorded holograms of polystyrene spheres suspended in water (Figure 4.2a,d). Both holograms we examined are 200×200 pixel regions cropped from larger holograms. The first hologram (Figure 4.2a) was of a 1-µm-diameter sphere and was originally recorded for an *in situ* viscometry experiment. We modeled this hologram using the Lorenz–Mie scattering solution[21], and we fit the solution to the data using five parameters: three spatial coordinates, the radius, and the index of refraction. The second hologram (Figure 4.2d) is of a cluster of six 1.3-µm-diameter particles that are nearly touching one another. It was originally recorded for the experiments described by Perry and coworkers[30]. We

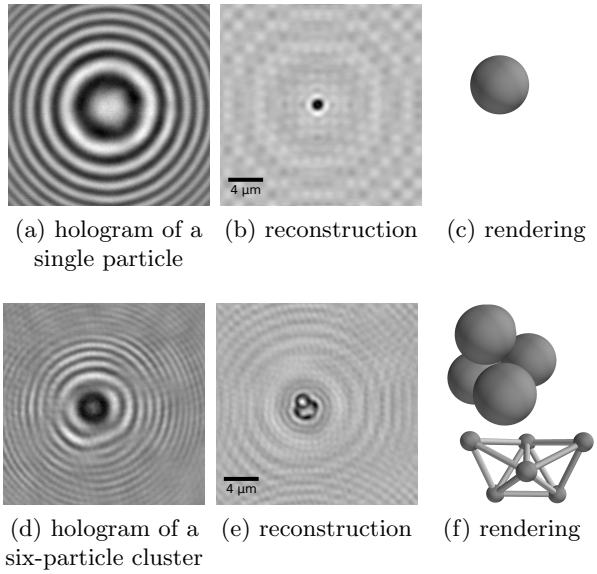

(a) hologram of a single particle (b) reconstruction (c) rendering

(d) hologram of a six-particle cluster (e) reconstruction (f) rendering

Figure 4.2: (a, d) The holograms we use in our study (200×200 pixels, 0.12 µm/pixel). (b, e) Numerical reconstructions showing a plane cutting through the center of the scatterers. (c, f) Renderings of the structures, based on fitting a scattering model to the holograms. Incident light propagation is into the page in all cases. The inset ball and stick model is rotated to show structure.

modeled this hologram using the multisphere superposition solution[88], and we fit the solution to the data using 18 parameters: three center coordinates for each particle. We fixed the index of each particle to 1.58 and the radius to 0.645 µm. To give the fitting algorithm more freedom to rearrange particles, we allowed the model to adopt configurations where the particles overlapped by up to 10 percent of the radius.

We measured the convergence properties and accuracy of our method as a function of the pixel fraction f and initial guess quality. The accuracy we report for the fitted parameters is relative to a set of standard values, which we calculated by fitting full frames and averaging the fit results over different initial guesses. The pixel fractions varied from $f = 1$, corresponding to the full hologram, to $f = 1 \times 10^{-3}$. For the single-

particle hologram, we chose 10 random subsets at each fraction, and we used an initial guess for z that differed by 0.2 µm from the standard z, corresponding to a displacement along the optical axis. For the six-particle hologram, we chose 600 random subsets at each fraction, and we used a range of initial guesses created by adding a Gaussian random variable to each standard center coordinate. To vary the guess quality, we increased the width of the Gaussian. We generated 50 initial guesses for each of 12 distributions whose standard deviations ranged from $\sigma = 1 \times 10^{-4}$ µm to $\sigma = 1$ µm. In all cases, we rejected guesses in which particles overlapped by more than 10 percent of their radius.

We used an implementation of the Levenberg–Marquardt algorithm from nmpfit, a Python adaptation of LMFIT from MINPACK-1[92]. We considered a fit to have converged if the sum of squares decreased by less than 10^{-10} between iterations, the normalized solution vector changed by less than 10^{-10} between iterations, or the cosine of the angle between the solution and and all columns of the Jacobian was smaller than 10^{-10}. If none of these conditions were met within 100 Levenberg–Marquardt iterations, we halted the minimization and considered the fit to have failed to converge.

All computations were performed on Harvard's Odyssey computing cluster. Each fit ran on a single core of a 2.67 GHz Intel Xeon X5650 processor. We used parallel processing only to run independent trials concurrently. Our implementation of random-subset fitting is available in our hologram processing code HoloPy[33], available as free and open-source software.

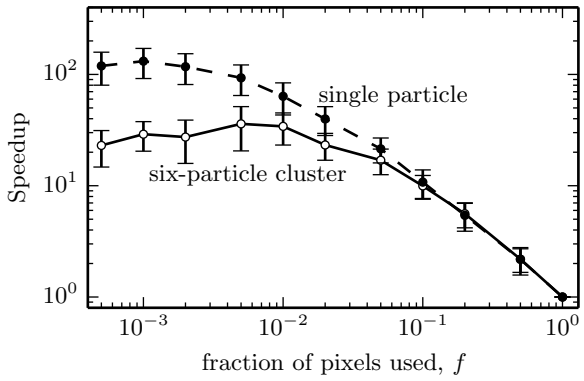

Figure 4.3: Fitting speedup as a function of subset fraction. The speedup is the time required to fit the full hologram divided by the time required to fit a random subset, using the same initial guess and convergence criteria. The single-particle fits used a guess offset by 0.2 µm along the axis perpendicular to the optical axis, and the six-particle cluster fits used guesses offset by a Gaussian random variable with $\sigma = 0.01$ µm added to each center coordinate of each particle. Error bars represent 99% confidence intervals.

4.3 Results

We find that random-subset fitting extracts particle parameters from holograms more than an order of magnitude faster than fitting the full frame. At pixel fractions near $f = 1$, the speedups we obtain increase linearly with decreasing pixel fraction. For both holograms, the speedups plateau at fractions smaller than $f = 0.01$ (Figure 4.3): the plateau is at 100× speedup for the single-particle hologram and 25× for the hologram of the six-particle cluster. The plateau occurs because the time required to compute scattering coefficients is independent of the pixel fraction. This overhead is larger for the multisphere superposition solution, explaining the lower plateau for the six-particle cluster.

We also find that fitting such small fractions does not require better initial guesses.

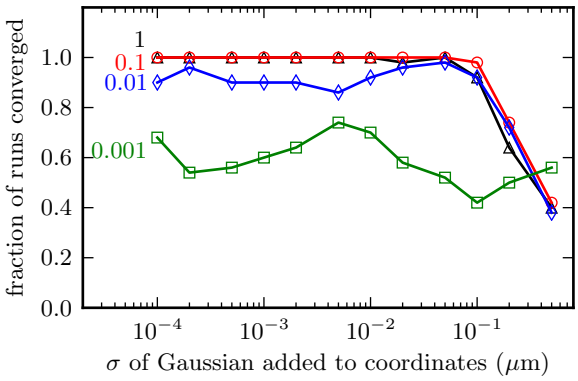

Figure 4.4: Fraction of six-particle cluster fits that converge, shown as a function of initial guess quality for subsets with $f = 1$ (black), $f = 0.1$ (red), $f = 0.01$ (blue), and $f = 0.001$ (green). The $f = 1$ and $f = 0.1$ lines overlap.

For the full hologram of the six-particle cluster, the fitting algorithm converges for initial guesses within about 100 nm (per coordinate) from the standard values, as shown in Figure 4.4. Fits converge with similar reliability for pixel fractions as low as $f = 0.01$. Reducing the fraction further reduces the probability of convergence, but there is little reason to use such small subsets given the diminishing returns on speedup.

Finally, we examine accuracy. In a simulated case with a known correct answer (Figure 4.5 dotted line) we find that we can recover the correct particle positions to within numerical roundoff error, even at subset fractions as low as $f = 1 \times 10^{-4}$ (which correspond to only 20 pixels). Thus, in the absence of noise, the fit accuracy is limited only by numerical precision and minimizer cutoff. With our experimentally-measured holograms, we find a larger deviation between the fitted parameters and the standard values at fractions smaller than the full frame $f = 1$ (Figure 4.5). However, the errors are comparable to the intrinsic measurement error, which is approximately 10^{-3} µm for

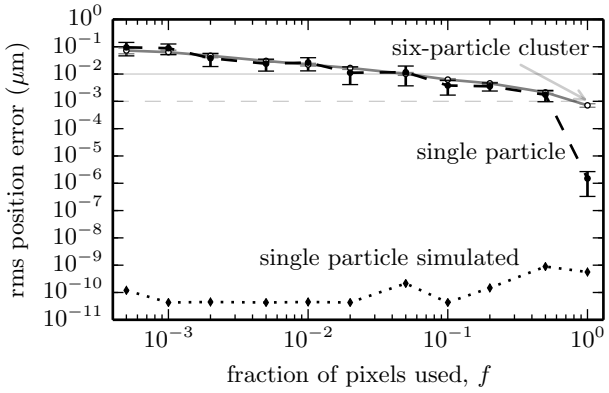

Figure 4.5: Per-particle fit error as a function of subset fraction. The horizontal lines show the expected intrinsic rms position error of fits to the single-particle hologram (dashed) and to the hologram of the six-particle cluster (solid). Error bars represent 99% confidence intervals. The rms position error in the $f = 1$ fit arises because we fit the same hologram multiple times using different initial guesses. Slightly different particle configurations yield model holograms that match the recorded one equally well, owing to noise in the recorded holograms.

single particles[4] and 10^{-2} µm for clusters[54]. Thus, we expect that fitting to subsets as small as $f = 0.1$ will not impact accuracy.

4.4 Discussion

Our results show that fitting scattering models to random subsets of holograms offers large speedups with few tradeoffs, either in convergence or accuracy. Two questions arise: How should it be used as part of a workflow for analyzing holograms? And why is random sampling so effective?

Given how simple it is to apply random sampling, there seems little reason not to use it when fitting holograms. Some experimentation may be required to find the appropriate subset fraction, which varies with the scatterer and experimental conditions. However,

we find that using ten percent of the pixels generally gives good results.

We have found the method to be particularly useful for analyzing time series of holograms quickly and accurately. We use a two stage procedure, the first stage of which is a serial "rough fit" of each hologram in the time series, using a fraction of 0.01 or 0.1. We also limit the minimization algorithm to two iterations. The rough fit completes in approximately 1/200th the time required to serially fit the full frames: we gain a factor of twenty by fitting to small subsets, and another factor of ten by limiting the number of iterations. In the second stage, we use the results from the first stage as initial guesses to fit the full holograms. Because these fits can be done in parallel, the speedup scales linearly with the number of nodes.

As to the second question, on the effectiveness of the technique, we offer only an intuitive justification here. Consider a linear transformation, operating on a hologram, that compresses all the information about particle parameters into a small number of basis elements. We expect this transformation to depend on the shape, size, and optical properties of the scatterer. Hence we do not—and quite possibly cannot—know either the transformation or the optimal basis. However, for certain classes of signals, it has been shown that random sampling allows efficient compression without knowledge of the optimal basis[89–91]. We cannot prove that a hologram is in this class of signal, but the success of our method suggests that it is close enough for practical purposes, at least for holograms of spherical particles or clusters of spheres. Future work might yield a more rigorous mathematical justification of this idea. Future work could also compare

in detail to downsampling, cropping, and other sampling other sampling methods.

4.5 Conclusions

We have shown that the random-subset technique speeds up fitting by a factor of twenty or more, with little effect on accuracy or convergence compared to fitting to the full hologram. The technique is simple to implement and works for different scattering models. It can be used in concert with GPU or CPU-based parallel processing, with other algorithmic improvements, or with faster implementations of scattering models. In our calculations we sampled from a uniform distribution of pixel locations, but non-uniform sampling may allow even greater speedups by weighting areas of a hologram most useful for constraining a fit. In future work, it will be particularly interesting to determine if there are non-uniform distributions that are effective for a wide variety of scatterers.

Because recovering precise 3D information from measured data is the most time-consuming step in holographic microscopy, the method we have demonstrated should allow faster turnaround in experiments and may enable new experiments that were previously infeasible. The ability to process large quantities of data more quickly will be particularly useful for studying rare events, such as nucleation, that can only be captured with high frame rates over long times.

4.6 Acknowledgments

This work was funded by the National Science Foundation (NSF) through grant no. DMR-1306410 and by the Harvard MRSEC through NSF grant no. DMR-0820484. Scattering calculations and hologram fitting were performed on the Odyssey cluster, managed by the FAS Science Division Research Computing Group at Harvard University.

We thank Ameya Agaskar for helpful discussion on the mathematical background of this technique.

Our nmpfit routine comes from STSCI_PYTHON, a product of the Space Telescope Science Institute, which is operated by AURA for NASA.

5

A Bayesian approach to analyzing holograms of colloidal particles

Digital in-line holography has emerged as a powerful tool for both tracking and characterizing colloidal particles. Such particles are a few hundred nanometers to a

few micrometers in diameter and dispersed in a fluid. In a typical tracking experiment, one measures the position of the particle in three dimensions as a function of time by analyzing a time-series of holograms; the particle might be subject to Brownian motion, interactions, or external forces. This type of measurement is important to a number of different experiments in soft-matter physics and fluid mechanics, including microrheology[93], the study of phase transitions[30], measuring interactions in colloidal systems[4,54], and particle image velocimetry[20,62]. In a characterization experiment, one measures the size, refractive index, and shape of individual colloidal particles. This type of experiment is important for quality control of formulations[24] and studies of polymerization kinetics[28], among other applications.

Many of these experiments have been enabled by a method of analyzing holograms based on scattering theory. A hologram is a two-dimensional interference pattern that encodes information about the three-dimensional (3D) positions and optical properties of a particle or set of particles. An inline hologram results from the interference of the scattered fields from the sample and the undiffracted beam, which acts as a reference wave (Figure 5.1). Traditionally, the 3D information is recovered through reconstruction—effectively, shining light back through the hologram to generate a real 3D image. However, when the particles are comparable to the wavelength, reconstruction produces a number of imaging artifacts[18]. An alternative to reconstruction was first demonstrated by Ovryn and Izen[19,20], who showed that inline holograms could be directly predicted from Lorenz–Mie scattering theory and a model for the propagation

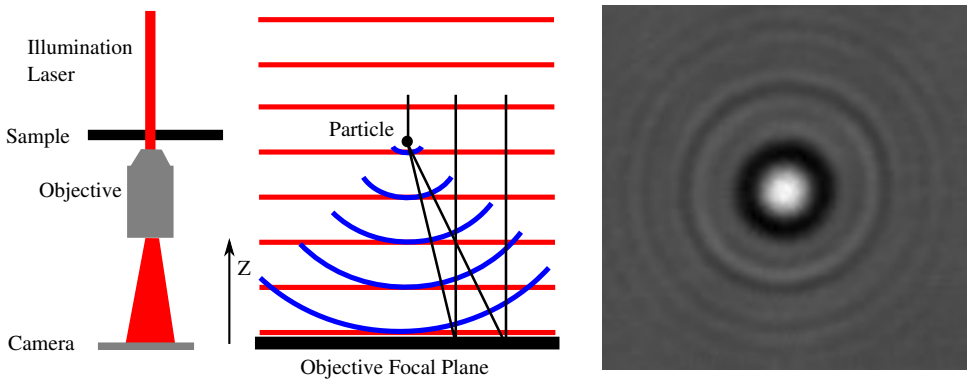

Figure 5.1: A typical optical setup used to record inline digital holograms. We image with a collimated laser beam and record the hologram on a CMOS camera. The coordinate z is defined as the distance from the particle to the focal plane of the objective ($z > 0$ if the particle is above (upstream of) the focal plane). x and y, not shown, represent the position of the particle in the plane normal to the optical axis. On the right, a measured, background-divided hologram of a 1-µm-diameter silica particle in water, taken using a 660 nm laser (hologram courtesy of Anna Wang).

and interference of the fields in the optical train of the microscope. Later work by Lee and coworkers[21], using a simpler model of the optical train, showed that one could fit a scattering model based on Lorenz–Mie theory to both track and characterize spherical colloidal particles. Further work from our research group has shown that other models, based on different scattering theories, can be fit to holograms of non-spherical particles[85] and clusters of spherical particles[30,31].

With these techniques, the optical field is never reconstructed; instead, a forward model—a model that can calculate a hologram based on scattering theory and propagation along the optical train— is fit to the data. In this method, which we will refer to as "least-squares fitting," an iterative minimization algorithm is used to find the parameters in the model that minimize the sum of the squared differences, pixel-by-pixel, between the measured hologram and the model. Although this process is computationally more

expensive than reconstruction, it yields much more accurate and precise measurements: for tracking experiments, analysis of the mean-square displacement shows that the uncertainty in the position can be on the order of a few nanometers or less[4].

However, there are two problems with the least-squares technique. First, it is difficult to obtain accurate uncertainties on parameters. In a characterization experiment, for example, one would like to know the uncertainty in the measured size or refractive index. The uncertainties on the fit parameters are critical not only for characterization experiments, but for any experiment in which the data will later be compared to a physical theory—for example an equation of motion for the particles—since the agreement with theory can only be judged when uncertainties are known. Second, the iterative minimization schemes used in least-squares fitting require good initial guesses for and constraints on all the parameters[30,54,85]. Determining the initial guesses and the constraints often requires a great deal of effort, particularly for more complex scattering models.

Here we demonstrate a Bayesian analysis method that overcomes these problems. Bayesian approaches to parameter estimation differ from classical frequentist approaches such as least-squares fitting in that we are allowed to speak of the probability distribution *of a parameter*. As in the least-squares technique, we start with a forward model, but instead of maximizing the likelihood, which is the probability of the data given parameter values, we maximize the posterior probability, which is the probability of the parameter values given the data. There are other Bayesian approaches to analyzing

holograms[94], but none that we are aware of use forward models based on scattering solutions.

There are two principle advantages of the Bayesian approach over least-squares fitting: first, any prior information about the particle, such as its position, size, or composition, can be formalized as a prior probability distribution, rather than implicitly encoded into initial guesses or constraints. Second, because probabilities are assigned to parameters, we can marginalize, or integrate out, the uncertainties on some of the parameters to obtain the uncertainties on the others. For example, in the characterization problem, we are interested in the uncertainty on the size or refractive index after taking into account the uncertainties on the positions. This type of analysis is not easily done using classical frequentist methods, and hence the parameters that are not of interest (but required by the model) are often called "nuisance parameters."

As an example of the utility of this technique, we show the results from a Bayesian analysis of a hologram of a 1-µm-diameter silica particle in water in Figure 5.2, compared to the results obtained from least-squares fitting. The Bayesian approach yields nearly the same "best fit" values for the parameters, but it also gives us much more information about the parameter uncertainties. From the Bayesian results, we see that there are strong correlations between the estimated particle refractive index, size, and axial position. The covariances between these parameters are difficult to extract in a least-squares approach. Because we have access to these covariances in the Bayesian approach, we can integrate them out, or "marginalize," to estimate the uncertainties on

any one parameter, as shown by the distributions along the diagonal.

The rest of this paper explains how we obtain these results and how we interpret them. In practice, the Bayesian approach involves three steps: specifying a forward model, specifying priors on all the parameters, and calculating the posterior probability distribution of all the parameters. The last step, calculating the posterior, is most effectively done using Markov-chain Monte Carlo (MCMC) techniques. In what follows, we limit our discussion to the measurement of a single microsphere, and we use a simple model for the propagation and interference along the optical train, equivalent to that of Lee and coworkers[21]. Nonetheless, the approach can easily be extended to more complex models of scattering or propagation. After describing the model and the Bayesian approach to calculating the posterior probability distribution, we apply our method to both the particle tracking problem and the particle characterization problem. We also show how to speed up the approach—and obviate the need for initial starting guesses for our MCMC chains—by a type of tempering based on choosing random subsets of pixels.

5.1 Background

Here we describe the forward model for hologram formation from a single spherical colloidal particle and the least-squares method used to fit this model to the data. The data are obtained from an in-line holographic setup as shown in Figure 5.1. We follow the work of Lee and coworkers[21] closely. In their model, a hologram **H** is computed as

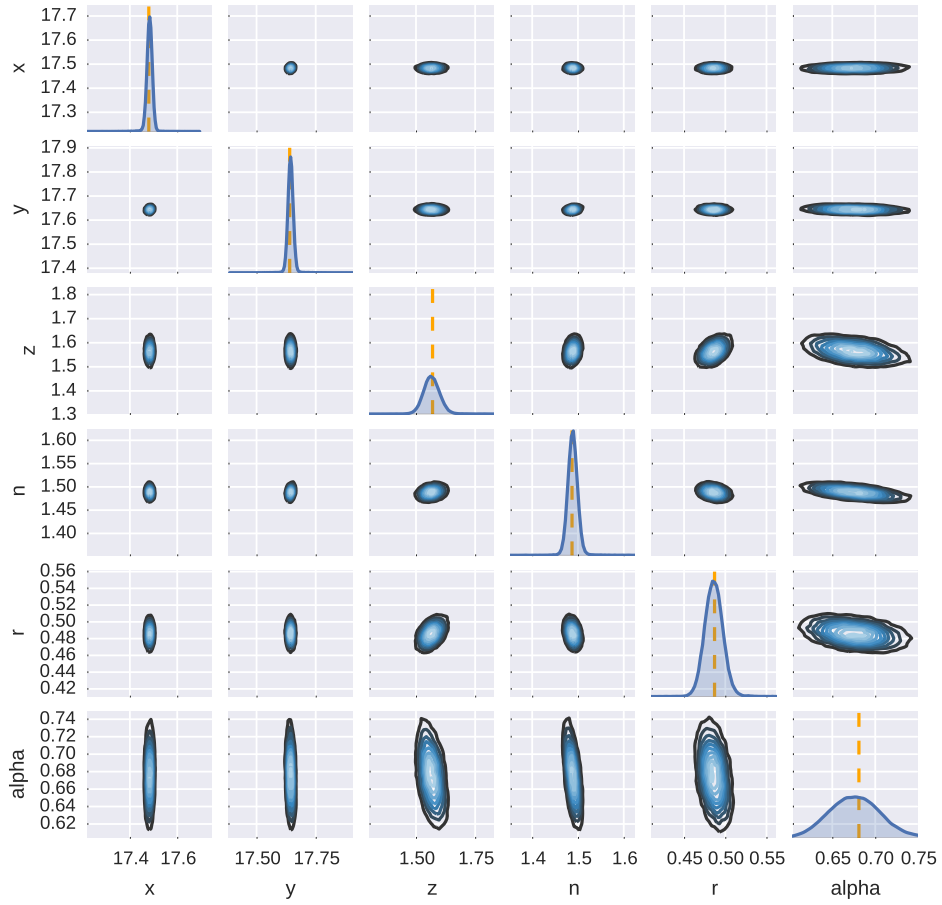

Figure 5.2: Comparison of results obtained from least-squares fitting using the Levenberg–Marquardt method, as described in Section 5.1, to those obtained from the Bayesian inference method described in this paper. Units of x, y, z and r are micrometers. The results from fitting are shown by the dashed orange lines. Fully marginalized distributions obtained from Markov-chain Monte Carlo (MCMC) sampling are shown along the diagonal, while the off-diagonal contour plots show the joint distributions of each pair of parameters, represented as Gaussian kernel density estimates. We explain the Bayesian approach and the MCMC technique later in the text.

a function of six parameters:

$$\mathbf{H}(x, y, z, n, r, \alpha) = |\alpha E_{\text{scat}}(x, y, z, r, n) + E_{\text{ref}}|^2 \tag{5.1}$$

where x, y, and z are the coordinates of the center of the particle, n is its (possibly complex) refractive index, r is its radius, and α is an auxiliary parameter that rescales the scattered field E_{scat}. E_{ref} is the reference field. We use boldface type for the hologram $\mathbf{H}$ to indicate that it is a matrix consisting of measurements of the intensity at each pixel (i, j) on our detector. The parameter α is needed to compensate for a number of different effects along the optical train and to produce satisfactory agreement between the model and measured holograms. Although a single scalar parameter is likely insufficient to account for all of these effects, the Bayesian analysis framework we develop below can easily be extended to more complex models, including ones that model the propagation along the optical train.

We divide the measured hologram $\mathbf{H}$ by a measured background, usually from a hologram captured without any particles in the field of view. The measured background should be very close to the intensity of the reference wave, $|E_{\text{ref}}|^2$, differing only by the light scattered by the particle of interest. We therefore approximate the background as the reference intensity to obtain an expression for the normalized hologram $\mathbf{h}$:

$$\mathbf{h} = \frac{\mathbf{H}(x, y, z, r, n, \alpha)}{|E_{\text{ref}}|^2} = \left| \frac{\alpha E_{\text{scat}}(x, y, z, r, n)}{E_{\text{ref}}} + 1 \right|^2. \tag{5.2}$$

To infer the parameters x, y, z, r, n, and α, we need a model for E_{scat} and an inference method. For a single sphere, the Lorenz–Mie theory can be used to calculate E_{scat}; other scattering theories[31,54,85] can be used for different types of particles or clusters of particles. In the inference method used by Lee and coworkers[21] and found in many subsequent papers, the parameters are inferred by minimizing the sum (over pixels) of the squared differences between model and data. Letting θ be the set of parameters $\{x, y, z, n, r, \alpha\}$, we have

$$\theta = \arg\min_{\theta} \chi^2 = \arg\min_{\theta} \sum_{i,j} \left[h_{ij} - h_{ij}^{M}(\theta)\right]^2 \tag{5.3}$$

where h_{ij} is the intensity measured at pixel (i, j) in the normalized recorded hologram and h_{ij}^{M} is the same for the hologram calculated from the forward model M. Using this least-squares formulation is equivalent to maximum-likelihood estimation when the noise in the measurements has a Gaussian distribution.

Because evaluating the Mie fields involves many computationally expensive special-function evaluations over a large number of pixels, several alternatives and variations on this method have been developed. Soulez and coworkers[95] used a local refinement approach to find particle positions more efficiently. Yevick and coworkers[96] used a machine-learning method employing a trained support-vector machine, instead of the full Lorenz–Mie solution, to infer the parameters. It is also possible to obtain particle positions from reconstructed volumes using deconvolution[97] or successive resampling[86].

In a recent paper our research group showed that good fits could be obtained by modeling only a small, randomly chosen subset of the pixels[98]. We later discuss how the random-subset technique can speed up the Bayesian approach as well.

Least-squares fitting is generally done using an iterative non-linear algorithm such as Levenberg–Marquardt[22], which requires the user to specify an initial guess for the parameters. Good initial guesses are required for the algorithm to converge. To keep the algorithm from wandering off into unphysical regions of parameter space, one can fix certain parameters or allow them to vary only over a certain range. One can also add arbitrary functions to the χ^2 function in Eq. (5.3) to penalize certain regions of parameter space. However, fixing parameters might not accurately reflect our knowledge of the system. We might have some prior information about the particle size (for example, from the manufacturer) but not trust it enough to set the size to a certain value. A penalty function can more accurately capture our state of knowledge, but as we shall see, Bayesian priors offer an elegant alternative to these various methods of limiting the parameter space.

5.2 Bayesian Approach

The problem of estimating parameters from a measured hologram and a forward model is well suited to Bayesian inference, which allows us to replace guesses, penalties, and constraints with a single prior probability function. Here we describe our approach using standard Bayesian terminology; for a more detailed discussion of Bayesian inference and

incorporation of priors, see chapter 3 of Gregory's book on Bayesian data analysis[99].

We are interested in the posterior probability $p(\theta|\mathbf{h}, M, I)$ of a set of parameters θ given the recorded hologram $\mathbf{h}$, the forward model M, and known prior information I. I might include information about the apparatus, such as the magnification, wavelength, and observation volume, and the fact that we know we are looking at a single sphere. Applying Bayes' rule, we can express the posterior in terms of a prior probability density $p(\theta|M, I)$, the likelihood function $p(\mathbf{h}|\theta, M, I)$, and a normalization factor $p(\mathbf{h}|M, I)$:

$$p(\theta|\mathbf{h}, M, I) = \frac{p(\theta|M, I)p(\mathbf{h}|\theta, M, I)}{p(\mathbf{h}|M, I)}. \tag{5.4}$$

Because the normalization factor $p(\mathbf{h}|M, I)$ has no effect on the parameters we measure or their uncertainties, we work with the unnormalized posterior probability distribution:

$$\begin{aligned} p'(\theta|M, \mathbf{h}, I) &= p(\theta|\mathbf{h}, M, I)p(\mathbf{h}|M, I) \\ &= p(\theta|M, I)p(\mathbf{h}|\theta, M, I). \end{aligned} \tag{5.5}$$

As shown by Eq. (5.5), we need to specify the prior $p(\theta|M, I)$ and the likelihood function $p(\mathbf{h}|\theta, M, I)$ to proceed.

5.2.1 Priors

The prior $p(\theta|M, I)$ is simply a mathematical statement of what we know about the parameters before doing the inference calculation. For example, priors for r and n might come from manufacturer's data sheets for commercially sourced particles. If we know the material composition of the particle, we can obtain a prior for n from a refractive index table and some estimated uncertainty. Similarly, we might obtain a prior probability distribution for r from the size distribution of the particles, as measured by the manufacturer or by some other technique, such as light scattering. In general, the prior distribution reflects our state of knowledge about a parameter, with wider distributions indicating greater uncertainty. We discuss specific choices of priors in Section 5.3.1.

5.2.2 Likelihood function

To evaluate the likelihood $p(\mathbf{h}|\theta, M, I)$, we must model the noise in the measured holograms. There are a variety of noise sources, including photon shot noise, electronic readout noise at the camera, and stray fringes from other scatterers in the optical train. We consider only the shot noise and readout noise in what follows. We then write the measured hologram as the sum of the model hologram plus noise:

$$h_{ij} = h_{ij}^{M} + u_{ij} \tag{5.6}$$

where the model hologram $\mathbf{h}^M$ (the elements of which are h_{ij}^M) is calculated from the forward model (Eq. (5.2)) and the Lorenz–Mie theory. The noise in the intensity at pixel (i, j) is represented by u_{ij}.

In writing this equation, we are assuming that the noise can be represented by a single additive term that characterizes pixel-by-pixel variations in the intensity. We now further assume that the noise $\mathbf{u}$ is delta correlated in space and time and that each u_{ij} is Gaussian distributed with a variance σ_{ij}^2 at each pixel. With these assumptions, we can write the likelihood as

$$p(\mathbf{h}|\theta, M, I) = \prod_{i,j} p(u_{ij}|\theta, M, I) \tag{5.7}$$

where

$$p(u_{ij}|\theta, M, I) = \frac{1}{\sqrt{2\pi}\sigma_{ij}} \exp\left\{ \frac{-\left[h_{ij} - h_{ij}^M(\theta)\right]^2}{2\sigma_{ij}^2} \right\}. \tag{5.8}$$

Combining Eq. (5.7) with Eq. (5.8), we obtain the full likelihood

$$p(\mathbf{h}|\theta, M, I) = \frac{1}{(2\pi)^{N/2} \prod_{i,j} \sigma_{ij}} \exp\left\{ -\sum_{i,j} \frac{\left[h_{ij} - h_{ij}^M(\theta)\right]^2}{2\sigma_{ij}^2} \right\} \tag{5.9}$$

where N is the total number of pixels in the hologram. We can estimate σ_{ij}^2 from a series of background holograms where no particles are present.

To justify this form of the likelihood, we first note that at the intensities we measure, the number of photons and readout electrons are large enough that the both shot

noise and readout noise are approximately Gaussian, and hence their sum is also Gaussian. Also, the non-linear least squares fitting methods used in previous work implicitly assume that the noise is Gaussian, so our approach is consistent with these methods. Other noise models that include the effect of correlations (that is, models for speckle) are a topic for future work.

5.2.3 Putting it all together

We recognize the term in the exponent of the likelihood in Eq. (5.9) as $\chi^2/2$, where

$$\chi^2(\theta) = \sum_{i,j} \frac{\left[h_{ij} - h_{ij}^M(\theta)\right]^2}{\sigma_{ij}^2}. \tag{5.10}$$

In this work we assume that the standard deviation of the noise at each pixel is identical and equal to a constant σ, so that

$$\prod_{i,j} \sigma_{ij} = \sigma^N. \tag{5.11}$$

Substituting the above into Eq. (5.5) yields an equation for the unnormalized posterior:

$$p'(\theta|, \mathbf{h}, M, I) = \frac{p(\theta|M, I)}{(2\pi)^{N/2}\sigma^N} \exp\left[-\frac{\chi^2(\theta)}{2}\right]. \tag{5.12}$$

Because the χ^2 term depends on the parameters θ through the forward model, it is difficult if not impossible to calculate the posterior directly from Eq. (5.12). We therefore

use a Markov-chain Monte Carlo technique as described in Section 5.3.1.

For the specific problems of particle tracking and particle characterization, we are interested only in subsets of the parameters. In particle tracking, we want to infer the position as a function of time. The position is described by one or more of the parameters x, y, and z along with their uncertainties, which can be obtained by marginalizing over n, r, and α (and over the position variables that are not of interest). In particle characterization, we want to infer the index of refraction n and radius r of the particle along with their associated uncertainties, which can be obtained by marginalizing over x, y, z, and α. Formally, marginalization involves integration of the posterior distribution over the parameters to be marginalized. The marginalized distribution for the particle tracking problem is

$$p'(x, y, z|\mathbf{h}, M, I) = \int dr\, dn\, d\alpha\, p'(x, y, z, r, n, \alpha|\mathbf{h}, M, I) \tag{5.13}$$

and that for the particle characterization problem is

$$p'(n, r|\mathbf{h}, M, I) = \int dx\, dy\, dz\, d\alpha\, p'(x, y, z, n, r, \alpha|\mathbf{h}, M, I). \tag{5.14}$$

In practice, we do not have to evaluate these integrals. As we describe in Section 5.3.1, we can sample the posterior using MCMC techniques and then marginalize by binning samples or by kernel density estimation.

5.2.4 Time-Series

In both particle tracking and particle characterization we work with time-series of holograms. There are two different categories of parameters within time-series: those that change over time and those that do not. In most cases, the particle radius r and refractive index n do not change as a function of time, whereas x, y, and z do. Therefore when calculating the marginalized posterior for the characterization problem, $p'(n, r|\mathbf{h}, M, I)$, we can use information from the entire time-series to minimize the uncertainty on the parameters. When computing $p'(x, y, z|\mathbf{h}, M, I)$ we analyze each frame individually.

There are at least two procedures to calculate the posterior $p'(n, r|\mathbf{h}, M, I)$ using information from the entire time-series. One is to evaluate it for a single frame and then use the result as the prior for the next frame. This procedure is similar to using a Kalman filter. It should yield the best estimate of n and r when the last frame is analyzed. The other procedure is to do a global calculation over the entire time-series. This procedure would involve expanding the model to include parameters for x, y, and z in each frame, then marginalizing all of these position parameters. Here we use the first approach because it is simpler to implement and computationally less demanding.

5.3 Results and discussion

5.3.1 Markov-chain Monte Carlo technique

Because Eq. (5.12) does not admit an analytic solution, we must calculate the posterior probability distribution numerically. Brute force evaluation of the posterior is prohibitively expensive owing to the six parameters. Thus we use a sampling technique based on a Markov-chain Monte Carlo algorithm.

Specifically, we use an affine-invariant ensemble sampler[100] as implemented in the Python library Emcee[101]. This sampler uses an ensemble of "walkers" to explore the posterior probability distribution. After a sufficient number of steps, the ensemble will converge to a steady-state distribution that is, by construction, equal to the posterior probability density. Thus, in the long-time limit, the method yields a set of samples directly from the posterior. Each sample consists of values for all six parameters in our forward model. We can plot these samples as in Figure 5.2 to visualize the joint distributions between the parameters, or we can marginalize simply by binning the samples as a function of only the parameters we wish to infer.

We demonstrate these techniques on holograms of 1-μm-diameter silica spheres in water (index 1.333), illuminated with a 660 nm laser. We assign priors as follows. For the refractive index n, we choose a Gaussian prior with a mean of 1.5 and a standard deviation of 0.1, chosen based on the typical variation of refractive index from particle to particle measured in holographic microscopy. Based on manufacturer's specifications,

we choose a Gaussian prior for the particle radius with mean of 0.5 µm and standard deviation of 0.05 µm. We estimate the in-plane particle position (x, y) using a Hough-transform based algorithm[62], and we assign Gaussian priors with standard deviations of 0.1 µm, chosen based on our prior experience with this algorithm. For α, we choose a Gaussian prior centered at 0.7 with width 0.05, based on prior experience from least-squares fitting. We have also used a uniform prior ranging from 0 to 1, with no significant effect on the results.

Choosing a prior distribution for z is the most difficult task; our knowledge of z is limited before we do the full inference calculation (indeed, one of the main goals of the model-based approach to holography is to make it possible to extract precise information about z). We discuss various choices for the prior on z below. In the end, we are able to use a uniform prior from 0 to 100 µm from the focal plane. We chose this width based on the sample height, which is on the order of 100 µm. The prior can be widened or narrowed with little effect on the final results.

We compute the noise (σ in Eq. 5.12) from a background image $\mathbf{b}$ as

$$\sigma = \frac{\mathrm{std}(\mathbf{b})}{\mathrm{mean}(\mathbf{b})}. \tag{5.15}$$

Typically, $\sigma \approx 0.1$ for holograms with no other particles present. For the results presented here, $\sigma = 0.119$.

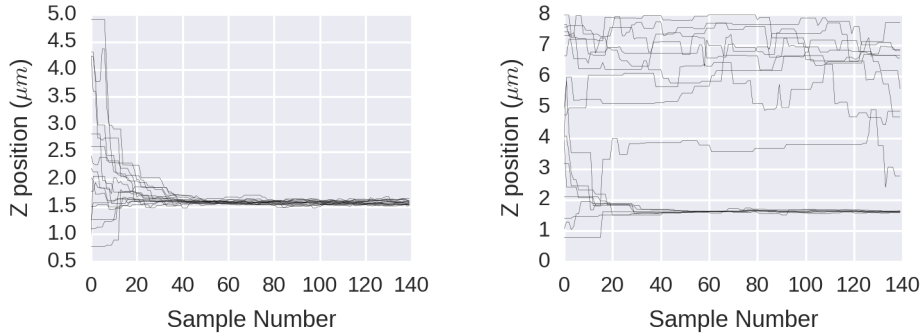

(a) Walkers initially distributed uniformly between 0 and 5 µm

(b) Walkers initially distributed uniformly between 0 and 8 µm

Figure 5.3: Plots of walker position as a function of MCMC time step. As can be seen from the plot on the right, walkers that are further away from the MAP position get stuck and do not reach steady state.

Equilibration of the MCMC chain

With these choices for priors, we run our MCMC algorithm and examine the period of time required to reach steady state, called the "burn-in" period. Evaluating the burn-in period is important because the samples obtained within it are not representative of the posterior and therefore must be discarded.

We find that the number of samples needed to reach steady state is strongly affected by the initial positions of the ensemble of walkers. Figure 5.3a shows that if the initial positions are not too far from the maximum *a posteriori* (MAP) estimate of z, the burn-in is quick. However, if the initial positions are slightly further from the MAP estimate, the MCMC walkers can get stuck in local minima of the posterior distribution and take a long time to equilibrate (Figure 5.3b).

Random-subset tempering

The slow equilibration poses a chicken-and-egg problem for the MCMC method: without a good guess for the position of the particle, the method does not yield a reliable MAP estimate in any reasonable time, but it is difficult to obtain a sufficiently precise guess (within a few micrometers of the MAP position) without knowing the MAP position ahead of time. One way to address this challenge is parallel tempering[102], in which several different MCMC chains are run in parallel at different "temperatures" (a parameter used to vary the sharpness of the probability distributions). Higher-temperature chains can more easily explore the entire probability distribution.

Inspired by this approach, we use a tempering scheme, but one that is simpler to implement and computationally less expensive. Our technique is based on previous work[98] showing that only a small, randomly selected subset of the pixels in a hologram are needed to extract information about the parameters in the model. We therefore start by choosing a small fraction $f = 0.0005$ of pixels (10 pixels out of 20,000 total in the hologram). We then initialize 500 walkers with positions sampled from the prior and run the MCMC algorithm for a short period (30 samples after burn-in, which takes 15 seconds or less on a modern CPU at these low fractions). We then calculate the marginalized probability distribution of z based on this MCMC chain, select a Gaussian distribution around the most likely value of z, and use this distribution to select the walker positions for the next iteration at higher fraction. We then repeat this procedure

at successively higher fractions. At each stage, we change only the initial distribution of walkers, and not the priors. For holograms of single spheres, we find that three preliminary stages (at $f = 0.00025$, $f = 0.00116$, and $f = 0.00539$) are sufficient to obtain a narrow distribution for z, which we then use for a final run at $f = 0.025$. We also find that the posterior distribution does not change significantly at fractions above $f = 0.025$ for the holograms we analyze in this paper.

To illustrate this approach, we analyze a hologram of a 1-µm-diameter polystyrene sphere in water (index 1.33). Figure 5.4 shows that there is a peak in the posterior probability of z even at the lowest pixel fraction. The probability distributions show that the fraction of pixels is analogous to inverse temperature in a tempering scheme; as the fraction increases, the peak in the posterior distribution becomes sharper, and the other modes subside. The prior and posterior distributions at each stage for a 1-µm-diameter silica sphere in water are shown in Figure 5.5. The plots show that one can start with a very wide prior in z (a range of 0 to 100 µm above the focal plane, representing nearly complete ignorance about the position of the particle within the sample volume), and refine the uncertainty in the position down to a fraction of the particle size.

This tempering approach allows us to start with very little information and—in under a minute—obtain an excellent starting point for a final MCMC calculation that requires only a short burn-in. We speed up the computation by doing this final run at a fraction $f = 0.025$, for which the calculation is about 40 times faster than that for the full

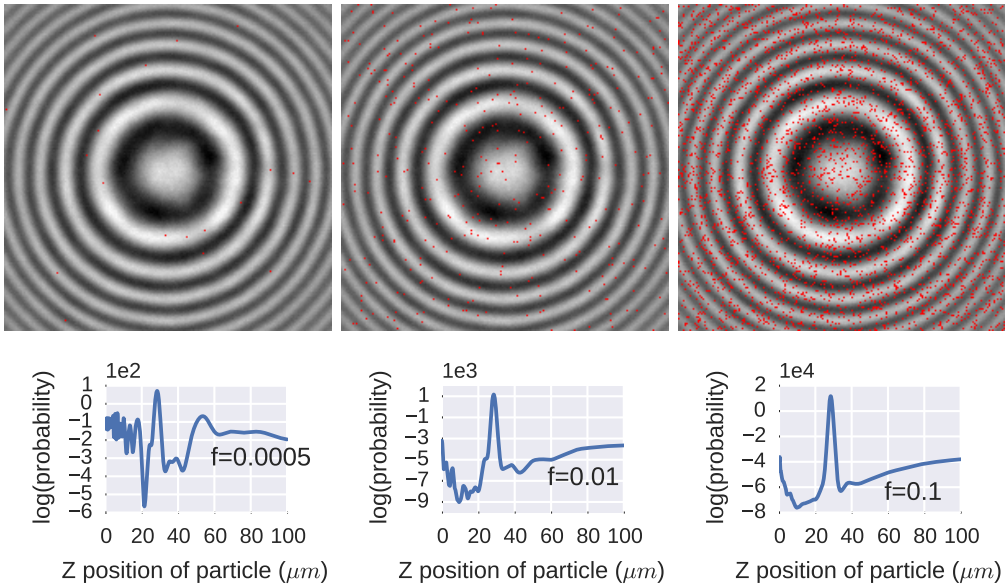

Figure 5.4: Comparison of posterior probability distributions of the parameter z, as obtained through MCMC calculations at three subset fractions. Top row shows a hologram of a 1-μm-diameter polystyrene sphere in water (same hologram in all three columns). The red pixels are the randomly chosen pixels used in the calculations. The bottom row shows the posterior probabilities at the different random-subset fractions f. The peak in the probability gets sharper as the fraction increases, but it is present even at the lowest fraction, which represents only 10 pixels in the hologram.

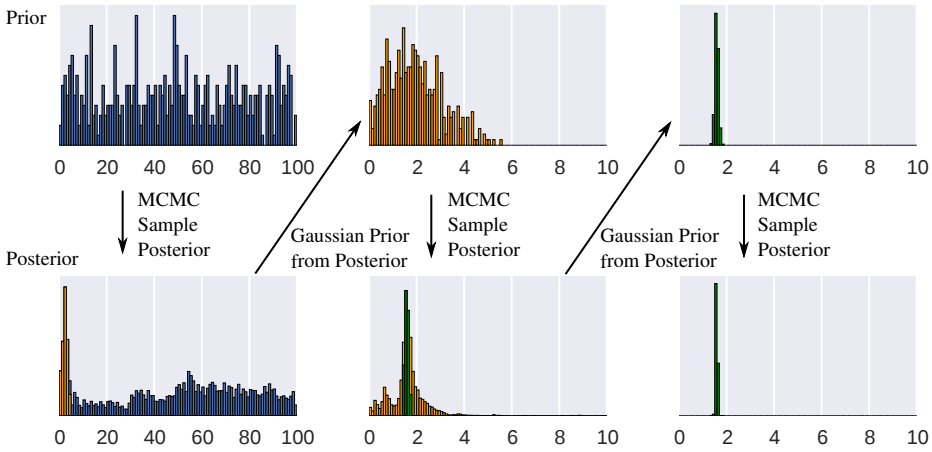

Figure 5.5: The random-subset tempering procedure involves refining position estimates at successively higher fractions, starting from coarse guesses. Top row shows the prior probability distributions of z for each fraction and bottom row the posterior, as obtained from MCMC ensemble sampling. The prior for the second and third stage are chosen from the peak in the posterior for the previous stage. The final posterior distribution is much narrower than the prior used for the first stage; note the change in scale in the horizontal axis, which shows the z position in micrometers.

hologram. We use 500 walkers in the final run. Under these conditions, it takes about 20 minutes on a modern desktop CPU to obtain 25000 samples, roughly 5000 of which are independent. Although the method takes longer than least-squares fitting, it does not require a good initial guess for z (or α). Thus there is less manual intervention required to accurately estimate the parameters, and the procedure can more easily be parallelized.

5.3.2 Particle tracking

With the tempering scheme in hand, we proceed to the particle-tracking problem. We present results from the analysis of a trajectory of a 1-µm silica particle, as described in

section 5.3.1. First, we note that the MAP estimates for the position, as obtained from a single frame of the trajectory, are close to those obtained from least-squares fitting, but with some key differences as shown in Figure 5.6, in which we have marginalized over n, r, and α. The MAP estimates, with boundaries of the 68% credibility interval shown in subscript and superscript, are $x = 17.48^{+0.024}_{-0.026}$ µm, $y = 17.64^{+0.027}_{-0.028}$ µm, and $z = 1.56^{+0.077}_{-0.072}$ µm. There is no appreciable covariance between x and y and only slight covariance between the in-plane coordinate (x, y) and the height z. By comparison, a least-squares fit yields $x = 17.483 \pm 0.0016$ µm, 17.640 ± 0.0026 µm, and $z = 1.569 \pm 0.0019$ µm, where the uncertainties represent one-sigma confidence intervals, as calculated from the parameter covariance matrix.

The least-squares fits yield parameters with smaller uncertainties, but this is because they do not account for the uncertainties in the other parameters. For example, the uncertainty in z is larger in the Bayesian estimate because z is correlated with n, r, and α; in the least-squares fit the uncertainty in z is calculated based on fixed (best-fit) values for these other parameters. The differences between the MAP estimates and the best-fit values likely also arise because the Bayesian approach accounts for the correlations. We therefore argue that the marginalized distributions offer more realistic estimates of the parameters and their uncertainties.

With the posterior distributions obtained for each frame in a time-series, we construct the trajectory of the particle along with the uncertainty by sequential MCMC calculations, in which we use the posterior from one frame in the time-series to inform

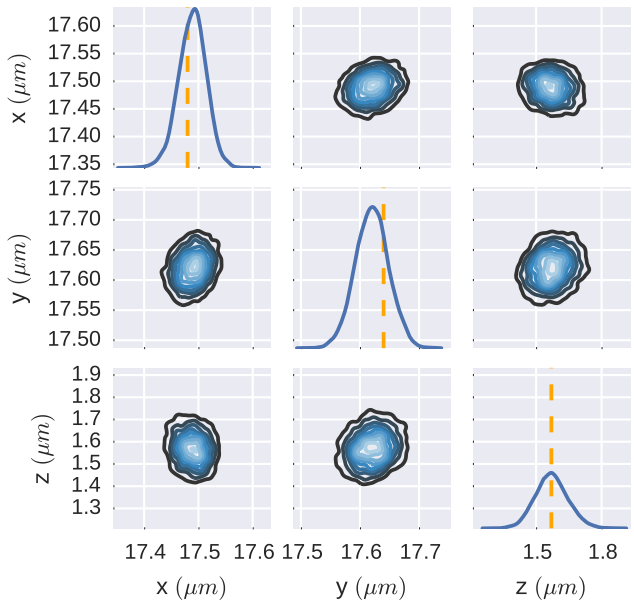

Figure 5.6: Results for particle position estimation from a single frame of a time-series. The marginalized posteriors for x, y, and z obtained from our MCMC scheme are shown along the diagonal, while the off-diagonal contour plots show the covariances (Gaussian kernel density estimates). Orange lines are results from least-squares fitting.

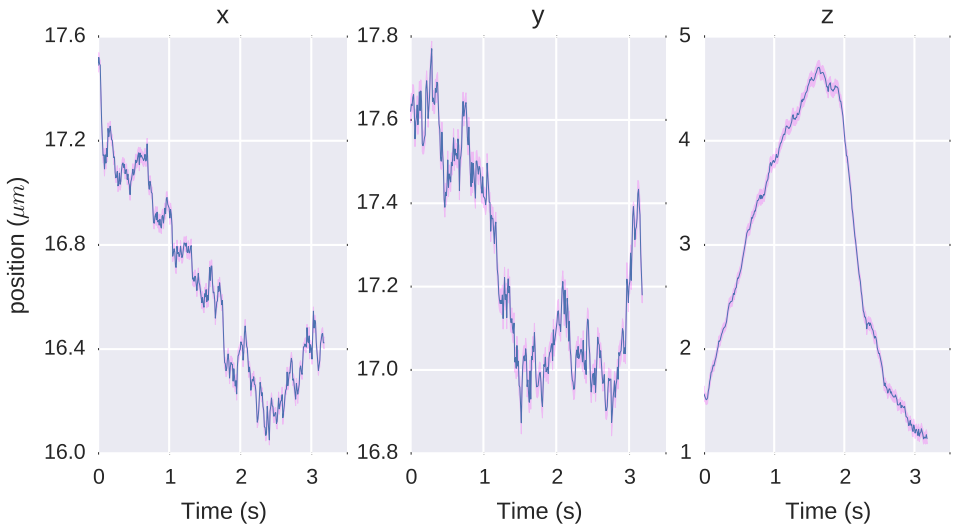

Figure 5.7: Particle trajectory inferred from analysis of a 300-frame time-series (3 seconds total). The blue curve shows the MAP position estimate at each frame, and the shaded magenta region indicates the uncertainty (68% credibility region).

priors for the subsequent frame. Specifically, we approximate the posterior from one frame as a Gaussian and use it as a prior for n, r, and α for the next frame. For x, y, z, we use a Gaussian prior centered on the MAP position of the previous frame and assign it a width of 0.1 µm to allow for diffusion. The results of our procedure are shown in Figure 5.7.

The uncertainty in the trajectory is important for hypothesis testing or model selection. In this data set the particle is subject to Brownian motion and forces from an optical trap and gravity. In other, similar experiments, one might wish to compare the actual trajectory to an equation of motion derived from a model of the forces. To quantitatively determine whether a particular equation of motion accurately models the data, one needs to know the uncertainty on the trajectory.

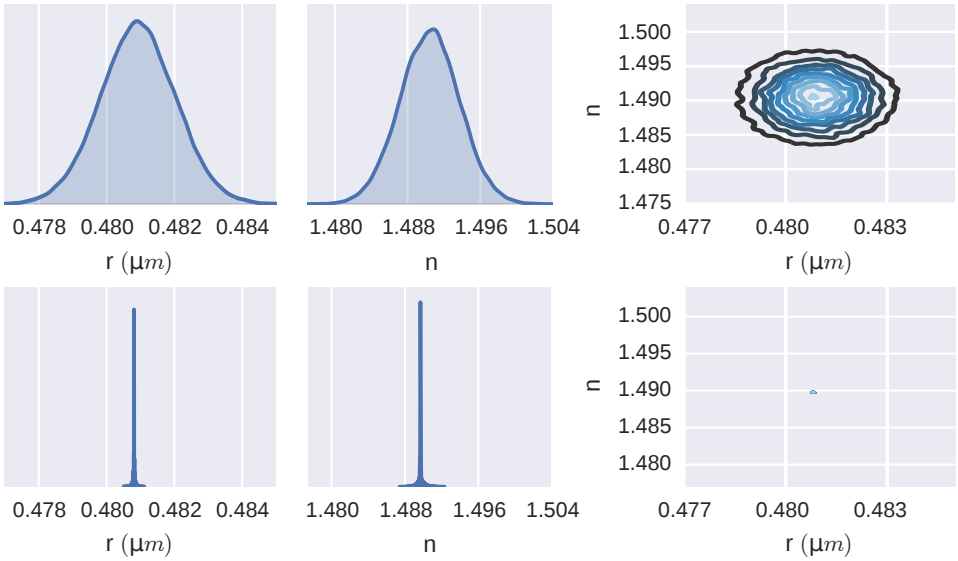

Figure 5.8: Posterior distributions of particle properties (refractive index n and radius r) inferred from analysis of one frame (top row) and ten successive frames (bottom row) of a time-series of holograms. Plots are Gaussian kernel density estimates of MCMC samples.

5.3.3 Particle characterization

We apply the same tempered MCMC scheme to the particle characterization problem, using the same data set as for the tracking analysis. We show the results from a single frame and those from 10 frames in Figure 5.8. For the 10-frame analysis, we use the posterior for one frame as the prior for the subsequent frame in the time-series, then we combine the MCMC samples from the 10 individual frames. The resulting posterior distribution is significantly narrower (by a factor of about 10 for n and 26 for r), thus yielding much better estimates for the refractive index and radius than could be obtained from a single frame.

The refinement of the particle parameters is important in applications such as particle

sizing. In holographic particle sizing, the size distribution is measured by combining measurements for individual particles. It is therefore important that the uncertainty on the each measured particle be as small as possible, so that narrow size distributions and multimodal distributions can be resolved. In the 10-frame measurement we show in Figure 5.8, the uncertainty relative to the MAP radius is 0.016%, which (as desired) is much smaller than the width of even a monodisperse colloidal size distribution.

5.4 Conclusions

The Bayesian approach we have demonstrated is well-suited to the problem of extracting information from holograms of colloidal particles. One reason is that we understand well how particles scatter light, and there are accurate ways to calculate the scattering. Another is that the technique captures essential details of the scattered field, such as the phase information that is encoded in the interference fringes. All of this means that accurate forward models for hologram formation can be developed from first principles. With a good forward model, MCMC schemes like the one we demonstrate here can be brought to bear on parameter estimation.

The main disadvantage of the Bayesian approach compared to least-squares fitting is that the computations take much longer. This disadvantage is outweighed, in our opinion, by the many advantages of the approach. First, prior information can be specified explicitly at the outset, and need not be spread over initial guesses, constraints, and penalty functions. Second, the tempered MCMC scheme we have developed does

not require accurate initial guesses, thus eliminating much of the manual effort required to tune the guesses. Third, the analysis yields the entire joint probability distribution of the parameters, so that the uncertainty can be quantified accurately and marginalized to account for uncertainties in parameters that are not of interest. Thus the method can be applied, as we have shown, to extract realistic uncertainties on the trajectory of a particle or on its properties, taking into account data over an entire time-series.

But we have not demonstrated what could be the most important advantage of the Bayesian approach: its adaptability to other, more complex forward models. In the future, we aim to apply this technique to more complex scattering models such as the superposition solution for multiple spheres, which allows us to infer the structure and dynamics of colloidal clusters[30]. We would also like to model correlated noise (speckle) in the holograms as well as other optical effects arising from propagation along the optical train of the microscope. A more detailed description of noise and imaging effects should allow us to model the hologram more accurately than the current formulation, which lumps these effects into a single parameter, α.

All of these extensions to the forward model will require more parameters, which is problematic for least-squares analysis but much less so for the Bayesian approach. In least-squares techniques, most of the computational time is spent evaluating the Hessian, which scales as the number of parameters squared. Furthermore, adding parameters complicates the chi-squared surface, producing local minima corresponding to unphysical solutions. It becomes more and more difficult to tune the least-squares algo-

rithm as the number of parameters increases, even when there is no risk of overfitting. In contrast, the MCMC techniques used in a Bayesian approach do not rely on gradients and are designed to handle models with hundreds or even thousands of parameters. Also, tempering methods allow one to explore the entire probability distribution, including local maxima. Thus we believe that tempered MCMC approaches like the one we describe in this paper will scale better to more complex models.

Finally, although the approach we describe is by no means simple, it is easy to use because it requires very little tuning. We therefore hope that it will be useful to other researchers considering the use of holographic microscopy. The subset tempering algorithm makes it straightforward for a scientist who is not experienced with holography to make measurements. To facilitate the adoption of these techniques, we have implemented them in our open-source hologram analysis library HoloPy[33].

Funding

We acknowledge support from the National Science Foundation through grant number DMR-1306410.

Acknowledgments

We thank Anna Wang for providing the raw holograms that we analyze in the paper.

6

Conclusion

This thesis shows a number of advances that should help holography to be useful to many more scientists working on many more problems. In Chapter 2 I demonstrated a simple, inexpensive, self-contained holographic microscope suitable for portable use and inside environments such as cell incubators. In Chapter 3 I showed how to use the discrete dipole approximation to analyze holograms of non-spherical scatterers, a

previously inaccessible regime. In Chapter 4 I showed an elegant technique for reducing the computation time required for analysis of holograms to a much more manageable duration. Finally, in Chapter 5 I introduced a Bayesian analysis paradigm which more naturally handles known prior information, allows marginalizing out nuisance parameters, and eliminates the need for manual initial guesses for parameter inference. Taken together, the advances in this thesis and my open source implementations of them in the software package HoloPy should allow a broader range of scientists to use high precision holographic measurements on a broader range of problems.

There is still work to be done in this field. Our scattering models still include a nuisance parameter, α, which should be replaced with physically motivated parameters obtained from a proper analysis of the holographic imaging system. Our ability to parameterize shapes still limits the objects we can analyze even with the DDA. My work with Ana Carpio and coworkers on topological techniques offers one route forward in this space[103], but this is still an open problem.

The DDA works well for small objects, but computation becomes prohibitive for larger objects. In order to work with holograms of objects the size of mammalian cells we will need another scattering theory, perhaps one based on Rayleigh-Gans or other models.

Shot noise from the reference wave also limits our ability to take holograms of small, weakly scattering objects. This situation can be improved by reducing the intensity of the reference wave relative to the scattered field from the object of interest. This can

be done by introducing an attenuation filter that selectively attenuates the reference field[104,105] or by using the reflection from a coverglass as the reference[106–110]. In some cases it also may be possible to increase the scattering signal of small objects by attaching dyes to the object whose resonant effect with incident light yields a signal out of proportion to their size. My colleague Aaron Goldfain is working on these approaches.

Holography works without any labeling of samples, giving contrast to refractive index differences similar to differential interference contrast. This is an advantage of holography over other three-dimensional imaging techniques like confocal microscopy that require labels. However, it is also a disadvantage because selective labeling of structures of interest is an extremely powerful tool for many investigations. Nonetheless, it may be possible to use labels in a two-color holographic instrument and see refractive or absorptive effects from attached dyes. If the dyes and wavelengths of incident light overlap, a dye should have a strong signal in only one channel[111].

Holography also is difficult to use on optically thick samples where very large amounts of light are scattered. Switching to an off axis-geometry may help, though at the cost of increased resolution requirements. For scenes without too much multiple scattering it might still be possible to model and infer parameters, but that will require efficient parameterizations, efficient scattering models, significant advances in available computing power, and probably new, more efficient inference techniques.

In spite of these limitations, holography has the potential to become an increasingly useful imaging and measurement tool. It is already useful for many problems,

particularly tracking three-dimensional dynamics at high speed. Additionally, because holography uses simple hardware and moves most of image formation and measurement into post-processing in software, it can easily benefit from advances in computing, and new techniques and improvements can be readily adopted.

References

[1] Nobel Media AB 2015, "All Nobel Prizes in Physics." http://www.nobelprize.org/nobel_prizes/physics/laureates/, 2016.

[2] D. Gabor, "A new microscopic principle," *Nature*, vol. 161, no. 4098, pp. 777–778, 1948.

[3] U. Schnars and W. P. Juptner, "Digital recording and numerical reconstruction of holograms," *Measurement Science and Technology*, vol. 13, no. 9, pp. 85–101, 2002.

[4] D. M. Kaz, R. McGorty, M. Mani, M. P. Brenner, and V. N. Manoharan, "Physical ageing of the contact line on colloidal particles at liquid interfaces," *Nature Materials*, vol. 11, pp. 138–142, Feb. 2012.

[5] S. K. Jericho, J. Garcia-Sucerquia, W. Xu, M. H. Jericho, and H. J. Kreuzer, "Submersible digital in-line holographic microscope," *Review of Scientific Instruments*, vol. 77, pp. 043706–10, Apr. 2006.

[6] N. Lewis, W. Xu, S. Jericho, H. Kreuzer, M. Jericho, and A. Cembella, "Swimming speed of three species of Alexandrium (Dinophyceae) as determined by digital in-line holography," *Phycologia*, vol. 45, pp. 61–70, Jan. 2006.

[7] J. Sheng, E. Malkiel, J. Katz, J. Adolf, R. Belas, and A. R. Place, "Digital holographic microscopy reveals prey-induced changes in swimming behavior of predatory dinoflagellates," *Proceedings of the National Academy of Sciences*, vol. 104, pp. 17512–17517, Oct. 2007.

[8] H. J. Kreuzer, M. J. Jericho, I. A. Meinertzhagen, and W. Xu, "Digital in-line holography with photons and electrons," *Journal of Physics: Condensed Matter*, vol. 13, no. 47, pp. 10729–10741, 2001.

[9] L. Repetto, E. Piano, and C. Pontiggia, "Lensless digital holographic microscope with light-emitting diode illumination," *Optics Letters*, vol. 29, no. 10, pp. 1132–1134, 2004.

[10] W. Xu, M. H. Jericho, H. J. Kreuzer, and I. A. Meinertzhagen, "Tracking particles in four dimensions with in-line holographic microscopy," *Optics Letters*, vol. 28, pp. 164–166, Feb. 2003.

[11] J. Sheng, E. Malkiel, and J. Katz, "Digital holographic microscope for measuring three-dimensional particle distributions and motions," *Applied Optics*, vol. 45, no. 16, pp. 3893–3901, 2006.

[12] S. Satake, T. Kunugi, K. Sato, T. Ito, H. Kanamori, and J. Taniguchi, "Measurements of 3D flow in a micro-pipe via micro digital holographic particle tracking velocimetry," *Measurement Science and Technology*, vol. 17, no. 7, pp. 1647–1651, 2006.

[13] Y.-S. Choi and S.-J. Lee, "Three-dimensional volumetric measurement of red blood cell motion using digital holographic microscopy," *Applied Optics*, vol. 48, pp. 2983–2990, June 2009.

[14] S. Seo, T. W. Su, D. K. Tseng, A. Erlinger, and A. Ozcan, "Lensfree holographic imaging for on-chip cytometry and diagnostics," *Lab on a Chip*, vol. 9, no. 6, pp. 777–787, 2009.

[15] W. Xu, M. H. Jericho, I. A. Meinertzhagen, and H. J. Kreuzer, "Digital in-line holography for biological applications," *Proceedings of the National Academy of Sciences*, vol. 98, no. 20, p. 11301, 2001.

[16] I. Moon, M. Daneshpanah, B. Javidi, and A. Stern, "Automated Three-Dimensional Identification and Tracking of Micro/Nanobiological Organisms by Computational Holographic Microscopy," *Proceedings of the IEEE*, vol. 97, pp. 990–1010, June 2009.

[17] F. Dubois, C. Yourassowsky, O. Monnom, J.-C. Legros, O. Debeir, P. Van Ham, R. Kiss, and C. Decaestecker, "Digital holographic microscopy for the three-dimensional dynamic analysis of in vitro cancer cell migration," *Journal of Biomedical Optics*, vol. 11, no. 5, p. 054032, 2006.

[18] Y. Pu and H. Meng, "Intrinsic aberrations due to Mie scattering in particle holography," *Journal of the Optical Society of America A*, vol. 20, pp. 1920–1932, Oct. 2003.

[19] B. Ovryn and S. H. Izen, "Imaging of transparent spheres through a planar interface using a high-numerical-aperture optical microscope," *Journal of the Optical Society of America A*, vol. 17, pp. 1202–1213, July 2000.

[20] B. Ovryn, "Three-dimensional forward scattering particle image velocimetry applied to a microscopic field-of-view," *Experiments in Fluids*, vol. 29, pp. S175–S184, Dec. 2000.

[21] S.-H. Lee, Y. Roichman, G.-R. Yi, S.-H. Kim, S.-M. Yang, A. van Blaaderen, P. van Oostrum, and D. G. Grier, "Characterizing and tracking single colloidal particles with video holographic microscopy," *Optics Express*, vol. 15, pp. 18275–18282, Dec. 2007.

[22] D. W. Marquardt, "An Algorithm for Least-Squares Estimation of Nonlinear Parameters," *Journal of the Society for Industrial and Applied Mathematics*, vol. 11, no. 2, pp. 431–441, 1963.

[23] H. W. Moyses, B. J. Krishnatreya, and D. G. Grier, "Robustness of Lorenz-Mie microscopy against defects in illumination," *Optics Express*, vol. 21, p. 5968, Mar. 2013.

[24] C. Wang, X. Zhong, D. B. Ruffner, A. Stutt, L. A. Philips, M. D. Ward, and D. G. Grier, "Holographic Characterization of Protein Aggregates," *Journal of Pharmaceutical Sciences*, vol. 105, pp. 1074–1085, Mar. 2016.

[25] M. Hannel, C. Middleton, and D. G. Grier, "Holographic characterization of imperfect colloidal spheres," *Applied Physics Letters*, vol. 107, p. 141905, Oct. 2015.

[26] F. C. Cheong, K. Xiao, D. J. Pine, and D. G. Grier, "Holographic characterization of individual colloidal spheres' porosities," *Soft Matter*, vol. 7, pp. 6816–6819, July 2011.

[27] C. Wang, H. W. Moyses, and D. G. Grier, "Stimulus-responsive colloidal sensors with fast holographic readout," *Applied Physics Letters*, vol. 107, p. 051903, Aug. 2015.

[28] C. Wang, H. Shpaisman, A. D. Hollingsworth, and D. G. Grier, "Celebrating Soft Matter's 10th Anniversary: Monitoring colloidal growth with holographic microscopy," *Soft Matter*, vol. 11, no. 6, pp. 1062–1066, 2015.

[29] H. Shpaisman, B. J. Krishnatreya, and D. G. Grier, "Holographic microrefractometer," *Applied Physics Letters*, vol. 101, p. 091102, Aug. 2012.

[30] R. W. Perry, G. Meng, T. G. Dimiduk, J. Fung, and V. N. Manoharan, "Real-space studies of the structure and dynamics of self-assembled colloidal clusters," *Faraday Discussions*, 2012.

[31] J. Fung, R. W. Perry, T. G. Dimiduk, and V. N. Manoharan, "Imaging Multiple Colloidal Particles by Fitting Electromagnetic Scattering Solutions to Digital Holograms," *Journal of Quantitative Spectroscopy and Radiative Transfer*, vol. 113, pp. 2482–2489, Dec. 2012.

[32] E. M. Purcell and C. R. Pennypacker, "Scattering and Absorption of Light by Nonspherical Dielectric Grains," *The Astrophysical Journal*, vol. 186, pp. 705–714, Dec. 1973.

[33] T. G. Dimiduk, J. Fung, R. W. Perry, and V. N. Manoharan, "HoloPy - Hologram Processing and Light Scattering in Python." https://github.com/manoharan-lab/holopy, 2016.

[34] J. B. Pawley, *Handbook of Biological Confocal Microscopy*. New York: Springer, 3rd ed., 2006.

[35] C. Depeursinge, T. Colomb, Y. Emery, J. Kühn, F. Charrière, B. Rappaz, and P. Marquet, "Digital holographic microscopy applied to life sciences," in *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, vol. 2007, pp. 6244–6247, 2007.

[36] F. Dubois, L. Joannes, and J. C. Legros, "Improved three-dimensional imaging with a digital holography microscope with a source of partial spatial coherence," *Applied Optics*, vol. 38, p. 34, 1999.

[37] Y. Emery, E. Cuche, T. Colomb, C. Depeursinge, B. Rappaz, P. Marquet, and P. Magistretti, "DHM (Digital Holography Microscope) for imaging cells," in *Journal of Physics: Conference Series*, vol. 61, pp. 1317–1321, 2007.

[38] C. Mann, L. Yu, and M. Kim, "Movies of cellular and sub-cellular motion by digital holographic microscopy," *BioMedical Engineering OnLine*, vol. 5, no. 1, p. 21, 2006.

[39] J. Goodman, *Introduction to Fourier optics*. Englewood Colo.: Roberts & Co., 3rd ed. ed., 2005.

[40] D. Coffin, "Decoding raw digital photos in Linux." http://www.cybercom.net/~dcoffin/dcraw/, 2010.

[41] T. Oliphant, P. Peterson, J. Eric, *et al.*, "SciPy: Open source scientific tools for Python," 2001.

[42] F. Pampaloni, E. G. Reynaud, and E. H. K. Stelzer, "The third dimension bridges the gap between cell culture and live tissue," *Nat. Rev. Mol. Cell Biol.*, vol. 8, pp. 839–845, Oct. 2007.

[43] S.-H. Lee and D. G. Grier, "Holographic microscopy of holographically trapped three-dimensional structures," *Optics Express*, vol. 15, pp. 1505–1512, Feb. 2007.

[44] D. A. Weitz, J. S. Huang, M. Y. Lin, and J. Sung, "Dynamics of Diffusion-Limited Kinetic Aggregation," *Physical Review Letters*, vol. 53, pp. 1657–1660, Oct. 1984.

[45] S. C. Glotzer, M. J. Solomon, and N. A. Kotov, "Self-assembly: From nanoscale to microscale colloids," *AIChE Journal*, vol. 50, no. 12, pp. 2978–2985, 2004.

[46] P. L. Biancaniello and J. C. Crocker, "Line optical tweezers instrument for measuring nanoscale interactions and kinetics," *Rev. Sci. Inst.*, vol. 77, p. 113702, 2006.

[47] D. T. Chen, E. R. Weeks, J. C. Crocker, M. F. Islam, R. Verma, J. Gruber, A. J. Levine, T. C. Lubensky, and A. G. Yodh, "Rheological Microscopy: Local Mechanical Properties from Microrheology," *Physical Review Letters*, vol. 90, p. 108301, Mar. 2003.

[48] S. M. Anthony, L. Hong, M. Kim, and S. Granick, "Single-Particle Colloid Tracking in Four Dimensions," *Langmuir*, vol. 22, pp. 9812–9815, Nov. 2006.

[49] G. L. Hunter, K. V. Edmond, M. T. Elsesser, and E. R. Weeks, "Tracking rotational diffusion of colloidal clusters," *Optics Express*, vol. 19, pp. 17189–17202, Aug. 2011.

[50] D. J. Kraft, R. Wittkowski, B. ten Hagen, K. V. Edmond, D. J. Pine, and H. Löwen, "Brownian motion and the hydrodynamic friction tensor for colloidal particles of complex shape," *Physical Review E*, vol. 88, p. 050301, Nov. 2013.

[51] T. M. Kreis, "Frequency analysis of digital holography with reconstruction by convolution," *Opt. Eng.*, vol. 41, no. 8, pp. 1829–1839, 2002.

[52] F. C. Cheong and D. G. Grier, "Rotational and translational diffusion of copper oxide nanorods measured with holographic video microscopy," *Opt. Express*, vol. 18, pp. 6555–6562, Mar. 2010.

[53] F. C. Cheong, B. J. Krishnatreya, and D. G. Grier, "Strategies for three-dimensional particle tracking with holographic video microscopy," *Optics Express*, vol. 18, pp. 13563–13573, June 2010.

[54] J. Fung, K. E. Martin, R. W. Perry, D. M. Kaz, R. McGorty, and V. N. Manoharan, "Measuring translational, rotational, and vibrational dynamics in colloids with digital holographic microscopy," *Optics Express*, vol. 19, p. 8051, Apr. 2011.

[55] J. Fung and V. N. Manoharan, "Holographic measurements of anisotropic three-dimensional diffusion of colloidal clusters," *Phys. Rev. E*, vol. 88, p. 020302, Aug. 2013.

[56] D. W. Mackowski and M. I. Mishchenko, "Calculation of the T matrix and the scattering matrix for ensembles of spheres," *Journal of the Optical Society of America A*, vol. 13, pp. 2266–2278, Nov. 1996.

[57] K. Chaudhary, Q. Chen, J. J. Juárez, S. Granick, and J. A. Lewis, "Janus Colloidal Matchsticks," *J. Am. Chem. Soc.*, vol. 134, pp. 12901–12903, Aug. 2012.

[58] A. Kuijk, A. van Blaaderen, and A. Imhof, "Synthesis of Monodisperse, Rodlike Silica Colloids with Tunable Aspect Ratio," *J. Am. Chem. Soc.*, vol. 133, pp. 2346–2349, Mar. 2011.

[59] W. Stöber, A. Fink, and E. Bohn, "Controlled growth of monodisperse silica spheres in the micron size range," *Journal of Colloid and Interface Science*, vol. 26, pp. 62–69, Jan. 1968.

[60] J. H. K. Song and I. Kretzschmar, "Assembled Surface-Anisotropic Colloids as a Template for a Multistage Catalytic Membrane Reactor," *ACS Appl. Mater. Interfaces*, vol. 1, pp. 1747–1754, Aug. 2009.

[61] M. A. Yurkin and A. G. Hoekstra, "The discrete-dipole-approximation code ADDA: Capabilities and known limitations," *Journal of Quantitative Spectroscopy and Radiative Transfer*, vol. 112, pp. 2234–2247, Sept. 2011.

[62] F. C. Cheong, B. S. Rémi Dreyfus, J. Amato-Grill, K. Xiao, L. Dixon, and D. G. Grier, "Flow visualization and flow cytometry with holographic video microscopy," *Optics Express*, vol. 17, p. 13071, July 2009.

[63] H. Flyvbjerg and H. G. Petersen, "Error estimates on averages of correlated data," *Journal of Chemical Physics*, vol. 91, p. 461, July 1989.

[64] I. Martchenko, H. Dietsch, C. Moitzi, and P. Schurtenberger, "Hydrodynamic Properties of Magnetic Nanoparticles with Tunable Shape Anisotropy: Prediction and Experimental Verification," *J. Phys. Chem. B*, vol. 115, pp. 14838–14845, Dec. 2011.

[65] T. Yoshizaki and H. Yamakawa, "Dynamics of spheroid-cylindrical molecules in dilute solution," *Journal of Chemical Physics*, vol. 72, p. 57, Jan. 1980.

[66] T. Norisuye, M. Motowoka, and H. Fujita, "Wormlike Chains Near the Rod Limit: Translational Friction Coefficient," *Macromolecules*, vol. 12, pp. 320–323, Mar. 1979.

[67] M. M. Tirado, C. L. Martínez, and J. G. de la Torre, "Comparison of theories for the translational and rotational diffusion coefficients of rod-like macromolecules. Application to short DNA fragments," *Journal of Chemical Physics*, vol. 81, p. 2047, Aug. 1984.

[68] H. Brenner, "Taylor dispersion in systems of sedimenting nonspherical brownian particles. I. Homogeneous, centrosymmetric, axisymmetric particles," *Journal of Colloid and Interface Science*, vol. 71, pp. 189–208, Sept. 1979.

[69] A. Einstein, *Investigations on the Theory of the Brownian Movement*. Dover, New York, 1956.

[70] P. Debye, *Polar molecules*. Dover, New York, 1929.

[71] M. Gittings and D. Saville, "The determination of hydrodynamic size and zeta potential from electrophoretic mobility and light scattering measurements," *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, vol. 141, pp. 111–117, Oct. 1998.

[72] J. E. Seebergh and J. C. Berg, "Evidence of a hairy layer at the surface of polystyrene latex particles," *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, vol. 100, pp. 139–153, July 1995.

[73] S. D'Agostino, P. P. Pompa, R. Chiuri, R. J. Phaneuf, D. G. Britti, R. Rinaldi, R. Cingolani, and F. Della Sala, "Enhanced fluorescence by metal nanospheres on metal substrates," *Opt. Lett.*, vol. 34, pp. 2381–2383, Aug. 2009.

[74] J. Fricks, L. Yao, T. C. Elston, and M. G. Forest, "Time-domain methods for diffusive transport in soft matter," *SIAM Journal on Applied Mathematics*, vol. 69, no. 5, pp. 1277–1308, 2009.

[75] B. J. Berne and R. Pecora, *Dynamic light scattering: with applications to chemistry, biology, and physics*. Wiley, 1976.

[76] P. M. Morse and H. Feshbach, *Methods of theoretical physics*. New York: McGraw-Hill, 1953.

[77] A. R. DiDonato, "Recurrence Relations for the Indefinite Integrals of the Associated Legendre Functions," *Mathematics of Computation*, vol. 38, p. 547, Apr. 1982.

[78] J. Saragosti, P. Silberzan, and A. Buguin, "Modeling E. coli Tumbles by Rotational Diffusion. Implications for Chemotaxis," *PLoS ONE*, vol. 7, p. e35412, Apr. 2012.

[79] D. A. Beard and T. Schlick, "Unbiased Rotational Moves for Rigid-Body Dynamics," *Biophysical Journal*, vol. 85, pp. 2973–2976, Nov. 2003.

[80] S. Yeom, I. Moon, and B. Javidi, "Real-Time 3-D Sensing, Visualization and Recognition of Dynamic Biological Microorganisms," *Proceedings of the IEEE*, vol. 94, no. 3, pp. 550–566, 2006.

[81] W. Xu, M. Jericho, I. Meinertzhagen, and H. Kreuzer, "Digital in-line holography of microspheres," *Applied Optics*, vol. 41, no. 25, pp. 5367–5375, 2002.

[82] S. Satake, T. Kunugi, K. Sato, T. Ito, and J. Taniguchi, "Three-Dimensional Flow Tracking in a Micro Channel with High Time Resolution Using Micro Digital-Holographic Particle-Tracking Velocimetry," *Optical Review*, vol. 12, pp. 442–444, Nov. 2005.

[83] J. P. Fugal, T. J. Schulz, and R. A. Shaw, "Practical methods for automated reconstruction and characterization of particles in digital in-line holograms," *Meas. Sci. Technol.*, vol. 20, no. 075501, p. 075501, 2009.

[84] L. Dixon, F. C. Cheong, and D. G. Grier, "Holographic deconvolution microscopy for high-resolution particle tracking," *Optics Express*, vol. 19, pp. 16410–16417, Aug. 2011.

[85] A. Wang, T. G. Dimiduk, J. Fung, S. Razavi, I. Kretzschmar, K. Chaudhary, and V. N. Manoharan, "Using the discrete dipole approximation and holographic microscopy to measure rotational dynamics of non-spherical colloidal particles," *Journal of Quantitative Spectroscopy and Radiative Transfer*, vol. 146, pp. 499–509, Oct. 2014.

[86] M. Seifi, C. Fournier, L. Denis, D. Chareyron, and J.-L. Marié, "Three-dimensional reconstruction of particle holograms: a fast and accurate multiscale approach," *Journal of the Optical Society of America A*, vol. 29, pp. 1808–1817, Sept. 2012.

[87] M. Seifi, C. Fournier, N. Grosjean, L. Méès, J.-L. Marié, and L. Denis, "Accurate 3D tracking and size measurement of evaporating droplets using in-line digital holography and "inverse problems" reconstruction approach," *Optics Express*, vol. 21, pp. 27964–27980, Nov. 2013.

[88] D. W. Mackowski, "Calculation of total cross sections of multiple-sphere clusters," *Journal of the Optical Society of America A*, vol. 11, no. 11, pp. 2851–2861, 1994.

[89] E. Candes and T. Tao, "Near-Optimal Signal Recovery From Random Projections: Universal Encoding Strategies?," *IEEE Transactions on Information Theory*, vol. 52, pp. 5406–5425, Dec. 2006.

[90] D. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[91] R. G. Baraniuk and M. B. Wakin, "Random Projections of Smooth Manifolds," *Foundations of Computational Mathematics*, vol. 9, pp. 51–77, Feb. 2009.

[92] C. B. Markwardt, "Non-linear Least-squares Fitting in IDL with MPFIT," in *Astronomical Data Analysis Software and Systems XVIII*, vol. 411, p. 251, 2009.

[93] F. C. Cheong, S. Duarte, S.-H. Lee, and D. G. Grier, "Holographic microrheology of polysaccharides from Streptococcus mutans biofilms," *Rheologica Acta*, vol. 48, pp. 109–115, Jan. 2009.

[94] L. Denis, D. Lorenz, E. Thiébaut, C. Fournier, and D. Trede, "Inline hologram reconstruction with sparsity constraints," *Optics Letters*, vol. 34, pp. 3475–3477, Nov. 2009.

[95] F. Soulez, L. Denis, C. Fournier, É. Thiébaut, and C. Goepfert, "Inverse-problem approach for particle digital holography: accurate location based on local optimization," *JOSA A*, vol. 24, pp. 1164–1171, Apr. 2007.

[96] A. Yevick, M. Hannel, and D. G. Grier, "Machine-learning approach to holographic particle characterization," *Optics Express*, vol. 22, pp. 26884–26890, Nov. 2014.

[97] T. Latychevskaia, F. Gehri, and H.-W. Fink, "Depth-resolved holographic reconstructions by three-dimensional deconvolution," *Optics Express*, vol. 18, pp. 22527–22544, Oct. 2010.

[98] T. G. Dimiduk, R. W. Perry, J. Fung, and V. N. Manoharan, "Random-subset fitting of digital holograms for fast three-dimensional particle tracking," *Applied Optics*, vol. 53, pp. G177–G183, Sept. 2014.

[99] P. Gregory, *Bayesian Logical Data Analysis for the Physical Sciences: A Comparative Approach with Mathematica® Support*. Cambridge University Press, Apr. 2005.

[100] J. Goodman and J. Weare, "Ensemble samplers with affine invariance," *Communications in Applied Mathematics and Computational Science*, vol. 5, pp. 65–80, Jan. 2010.

[101] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, "emcee: The MCMC Hammer," *Publications of the Astronomical Society of the Pacific*, vol. 125, pp. 306–312, Mar. 2013.

[102] D. J. Earl and M. W. Deem, "Parallel tempering: Theory, applications, and new perspectives," *Physical Chemistry Chemical Physics (Incorporating Faraday Transactions)*, vol. 7, p. 3910, 2005.

[103] A. Carpio, T. G. Dimiduk, M. L. Rapún, and V. Selgas, "Noninvasive imaging of 3D micro and nanostructures by topological methods," *SIAM Journal on Imaging Sciences*, In Press.

[104] Q. D. Pham, Y. Kusumi, S. Hasegawa, and Y. Hayasaki, "Digital holographic microscope with low-frequency attenuation filter for position measurement of a nanoparticle," *Optics Letters*, vol. 37, p. 4119, Oct. 2012.

[105] K. Goto and Y. Hayasaki, "Three-dimensional motion detection of a 20-nm gold nanoparticle using twilight-field digital holography with coherence regulation," *Optics Letters*, vol. 40, p. 3344, July 2015.

[106] V. Jacobsen, P. Stoller, C. Brunner, V. Vogel, and V. Sandoghdar, "Interferometric optical detection and tracking of very small gold nanoparticles at a water-glass interface," *Optics Express*, vol. 14, no. 1, p. 405, 2006.

[107] J. Ortega-Arroyo and P. Kukura, "Interferometric scattering microscopy (iSCAT): new frontiers in ultrafast and ultrasensitive optical microscopy," *Physical Chemistry Chemical Physics*, vol. 14, no. 45, pp. 15625–15636, 2012.

[108] M. Piliarik and V. Sandoghdar, "Direct optical sensing of single unlabelled proteins and super-resolution imaging of their binding sites," *Nature Communications*, vol. 5, p. 4495, July 2014.

[109] P. Kukura, H. Ewers, C. Müller, A. Renn, A. Helenius, and V. Sandoghdar, "High-speed nanoscopic tracking of the position and orientation of a single virus," *Nature Methods*, vol. 6, pp. 923–927, Dec. 2009.

[110] A. M. Goldfain, R. F. Garmann, Y. Jin, Y. Lahini, and V. N. Manoharan, "Dynamic Measurements of the Position, Orientation, and DNA Content of Individual Unlabeled Bacteriophages," *The Journal of Physical Chemistry B*, vol. 120, pp. 6130–6138, July 2016.

[111] J. L. Nadeau, Y. B. Cho, and C. A. Lindensmith, "Use of dyes to increase phase contrast for biological holographic microscopy," *Optics Letters*, vol. 40, p. 4114, Sept. 2015.

This thesis was typeset using LaTeX, originally developed by Leslie Lamport and based on Donald Knuth's TeX. A template that can be used to format a PhD dissertation with this look & feel has been released under the permissive AGPL license, and can be found online at github.com/suchow/Dissertate or from its lead author, Jordan Suchow, at suchow@post.harvard.edu.