Studying Electronic Textures with Coherent Lensless Imaging

by

Abraham Lewis Levitan

B.S., Olin College of Engineering (2016)

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2023

Author . . . .
Department of Physics
April 24, 2023

Certified by . . . .
Riccardo Comin
Class of 1947 Career Development Associate Professor of Physics
Thesis Supervisor

Accepted by . . . .
Lindley Winslow
Associate Department Head of Physics

Studying Electronic Textures with Coherent Lensless Imaging

by

Abraham Lewis Levitan

Submitted to the Department of Physics
on April 24, 2023, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

X-ray microscopes have opened our collective eyes to the richness of nanoscale texture in systems such as correlated and quantum materials. These microscopes draw their power from the combination of short wavelengths, which provide high resolution, and interaction with atomic resonances, which makes them sensitive to subtle changes in electronic structure. However, x-ray microscopy remains an area where the main limits are technological rather than fundamental. Therefore, major progress is still possible with methodological improvements.

In the past twenty years, research has exploded into the use of coherent x-ray light to improve the quality and resolution of x-ray microscopes. In many cases, using coherent light makes it possible to remove the objective lens in a microscope, replacing it with an algorithmic analysis of the direct scattering data. This can increase the quality and resolution of the resulting quantitative images.

In this thesis, I present the results from a collection of projects aimed at using coherent imaging methods to study the real-space texture of electronic phases of matter with soft x-ray light. I first discuss the methods we developed and implemented to counteract the experimental errors that we found to be ubiquitous in our data, focusing on ptychography, the most commonly used lensless imaging method. Then, I turn to the development of an entirely new single-shot lensless imaging method, randomized probe imaging (RPI).

RPI has proven to be reliable and robust across a broad range of scenarios. The remainder of the thesis is devoted to applications of RPI at a free electron laser and a synchrotron. Also reported are further projects designed to improve the method, as well as attempts to expand our understanding of the mechanisms behind it and its limitations. I sincerely hope that the availability of RPI will help bring x-ray imaging to a broader group of scientists and lead to a better understanding of the nanoscale details of electronic texture.

Thesis Supervisor: Riccardo Comin
Title: Class of 1947 Career Development Associate Professor of Physics

Acknowledgments

I owe a huge debt of gratitude to my advisor, Riccardo Comin, who gave me leeway to explore my wilder ideas, and supported them with his expertise, time, and funding. He also brought me back to earth and focused my attention when needed, without which I may never have brought any of my projects to a conclusion!

Many of the ideas in this thesis also rest on everything I learned even before beginning my Ph.D. from the scientists I worked with in Siegfried Glenzer's group at SLAC: my advisor Luke Fletcher, who was a great source of professional and technical advice; Michael Macdonald, who started developing the project I worked on while there; Jerry Hastings, who was always happy to point out that what I was thinking of had already been done in the the 1970s; and my officemate Peihao Sun, the source of dozens of interesting discussions.

Working closely with David Shapiro at his beamline of the Advanced Light Source this past year, I learned more about the operation of x-ray microscopes than I had in the previous five years, and learned a tremendous amount about coherent imaging and ptychography. Nearly all of the work described in the chapter on further development of RPI would not have existed without David. My additional collaboration with Sophie Morley and Sujoy Roy also made this a year of broad-based learning and growth.

George Barbastathis and Zhen Guo, collaborators on a project to bring machine learning methods into the analysis of RPI data, also brought a whole new realm of science into my work, and for that I am grateful. I thank George especially for his scientific rigor and willingness to think deeply and have serious discussions about basic principles.

Nuh Gedik, my academic advisor and second committee member, gave me strong advice on how to navigate MIT and, eventually, my career post-MIT. Many thanks to Kiyoshi Masui for stepping in to serve on the committee with such late notice, which was even more appreciated given the surprising connections between his work and mine, and his thoughtful comments on this project.

I also thank everyone who gave me comments on early versions of my paper on randomized probe imaging, in particular Stephan Hruszkewycz and Manuel Guizar Sicairos.

I am thankful to have worked with everyone in the photon scattering group at MIT, who made that work such an enriching experience. Mingu and Jiarui, who were the only grad students in the lab when I joined, already made the lab feel like a vibrant place. Johnny and Zhihai, and then Connor, Qian, and Luiz, and Jiaruo, Shua, and Yi who came after, I enjoyed so much getting to know and learn from all of you. Most importantly, without Kahraman Keskinbora, who contributed more than anyone else to the work on RPI, the paper would have taken twice as long to create and been half as good.

I especially thank Anastasia Kutakh for her work on a project involving polarization-sensitive ptychography; it was a joy to work with you and see your progress over the timescale of the project.

Of course, I have to thank Monica Wolf and Gerry Miller for keeping the CMX floor of Building 13 running as smoothly as possible.

To my friends in MIT physics, Bryan Fichera and Jeff Krupa especially, for making my time at MIT so enjoyable, I can't thank you enough. And to the rest of my Cambridge friends, Connor and Paul, Simon, and Brenna, and especially Zach, thanks for living out the pandemic with me and keeping me sane through five years of grad school!

Thanks to my cubicle-mates from the Advanced Light Source, Michelle and Cissy, for doing the same for my sixth year spent at the ALS, and for making building 80 the most fun place to work at LBL.

To Mom, Dad, and Rebecca, thanks for instilling in me the love of discovery and being unconditionally supportive of my choices and work.

And finally, to Wanyi, for supporting me, helping me with presentations across all six years, being in Cambridge and then in the San Francisco area with me, and always supporting my work and believing in me. None of these projects would have been completed without the strength I drew from that.

Contents

List of Figures

List of Algorithms

Chapter 1

Introduction

From Galileo's discovery of the moons of Jupiter with the telescope [1], to Van Leeuwenhoek's discovery of the cellular nature of life with the light microscope [2], to Rosalind Franklin's first observation of the x-ray diffraction pattern of DNA [3], the key to the discovery of many facets of nature has simply been the ability to see them. Improving our ability to visualize the world around us has driven scientific progress for as long as the concept of a scientific method has existed.

Advances in light microscopy made diffraction-limited spatial resolutions in the hundreds of nanometers easily achievable, and the advent of the electron microscope and scanning tunneling microscope allowed scientists to see down to the single-atom level. Simultaneously, the growing importance of nanoscale phenomena to science and engineering have made access to nanometer-level resolution increasingly important.

This progress in nanoscale engineering is now driving a massive transition in every field of science, as the cost of computation falls and the computational power available to the average scientist grows exponentially. Access to computation has supported multiple recent revolutions in microscopy, from the advent of cryo-electron microscopy [4] to the development of computed tomography [5].

However, yet another revolution in computational imaging is on the cusp of occurring, after nearly 50 years of steady progress. This is the wholesale replacement of lenses with a combination of coherent light and algorithmic analysis, under the name of coherent lensless imaging. By 2010, the highest resolution lensless x-ray images were already well ahead of lens-based x-ray imaging methods [6]. In 2018, the same thing happened for electron microscopy [7].

Coherent lensless imaging methods have become ubiquitous in imaging research, but have not yet come into widespread use due the computational demands of the phase retrieval algorithms needed to analyze the data. It is still cheaper and easier to buy a good microscope than to pay for a cheap microscope, an expensive computer, and the imaging expert needed to operate them both. However, this will change as the price of computation drops further and data analysis programs mature. Furthermore, lensless imaging methods measure both the absorption and phase delay imparted on light as it traverses a sample, information which is not routinely accessed with standard microscopy. Because of the quality and (eventual) cost advantages, it seems likely that in ten or twenty years, every lab's bench will have at least one microscope driven by lensless imaging, enabling all kinds of new experiments.

1.1 X-ray Microscopy

As mentioned above, much of the original research into coherent lensless imaging focused on applications in x-ray microscopy. This is because for x-ray microscopy the gap between what is possible in principle and what is possible in practice has always been especially large.

Regarding what is possible in principle: an efficient, diffraction-limited x-ray microscope with high numerical aperture could take images with sub-nanometer resolution, with sensitivity to chemical and electronic information through interaction with atomic resonances. However, in practice the optics needed to build such a microscope don't exist. State-of-the-art optics cannot reach diffraction-limited performance for resolutions beyond tens of nanometers, and rarely have a photon throughput above 10%.

Removing lenses from the equation gets rid of that limitation, and leaves only the fundamental balance of imaging contrast with sample damage. As we will discuss next, many fields of science can benefit from x-ray imaging, and the advances in resolution and efficiency that lensless imaging will make possible can have an impact across all these fields.

The simplest and possibly most compelling use of an x-ray microscope is as a miniature version of the x-ray projection cameras that are commonplace in medicine. We get medical x-ray images taken when we want to see inside our bodies, because the x-rays penetrate through us and reveal their inner structure. Similarly, we often want to image, for example, buried structures in integrated circuits [8]. Visible light will reflect off the top layers, and electrons are strongly absorbed by all materials, but x-rays sail right through.

Furthermore, individual x-ray photons carry enough energy to ionize the core shell electrons of atoms. Every atom has a related set of ionization energies that define the threshold energies for this ionization process to occur. When the energy of an x-ray matches one of these threshold energies, it interacts resonantly with the atom and is scattered much more efficiently. Consequently, carefully selecting the wavelength of the light within an x-ray microscope can highlight the locations of specific atomic species.

Moreover, the detailed dependence of this interaction on the wavelength of the incoming x-rays contains information about the electronic configuration of the atom and its local environment. This can reveal chemical information such as the oxidation state [9] and the spin polarization of the valence electrons [10]. Combined with imaging, this makes it possible to study the spatial dependence of oxidation states within the nanoparticles of a battery material[9], or take photographs of magnetization domains in thin films [11]

Finally, the wavelength of x-ray light is short enough to diffract from the atomic lattice in many materials. At atomic resonances, it is also possible to diffract off of longer-scale perturbations of the electronic structure such as charge density waves, spin waves, and antiferromagnetic order [10]. These subtle, spatially-modulated electronic orders are often extremely disordered at intermediate length scales down to the nanoscale, full of strained regions and defects. High resolution imaging of the diffraction can reveal the details of these nanoscale textures [12], providing one of the most powerful tools available for understanding their relationships to one another and other causes of nanoscale heterogeneity.

Each of these areas benefits from the resolution and efficiency advantages afforded by coherent lensless imaging, and depending on the specific application is either now routinely accomplished with coherent methods or beginning to take advantage of their benefits. The following chapter will discuss the obvious next question: What are these methods, and how do they work?

Chapter 2

Coherent Lensless Imaging

Coherent lensless imaging is a broad field encompassing many experimental methods, so we begin by discussing the commonalities that bind the field. What ties all these methods together is a reliance on coherent light. Coherent light, in the context of this research field, is monochromatic light that can be fully described by a single, time-independent, complex-valued vector electric field:

$$\vec{E}(\vec{r}, t) = \mathrm{Re}\left[\vec{E}_0(\vec{r})e^{-i\omega t}\right]. \tag{2.1}$$

Typically, a further approximation is introduced, that the light is uniformly polarized. In this case, only one polarization component $\vec{\epsilon}$ needs to be considered, and we can write the electric field as

$$\vec{E}(\vec{r}, t) = \vec{\epsilon}\mathrm{Re}\left[\psi(\vec{r})e^{-i\omega t}\right]. \tag{2.2}$$

In the remainder of this work, we will proceed from the above complex scalar approximation for coherent light.

It might seem as though all monochromatic light should be coherent, and formally that is true. However, the period of optical light is roughly a femtosecond, and measurements can easily last for seconds. Consequently, optical light whose bandwidth $\frac{\Delta E}{E}$ is less than 10^{-15} will necessarily be stable over a period of seconds, but this is far, far, far more monochromatic than we typically mean when we say "monochromatic". For most coherent lensless imaging experiments, $\frac{\Delta E}{E} < 10^{-4}$ is sufficient to be practically indistinguishable from fully monochromatic light.

Most light sources are simply not stable on length scales of seconds. This is the basic reason that most light, even monochromatic light, is incoherent. For example, the filament in an incandescent light bulb is emitting light essentially randomly from various locations along the filament, each location producing a different propagating electromagnetic field. Even after filtering the light by wavelength to produce a "monochromatic" source, the resulting electric field will still vary dramatically over most experimentally relevant timescales. Therefore we cannot simply assume it is well described by one stationary state of the electric field, and have to resort to a statistical description. Coherent light is light that is stable enough to not require a statistical description.

The most recognizable property of coherent light is interference. This is typified by the classic double-slit experiment, wherein a wavefield of coherent light passes through two nearby slits of light. This light is then projected on a camera sensor. The camera sensor measures the average power contained within the wavefield as a function of the pixel location $\vec{r}$,

$$\mathbf{I}(\vec{r}) = \frac{1}{2}|\psi(\vec{r})|^2. \tag{2.3}$$

For the remainder of this thesis, we will omit the factor of $\frac{1}{2}$, as neither measured intensity values nor scalar wavefields are tied to fundamental units in any of the experiments we discuss.

The important thing to recognize is that the camera records the magnitude of the wavefield, but not its phase. Still, the relative phase differences between the light passing through the two slits generates a pattern of light and dark fringes in this intensity image, as seen in Figure 2-1. In this way, interference between the light emerging from the two slits turns phase information into intensity information.

The orientation and spacing of the fringes in this diffraction pattern is uniquely determined by the spatial structure of the slits. In other words, even though we

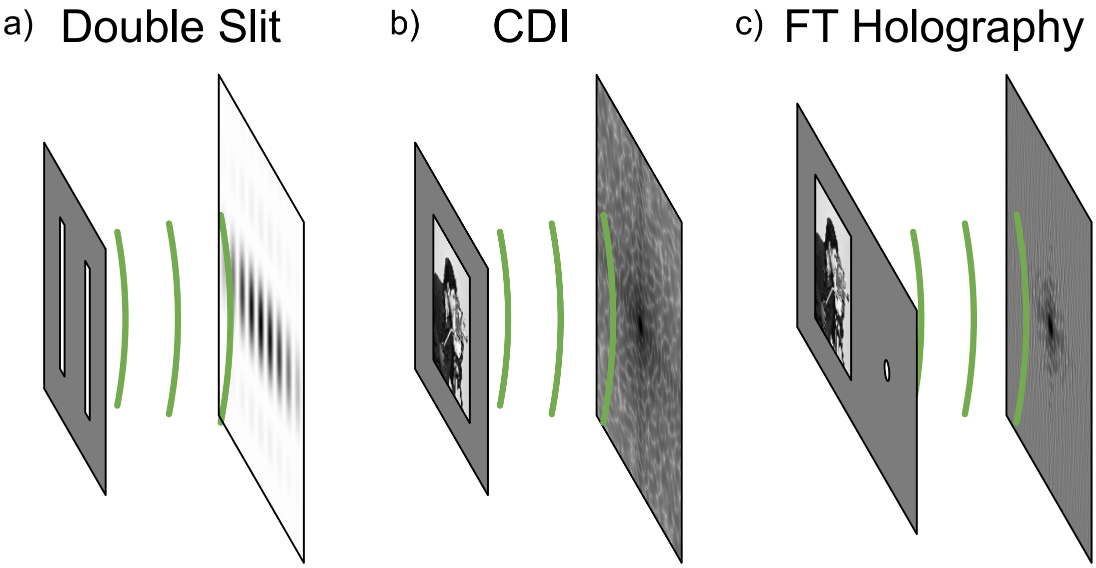

Figure 2-1: **Classic Experiments with Coherence**. (a) A visualization of the double-slit experiment, where two neighboring slits produce a periodic pattern of fringes. (b) A CDI experiment, where the pair of slits is replaced with an arbitrary image. The diffraction pattern becomes much more complex, but in principle the details of the image are still recoverable from the diffraction pattern alone. (c) A visualization of a Fourier transform holography (FTH) experiment. The additional reference hole introduces more fringes to the measured diffraction pattern, enabling a one step reconstruction.

aren't sensitive to the phase of the light impinging on the camera, we can still learn almost everything there is to know about the structure that diffracted the light. This hints at the essential question: how complicated can we make the pattern of slits, dots, or whatever, before we can no longer figure out what the arrangement was from the diffraction pattern alone? The answer turns out to be: with a few caveats, as complicated as you want!

The idea is shown in Figure 2-1: one measures the diffracted intensity from an object illuminated with coherent light, and then uses algorithms to figure out what that object had to be. This idea is known as coherent diffraction imaging (CDI) and is conceptually the simplest method in the world of coherent lensless imaging. We will therefore turn to it first.

2.1 Coherent Diffractive Imaging

If the screen is far enough away from the slit, the pattern of light projected on it has an interesting property, within the paraxial approximation [13]. The magnitude of the wavefield at the plane of the screen will be related to the magnitude of the Fourier transform of the wavefield that emerged from the object.

$$\begin{aligned}
|\psi(x, y; z)| &\propto \left| \int dx' dy' e^{i(k_x x' + k_y y')} \psi(x', y', 0) \right| \\
k_x &= \frac{2\pi}{\lambda z} x \\
k_y &= \frac{2\pi}{\lambda z} y
\end{aligned} \tag{2.4}$$

This reformulates our problem as the recovery of an arbitrary complex-valued image (the *object*) from the magnitude of its Fourier transform. On the face of it, this sounds impossible. Surprisingly, this *inverse problem* nearly always has a unique answer, up to a few trivial degeneracies such as a global phase factor $e^{i\phi}$. The essential condition is that the wavefield emerging from the object drops to 0 beyond some well-defined finite region of space. This condition is known as the *finite support constraint.*

That this problem can be solved, and why, appears to first have been noted by David Sayre in 1952 [14], a few years after Shannon's work on the Nyquist-Shannon theorem was published [15]. The critical detail is that a finitely supported function has a Fourier transform that can be fully parameterized by discrete samples taken at the Nyquist rate. Consequently, there is redundant information within the continuous electric field distribution at the detector plane, because the full continuous distribution can be recovered from a sparsely spaced set of samples.

This excess information is not so redundant when we remove the phases and begin to study the intensity distribution. Due to the convolution theorem, the intensity pattern is the Fourier transform of the object function's autocorrelation. This autocorrelation is supported on a region with roughly twice the linear dimension of the original object function's support.

This means that, to reconstruct the full distribution of diffracted intensity, we need to sample it at most at twice the Nyquist rate of the complex-valued wavefield. The information from these additional samples was redundant when we knew the phases, but it is no longer redundant after we've lost them. This information hiding in the additional sampling points is what makes phase retrieval possible, and also why we need to ensure that the object function has finite support.

Of course, this is a plausibility argument, and we have not yet proven anything. However, it is understood that in essentially every realistic case the phase retrieval problem of CDI does indeed have a unique solution in the absence of noise [16] in 2 or higher dimensions. Furthermore, there are strong indications that the solution remains reasonably well conditioned in the presence of noise [17].

The first practical algorithms for solving the phase retrieval problem of CDI were developed by Fienup in the 1980s [18, 19], and CDI has been used for x-ray based imaging starting in 1999 [20]. It continues to be used today, typically in situations where the samples are naturally isolated, such as viruses [21] or soot particles [22].

Unfortunately, the stability and accuracy of the retrieval step has continued to pose problems. It is rarely possible to obtain reliable reconstructions with a field of view more than a few tens of pixels across, and the difficulty of finding a unique solution increases as the images get larger. It is not entirely clear whether the issues are related to the underlying structure of the problem or to the inadequacies of current reconstruction algorithms. Nonetheless, these issues have led many researchers to focus on alternate methods with more stability and reliability.

2.2 Holography

If CDI data were easy to analyze, then it would likely still be the go-to method for x-ray based lensless experiments everywhere. However, as we have discussed, it is often anything but easy to analyze. It is for this reason that Fourier transform holography (FTH) [23] has become a popular alternative. The concept of FTH is simple: drill a hole next to your sample, and measure the diffraction pattern directly, just like you would with CDI. The simple act of drilling the hole enables a simple, direct reconstruction method to be used.

As discussed before, the intensity pattern is the Fourier transform of the spatial autocorrelation of the object function with itself. If the reference hole drilled next to the sample window is sufficiently far away, then this autocorrelation includes a term for the cross-correlation between the object and reference, i.e. a slightly blurred image of the sample. Simply taking the inverse Fourier transform of the diffraction pattern is enough to recover this image [23].

This method, however, is limited in two significant ways. First, just like CDI, it only works on small, isolated samples, and large samples need to be physically masked off before they can be studied with FTH. Second, the resolution of the final reconstruction is limited by the diameter of the reference hole, so it can't be used for the absolute highest resolution work.

These limitations have not stopped FTH from becoming a dominant method for dynamic imaging of thin films [24, 25] and other structures that can easily be masked. However, the frustration of dealing with limited sample sizes has led to the exploding popularity of the method I will discuss next, ptychography.

2.3 Ptychography

While CDI and FTH are both single-frame methods that reconstruct an image from a single exposure's worth of data, ptychography fundamentally relies on collecting multiple frames of data. This is a gift and a curse. A gift, because the multiple frames of data expand the field of view and make it more reliable. A curse because it can't capture fast events with single flashes of light, and because precise mechanical motion is needed to make it work.

The basic setup of a ptychography experiment is shown in Figure 2-2. Instead of using a small, isolated or masked sample, a large extended sample interacts with a (usually) focused beam of coherent light. The diffraction pattern formed by this process would be difficult or impossible to analyze by CDI on its own, likely producing unstable results as discussed above.

Instead, a collection of typically hundreds or even thousands of diffraction patterns is taken as the beam is scanned across the sample. Because the position of the illuminated region in each diffraction pattern is different, each diffraction pattern essentially contains information about the mutual interference within a different collection of sample regions. Scanning with a small step, to ensure that each part of the sample is included in a few different diffraction patterns (the *overlap* constraint), creates highly redundant data which is amenable to a stable reconstruction.

The concept of ptychography appears to have first been discussed by Walter Hoppe in a series of three 1969 papers [26–28]. However, because these papers have never been translated to English, I am not personally familiar with the details of his proposal. The name itself was coined by Hergel and Hoppe one year later [29]. By 1982, it had been recognized that scanning transmission electron microscopes were very close to an ideal setup for taking ptychography data [30, 31]. This led to the earliest experimental electron ptychography, led by John Rodenburg throughout the 1990s. This started with the proposal for Wigner distribution deconvolution [32], the first ptychography algorithm still recognizable to modern researchers, and subsequent demonstrations of the method with optical light [33] and electrons [34].

The first hints of modern ptychography appeared in 2004 when Faulkner and Rodenburg began to build on CDI algorithms by designing and simulating the performance of an iterative algorithm (rather than the earlier direct algorithm) for ptychography reconstructions [35]. The key advance here was the realization that the diffraction patterns need not be sampled at the resolution of the final image. Previous algorithms had relied on this fine sampling to work, and removing that constraint made ptychography much more practical. This was followed in 2008 by Manuel Guizar-Sicairos and James Fienup's work showing how the illumination and object could be jointly retrieved [36], and the nearly simultaneous publication of an experimental demonstration of the same idea with x-rays by a different group [37].

From there, the field burgeoned, and by the present day over a thousand studies have been published, extending upon and using the principle of ptychography. A diverse collection of iterative algorithms for reconstructing ptychography data have been proposed since 2008, including the extended Ptychographical Iterative Engine (ePIE) [38], difference map [39], maximum-likelihood refinement [40], as well as many variations that correct for sources of experimental error, which will be discussed more in Chapter 3.

Although the details vary, nearly all the algorithms make the same set of basic assumptions. The physical experiment is modeled as the interaction of a probe wavefield $\mathbf{P}(x, y)$ with a complex-valued object function $\mathbf{O}(x, y)$, producing an exit wave $\mathbf{E}(x, y)$:

$$\mathbf{E}(x, y) = \mathbf{P}(x, y)\mathbf{O}(x, y). \tag{2.5}$$

This model is derived from the linearity of light propagation and the assumption of *locality*. Locality here means that the properties of the object at a point x, y only affect the wavefield at that point. This is a valid assumption when the object is thin, and there is no opportunity for light at one point on the object to be scattered to a different region of the object. The linearity of light propagation then constrains the interaction to take on this form.

In less abstract terms, the object function is related to an integral through the index of refraction of the sample:

$$\mathbf{O}(x, y) = \exp\left(i \int_{z=0}^{t} dz\, n(x, y, z)\, k \right). \tag{2.6}$$

The exit wave propagates to a detector, which is typically placed in the far field, satisfying the Fraunhofer diffraction condition as encoded by Equation (2.4). This diffraction is measured on an array of pixels, with some spacing between neighboring pixels p and some overall number of pixel $M \times N$.

In physical reality, the probe, object, and exit wave are continuous functions of x and y. However, to be represented computationally, they must be stored as 2D images, sampled at a finite spacing and over a finite size. Usually, that spacing and size is chosen so that the discrete Fourier transform of the 2D images maps, pixel-for-pixel, onto the detector image. This spacing is $\frac{\lambda d}{Mp}$.

The reconstruction algorithms operate on these discrete approximations. Generally, the algorithms begin with a sensible guess of the object and probe, say a uniform object $\mathbf{O}(x, y) = 1$ and a probe that roughly matches the expected probe, such as the focus of a known zone plate lens or an aperture with a particular shape and size.

From here, the diffracted intensity patterns at one or more probe positions are simulated and compared to the measured data. Information from the measured data is included, for example by replacing the simulated Fourier components of the exit wave with the measured ones, or by calculating a gradient of a loss function. Then, the exit wave is returned to real space, and another update is applied to bring the guess of the object and probe more in line with this new information. This process is continued until both the object and probe converge to stable values, producing both a measurement of the sample's properties and a measurement of the illumination's structure.

Because of its reliability and scalability, ptychography has become a workhorse imaging method relied upon even by many non-experts for scientific investigations, and is likely to become the standard method for x-ray microscopy in the near future. However, the simplified discussion above misses a big part of the challenge of ptychography—overcoming the gap between the simplified models and the real world. In the next chapter, we will discuss in detail modern perspective on how to perform reconstructions on ptychography data that take these issues into account explicitly.

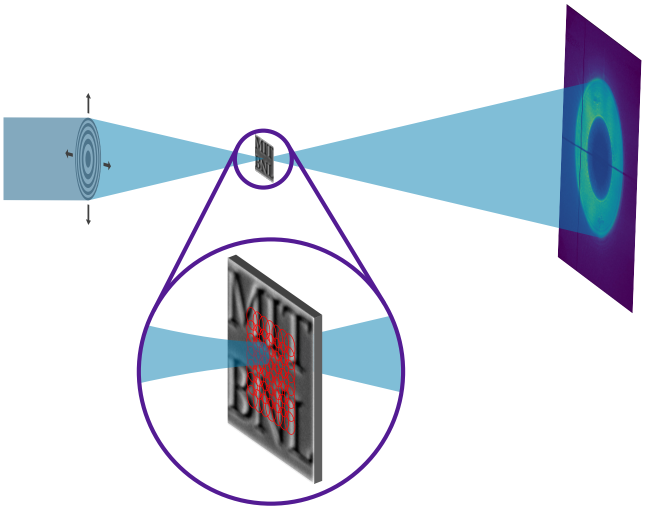

Figure 2-2: **Ptychography**. The basic concept behind ptychography. A small beam of light, or *probe*, is scanned across a sample, or *object*. The diffraction patterns that form on the detector are saved for each position, and the overlap between the neighboring illuminated regions is used as an extra constraint to lead to more reliable convergence.

Chapter 3

Models for Phase Retrieval with Automatic Differentiation

3.1 Introduction

One of the major problems facing coherent imaging methods is the difference between the idealized settings that canonical algorithms rely on and the reality of a physical experiment, with all its imperfections. A sizable literature has consequently sprung up to deal with the various types of errors and imperfections that can exist in ptychography data. These corrections typically work by modeling the cause of the error and recovering a set of parameters which characterize it in addition to the parameters of interest.

Methods have been developed that extend most of the standard algorithms to correct for everything from inaccurate position information [41, 42] to partial coherence [43]. However, combining these corrections is cumbersome, because it is not always obvious how two different adjustments to the same base algorithm should be blended together.

A recently popularized solution to this problem is automatic differentiation. As we discuss below, using automatic differentiation-based algorithms makes it simple to add many standard corrections to a ptychography algorithm. More importantly, blending multiple corrections is much more straightforward within this framework,

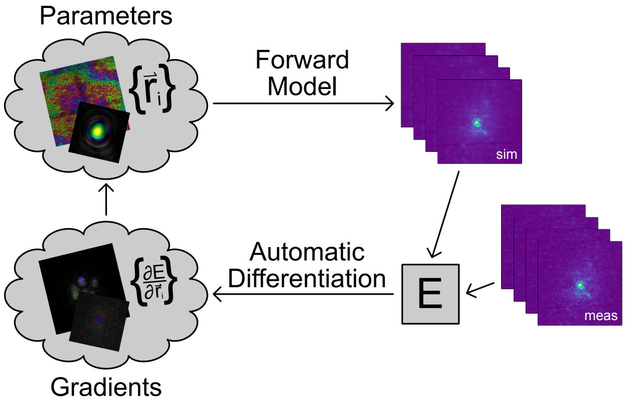

Figure 3-1: The flow of data in a generic automatic differentiation based algorithm for ptychography

because one only needs to consider how to simulate the source of error—not how to define an algorithm to recover it. As part of my thesis research, I developed a library of correction models, some based on existing work and some original, to enable ptychography reconstructions within the group.

3.2 Automatic Differentiation

Automatic differentiation is a computational method that traces the chain of mathematical operations occurring in a script or function. As it does this, it builds up a secondary function capable of differentiating the output of the computation with respect to the inputs [44]. This secondary function is defined using the same types of primitives as the original function, and in simple cases can execute just as quickly. Specifically: automatic differentiation is not a finite difference method. With automatic differentiation, derivatives are computed exactly and analytically.

The fundamental idea essentially boils down to cleverly identifying a synergy between the chain rule for differentiation and the structure of computer programs. The method has been around in various forms since at least the 1960s [45]. Its popularity has exploded in recent years because of its application to training deep neural networks. As a result, many libraries for high-performance computing now have comprehensive support for it.

Automatic differentiation was first applied to ptychography in two studies from 2018 [46] and 2019 [47]. The basic concept is illustrated in Figure 3-1. First, a set of parameters to be recovered from the data is defined. Then, a *forward model* is written, in the form of a function that maps those parameters onto a simulated set of diffraction patterns $\{\mathbf{I}_j^{\text{sim}}\}$. The prototypical forward model will be given in Equations (3.1) - (3.3). Importantly, this function must be *differentiable* in its inputs with respect to its outputs.

Next, an error metric or *loss function* $\mathcal{L}$ is defined, which quantifies the difference between the simulated and measured diffraction patterns. Many error metrics can be used, but the most common are the negative log-likelihood function associated with the dominant error model on the detector [40] or the euclidean distance between the square-root intensities, i.e. Equation (3.4). This latter model is less mathematically elegant, but empirically is robust to most detector noise models [40].

Finally, a gradient descent algorithm is used to find the minimum of this loss function. The gradients themselves are calculated using automatic differentiation, and many standard algorithms exist for converting these gradients into steps. With a well-tuned step-size, even simple stochastic gradient descent [47] is sufficient. However, tuning step sizes can quickly become difficult when error models are introduced. This is because the new parameters, e.g., the detector background, can have gradients that are orders of magnitude larger or smaller than the gradients of the object pixels.

In practice, the Adam algorithm [48] is used most frequently, and it is robust across a wide range of settings, so long as care is taken to ensure that the parameters being optimized over are of order unity. In many settings, second order quasi-Newton algorithms like L-BFGS [49] also have good performance, and there has been recent work exploring algorithms specifically designed for this setting [50].

The result of this cycle of steps is an easy to implement and flexible algorithm because, in most cases, making an adjustment is as simple as updating the forward model. In the following sections, we discuss a variety of forward models which account for commonly encountered sources of error.

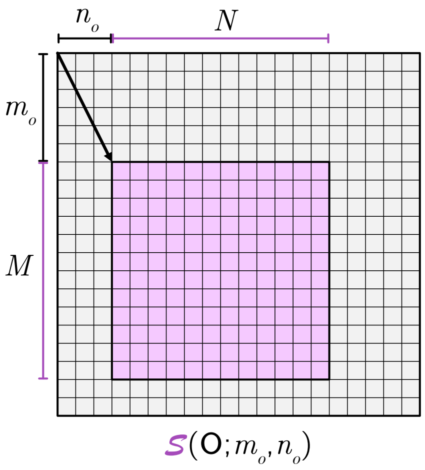

Figure 3-2: A diagram showing the operation of the shift and clip function $\mathcal{S}$. The $M \times N$ region in lavender is extracted from the larger object array $\mathbf{O}$, at the shift m_0, n_0.

3.3 The basic forward model

The simplest forward model upon which all others are built simulates the diffracted intensity pattern $\mathbf{I}$ sampled on a $M \times N$ pixel detector with pixel pitch p at a distance d from the sample. The simulation uses two complex-valued images, the object array $\mathbf{O}$ and probe array $\mathbf{P}$. The object and probe arrays are defined as sampling the real-space structure of the sample and illumination, respectively, with a pixel pitch of $\frac{\lambda d}{Mp} \times \frac{\lambda d}{Np}$, as described in Section 2.3. From these components, the diffraction from each frame with index j is simulated as follows:

$$\mathbf{E}_j = \mathbf{P}\,\mathcal{S}(\mathbf{O}; \lfloor x_j \rceil, \lfloor y_j \rceil) \tag{3.1}$$

$$\tilde{\mathbf{E}}_j = \mathcal{F}\{\mathbf{E}_j\} \tag{3.2}$$

$$\mathbf{I}^{\text{sim}}_{j,mn} = |\tilde{\mathbf{E}}_{j,mn}|^2 \tag{3.3}$$

$$\mathcal{L} = \sum_j \sum_{mn} \left[\sqrt{\mathbf{I}^{\text{sim}}_{j,mn}} - \sqrt{\mathbf{I}^{\text{meas}}_{j,mn}} \right]^2, \tag{3.4}$$

where $\mathbf{P}$ is the probe array, whose dimensions $M \times N$ match the detector's, and $\mathbf{O}$ is the larger object array. $\lfloor x \rceil$ refers to rounding x to the nearest integer, and (x_j, y_j)

is specified in units of pixels. $\mathcal{S}(\boldsymbol{O}; \lfloor x_j \rceil, \lfloor y_j \rceil)$, the shift and clip function, extracts an $M \times N$ region of the object at the specified location:

$$\mathcal{S}(\mathbf{O}; m_0, n_0)_{mn} = \mathbf{O}_{m_0+m, n_0+n}, \quad m \in \mathbb{Z}_M, n \in \mathbb{Z}_N. \tag{3.5}$$

This operation is shown graphically in Figure 3-2. The 2D discrete Fourier transform $\mathcal{F}$ and inverse 2D discrete Fourier transform $\mathcal{F}^{-1}$ are defined here as:

$$\begin{aligned}
\mathcal{F}\{\mathbf{A}\}_{mn} &= \frac{1}{\sqrt{MN}} \sum_{\mu\nu} \mathbf{A}_{\mu\nu} e^{-2\pi i \left(\frac{(m - \lfloor \frac{M}{2} \rfloor)(\mu - \lfloor \frac{M}{2} \rfloor)}{M} + \frac{(n - \lfloor \frac{N}{2} \rfloor)(\nu - \lfloor \frac{N}{2} \rfloor)}{N} \right)} \\
\mathcal{F}^{-1}\{\mathbf{A}\}_{mn} &= \frac{1}{\sqrt{MN}} \sum_{\mu\nu} \mathbf{A}_{\mu\nu} e^{2\pi i \left(\frac{(m - \lfloor \frac{M}{2} \rfloor)(\mu - \lfloor \frac{M}{2} \rfloor)}{M} + \frac{(n - \lfloor \frac{N}{2} \rfloor)(\nu - \lfloor \frac{N}{2} \rfloor)}{N} \right)}.
\end{aligned}$$

This definition defines the origin of coordinates at the center of the image, rather than at the corner of the image as is standard. This setting maps naturally onto the physical model of light propagation, where the low-frequency data winds up on the center of a physical detector.

Equations (3.1) - (3.4) define the basic mapping from the input parameters to the single objective (loss) function we will attempt to minimize. Here I have broken it out into four conceptual steps: the interaction model (3.1); propagation model (3.2); measurement model (3.3); and loss model (3.4).

Separating out these stages of the model is worthwhile because many corrections we will introduce only affect one or a few of the stages. For example, as we will see in Section 3.6, introducing position correction only requires an update to the interaction model. Similarly, introducing background estimation (Section 3.7) only causes a change to the measurement model.

The following sections explore various corrections that account for ways in which reality deviates from this simplified model, but it is important to note that this simple model is already in many cases sufficient to obtain state-of-the-art reconstructions.

As has been noted elsewhere, the gradients $\frac{\partial \mathcal{L}}{\partial \mathbf{O}_{mn}}$ and $\frac{\partial \mathcal{L}}{\partial \mathbf{P}_{mn}}$ are closely related to the ptychographic iterative engine (PIE) [51] and extended PIE (ePIE) [38] update steps [47], which are among the oldest and simplest algorithms for ptychography. The ePIE algorithm has been used to produce some of the most impressive results in the field, most recently the first successful modern-style electron ptychography [7]. Thus, before we dive into a detailed discussion, it is important to remember the following rule of thumb:

Adage 3.1

If your ptychography data doesn't converge at all with ePIE, it probably won't converge with a fancy algorithm either.

With that said, introducing error corrections into the reconstruction algorithm can improve the quality of the reconstruction significantly, and it is often worthwhile even if the errors in question are small. However, these models, however comprehensive, cannot turn bad data into good data. The existence of the error correction models is not a license to take bad ptychography data.

3.4 Correction for incoherent illumination

Traditionally, ptychography and coherent diffraction imaging (CDI) assume that the illumination on the object is fully coherent and monochromatic. This is a good approximation for many sources of light, but is never truly realized in an actual experiment. Common causes of a lack of perfect coherence include:

- Partially coherent light sources
- Vibrating optics
- Fly scans (collecting data with a continuously-moving probe)
- A lack of sufficiently high monochromaticity

Luckily, all these sources of incoherence (with the exception of some of the effects of non-monochromaticity [52]) can be treated using the same framework, which is nowadays ubiquitous. This is the multi-mode model, which was first explored by Pierre Thibault and Manuel Guizar-Sicairos [43].

$$\mathbf{E}_j^k = \mathbf{P}^k \cdot \mathcal{S}\left(\boldsymbol{O}; \lfloor x_j \rceil, \lfloor y_j \rceil\right) \tag{3.6}$$

$$\tilde{\mathbf{E}}_j^k = \mathcal{F}\{\mathbf{E}_j^k\} \tag{3.7}$$

$$\mathbf{I}_{j,nm}^{\text{sim}} = \sum_{k=1}^{K} |\tilde{\mathbf{E}}_{j,nm}^k|^2 \tag{3.8}$$

In this model, the illumination is defined by a set of K wavefunctions $\boldsymbol{P}^k$, known as coherent modes. Each mode interacts with the sample separately and propagates to the detector separately, where their intensities add incoherently.

Conceptually, this refers to a time-varying illumination wandering through the defined probe modes. From a fundamental optics perspective, this list of modes is a truncated eigendecomposition of the mutual coherence function of the illuminated light, as explored more in Appendix B.

Most ptychography algorithms have been extended to the multiple mode case, including ePIE [43], difference map [43], maximum likelihood [43, 53], and automatic differentiation [46].

In most implementations, mutual orthogonality of the probe modes is explicitly enforced. However, in automatic differentiation, we typically allow the modes to be non-orthogonal during reconstruction [46]. This doesn't affect the validity of the results, as any set of three non-orthogonal modes can always be converted to an equivalent set of orthogonal modes.

The number of modes K is related to the overall degree of coherence of the incoming light, and in fact fully incoherent light can be thought of as a probe with many modes contributing meaningfully. A typical reconstruction from good quality data and a coherent light source will have one or several dominant modes, followed by a rapidly decaying tail of weaker modes. One rarely sees reconstructions performed with more than 10 modes, because if that many modes are really required, the light is likely to be insufficiently coherent to produce strong speckle contrast.

3.5 Correction for unstable illumination

While the model above corrects for fluctuations in the illumination occurring faster than the exposure timescale, it is also common to find that the illumination function changes subtly between exposures. This can be due to slow vibration, drifts, the use of an intrinsically unstable source such as a high harmonic generation (HHG) laser [54] or free electron laser (FEL) [55], or the physical motion of optics.

In this case, we can relax the constraint that the probe at each exposure is precisely the same as the probe at every other exposure. We can accomplish this elegantly by recognizing that the multi-mode model described above is successful because, in many cases, the illumination only "wanders" around a low-dimensional subspace of dimension K. This suggests that we can allow the probe at each exposure to look like any collection of K modes, each mode consisting of a linear combination of modes from a (potentially larger) subspace with dimension L using the following model:

$$\mathbf{P}_j^k = \sum_l \alpha_j^{kl} \mathbf{\Psi}^l. \tag{3.9}$$

This model allows for each exposure to contain K probe modes, and for the probe to explore L modes throughout the full data collection. It is closely related to orthogonal probe relaxation ptychography (OPRP) [54], which defines an explicit reconstruction algorithm for the case where $K = 1$ but $L \neq 1$. It is also related to the model explored in [55], where the authors force the matrix α_j^{kl} to be diagonal in k, l for each frame j. Finally, if α_j^{kl} is set to the identity matrix for each j, the standard multi-mode model for partially coherent light is recovered.

3.6 Correction for position inaccuracy

With the possible exception of partial coherence, the most discussed source of error in ptychography is inaccurate information about the position of the probe. This kind of error is ubiquitous in real experiments and correcting for it pays off handsomely because positioning errors do not degrade the quality of the diffraction patterns themselves. Simply put, if one only knew the actual positions, the dataset wouldn't be degraded at all. In fact, due to the raster grid pathology that affects data collected from scan points on a Bravais lattice [56], sometimes a dataset with imprecise positions is better than it would have been if the probe had been placed accurately.

There are a variety of methods in common use for position correction, falling broadly into two categories. The first category is exploration-based approaches, which use an explicit method to generate suggested position updates between iterations. The canonical example of a method in this realm is position annealing [41], which generates a set of random trial positions and accepts the one that most improves the loss metric. Another method uses the small offset present in the per-pattern probe update to propose an update step for the position [57]

The methods in the second category use the diffraction patterns themselves to determine the position update at each iteration, most popularly by calculating the gradient of a loss function with respect to the probe positions, usually following a calculation of the gradient done by hand [42]. However, in order for this gradient to be well-defined, the interaction model must be differentiable in $\vec{r}$, which Equation (3.1) is manifestly not.

To rectify this, we need to introduce a method for shifting either the probe or the object to subpixel accuracy. The natural choice is to apply a shift to the probe using a linear phase ramp in the Fourier domain, as follows:

$$\mathbf{P}'_{j,mn} = \mathcal{F}^{-1}\left\{\exp\left(2\pi i\left(\frac{\mu(x_k - \lfloor x_k \rceil)}{M} + \frac{\nu(y_k - \lfloor y_k \rceil)}{N}\right)\right)\mathcal{F}\{\mathbf{P}\}_{\mu\nu}\right\} \tag{3.10}$$

This model selects the closest region of the object function $\boldsymbol{O}$ to extract, and

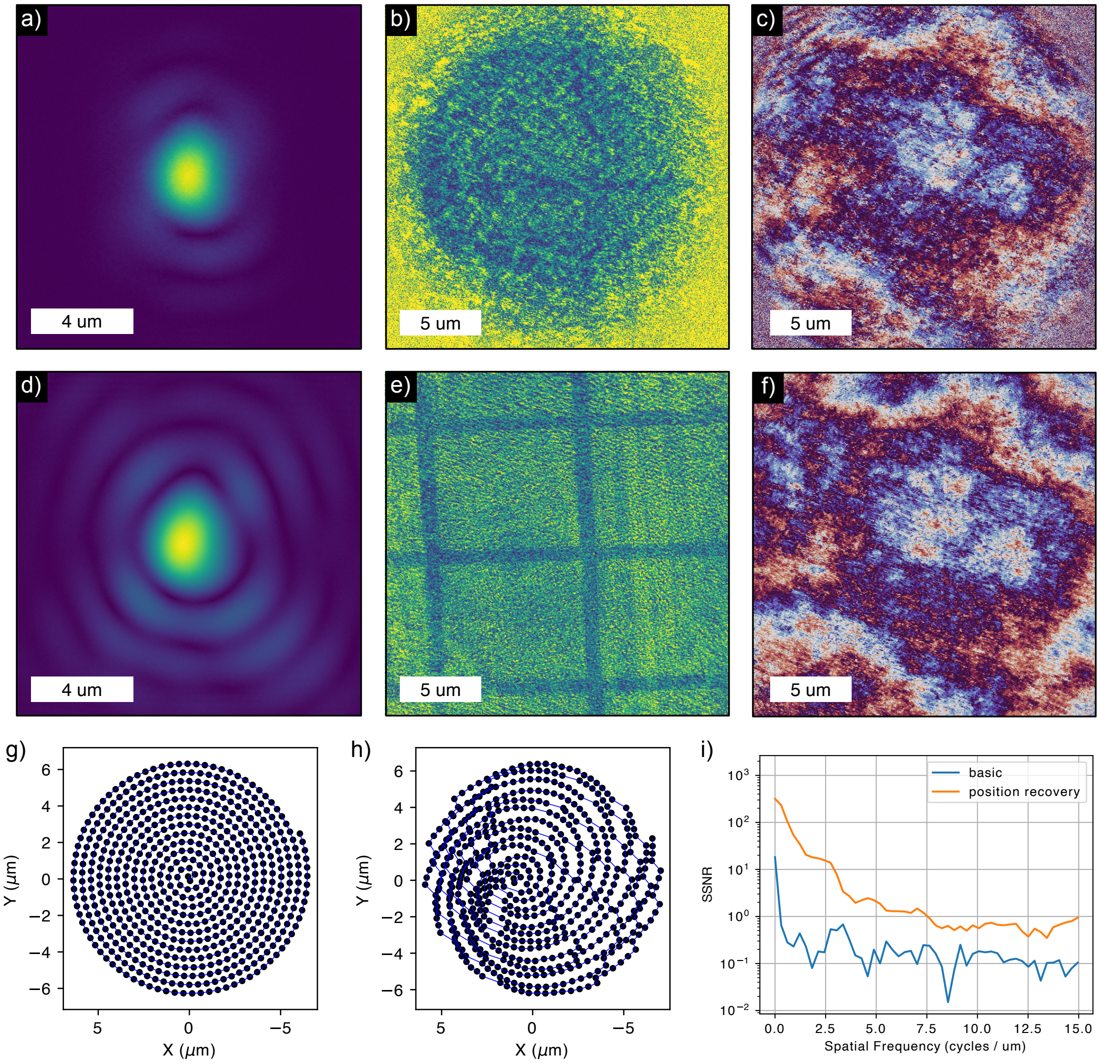

Figure 3-3: **Role of Position Correction**. An example showing the power of position correction, using resonant soft x-ray reflection data collected at the Ni L3 edge from a thin film sample of NdNiO3. (a) The amplitude of the probe function, (b) the amplitude of the object function, and (c) the phase of the object function, using the nominal probe translations without position correction. (d), (e), and (f) the same, but after including the position correction model in Equation (3.10). (g) the original positions, and (h) the corrected ones. (i) comparison of the spectral signal-to-noise ratio (SSNR) associated with the resulting object reconstructions.

then performs a small subpixel shift of the probe to adjust the position to subpixel accuracy. This method of performing the shift is equivalent to a sinc-interpolation with periodic boundary conditions, an appropriate method for typical ptychography probe functions because they often can be well approximated as band limited.

Although there are still discrete changes in the calculated exit wave when the rounded positions $\lfloor x_k \rceil$ and $\lfloor y_k \rceil$ jump across a step, the forward model remains differentiable because the rounding operation is not included in the computational graph and the derivative is always well-posed when $\lfloor x_k \rceil$ and $\lfloor y_k \rceil$ are considered to be fixed. As Figure 3-3 shows, this model can produce a dramatic improvement in reconstruction quality on otherwise good data with positioning errors.

3.7 Corrections for other common errors

A variety of small, but common, sources of error plague many ptychography experiments. Here we treat them simply:

3.7.1 Detector background

To offset a detector background, alter the measurement model:

$$\mathbf{I}^{\text{sim}}_{j,mn} = |\tilde{\mathbf{E}}_{j,mn}|^2 + \mathbf{B}^2_{mn}. \tag{3.11}$$

The definition of the parameter $\mathbf{B}$ as the square-root of the background is a common trick used for non-negative parameters. The background is usually constrained to be non-negative because negative backgrounds can produce negative simulated intensities, causing issues for the standard mean-squared-amplitude-error loss model and the Poisson negative log-likelihood metric.

3.7.2 Probe translation

It is often the case that the optics which produce the probe are scanned across the sample, rather than the canonical situation of a sample scanning across a fixed beam.

When this is the case, the diffraction pattern will shift slightly on the detector as the probe moves. This becomes a genuine issue when the scanned field of view approaches an appreciable fraction of the size of a detector pixel. Because pixels on most x-ray detectors are a few tens of microns wide, this effect can start to be noticeable for scan windows even as small as a few microns. This can be fixed by introducing a small correction to the exit wave:

$$\mathbf{E}'_{k,mn} = \exp\left(2\pi i\left(\frac{m(x_k - \bar{x})}{Md} + \frac{n(y_k - \bar{y})}{Nd}\right)\right)\mathbf{E} \tag{3.12}$$

Where d is the size of an individual pixel on the detector. This linear phase ramp, applied to the exit wave in real space, creates a subpixel shift in the Fourier-space wavefield at the detector plane. This is parallel to the same approach used to introduce subpixel shifts to the probe's position in real space. If uncorrected, this effect will cause a quadratic phase structure to appear in the reconstruction, which can start to degrade the quality of the reconstruction if the probe is large enough for the phase curvature to make an appreciable difference at the per-diffraction-pattern level.

3.7.3 Large probes

Sometimes, one is forced to use an illumination that doesn't fit into the finite-sized array that stores the probe function. In this case, it is possible to increase the size of the probe array by simulating wavefields on a finer grid of pixels than the actual sampling at the detector plane [58].

To introduce this into automatic differentiation, we create a probe with $aM \times aN$ pixels, and then simply modify the measurement model

$$\mathbf{I}'^{\text{sim}}_{j;m,n} = \sum_{\mu,\nu=1}^{a} \mathbf{I}^{\text{sim}}_{j;am+\mu,an+\nu} \tag{3.13}$$

so that each pixel on the detector is compared to the sum of an $a \times a$ region of the simulated wavefield at the detector plane. Running a reconstruction with this

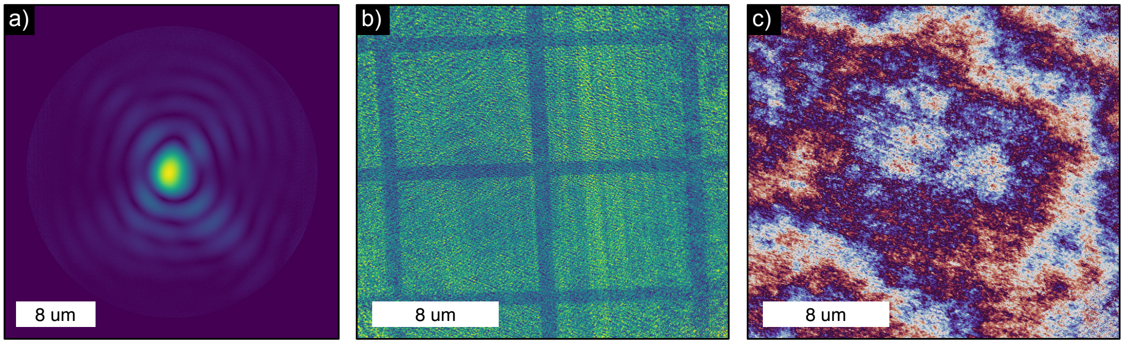

Figure 3-4: **Further Small Corrections**. A reconstruction from the same data used in Figure 3-3, but including a correction for the detector saturation, doubling the size of the simulated probe, and the unstable illumination model to allow for varying illumination intensity. (a) The amplitude of the probe function, (b) the amplitude of the object function, and (c) the phase of the object function.

forward model can then recover the structure of the probe over a larger area, at the cost of setting up a more poorly conditioned inverse problem. An example of the improvement in quality that can be achieved in a typical case is shown in Figure 3-4.

3.7.4 Superresolution

Sometimes the signal rate is so high, even out at the edge of the detector, that one has hope of reconstructing the object to a higher resolution than the naïve pixel size in the detector conjugate coordinate space [58]. This can be set up by simply simulating $\tilde{\mathbf{E}}$ on a larger array than the actual detector images and cropping it when simulating the intensities. This results in a correspondingly smaller pixel size in real space for both the probe and object images. I used this method in the reconstructions from the data in Chapter 5, due to the high signal rates.

However, caution needs to be taken because the high-frequency components may not be actually constrained by the data. Specifically, the probe needs to contain frequencies high enough to map the highest frequency components in the object back to the detector field of view. This can be rephrased as a the traditional limitation that

$$\mathrm{NA}_{\mathrm{total}} = \mathrm{NA}_{\mathrm{detector}} + \mathrm{NA}_{\mathrm{illumination}} \tag{3.14}$$

Usually, superresolution is a hopeless endeavor unless the signal is extraordinarily high, and the high-frequency components will become dominated by noise before they reach the ultimate limit. There's also not much point to it, because simply moving the detector closer to the sample will capture a higher numerical aperture. Perhaps in some far-off future, numerical apertures for extreme ultraviolet (EUV) or soft x-ray experiments will approach 1, and in that case super-resolution imaging may become valuable.

3.7.5 Finite Pixels

When the diffracted wavefield imaged on the detector has structure on a length scale near the pixel length scale, the effect of the finite pixel size can become relevant. The naïve model simply assumes that each pixel is sampling an infinitely small region of the wavefield at a particular location. In contrast, the physical pixel integrates over a specifically shaped area, typically a square region with edge length equal to the pixel pitch.

We can account for this difference explicitly by ensuring that we accurately sample the intensity distribution at the detector plane, and then convolving that distribution with a pixel shape function. Formally, this is accomplished with:

$$\tilde{\mathbf{E}}_j = \mathcal{F}\{\mathrm{Pad}(\mathbf{E}, M, N)\} \tag{3.15}$$

$$\mathbf{I}_j^{\mathrm{sim}} = |\tilde{\mathbf{E}}_j|^2 \tag{3.16}$$

$$\mathbf{I}_j'^{\mathrm{sim}} = \mathcal{F}\{\mathbf{H}_{mn}\,\mathcal{F}^{-1}\{\mathbf{I}_j\}_{mn}\} \tag{3.17}$$

$$\mathbf{I}_{j;m,n}''^{\mathrm{sim}} = \mathbf{I}'_{j;2m,2n}. \tag{3.18}$$

Where the function $\mathrm{Pad}(\mathbf{E}, M, N)$ above pads the exit wave array with zeroes to double its size. $\mathbf{H}_{mn}$ is defined as the Fourier transform of the pixel's sensitivity function, which can be calculated analytically beforehand. For the simple case of a square pixel with side length matching the pitch, this is:

$$\mathbf{H}_{mn} = \text{sinc}(m - M)\,\text{sinc}(n - N), \tag{3.19}$$

or, in other words, a pair of sinc functions that fall to their first zeroes right at the edge of the simulated probe field of view. This correction is especially valuable—although also especially computationally expensive—when using large probe that already requires oversampling on the detector.

3.7.6 Saturation

Because of the large dynamic range present in typical ptychography datasets, detector saturation is a perennial problem. Fortunately, provided that the saturated pixels don't affect their neighbors (as they, unfortunately, often do), it can be easily corrected for with the following addition to the measurement model:

$$\mathbf{I}'^{\text{sim}}_{j;mn} = \min(\mathbf{I}^{\text{sim}}_{j;nm}, s) \tag{3.20}$$

Where s refers to the saturation threshold. This model effectively prevents the error term arising from any saturated pixels from propagating backward through to the object gradients. This works remarkably well, but it does create the risk that if the initial simulation is fully saturated, there will be no gradients propagated whatsoever and the reconstruction will stagnate. For this reason, it is preferable to mask off saturated pixel based on $\mathbf{I}^{\text{meas}}$, but this trick will indeed work in a pinch. This model was also included in the reconstructions feeding into Figure 3-4.

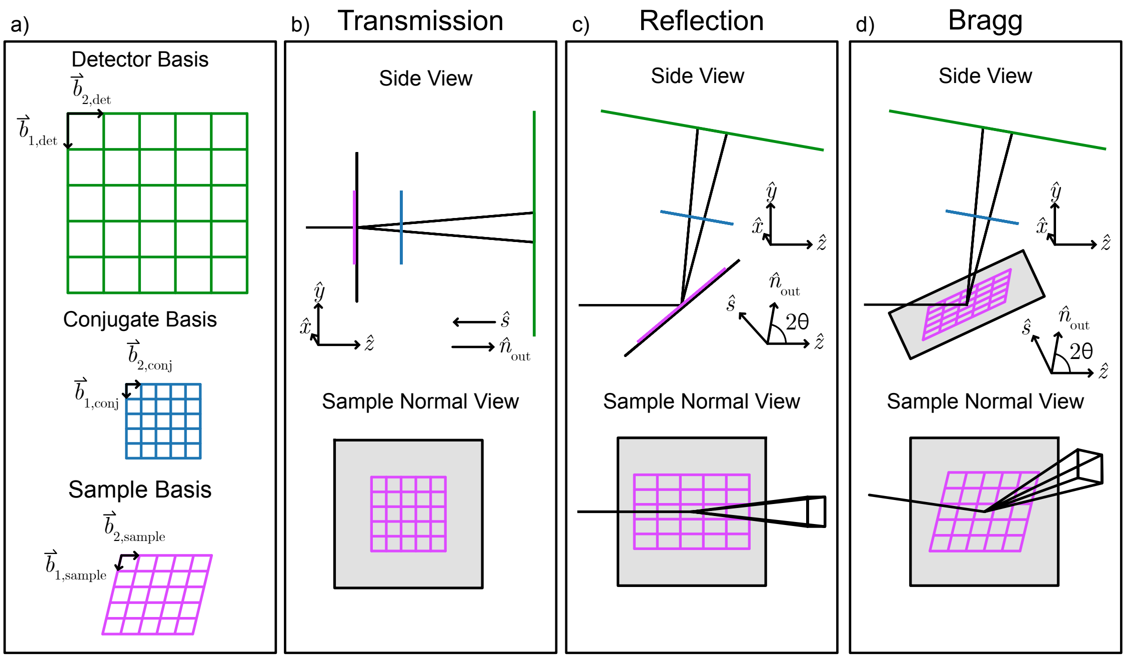

Figure 3-5: **Projection geometry**. (a) The definition of the basis vectors for the various coordinate systems for stored arrays: detector, detector conjugate, and sample-plane. (b) The geometry for transmission ptychography experiments. (c) The specular reflection geometry. (d) The general Bragg geometry, where the Bragg vector need not be normal to the sample surface.

3.8 Corrections for the reflection and Bragg geometries, as well as high NA

So far, we have considered ptychography data collected in transmission through a sample that is perpendicular to the propagation direction of incoming light. It is also possible to collect ptychography data in reflection mode [59, 60], and it is even possible to collect ptychography from light scattered into a Bragg diffraction peak [61, 62]. Bragg ptychography data is often reconstructed in 3 dimensions, but when studying sufficiently thin films it is often possible to do a Bragg mode reconstruction of a 2D object [63].

We can treat both the reflection and 2D Bragg cases with a common framework. Without loss of generality, we consider a canonical scenario where the illuminating light propagates along the positive z axis, strikes a sample with surface normal vector

$\hat{s}$, and reflects off at an angle 2θ from the surface, as shown in Figure 3-5. This defines an outward direction

$$\hat{n}_{\text{out}} = [0, \sin(2\theta), \cos(2\theta)] \tag{3.21}$$

We assume that the center of the detector is along a line defined by $\hat{n}_{\text{out}}$, and the detector face is perpendicular to $\hat{n}_{\text{out}}$. This gives us the detector basis vectors (see Figure 3-5):

$$\boldsymbol{B}_{\text{det}} = \begin{pmatrix} 0 & p \\ p\cos(2\theta) & 0 \\ p\sin(2\theta) & 0 \end{pmatrix} = \begin{pmatrix} \vec{b}_{1,\text{det}} & \vec{b}_{1,\text{det}} \end{pmatrix}. \tag{3.22}$$

Our first task is to understand how an exit wave leaving the plane of the sample will ultimately wind up propagating to the detector. We start by considering the detector conjugate coordinate system parallel to the plane of the detector, shown in blue in Figure 3-5. We find a square grid of points with basis:

$$\boldsymbol{B}_{\text{conj}} = \begin{pmatrix} 0 & \frac{\lambda d}{Mp} \\ \frac{\lambda d}{Mp}\cos(2\theta) & 0 \\ \frac{\lambda d}{Mp}\sin(2\theta) & 0 \end{pmatrix} = \begin{pmatrix} \vec{b}_{1,\text{conj}} & \vec{b}_{1,\text{conj}} \end{pmatrix}. \tag{3.23}$$

This is just the normal grid of points of spacing $\frac{\lambda d}{Mp}$, but tilted to stay parallel to the detector plane. Any wavefield sampled at this plane will have the standard Fourier transform relationship with the detector.

To zeroth order, we can approximate the diffraction from the tilted sample by projecting the exit wave at the sample surface along $\hat{n}_{\text{out}}$ onto the detector conjugate coordinate grid. This approximation becomes valid when the projection distance required is less than the depth of focus corresponding to the detector's numerical aperture. Our task is now to find the sampling points at the sample plane that project back onto the sampling points in the detector conjugate coordinate system. We parameterize the sample basis as:

$$\boldsymbol{B}_{\text{sample}} = \begin{pmatrix} \vec{b}_{1,\text{sample}} & \vec{b}_{1,\text{sample}} \end{pmatrix} . \tag{3.24}$$

where $\vec{b}_{1,\text{sample}}$ and $\vec{b}_{1,\text{sample}}$ are related to the basis vectors $\vec{b}_{1,\text{det}}$ and $\vec{b}_{2,\text{det}}$ through the projection shown in Figure 3-5. The projection is calculated by setting up a set of linear equations. First, we know that the component of the projected vectors along the surface normal direction is 0, so:

$$\hat{s} \cdot \vec{b}_{i,\text{sample}} = 0 \tag{3.25}$$

And secondly, we know that the re-projection of the sample-plane vectors back onto the detector conjugate plane will recover the original detector conjugate coordinate system:

$$(\boldsymbol{I} - \hat{n}_{\text{out}} \otimes \hat{n}_{\text{out}})\vec{b}_{i,\text{sample}} = \vec{b}_{i,\text{conj}} \tag{3.26}$$

Where $\otimes$ is the outer product. Combining these and solving with a pseudoinverse, we find that:

$$\vec{b}_{i,\text{sample}} = \text{pinv}\left[\begin{pmatrix} \boldsymbol{I} - \hat{n}_{\text{out}} \otimes \hat{n}_{\text{out}} \\ \hat{s} \end{pmatrix}\right] \begin{pmatrix} \vec{b}_{i,\text{conj}} \\ 0 \end{pmatrix}, \tag{3.27}$$

with pinv as the Moore-Penrose pseudoinverse.

To zeroth order, we can now store an object array related to the detector plane through a discrete Fourier transform, just as we do with standard transmission ptychography. The only difference is that the meaning of the pixels has changed. Now, each pixel refers to a specific location on the sample plane defined by the basis vectors above.

When the depth of focus of the detector NA is sufficiently the only correction which actually needs to be applied to account for the reflection or Bragg geometry is the modified mapping between the sample translation in real-space coordinates and the translation in pixel-space coordinates. If the translation occurs along the plane of the sample, then the corrected mapping can be accomplished using the standard method of multiplying by the pseudoinverse of the basis vectors:

$$\begin{pmatrix} x_j \\ y_j \end{pmatrix} = \mathrm{pinv}[\boldsymbol{B}_{\text{sample}}]\vec{t}_j. \tag{3.28}$$

However, if the translations aren't entirely within the plane of the sample, then they first need to be projected into the plane of the sample:

$$\vec{t}_i = \mathrm{pinv}\left[\begin{pmatrix} \boldsymbol{I} - \hat{z} \otimes \hat{z} \\ \hat{s} \end{pmatrix}\right] \begin{pmatrix} \vec{t}_i \\ 0 \end{pmatrix}. \tag{3.29}$$

Here, as a reminder, the incoming light is defined to propagate along the positive $\hat{z}$ direction.

In the simplest case, this geometric transformation is all that is needed to enable Bragg mode ptychography. Helpfully, in the case of reflection-mode ptychography, all the transformations cancel out. Consequently, in many cases a reflection-mode ptychography dataset can be treated as though it were a transmission-mode dataset with the motion along the y axis inverted. In fact, the data shown in Figures 3-3 and 3-4 was taken in the specular reflection geometry with $2\theta = 114^\circ$. The reconstructions shown were run with no correction for the reflection geometry other than flipping the vertical axis of the diffraction pattern images. This means that the images shown in Figures 3-3 and 3-4 are actually images of the sample surface, viewed from the perspective of the detector.

Using the projections described above, we could also transform these images into the natural plane of the detector, as is done with the results in Figure 3-6. However, those results include the correction for sample tilt below, which is needed when the pure-projection approximation breaks down.

Because the sample and detector planes are no longer parallel, the calculation of the diffraction patterns as the discrete Fourier transform of the real-space array can quickly break down.

Specifically, each pixel on the diffraction pattern is still related to the inner product of the object function with a complex exponential. However, the wavevectors sampled

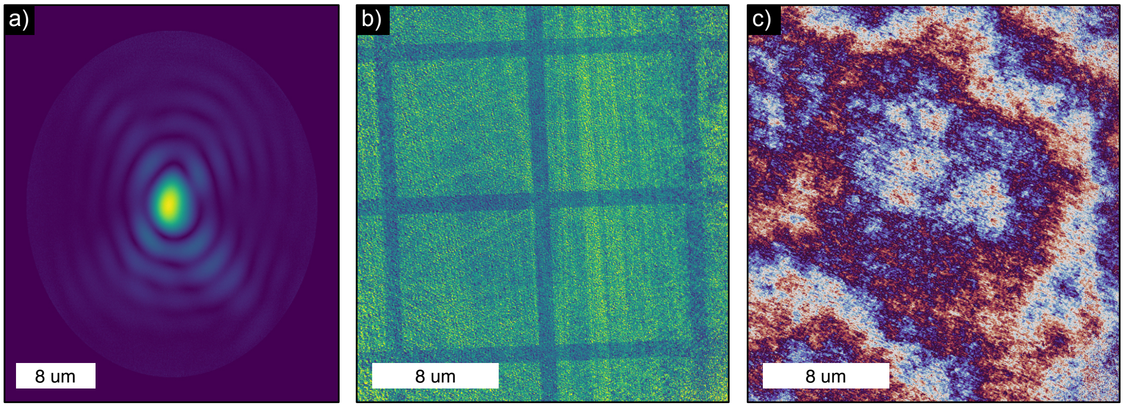

Figure 3-6: **Reflection Geometry Correction**. A reconstruction from the same data used in Figure 3-3 and 3-4, but now also correcting for the probe propagation with the optic translations, and the effect of the tilt between the sample and detector planes. Note that, from this view, the square grid of gold lines on the sample appears with the correct aspect ratio.

by the detector deviate from the perfect grid associated with the discrete Fourier transform.

We can correct for this by explicitly calculating the points in Fourier space sampled by the detector pixels, and then sampling the discrete Fourier transform at these locations. The essential calculation is determining those points. A given detector pixel at location (m, n), defined relative to the center of the detector, samples light with the wavevector

$$\vec{k}(m,n) = \frac{2\pi}{\lambda} \frac{d\hat{n}_{\text{out}} + m\vec{b}_{1,\text{det}} + n\vec{b}_{2,\text{det}}}{|d\hat{n}_{\text{out}} + m\vec{b}_{1,\text{det}} + n\vec{b}_{2,\text{det}}|} \tag{3.30}$$

where d is the sample-detector distance. However, these wavevectors have a problem, because they still contain the carrier frequency associated with light traveling in the direction $\hat{n}_{\text{out}}$. This ignores the fact that, when working in the reflection or Bragg mode, there is a phase-matching process going on, either to satisfy the specular reflection condition or the Bragg condition. In other words, the sample has some high frequency components which create the high-angle scattering we're measuring, we're interested in the low frequency modulations around the central diffraction or reflection wavevector.

The simplest way to remove this carrier frequency is to simply subtract off the frequency associated with $\hat{n}_{\text{out}}$:

$$\vec{k}'(m,n) = \vec{k}(m,n) - \frac{2\pi}{\lambda}\hat{n}_{\text{out}}. \tag{3.31}$$

Finally, we transform this back into the pixel-space basis of the Fourier transform array. We can do this using multiplication with the appropriately scaled real-space basis of the exit wave array:

$$\begin{aligned} m_{\text{out}}(m,n) &= \frac{M}{2\pi}\vec{B}_{1,\text{sample}} \cdot \vec{k}'(m,n) \\ n_{\text{out}}(m,n) &= \frac{N}{2\pi}\vec{B}_{2,\text{sample}} \cdot \vec{k}'(m,n). \end{aligned} \tag{3.32}$$

This sampling should ideally be done on the final simulated diffraction pattern, and in practice a linear interpolation appears to be sufficient.

This correction is appropriate for both reflection-mode and Bragg-mode, but notably it also is applicable to transmission-mode measurements. In this case, it naturally corrects for the deviations from the paraxial approximation that occur with high numerical aperture detectors. Consequently, this correction can also be dropped in for transmission-mode ptychography to correct for the effects of high-NA detectors.

Chapter 4

Single-frame far-field diffractive imaging with randomized illumination

4.1 Preface

The following chapter is based on a publication that originally appeared in *Optics Express* [64], modified to fit the format of this thesis. It also includes a new section at the end with further discussion highlighting progress in our understanding since the paper's publication.

The paper was coauthored by myself, Kahraman Keskinbora, Umut T. Sanli, Markus Weigand, and Riccardo Comin, copyright is held by Optica Publishing Group. I initiated the line of inquiry, developed the reconstruction algorithm, performed the optical experiment, designed all the optics, and analyzed the data for all experiments. I was not responsible for the manufacturing of the x-ray optics, nor was I involved in performing the x-ray experiment. See the published version [64] for full affiliation and funding information.

4.2 Abstract

We introduce a single-frame diffractive imaging method called Randomized Probe Imaging (RPI). In RPI, a sample is illuminated by a structured probe field containing

speckles smaller than the sample's typical feature size. Quantitative amplitude and phase images are then reconstructed from the resulting far-field diffraction pattern. The experimental geometry of RPI is straightforward to implement, requires no near-field optics, and is applicable to extended samples. When the resulting data are analyzed with a complimentary algorithm, reliable reconstructions which are robust to missing data are achieved. To realize these benefits, a resolution limit associated with the numerical aperture of the probe-forming optics is imposed. RPI therefore offers an attractive modality for quantitative X-ray phase imaging when temporal resolution and reliability are critical but spatial resolution in the tens of nanometers is sufficient. We discuss the method, introduce a reconstruction algorithm, and present two proof-of-concept experiments: one using visible light, and one using soft X-rays.

4.3 Introduction

Diffractive imaging refers to a collection of computational imaging techniques that reconstruct quantitative amplitude and phase images directly from diffraction patterns [65, 66]. Twenty years after the first demonstrations, [20, 67–69] this methodology has spurred a wave of technical and scientific innovations and motivated the development of a new generation of advanced X-ray light sources. Today, diffractive imaging offers unprecedented opportunities for quantitative phase imaging using both single-frame and multi-frame methods.

In general, single-frame methods can perform time-resolved imaging of non-reproducible dynamics but multi-frame methods are more flexible and reliable. The comparison between Coherent Diffractive Imaging (CDI) and ptychography exemplifies this compromise. CDI, like most single-frame methods, uses small samples with a sharply defined boundary. The boundary is typically imposed by an opaque mask deposited on the sample [66]. This requirement ultimately restricts what systems can be studied and predefines the field of view. In contrast, ptychography uses multiple diffraction patterns from overlapping regions of a sample to improve the reconstruction's reliability and allow extended samples to be imaged [36–39, 51, 70, 71]. This trade-off has driven an enduring search for single-frame imaging methods which retain the reliability and flexibility of ptychography. Here we implement an approach to tackling this challenge which synthesizes two ingredients from different corners of the diffractive imaging world.

The first ingredient is Band-Limited Random (BLR) illumination, a variety of structured illumination which has been explored in the context of ptychographic imaging [72–76]. Our choice to use this illumination arises from the longstanding observation that the reliability of a phase retrieval problem is typically improved by the presence of high-frequency phase structures [77]. This idea has led to a number of proposed diffractive imaging methods which use randomized illumination [78–81] or a quadratic phase ramp as in Fresnel CDI[82–84]. The concept of using randomized illumination is also closely related to Coherent Modulation Imaging (CMI) [85–88], a method which imprints randomized high-frequency structures on the wavefield with a mask placed downstream from the sample. In this work we point out that there is strong reason to believe that, when the object function is band-limited to a sufficiently low frequency, diffraction data resulting from many of these methods can support a unique solution to the phase retrieval problem even without a finite support constraint.

This observation leads naturally to the second ingredient, a band-limiting constraint applied to the object function during reconstruction. This constraint is loosely inspired by single-shot ptychography algorithms [89–92]. In single shot ptychography, an entire ptychography dataset is taken in a single frame by converging a grid of probes onto the sample and imaging all the diffraction patterns simultaneously. To avoid interference between neighboring diffraction patterns, the Fourier transform of the object must fall to zero beyond some band-limiting frequency. To acknowledge this band-limiting requirement on the underlying object function, single-shot ptychography algorithms explicitly impose a band-limiting constraint on the object via a judicious choice of the pixel pitch in the object array. These algorithms can produce remarkably robust reconstructions, with reliability in the same class as ptychography. It therefore seems plausible that a similar constraint could also pay dividends when analyzing data arising from experiments, such as those discussed above, where the diffraction data alone provides a particularly strong constraint when the object is band-limited.

Randomized Probe Imaging (RPI) synthesizes these two ingredients by applying a band-limiting constraint to diffraction data collected under BLR illumination. The exact structure of the illumination is found with an initial ptychography scan used for calibration. As we show, this combination provides enough information to solve the phase retrieval problem without the need for either a finite support constraint or multiple diffraction patterns. The result is a reliable single-frame (and therefore potentially single-shot) technique which can be easily implemented at most ptychography beamlines by simply replacing an optic. As a bonus, RPI is applicable to extended samples and is robust to missing regions of data, such as those caused by dead pixels or the presence of beamstops. RPI is therefore an attractive alternative to holography, CDI, and CMI in a variety of situations, such as x-ray microscopy applications where reliability is paramount, sample environments are bulky and a resolution limit in the tens of nanometers is acceptable.

4.4 Randomized probe imaging

4.4.1 Basic Principles

The experimental geometry at the heart of RPI, outlined in Figure 4-1a, is the same as a geometry commonly used for ptychography. In fact, the first step of any RPI experiment is to collect a ptychography dataset to solve for the illumination wavefield. A probe-forming diffractive hologram known as a Randomized Zone Plate (RZP) focuses light of a wavelength λ to a focal spot over a focal distance $f(\lambda)$. An Order Selecting Aperture (OSA) is used in conjunction with a central beamstop on the probe-forming optic to block all but the intended diffractive order from reaching the sample.

The only nonstandard requirement in our geometry is the use of an RZP as the

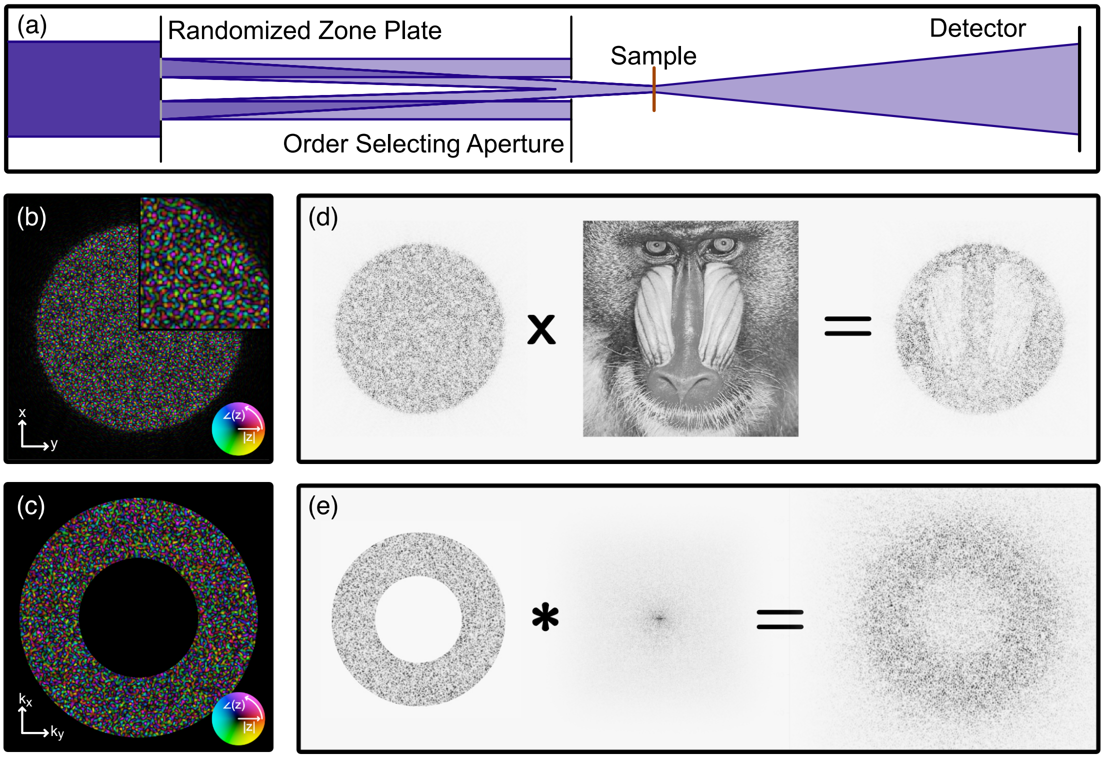

Figure 4-1: **Conceptual Model**. (a), A diagram outlining the experimental geometry. (b), A typical band-limited random focal spot in real space, with an inset showing the details of the finely structured amplitude and phase. (c), The corresponding representation in Fourier space. (d), The interaction model in real space, where the incident light $P(x, y)$ is multiplied by the object $O(x, y)$ to find the exit wave. (e), The same interaction model in Fourier space, where the incident illumination's Fourier transform $\tilde{P}(k_x, k_y)$ is convolved with the Fourier transform of the object $\tilde{O}(k_x, k_y)$ to find the diffraction pattern.

probe-forming optic in place of the traditional Fresnel zone plate or pinhole. As shown in Figures 4-1b and 4-1c and further explored in Figure A-2, distortions in this zone plate broaden the focal spot and fill it with speckles. The focal spot diameter can be controlled over a wide range during the optic design process, as discussed in Appendix A.2. The diameter of this speckle-filled focal spot ultimately determines the field of view of the RPI reconstructions. The speckle size, related to the numerical aperture of the zone plate, sets the resolution element. The effect of this randomized wavefield on the final diffraction pattern is very similar to the effect of the phase modulator used in CMI.

This BLR probe wavefield $P(x, y)$ then interacts with a thin sample described by an object function $O(x, y)$. The resulting exit wave $E(x, y) = P(x, y)O(x, y)$ propagates into the far field where its intensity is imaged onto an area detector (Fig 4-1d). The goal of RPI is to reconstruct a discrete representation of the object $O_{ij} = O(i\Delta_x, j\Delta_x)$, using only the measured intensities $I_{ij} = |\tilde{E}(i\Delta_k, j\Delta_k)|^2$ and knowledge of the incident probe wavefield P_{ij}. In this notation, $\tilde{E}(k_x, k_y)$ refers to $\mathcal{F}\{E(x, y)\}$, the Fourier transform of $E(x, y)$ and Δ_x, Δ_k are step sizes in real space and Fourier space, respectively.

To understand how a reconstruction of O_{ij} is possible from a single diffraction pattern without the use of a finite support constraint, we consider the case where $O(x, y)$ is band-limited to a maximum frequency k_o which is smaller than the highest frequency contained in the probe k_p. As the visualization of this process in Figure 4-1e makes clear, when $k_o \ll k_p$, the final diffraction pattern will occupy a much larger region of Fourier space than the object itself does. As a result, the measured diffraction pattern can easily contain more measurements than there are independent degrees of freedom in a band-limited object, leading to a potentially well-defined inverse problem on the space of band-limited objects. This suggests the reconstruction strategy that we eventually pursue, where we impose an explicit band-limiting constraint on the object in lieu of a finite support constraint.

We additionally note that the map between the complex-valued object and the complex-valued wavefield at the detector plane is conveniently linear. As a result, the inverse problem can be expressed as a specific case of the generic phase retrieval problem [93]. This connection lets us propose a limit on how high the band-limiting frequency k_o of the original object can be as compared to the probe's maximum frequency k_p before the inverse problem becomes fundamentally unstable. We note that for phase retrieval to be well posed, the number of intensity measurements (which scales with the area of the diffraction pattern) must typically be greater than four times the number of complex parameters in the object (which scales with the object's support in Fourier space) [94, 95]. In Appendix A.5, we show that in the context of our reconstruction algorithm, this requirement limits the applicability of RPI to objects with resolution ratios $R = \frac{k_o}{k_p}$ below approximately 0.94. Later, we demonstrate experimentally that reconstructions on objects with at least $R \leq 0.6$ are feasible in practice.

A practical advantage of this method is that the raw diffraction pattern itself provides convincing evidence that the underlying object meets the band-limiting requirement. This is because the diffraction pattern is derived from the convolution of the Fourier representations of the object and probe. The speckles in the diffraction pattern will therefore generically fill a region defined by the binary dilation of the probe's support in Fourier space with the object's. As a result, for a BLR probe like that shown in Figure 4-1c, $k_p + k_o$ can be estimated as the highest frequency at which the measured intensity remains above the noise floor. In other words, if the true object does not satisfy the band-limiting requirement, the diffraction pattern itself will almost always reveal this via diffracted intensity at frequencies which are simply too high to be accessed by a sufficiently band-limited sample. This fact further suggests that reconstructions of band-limited objects may benefit from randomized illumination even if no explicit band-limiting constraint is applied - as is the case in CMI. This idea, and it's limitations, are explored further in our numerical experiments.

Finally, a real-space view of the probe's structure further motivates the use of an explicit band-limiting constraint to perform reconstructions from RPI data. As seen in Figure 4-1b, a fine mesh of zeros interpenetrates the BLR illumination - a feature not shared with CMI modulators. Attempting to reconstruct an object at

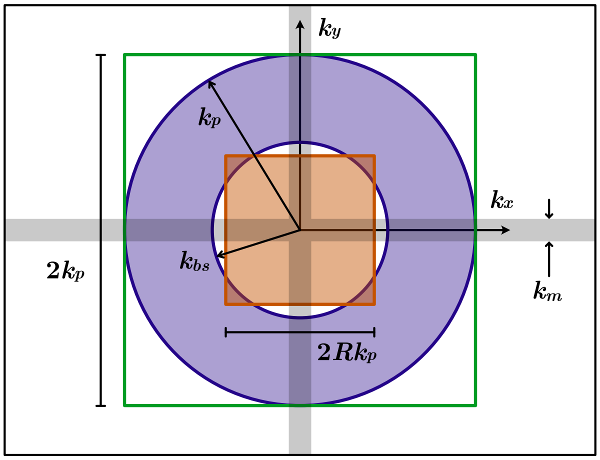

Figure 4-2: **Reciprocal Space Diagram**. The simulation geometry for our numerical experiments, as viewed in Fourier space. The purple region shows the extent of the probe function, bounded on the outside at a frequency k_p and bounded on the inside by the extent of the central beamstop, k_{bs}. The orange region denotes the band-limited object, and the region bounded in green shows the extent of the noise introduced outside the band-limiting constraint. The grey shaded region is the extent of the simulated detector "dead zone".

a finer resolution than the speckle size will therefore inevitably lead to pixels in the reconstructed object which are poorly constrained simply because they are weakly illuminated. If we explicitly apply a band-limiting constraint with $R < 1$, every pixel in the object will be illuminated by at least one speckle. In this way, the band-limiting constraint alleviates this issue of weakly constrained modes at the same time as it dramatically improves the reliability of reconstructions. The limitations imposed on RPI by this issue, and their relationship to the beamstop diameter, are explored further in Appendix A.6.

4.4.2 Reconstruction Algorithm

Our reconstruction algorithm is based on Automatic Differentiation Ptychography [46, 47], using the forward model defined in Equation (4.1). This model incorporates the band-limiting constraint by taking a low-resolution representation of the unknown object, O'_{kl}, as input and using it in conjunction with a higher-resolution known probe P_{ij} defined on the standard detector conjugate coordinate grid used in CDI and ptychography:

$$\tilde{E}_{ij} = \mathcal{F}\{P_{ij}\mathcal{F}^{-1}\{\mathrm{pad}(\mathcal{F}\{O'_{kl}\})\}_{ij}\}. \tag{4.1}$$

The first step in the model is to upsample the low-resolution object O'_{kl} by padding it with zeros in Fourier space. We then directly multiply the upsampled object by P_{ij}, the high-resolution representation of the probe, and Fourier transform the resulting exit wave to propagate it to the detector plane. The size of the low-resolution object array O'_{kl} is chosen ahead of time based on the measured radius of the freely propagated probe's diffraction pattern and a target resolution ratio R_{rec}.

To perform a reconstruction, we start with an initial guess of the object function and use Equation (4.1) to simulate the corresponding diffraction pattern. Next, we calculate the normalized mean squared error between the measured diffraction amplitudes and a simulated diffraction pattern (including a known detector background):

$$L = \frac{1}{\sum_{ij} I_{ij}} \sum_{ij} \left(\sqrt{|\tilde{E}_{ij}|^2 + B_{ij}} - \sqrt{I_{ij}} \right)^2 . \tag{4.2}$$

The quantity L is referred to as the diffraction loss. We then use automatic differentiation to calculate the Wirtinger derivative of L with respect to O'_{kl} and feed those derivatives into an update for the object guess O'_{kl} using the Adam algorithm [48]. This process is repeated iteratively until the object converges. Remarkably, we were able to use the same set of algorithmic parameters across all our numerical experiments as well as both experimental demonstrations, indicating the robustness of this reconstruction method to variations in experimental conditions.

Finally, we note that because the forward model simply defines a special case of the generic phase retrieval problem, other phase retrieval algorithms originally designed for geometries such as CDI and CMI can likely be modified to include a band-limiting constraint without affecting their validity. Potentially applicable families of algorithms include those designed to solve the generic phase retrieval problem, such as Wirtinger Flow [96] and SketchyGCM [97]. In addition, algorithms such as hybrid input-output [18] and difference map [98] that focus on finding solutions

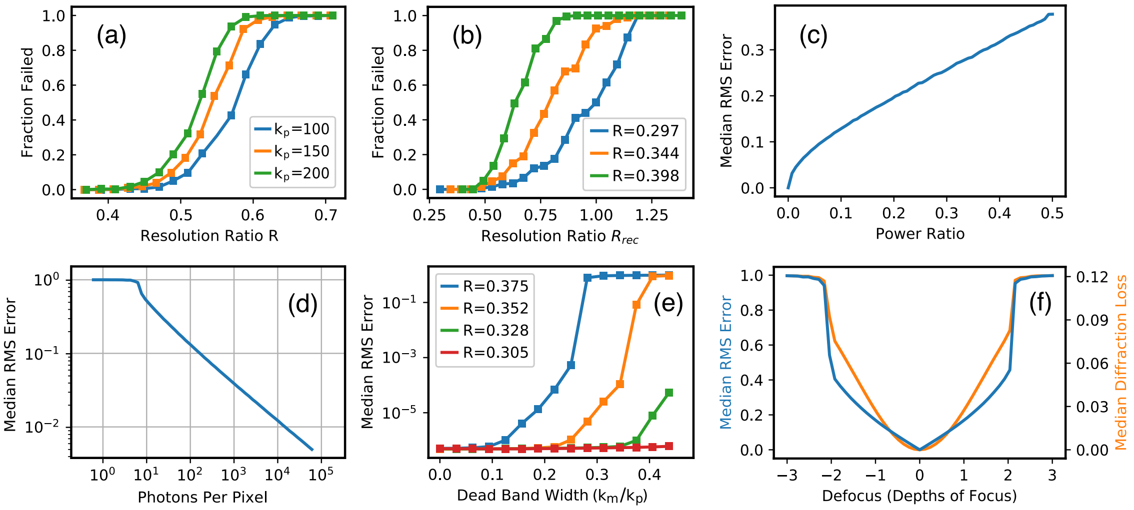

Figure 4-3: **Numerical Results**. (a), The dependence of failure probability of a reconstruction on the true object's resolution ratio R for various k_p in pixels. (b), The dependence of failure probability on the resolution ratio R_{rec} of the applied band-limiting constraint, for diffraction from underlying objects with various R. (c)-(e), The dependence of the median image RMS error ϵ (Eq. (4.3)) on (c), the ratio of power outside the band-limiting constraint to power inside; (d), the number of photons per illuminated pixel in the reconstructed object; and (e), the size of the cross-shaped dead region. (f), The dependence of the normalized diffraction RMS loss L (Eq. (4.2)) and the normalized image RMS error ϵ (Eq. (4.3)) on the error in the assumed propagation state of the probe.

within the intersection of constraint sets are likely to be applicable, with the sets in this case generated by the band-limiting and detector intensity constraints. Although we pursued an Automatic Differentiation based approach here to provide the most flexibility while exploring this concept, it is likely the case that the use of tailored reconstruction algorithms such as these, designed specifically with the phase retrieval problem in mind, will lead to significant gains in computational efficiency.

4.5 Results

4.5.1 Numerical Results

Although the previous discussion demonstrates that RPI is valid for sufficiently small k_o, the precise region of parameter space where the inverse problem is well posed remains poorly defined. We therefore performed a collection of numerical experiments to clarify under what circumstances RPI reconstructions succeed, as well as to understand the impact of various noise sources on the quality and reliability of RPI reconstructions.

In the following numerical experiments, we used an implementation of our reconstruction algorithm developed in PyTorch [99], which is available from the authors upon request. All object guesses were initialized with Gaussian random real and imaginary parts, $\sigma = 1$. The Adam reconstructions used $\alpha = 0.4$, $\beta_1 = 0.9$, and $\beta_2 = 0.999$. For each reconstruction, 1000 epochs of Adam were run, lowering α by a factor of 0.1 whenever the diffraction loss L (Eq. (4.2)) failed to improve after 10 iterations. Reconstructions were ended early if the loss fell below 1×10^{-9} or α fell below 1×10^{-4}. The reported object error in each case is calculated from a region of the reconstructed object which falls entirely within the illuminated region. All reconstructions were run with a full detector, i.e. $k_m = 0$ from Figure 4-2, except those designed to uncover the impact of a finite missing region. In addition, unless otherwise noted, the beamstop diameter in Fourier space k_{bs} is set to half the probe extent k_p, a typical choice which leads to good separation of diffraction orders in most common experimental geometries.

We first determined the maximum achievable resolution ratio $R = \frac{k_o}{k_p}$ in a noise-free experiment. To do this, we simulated diffraction from random objects with Gaussian-distributed real and imaginary parts, band-limited to different resolution ratios R. BLR illumination with $k_p = \{100, 150, 200\}$ pixels was used in the geometry described in Figure 4-2 and 1000 RPI reconstructions were attempted for each diffraction pattern, using a target resolution ratio $R_{rec} = R$. Information on typical reconstruction results and the full set of parameters are found in Appendix A.1. We monitored the success of reconstructions using the normalized Root Mean Squared (RMS) error ϵ of the reconstructed images R_{kl} as compared to the ground truth O'_{kl}:

$$\epsilon = \frac{1}{\sqrt{\sum |O'|^2_{kl}}} \sqrt{\sum |O'_{kl} - \gamma R_{kl}|^2}, \tag{4.3}$$

$$\gamma = \frac{\sum O'_{kl} R^{\dagger}_{kl}}{\sum |R_{kl}|^2}, \tag{4.4}$$

where $R^{\dagger}_{kl}$ represents the complex conjugate of R_{kl}. We find that reconstructions typically converge to machine precision or entirely fail within the allotted number of iterations, as seen in Figure A-1. We therefore classify a reconstruction as successful when $\epsilon < 0.1\%$ and unsuccessful otherwise. As shown in Figure 4-3a, when $R < 0.4$, the reconstructions are virtually guaranteed to succeed, whereas when $R > 0.6$, our algorithm becomes almost completely ineffectual. This should be compared to the theoretical limit of $R \approx 0.94$.

We next considered how tight the band-limiting constraint needs to be, compared to the true band-limited frequency of the underlying object. This is a critical question because it addresses the extent to which reconstructions can succeed on data from band-limited objects, even if a loose band-limiting constraint (or no constraint) is applied. This is relevant because of the close connections between the data generated in RPI experiments and data arising from experiments, such as CMI, that also make use of high-frequency randomized phase structures. Viewed from the lens of RPI, CMI reconstructions of objects with band-limited spectra are analogous to RPI reconstructions where no band-limiting constraint has been applied. Despite that, the reconstruction may still benefit - particularly because the band-limited nature of the object is hinted at by the diffraction pattern itself. In this way, information about the object's support in Fourier space could potentially be available to the reconstruction algorithm implicitly, through the intensity constraint at the detector plane.

To separate the value of the explicitly imposed band-limiting constraint from the effect of the band-limiting frequency of the underlying object, we simulated diffraction from a set of ground truth objects which were band-limited to $R = \{0.3, 0.35, 0.4\}$. We next ran reconstructions with the reconstructed object band-limited to values of R_{rec} ranging from the true R of the original object up to $R_{rec} \approx 1.25$. We then considered the success or failure of each reconstruction using the same criteria as above. The results in Figure 4-3b show that, especially for objects with a low R, the band-limiting constraint can be significantly relaxed without affecting the reliability of the reconstructions. However, especially for objects with higher values of R, relaxing the constraint too far leads to a quickly increasing probability of failure.

This experiment supports the view that from a practical standpoint, the band-limiting constraint has a clear impact on the reliability of actual reconstructions. It also suggests that once the maximum stable resolution ratio is determined for a particular experimental system, it is safe to run all reconstructions on that system at the maximum resolution. Finally, it indicates that under favorable enough circumstances it is possible to perform RPI reconstructions using only the band-limiting constraint naturally imposed by the discretization of the detector conjugate coordinate space. This provides a strong link with CMI and provides an alternate perspective on the mechanism underlying CMI reconstructions.

We next considered the robustness of RPI to objects which contain some spectral weight beyond the band-limiting constraint imposed by the reconstruction algorithm. To do this, we divided reciprocal space into one zone within the band-limiting frequency $k_o = Rk_p$ and a second zone extending to k_p (Fig 4-2). We then simulated diffraction from randomly generated objects normalized such that the power in the outer zone was a specified fraction of that within the inner one. Reconstructions were performed with $R_{rec} = R$, and the ground truth was taken as the low-resolution representation of the object consisting of only frequency components within the band-limiting constraint. The results in Figure 4-3c demonstrate that even if the band-limiting constraint is not perfectly satisfied, for many applications the resulting error is tolerable provided the object's dominant length scale is above the pixel size of the low resolution object.

We then studied how efficiently RPI uses the information in shot noise limited experiments. We compared the RMS error from a series of reconstructions against the number of simulated photons per illuminated pixel in the low-resolution object. We find that information starts to be extracted at signal levels close to 10 photons per pixel, with the normalized RMS error ϵ falling below 1% by the time signal levels reach $20,000$ photons per pixel (Fig 4-3d).

We subsequently examined the impact of missing data - for example, due to damaged regions or missing regions in a segmented detector. We modeled this situation by masking off a cross-shaped region of pixels of width k_m (Fig 4-2) representing a typical detector dead zone. We then performed reconstructions at several resolution ratios R while varying k_m. The results confirm that the major effect of missing data is simply to lower the maximum stable resolution ratio R (Fig 4-3e). This behavior is similar to that of CMI as described in [85], but contrasts with traditional CDI and FCDI, where missing data lead to poorly constrained regions of Fourier space and real space, respectively, as discussed in [100].

Finally, we investigated the effect of poor alignment between the object plane and the plane at which $P(x, y)$ is known. We find that, when no correction is applied to account for the probe's propagation, reconstructions succeed within roughly 2 depths of focus ($\mathrm{DOF} = \frac{\lambda}{2\mathrm{NA}^2}$) (Fig 4-3f). Furthermore, within this region the observable diffraction loss L is correlated with the true RMS error ϵ of the underlying image. This means that calibration of the probe defocus state is possible from each individual diffraction pattern, removing the need for precise shot-to-shot alignment.

4.5.2 Optical Experiment

To empirically validate RPI, we performed an experiment with a 532 nm laser (Thorlabs CPS532) in a table-top setup. The experimental apparatus, shown in Figure 4-4a, is a rough scale model of a typical scanning transmission X-ray microscope. The initial coherent wavefront was prepared using a beam expander with a 5 μm spatial filter. This light then illuminates a 2 cm diameter RZP (Fig 4-4b) prepared via photolithography of a chrome-on-glass photomask. The RZP focused the light to a 4 mm diameter spot at a focal distance of 33 cm, using a design generated via the

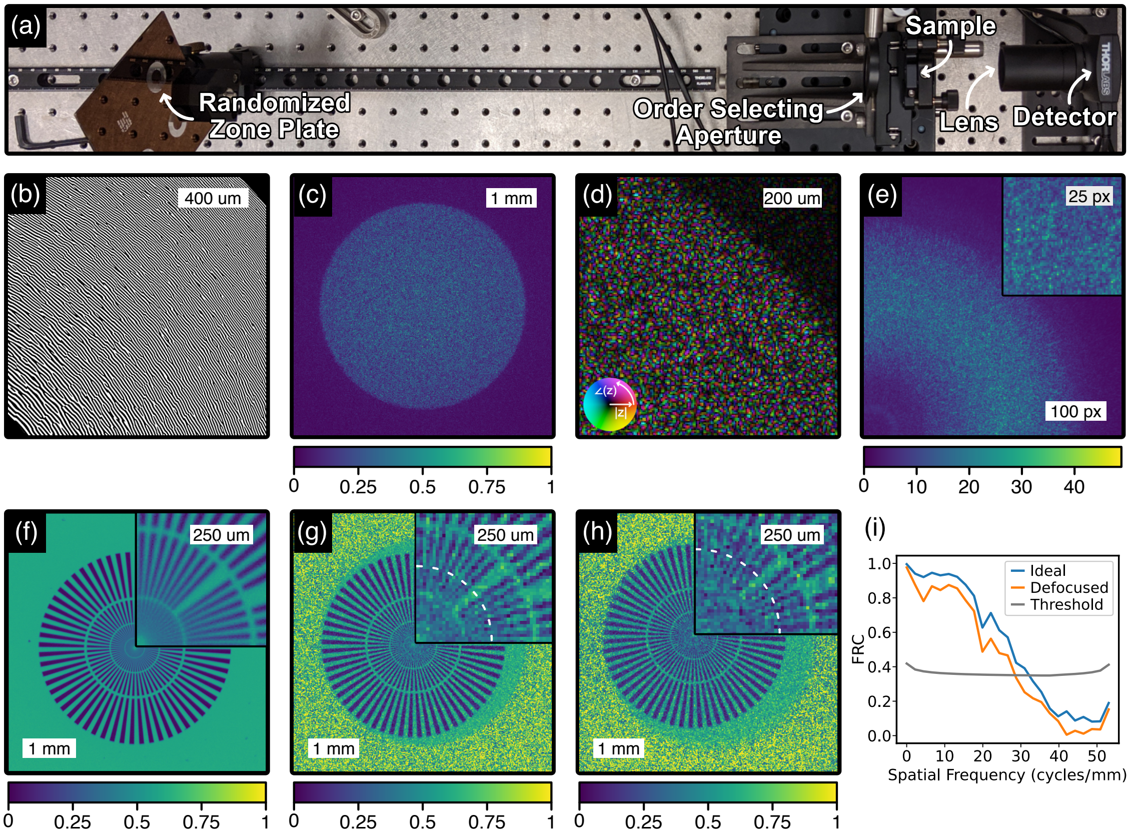

Figure 4-4: **Optical Demonstration**. (a), A top-down view of our optical test bench. (b), A section of the design pattern for the diffractive optic used in this experiment. (c), The amplitude of the incident probe, retrieved via ptychography. (d), A detailed view of the amplitude and phase of the recovered probe. (e), A segment of the square root of the measured diffraction pattern on the detector, in $\sqrt{\text{ADU}}$. Inset shows a further magnified region. (f), The amplitude of the test object, retrieved via ptychography. (g), The same test object's amplitude, retrieved from a single diffraction pattern via RPI. The dotted line shows the radius at which the pitch of the Siemens star matches the resolution calculated via FRC. (h), The amplitude of an RPI reconstruction from the same object after it was intentionally moved from the probe focus. (i), The Fourier ring correlation calculated between the ptychography and RPI reconstructions. All reconstructed amplitudes are reported in arbitrary units.

"ZP_0" method described in [75] and discussed in Appendix A.2. Higher diffractive orders were removed through the combination of a 1 cm beamstop integrated into the RZP and an iris placed approximately 2 cm upstream from the sample acting as an OSA. The diameter of the iris was set to the smallest diameter at which no change in the freely-propagated first order diffraction pattern could be detected. A Siemens star test target was scanned through this focal spot to provide ptychography and RPI data.

The diffraction pattern was imaged on a monochrome camera (Thorlabs DCC1545M) with 5 μm pixels in a $2f$ geometry using an achromatic doublet (Thorlabs AC254-150-A) with a 50 mm focal length. The freely propagated probe fills roughly $\frac{2}{3}$ of the detector in this geometry. For the calibration ptychography scan, we collected a 20x20 grid of diffraction patterns in 200 μm steps (Figs 4-4c-4-4f). Ptychography was performed on this data using Automatic Differentiation ptychography with a background correction, and the retrieved background was input into the RPI reconstructions.

We then performed an RPI reconstruction using a single diffraction pattern that was withheld from the calibration ptychography grid. The RPI reconstruction (Fig 4-4g) was performed using a 400x400 pixel object, corresponding to $R_{rec} \approx 0.6$ and a pixel size of 12.8 μm. The parameters for the reconstruction algorithm were chosen to match those used in our numerical experiments, with the best result from a pool of 20 random initializations reported in each case. The final single-frame full-pitch resolution of 31 μm was calculated via a Fourier Ring Correlation (FRC) (Fig 4-4i) between the RPI and ptychographic reconstructions at a threshold signal to noise ratio of 1 [101], as discussed in Appendix A.3. This shows that studies of objects with resolution ratios of $R > 0.5$ are possible in real experiments, with actual resolution nearly reaching the limit imposed by the band-limiting constraint.

Finally, to demonstrate the robustness of RPI we waited 4 hours, power-cycled the illuminating laser, and defocused the probe by several millimeters before collecting an additional diffraction pattern. We then performed an ensemble of RPI reconstructions with computationally defocused probes and choose the result with minimum diffraction error. The retrieved defocus corresponded to a shift of the sample by 3.4

mm along the propagation direction, roughly 12 times the depth of focus of the zone plate. This reconstruction (Fig 4-4h) converged successfully but was found to have a slightly degraded full-pitch resolution of 35 μm.

These experiments show that RPI is straightforward to implement in a typical ptychography geometry. In addition, we show that the illumination function can be chosen to fill a high proportion of the natural detector conjugate space - in this case, almost 80%, without major issues. In other words, although the reconstructed object's resolution is lowered by the need to measure a larger region in Fourier space than that occupied by the band-limited object, there is a corresponding increase in the field of view because the requirement to sample the intensity distribution in Fourier space at the Nyquist rate (typically referred to as "oversampling" in the CDI literature) is relaxed. This leads to a final space-bandwidth product of the reconstructed object which is roughly equivalent to that of a standard CDI experiment using the same detector geometry.

4.5.3 Soft X-ray Experiment

Subsequently, we performed an experiment at the MAXYMUS beamline of BESSY II to demonstrate that RPI is applicable to X-ray microscopy. A 60 μm diameter RZP was designed to focus 707 eV light to a 2.6 μm diameter focal spot over a focal length of 1.4 mm, with a numerical aperture of 0.021. Higher diffraction orders were filtered using a 30 μm beamstop on the RZP and a 15 μm OSA placed less than 385 μm upstream from the sample. Scattered light was collected on a 264×264 CCD detector with 48 μm pixels [102] placed 17 cm downstream from the sample.

We imaged a commercially available Siemens Star test sample (Carl Zeiss AG, Germany) made of 180 nm thick Au with a minimum feature size of 30 nm, as well as a ferromagnetic Fe/Gd multilayer prepared by sputtering 50 or 70 alternating layers of 3.6 Å thick Fe and Gd. All imaging was done at the Fe L_3-edge using circularly polarized illumination from the third undulator harmonic. The exit slits were set to 35×35 μm, and a 200 ms dwell time was used.

Typical raw diffraction patterns (Figs 4-5a, 4-5b) show many of the issues common

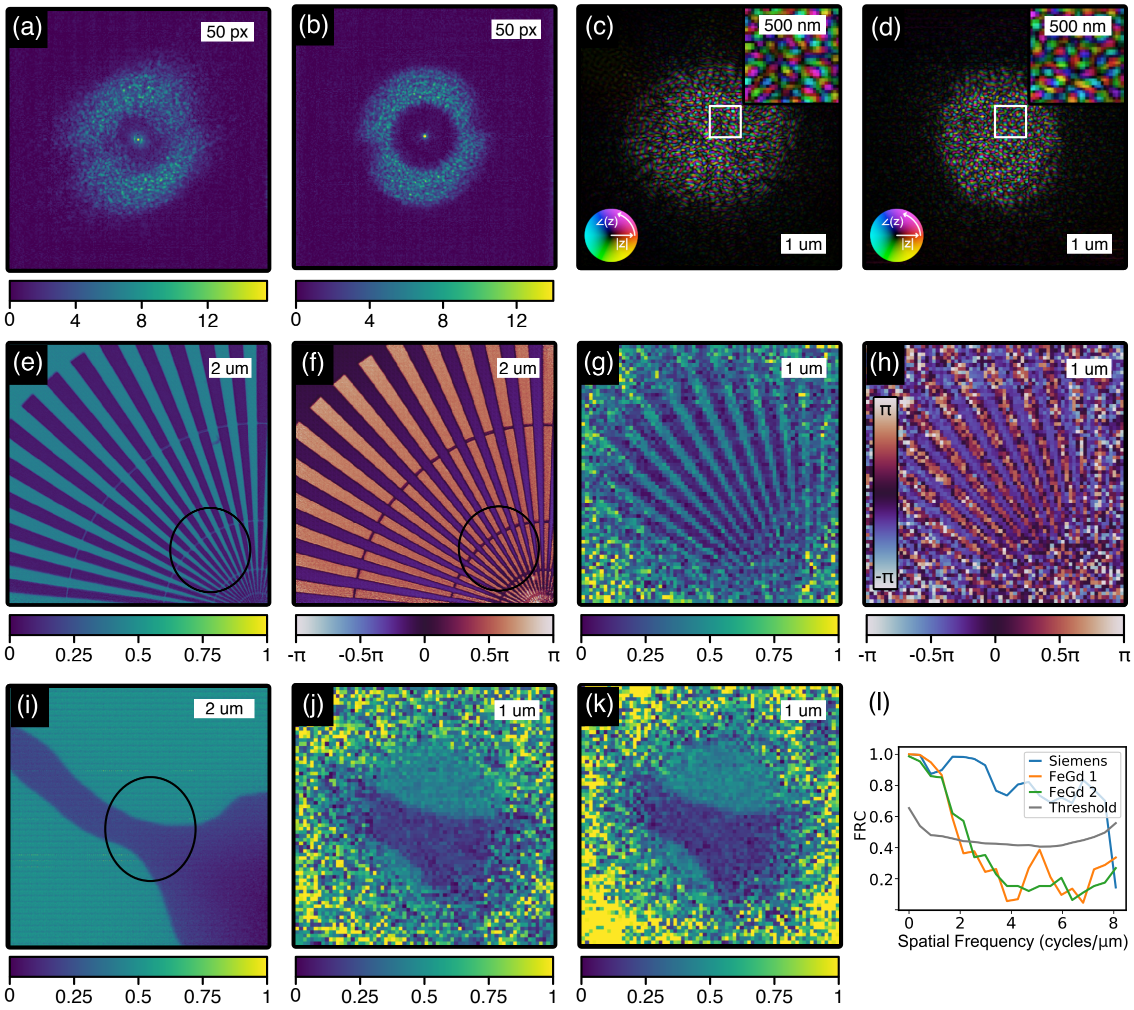

Figure 4-5: **Soft X-ray Demonstration**. (a),(b), The square root of a diffraction pattern from the Siemens Star and FeGd samples, in $\sqrt{\text{ADU}}$. (c),(d), The probe as reconstructed from the Siemens Star and FeGd ptychography. Inset shows the detailed speckle structure. (e),(f), The amplitude and phase of the ptychographic reconstruction of the Siemens Star. (g),(h), The amplitude and phase of an RPI reconstruction of a portion of the Siemens Star. (i) The amplitude of the ptychographic reconstruction of FeGd, showing three ferromagnetic domains. (j), The amplitude of an RPI reconstruction of a region of the Fe/Gd Sample. (k), The amplitude of an RPI reconstruction from the same diffraction pattern, using the probe reconstructed from the Siemens Star. (l), FRCs calculated for the various RPI reconstructions. All amplitudes are reported in arbitrary units.

to coherent X-ray scattering experiments. A strong zeroth-order component from the zone plate is seen, which was masked off in our RPI reconstructions. The data also reveals damaged regions on the optic and astigmatism is visible in the reconstructed focal spots (Figs 4-5c, 4-5d). In addition, the ptychography reconstruction from the Fe/Gd multilayer (Fig. 4-5i) is affected by raster grid pathology.

Despite these shortcomings, single-frame RPI reconstructions were successful on both the Siemens star (Figs 4-5g, 4-5h) and Fe/Gd multilayer (Fig. 4-5j). The reconstructions were obtained on a 70 × 70 pixel object with a pixel pitch of 83 nm, corresponding to $R_{rec} \approx 0.5$. In both cases, the diffraction pattern used for RPI was withheld from the calibration ptychography reconstruction. The resolution was estimated via a FRC between the RPI and ptychography reconstructions (Fig. 4-5l). We find that the Siemens star, which contains strong high-frequency components, is reconstructed to the resolution of the pixel pitch at a signal-to-noise threshold of 1, while the Fe/Gd sample is reconstructed at a full-pitch resolution of 510 nm.

As a final demonstration, we performed a RPI reconstruction of the FeGd sample using a numerically propagated probe reconstructed from the Siemens star ptychography data (Fig. 4-5k). We found that the probe's focal spot had shifted in the propagation direction by roughly 40 μm between samples, roughly 20 times the depth of focus of the zone plate. Ultimately, the retrieved full-pitch resolution of 420 nm was actually negligibly improved.

These experiments again demonstrate the resiliency of RPI. Single-frame reconstructions succeeded despite serious aberrations in the optics and typical data quality issues. We also showed that the calibration ptychography and RPI reconstructions can be performed on different samples, with potential consequences for the ultimate resolution achieved. Therefore, only one X-ray pulse is required to interact with the sample of interest, potentially enabling “diffract before destroy” experiments.

4.6 Discussion

We have demonstrated a method for single-frame quantitative lensless amplitude and phase imaging that is reliable, straightforward to implement, robust to missing data, and applicable to extended samples. As a result, RPI complements and extends on many existing diffractive and traditional imaging methods. We therefore envision a variety of scenarios where RPI can become a valuable addition to the experimentalist's arsenal of techniques.

The most immediate use case we envision for RPI is as an alternative to off-axis holography, particularly in the soft X-ray regime. We note that while the resolution limits for both off-axis holography and RPI are comparable, in off-axis holography the signal rate diminishes as the size of the reference hole shrinks. Because RPI can make much more efficient use of the available flux, it is a preferable option for many experiments which would otherwise be performed using soft X-ray off-axis holography. As a bonus, because RPI doesn't require the deposition of a mask or the drilling of a reference hole, it is possible to study many systems with RPI which could not be prepared as samples appropriate for soft X-ray holography.

RPI may also have value as an alternative to Transmission X-ray Microscopy (TXM), which it has two major advantages over. First, RPI uses every photon that hits the sample, while transmission X-ray microscopy uses an inefficient X-ray optic between the sample and detector. As light sources get brighter and brighter, sample damage thresholds increasingly limit the resolution of X-ray imaging experiments. In such a situation, moving an inefficient optic from downstream to upstream of the sample can be extremely valuable. The second advantage is that, whereas optical defects lead to aberrations in TXM images, all such aberrations are naturally corrected for in RPI during the probe calibration step. This could potentially lead to cheaper optics, larger optics, or even higher resolution optics than could be produced for a TXM experiment.

Similarly, RPI is an attractive alternative to CMI in many situations. Importantly, the experimental setup for RPI obviates the need for downstream optics in the sample's near-field region. This enables experiments in bulkier sample environments and removes a challenging alignment step. In addition, the RZP concentrates incident light onto the sample, allowing for much brighter illumination at the sample than can be achieved in a pinhole-based CMI geometry. The trade-off is the inevitable resolution limit that results from the truly band-limited nature of BLR illumination.

A second valuable application of RPI is as a complementary imaging method at transmission ptychography beamlines. In such a setup, users could switch between a high spatial resolution mode (ptychography) and a high time resolution mode (RPI) at will. This has obvious value - for example, allowing researchers to explore several regions of interest before commencing a dynamic imaging study, or making it possible to quickly study damage mechanisms in a material using RPI before commencing a static ptychography measurement. Although the large beam waist would prevent researchers from collecting fluorescence maps along with ptychography data, in many situations this trade-off would be valuable.

A third application of RPI is as a unique time-resolved quantitative imaging method for studying unpatterned thin films with soft X-rays in the reflection and Bragg geometries. Currently, no single-frame techniques exist for this kind of study. Because RPI is reference-free but much more reliable on highly textured samples, it is particularly well suited to this type of experiment provided that the sample itself is thin enough to ensure that the full angular frequency spectrum of the illumination can be diffracted.

Last, we propose a few avenues for improvement of RPI. First, our algorithm is tailored for flexibility. As a result, there are almost certainly gains in computational efficiency to be realized by porting other phase retrieval algorithms to this context. Second, it is possible that algorithmic improvements can push R closer to the theoretical limit, improving the potential resolution of RPI experiments. Third, modern fabrication methods are capable of producing RZPs with higher numerical apertures than that used in our proof-of-concept experiment, further improving resolution. Finally, we note that the spatial resolution is limited by the numerical aperture of the optics, not the quality of the focal spot. This motivates the design of X-ray optics, such as those using higher diffraction orders, which are capable of generating BLR radiation at higher numerical apertures than would be possible if aberrations were a major concern.

4.7 Conclusion

We have demonstrated a reliable single-frame lensless X-ray imaging method that can be easily commissioned anywhere ptychography is done. As a result, it makes single-frame and time-resolved quantitative X-ray phase contrast imaging simpler, more reliable, and more accessible. RPI is likely to be applicable in fields where holography, TXM, 2D CDI, Fresnel CDI, ptychography, and X-ray photon correlation spectroscopy are used. Because of its flexibility and robustness, it will potentially enable a new generation of time-resolved studies on samples ranging from correlated materials and magnetic devices to soft matter, biological systems, and beyond.

4.8 Postface

Following the publication of the paper, we made significant progress in randomized probe imaging (RPI) reconstruction algorithms, developing an algorithm that routinely reconstructs noiseless data out to $R = 1$ and, occasionally, beyond. This highlights another area of progress. In the above analysis of the theoretical maximum resolution ratio R, we ignored the role of finite support in real space. Further experience has made it evident that properly sampling the diffracted intensity at the Nyquist frequency (instead of half the Nyqyist frequency as assumed above) generates additional measurements. This makes convergence theoretically possible even when $R > 1$. However, the inverse problem rapidly becomes ill-conditioned due to illumination nonuniformity as R passes through 1.

The upshot is that we can now confidently state that $R = 1$ is a viable target for generic implementations of RPI. Further detail regarding the algorithmic improvements and improved understanding of oversampling is given in Chapter 6.

Chapter 5

Single-Shot Randomized Probe Imaging at a Free Electron Laser Source

The following chapter is based on work intended for standalone publication. Flavio Capotondi, Emanuele Pedersoli, Mattero Pancaldi, and Kahraman Keskinbora were all deeply involved in the data collection, and Kahraman Keskinbora fabricated the optics used for this experiment. Riccardo Comin was involved in a supervisory role.

5.1 Abstract

Nanoscale imaging with single free electron laser (FEL) pulses promises to revolutionize our understanding of the dynamics of heterogeneous phases of matter. Short-wavelength radiation can visualize structures on the nanoscale, while atomic resonances give access to electronic degrees of freedom. Unfortunately, existing methods for quantitative single-shot imaging at FELs sources struggle to image large areas with nanoscale resolution, i.e. reach a high space-bandwidth product. We report on an FEL-based implementation of a method, randomized probe imaging (RPI), that addresses these difficulties in a scalable way. We demonstrate repeatable single-shot reconstructions with a space-bandwidth product of 25,000 at a 400 nm full-pitch res-

olution, free of visual artifacts. RPI's well understood limitations promise an easy path to further improvements along both axes, leading to revolutionary applications in areas such as the study of shock physics, magnetic devices, and correlated electronic materials.

5.2 Introduction

Since the first demonstration in 2006 [69], the femtosecond pulses of x-ray and extreme ultraviolet (EUV) light produced by free electron lasers (FELs) have been used to image an astonishing variety of phenomena. Researchers have used FELs to study the morphology of metallic nanoparticles [22, 103, 104], soot [105, 106], stick figure drawings [69, 107], helium nanodroplets [108], battery electrolytes [109], viruses [21, 110], organelles [111–113], and even whole bacterial cells [114–116]. Because the flashes of light are quick enough to partially outrun the resulting damage, it is possible to reach higher signal-to-noise ratios and higher resolution than could be achieved using continuous wave illumination.

Hard x-ray FELs have enabled unprecedented visibility into the mechanical properties of matter undergoing rapid changes, by capturing snapshots of shock-wave propagation [117–120], cavitation bubbles [121], liquid microjets [122], and exploding liquid-filled capillaries [123]. The well-established interaction of hard x-rays with matter means the images can act as a quantitative measurement of local sample properties [121].

At lower energies, the resonant coupling of XUV and soft x-ray light to the electronic degrees of freedom has made it possible to study magnetism in thin films [24, 25], including picosecond-scale dynamics [124]. At these wavelengths, single-shot imaging [125] using FELs uniquely enables studies of non-repeatable dynamics.

This research frontier relies not only on the technical advancements behind FELs, but also on computational lensless imaging methods [65] like coherent diffraction imaging (CDI) [19, 20, 69]; Fourier transform holography (FTH) [23, 126] and its derivatives [127–129]; in-line holography [130, 131] and Fresnel/keyhole CDI [83, 84];

as well as single-shot ptychography [90, 132]. These methods use interference effects to computationally reconstruct quantitative images of the absorption and phase delay experienced by light as it traverses the sample.

Despite this progress a web of frustrating limitations hold back the computational imaging methods these studies rely on. Of these methods, x-ray CDI can achieve the highest resolution, better than 10 nm under ideal circumstances [22]. However, CDI becomes unreliable as the space-bandwidth product of the imaged region grows and the sample structure becomes more complex [110]. Reconstructions of large samples will consequently often fail to converge, or worse, produce plausible-looking but incorrect images.

At slightly lower resolutions in the tens of nanometers, both CDI and FTH can be used, but due to the stability issues with CDI, methods based on FTH dominate. Reference [110] exemplifies the lengths to which researchers will go to to introduce a reference and avoid CDI's stability issues. In this experiment, an entire second beam of reference particles was aligned to intersect with a main beam of sample particles, with data used only when a particle from both beams fell within the field of view.

Unfortunately, both FTH and CDI only work on isolated samples with typical max dimensions of a few microns. To study extended samples such as bulk crystals and thin films, a mask has to be lithographically patterned. This limits the field of view to a predefined region of interest. In some cases, such preparation is not even possible: consider the case of an exploding glass capillary [123]. Consequently, hard x-ray experiments on larger-scale samples have turned to a third family of methods that use diverging illumination to set up an effective Fresnel-regime propagation from the sample to the detector. Depending on the details of the reconstruction algorithm and the author's predilection, methods in this family have been called in-line holography [130, 131], near-field holography [121], propagation-based phase contrast imaging [122], Fresnel CDI [84], and keyhole CDI [83].

These methods are easy to use and do not require sample masks. However, in every FEL-based experiment of this kind performed to date, the analysis was stabilized by assuming that the sample's phase and optical density were linearly related [121–123, 133]. This is not always a valid assumption, especially for samples which contain multiple materials, and for images taken at longer wavelengths or near resonances. In many cases the raw detector images were simply used without phase retrieval [118–120, 134, 135] due to the difficulty of achieving a stable reconstruction.

Even when they are valid, it is well known that these methods are prone to artifacts, which often present as lines running parallel to sharp edge features [122, 133]. The supplement to [122] highlights the additional susceptibility of this approach to fixed pattern noise arising from a poorly characterized illumination. This has precluded quantitative analysis of 2D results, and researchers have relied on dimension reduction [121] or comparison to simulated data [136] when quantitative accuracy is needed.

Consequently, there is an acute need for a fresh, versatile approach to quantitative single-shot FEL imaging that is reliably free of artifacts, works on samples that are not isolated, and works even when the region of interest is large and there is no simple a priori relationship between the amplitude and phase channels. In this work, we discuss how randomized probe imaging (RPI), a recently developed method for quantitative single-frame imaging [64], can fill this niche. We implement RPI at an FEL source for the first time, showcasing its potential to simplify quantitative phase imaging and make it accessible to non-expert researchers.

5.3 Methods

At the DiProI beamline of the FERMI FEL [137], we illuminated an 800 µm diameter gold on silicon nitride randomized zone plate (RZP) with an outer zone width of 200 nm (Figure 5-1) using 60 eV light. The optic design strategy is the same as described in [64], and the optic was fabricated using focused ion beam lithography.

This formed a circular illumination spot on the sample with a diameter of 40 µm. Using this focal spot, we took a 300-exposure ptychography scan on a 64-spoke Siemens star test sample. The diffraction patterns were captured on a 2048 × 2048 CCD detector (Princeton Instruments MTE2048B) with 13.5 µm pixels placed 135 mm

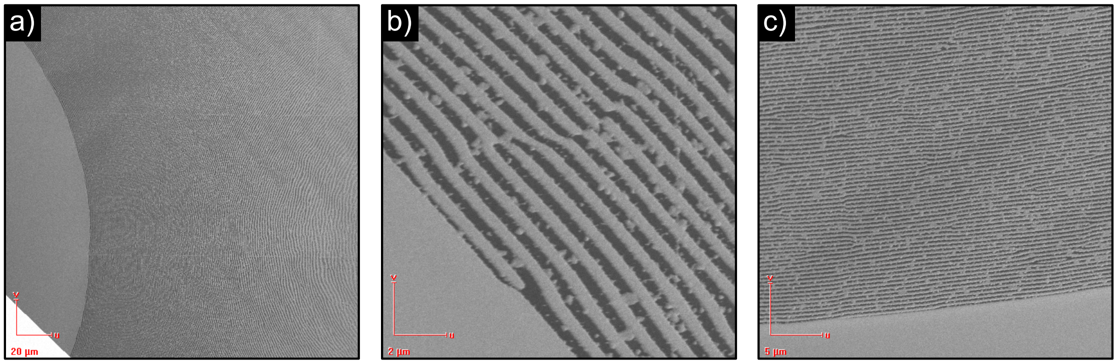

Figure 5-1: **Optics**. (a) A zoomed out view, from a scanning electron microscope, of an inner section of the zone plate. This images shows a Moire pattern generated by the randomized zones and the scan pattern of the microscope. (b) A closeup view of generic zones near the inner edge. (c) A closeup view of generic zones near the outer edge.

downstream from the sample. The pixels were binned into 2 × 2 regions. Each exposure captured approximately $6.9e6$ photons total. Accounting for absorption in the sample and the detector quantum efficiency, this corresponds to an estimated illumination fluence of $50\,\mu\mathrm{J\,cm^{-1}}$.

The ptychography data was reconstructed using an automatic differentiation-based superresolution reconstruction algorithm similar to the one described in [58], with three probe modes. Following characterization of the object and probe, we collected 128 further single-shot diffraction patterns from the same test target for analysis with RPI.

The probe and detector background calibration extracted from the ptychography reconstruction were used to perform RPI on all exposures from the subsequent single-shot dataset. The reconstructions were performed on 650 × 650 pixel objects with a pixel size of 160 µm, a nominal resolution ratio $R_{rec} = 1.25$. An automatic differentiation based algorithm with two object modes was used for these RPI reconstructions.

5.4 Results

We first characterized the reliability of the ptychography reconstructions of the structural test target, shown in Figure 5-2. As expected, the probe covers a roughly 40 µm

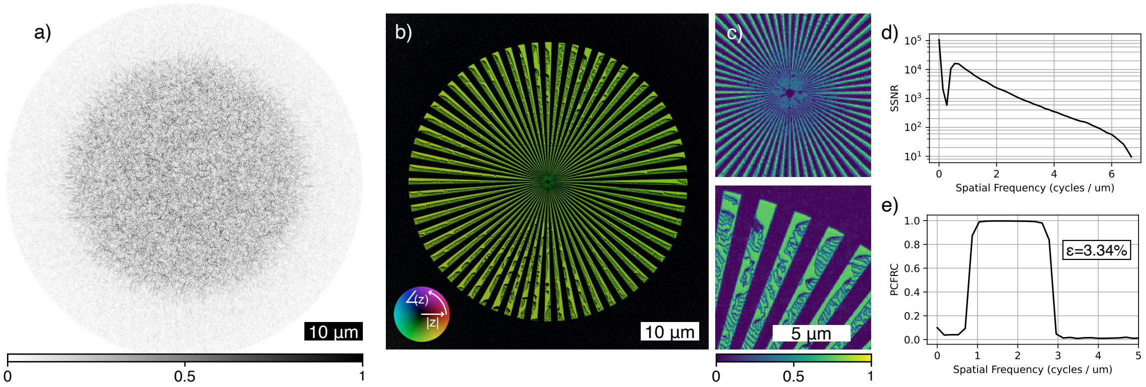

Figure 5-2: **Ptychography Calibration**. (a) The amplitude of the top reconstructed probe mode. (b) The amplitude and phase of the reconstructed Siemens star. (c) Detailed views of two regions of the reconstruction. (d) The spectral signal-to-noise ratio (SSNR) of the reconstructed object estimated via FRC. (e) The PCFRC of the reconstructed probe.

diameter circular focal spot. The Fourier ring correlation (FRC) calculated between two reconstructions, each using 50% of the exposures, indicated that the object was reconstructed at the pixel-by-pixel level for a pixel pitch of of 73 nm. This corresponds to a full-pitch resolution of 146 nm. The partially coherent Fourier ring correlation (PCFRC) [138] of the probe indicated reliable reconstruction within the relevant frequency band (Figure 5-2d). The normalized amplitude-minimized partially coherent mean squared error (PCMSE) [138] between the two reconstructed probes was 3.34%, suggesting that the accuracy of our single-frame RPI reconstructions will likely be limited to a few percent, as our probe calibration is only reliable to that level.

The single-shot RPI reconstructions are shown in Figure 5-3. We compared each single-shot RPI reconstruction with downsampled ptychography data using the FRC, a.k.a. the Fourier cross resolution (FCR), because we are comparing a noisy signal to an ostensible ground truth [139]. We then converted the FCR curves to SSNR [139], as shown in Figure 5-3f. Comparison with the half-bit threshold of SSNR $= 0.4142$ [101], commonly used to assess the resolution of lensless imaging data [77, 140–142], showed than only two of 128 reconstructions failed to surpass a resolution of 400 nm, $R = 1$. We therefore claim a robust resolution of 400 nm, which corresponds to our target

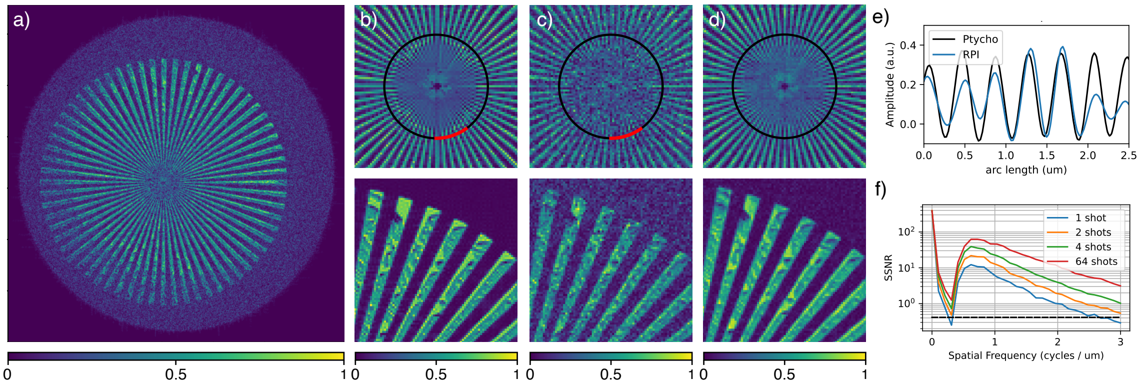

Figure 5-3: **RPI Reconstruction**. (a) A typical single-shot reconstruction. (b) A closeup of the calibration ptychography reconstruction, downsampled to match the resolution of RPI. (c) The same closeup regions of the single-shot RPI reconstruction, for comparison. (d) The mean of 64 subsequent single-shot reconstructions. (e) A circular cut from the single-shot and downsampled ptychography reconstruction, extracted at the radius marked on (b-d). At this radius, the spokes have a periodicity of 200 nm, or $R = 1$. (f) The SSNR of the objects, estimated via the FRC.

RPI resolution ratio $R = 1$. A sinc-interpolated linecut along an arc intersecting the spokes at a 400 nm pitch (Figure 5-3e) confirms that structures at this resolution are resolved.

The space-bandwidth product is the number of individually resolved resolution elements within our reconstruction, defined as the product of the accessible area in real space and reciprocal space:

$$\mathrm{SBP} = A_{\mathrm{real}} A_{\mathrm{Fourier}} = \pi^2 \frac{r^2}{\mathrm{res}^2}. \tag{5.1}$$

The uniformly illuminated 40 µm diameter focal spot and maximum resolved spatial frequency of 2.5 cycles per µm result in a space-bandwidth product of 2.5×10^4.

Aligning the stack of 200 RPI images to account for the shot-to-shot probe instability of roughly ±0.15 µm, we then studied how increased dose improved the SSNR of the resulting reconstructions. The improvement in SSNR grows linearly (as expected for uncorrelated noise) as we combine the first 4 images, but then the growth rapidly slows, although the signal strength continues to grow as we average more

RPI reconstructions. This likely indicates that other effects, such as the presence of unaccounted-for high frequency structures in the sample or issues with imperfect image alignment, become relevant.

Crucially, we find that the synthesized high-dose images are visually free of imaging artifacts. This contrasts with in-line holography, where even under bright illumination conditions the resulting reconstructions reveal artifacts [120, 122]. Furthermore, the visual quality of the single-shot images highlights that RPI handles noise gracefully, producing noisy images rather than introducing artifacts on larger length scales that could be mistaken for real features.

5.5 Discussion

The single-shot images we produced are on the Pareto frontier of resolution and space-bandwidth product for quantitative phase imaging, with higher resolution than any previously reported single-shot in-line holography reconstructions and a higher space-bandwidth product than any work produced with CDI. Moreover, only one publication has ever reported a quantitative single-shot FEL reconstruction of a decoupled phase and amplitude object with a higher space-bandwidth product than we achieve here [143].

There is also a clear path to improving the resolution of RPI at FELs, down to the minimum outer-zone widths which x-ray zone plate optics can be manufactured at. Synchrotron-based RPI has already achieved a resolution of 160 nm, and optics with 25 nm zones can be bought commercially.

Equally important is RPI's decoupling of the sample and experiment, allowing for a clean separation of concerns between the user and the beamline scientists and not forcing the user to become an expert at computational x-ray imaging. Furthermore, because RPI doesn't rely on sample sparsity, it is feasible to study complicated samples that are densely packed with relevant features. Finally, because the burden of optic and algorithm design can be shared across most experiments and samples, it is possible to maintain a setup that just "works like a camera" from the perspective of a user.

In conclusion, we have demonstrated that RPI is a practical choice for imaging at FEL light sources, and compares favorably to a wide variety of methods across a broad swath of parameter space. We expect that its simplicity and compatibility with unmasked samples will make it an attractive option for many experiments that otherwise would have used FTH, CDI, or in-line holography.

Chapter 6

Further development of RPI

Following the publication of the paper describing randomized probe imaging (RPI) [64], I pursued several projects to improve RPI. Primarily I worked to develop better reconstruction algorithms and to improve the design process for randomized zone plates (RZPs). This activity centered around beamline 7.0.1.2 of the advanced light source (ALS), where I worked under two fellowships from April 2022 to April 2023. In addition to developing improvements for RPI in general, I had a secondary goal to create and improve a specific implementation of RPI at beamline 7.0.1.2 that can be used by scientists who are not imaging specialists.

In this chapter, I report on the main products of this research. First, I discuss an improvement to the reconstruction algorithm for RPI, which enabled it to work at much higher resolution ratios R. Then, I discuss an alternate perspective on the analysis of the maximum possible resolution ratio R, expanding on the discussion in Appendix A.5. The new analysis accounts for the possibility of properly Nyquist sampling the diffraction pattern. Next, I discuss an overhaul of the optic design process that produces zone plates with much cleaner focal spots. Finally, I discuss progress toward implementing RPI at the ALS. I show data from structural and magnetic samples and explore the design of an entirely new algorithm for live RPI reconstructions.

6.1 Improving the algorithms

In reference [64], we identified one question for further research whose importance dominated everything else: *Can we push the practically achievable resolution ratio* R *to its limiting value of* ≈ 1*?*

Surprisingly, a very simple change to the algorithm achieved this goal, and beyond. This is the introduction of *multiple object modes* into the forward model. These additional modes are not needed to describe the final result, but introducing them opens up new intermediate states through which the reconstruction can proceed. This appears to help the reconstruction avoid stagnation, improving the maximum achievable resolution even beyond $R = 1$.

Formally, multiple modes are introduced by generalizing the basic interaction model from (4.1) to:

$$\begin{aligned} I_{ij} &= \sum_{\alpha,\beta=1}^{L_p,L_o} |\tilde{E}_{ij}^{\alpha,\beta}|^2 \\ \tilde{E}_{ij}^{\alpha,\beta} &= \mathcal{F}\{P_{ij}^{\alpha}\mathcal{F}^{-1}\{\text{pad}(\mathcal{F}\{O_{kl}^{\beta}\})\}_{ij}\}. \end{aligned} \tag{6.1}$$

Where α indexes the modes of the probe (from 1 to L_p) and β indexes the modes of the object (from 1 to L_o).

This model is sketched in Figure 6-1 for the simple case of one probe mode. Typical real-world objects consist of only a single mode. We enforce this constraint by adding a small penalty to power contained within modes 2 to L_o that slightly biases the reconstruction toward the top mode:

$$\text{Reg} = \sum_{\beta=2}^{m} \sum_{i,j} \gamma |O_{ij}^{\beta}|^2. \tag{6.2}$$

Here, γ is a parameter controlling the strength of the regularization. This parameter is set low enough to allow the reconstruction to explore multiple-mode configurations, but high enough to push most of the power in the final reconstruction into

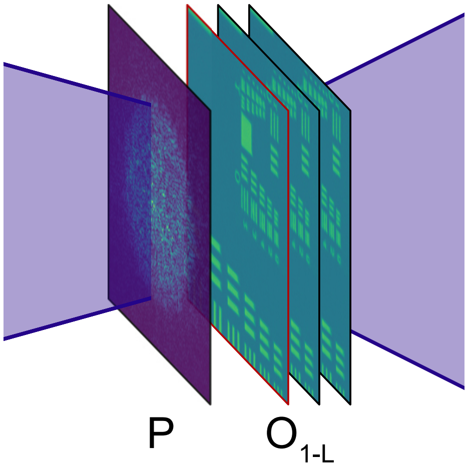

Figure 6-1: **Multi Mode Model**. The RPI interaction is simulated between the probe and multiple object functions $\mathbf{O}_{1-L}$, and the resulting diffraction patterns are added up at the detector plane. Typically, the final result can be captured by a single mode, so the reconstruction is slightly biased toward first mode (outlined in red), which is returned as the final result.

the first object mode.

Figure 6-2a shows the consequences of this change, using a simulation setup comparable to the simulation of Figure 4-3a. Pure random noise objects were created at various resolution ratios R, with each pixel drawn from the complex normal distribution. k_p was set to 50 pixels and reconstructions were attempted using $\gamma = 0.001$ and 500 iterations of L-BFGS. We used a probe with a diameter in real space of half the detector conjugate field of view, as discussed in Section 6.2. The failure probability was determined as a function of R, using reconstructions with R_{rec} matched to the true resolution ratio of the object. The single-mode reconstructions rapidly begin to fail at a resolution ratio around 0.8 in this case, whereas the addition of even a single extra mode leads to successful reconstructions with a probability over 50% even at $R = 1.2$.

However, there is no free lunch, because the non-uniformity in the speckle-filled illumination still causes issues at high resolution. Figure 6-2b shows the consequence at high R: an increasing sensitivity to noise. This graph shows the result of 500 reconstructions per R value, with simulated Poisson noise at $10,000$ photons per

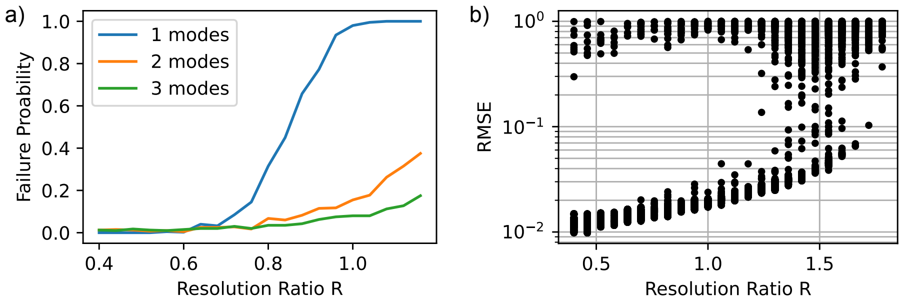

Figure 6-2: **Numerical Tests**. (a) The dependence of failure probability, defined as the chance of a reconstruction resulting in $> 1\%$ RMSE, on resolution ratio, from 500 simulated reconstructions. (b) The RMSE of reconstructions with $10,000$ photons per pixel in the reconstructed object, at various resolution ratios. Each resolution ratio shows 500 individual reconstruction attempts.

reconstructed pixel. As R increases, the root mean squared error (RMSE) of the successful reconstructions increases. This suggests that the condition number of the inverse problem is growing rapidly as R increases.

By $R = 1.75$, the per-pixel error in the final reconstructions is already $10\times$ the error resulting from data with the same per-output-pixel signal level using $R = 0.5$. This means that $100\times$ as many photons would be needed to reach the same quality level on the same object, and the results will become extremely sensitive to even small systematic errors. Consequently, $R = 1$ remains a practical choice for experiment design, and there is not a major advantage to developing algorithms capable of operation at significantly higher R.

The obvious questions at this point is, why does this work? First, I will note that the introduction of multiple modes for object reconstructions is not new, and in fact was discussed in detail by Pierre Thibault in his original paper on multi-mode ptychography [43]. What has not been fully appreciated by the ptychography community is the extent to which multi-mode reconstructions can be more stable, even in situations where the true result is fully coherent.

This is because the introduction of multiple object modes has a close connection to a convex relaxation of the generic phase retrieval problem known as PhaseLift [93]. The standard formulation of phase retrieval considers minimization of the loss function $\mathcal{L}$ under the constraint of perfect coherence. PhaseLift reformulates this as the minimization of the rank of an incoherent object's density matrix while enforcing the measured Fourier transform magnitude. In other words, PhaseLift hunts for the most coherent object that is still fully consistent with the measured data, rather than hunting for the fully coherent object that is most consistent with the measured data.

The reformulated optimization problem not only produces the same result as the original optimization problem in an idealized setting, but it is also convex [93] and therefore comes with convergence guarantees. Roughly speaking, the lifted problem is convex because the space of all object density matrices is much higher-dimensional than the embedded manifold of fully coherent objects consistent with the measured data. This embedding has the topology of a torus [144] and, in the space of density matrices, forms a non-convex set. Although the final solution to the optimization problem lies within this manifold, the relaxed problem achieves convexity by allowing the iterates to move through the larger convex space.

We could then ask, how incoherent do we really need to let the object be if we want to open up these new paths? In the case of the generic phase retrieval problem, the answer turns out to be, not very incoherent! This issue was explored by [97], where they found that even for large problems, a small (order unity) number of modes suffices to enable robust reconstructions. This principle appears to extend to RPI.

The outcome of this work is a simple and reliable algorithm for RPI reconstructions, which we have used nearly universally following the original publication. This algorithm is also based on automatic differentiation, and therefore can be easily modified to account for common corrections, such as a reflection-mode geometry or pixel oversampling on the detector. Its superior stability means that the limitation on resolution is now nearly always determined by the point at which the noise in the final reconstruction becomes intolerable. The only situation in which we have used an alternate approach is for live reconstructions, described in section 6.6, where it is preferable to use an explicit algorithm designed for speed at all costs is.

6.2 Improved understanding of stability

Notably, the results in Figure 6-2 show reliable reconstructions occurring even beyond the $R = 1$ threshold defined in our original paper on RPI [64] as the theoretical limit. What gives?

In the original paper, we considered a setup where the diffracted intensity pattern of our object was sampled at half its Nyquist frequency. This is because we set the object up without any finite support constraint, assuming that it filled the detector conjugate coordinate space exactly. We used this setting to demonstrate clearly that RPI does not rely on a finite support constraint, highlighting the difference between RPI and coherent diffraction imaging (CDI). In doing so, though, we ignored an important point: the illumination actually does provide a reasonably clean support constraint, even if we don't use it in the traditional way.

Although the exit wave does not have a pixel-level sharp cutoff (a requirement for good CDI data), it is still true that the finite size of the illuminated region creates the possibility of properly sampling the diffracted intensity at twice the Nyquist rate of the wavefield, as discussed in Section 2.1.

Consequently, it must be recognized that *both* the expansion of the diffraction pattern by convolution with the probe [64] and the "oversampling" in real space achievable by properly sampling the diffracted intensity can contribute to the all-important ratio of independent measurements to free parameters.

Here, I consider how the ratio of independent measurements to free parameters varies both as a function of the resolution ratio R in reciprocal space [64] and an "overfilling ratio" in real-space, which I define as the ratio of the detector conjugate field of view to the diameter of a circular RPI focal spot $D = \frac{d_{\text{fov}}}{d_{\text{spot}}}$.

I quantify this measurement ratio as the ratio of the space-bandwidth product [145, 146] of the measured intensity data to the space-bandwidth product of the reconstructed complex-valued objects. The space-bandwidth product is a measure of the number of independent parameters that can be encoded in a signal [145, 146], and is calculated by multiplying the area the signal occupies in real and Fourier space.

For this calculation, we consider the canonical case of a circular diffractive optic producing a circular focal spot, producing data that is then analyzed using a square detector with square pixels. For simplicity, we take the low resolution object to be band-limited to a circular region in Fourier space, rather than a square region. This deviates slightly from the typical practice of band-limiting to a square region of Fourier space.

In this scenario, the space-bandwidth product of the signal to be reconstructed, the object function $\mathbf{O}$, is:

$$\mathrm{SPB}_{\mathrm{obj}} = (\frac{\pi}{4} d_{\mathrm{spot}}^2)(R^2 \pi \vec{k}_{\mathrm{max}}^2). \tag{6.3}$$

In this equation, d is the focal spot's diameter in real space and R is the resolution ratio. The space-bandwidth product of the measured real-valued intensity pattern is:

$$\mathrm{SPB}_{\mathrm{Data}} = \mathrm{overlap}(d_{\mathrm{spot}}, d_{\mathrm{fov}})((R+1)^2 \pi \vec{k}_{\mathrm{max}}^2). \tag{6.4}$$

$\mathrm{overlap}(d_{\mathrm{spot}}, d_{\mathrm{fov}})$ is defined as the area within the intersection of a circle of diameter $2d_{\mathrm{spot}}$ and a square field of view of side length d_{fov}, as below:

$$\mathrm{overlap}(d_{\mathrm{spot}}, d_{\mathrm{fov}}) = \begin{cases} d_{\mathrm{fov}}^2 & d_{\mathrm{fov}} \le \sqrt{2} d_{\mathrm{spot}} \\ \pi d_{\mathrm{spot}}^2 - 2(\theta - \sin(\theta)) d_{\mathrm{spot}}^2 & \sqrt{2} d_{\mathrm{spot}} < d_{\mathrm{fov}} \le 2 d_{\mathrm{spot}} \\ \pi d_{\mathrm{spot}}^2 & 2 d_{\mathrm{spot}} < d_{\mathrm{fov}} \end{cases} \tag{6.5}$$

$$\theta = 2 \cos^{-1} \left(\frac{d_{\mathrm{fov}}}{2 d_{\mathrm{spot}}} \right) \tag{6.6}$$

The circle with diameter $2d_{\mathrm{spot}}$ is the support of the exit wave's autocorrelation, which has twice the diameter of the focal spot.

The ratio of these two space-bandwidth products is shown in Figure 6-3, as a function of R and D. The $4n - 4$ conjecture [95] suggests that in generic cases, the ratio of measurements to complex-valued free parameters must be at least 4 in order to support a unique solution in the noise-free case. This is therefore a reasonable

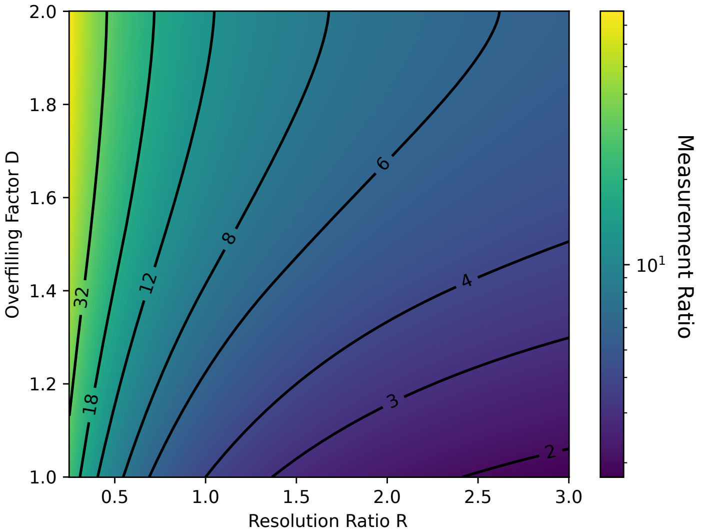

Figure 6-3: **Measurement Ratio**. The dependence of the measurement ratio defined in Equations (6.3) and (6.4) on the resolution ratio R and overfilling ratio D. Note that in the originally considered case with $D = 1$ the ratio drops below 4 at $R = 1$, but with higher oversampling ratios, enough measurements can be collected even when $R > 1$.

lower limit for the amount of data needed to support a reconstruction in the best case.

Just as in traditional CDI, once we acknowledge the presence of a finite support constraint, reconstructions should be possible in principle out to any resolution. However, as we discussed in Section 6.1, noise becomes an issue at high R. Nonetheless, as supported by the results in Figure 6-2, it is possible to capture routine reconstructions at $R = 1$.

These results change a few key aspects of how RPI should be understood. First, there is not a sharp cutoff at $R = 1$. Instead, the reconstructions gradually become less well-posed as R crosses through 1, rapidly producing noise-filled results as R grows past 1. Second the key to enabling reliable work around $R = 1$ is using focal spots that are small enough to produce data that can be properly sampled. Therefore, when designing optics for any experiment, it should be standard practice to work near $D = 2$, rather than the originally proposed $D = 1$.

Adage 6.1

Design randomized zone plates so the detector conjugate field of view is twice the focal spot diameter

6.3 Amplitude control for optic designs

One major issue that we encountered in our initial work with RPI was the difficulty of confining the illumination's power within a small, well-defined spot. The simple design strategy we used [64] led to sharp-edged and even focal spots, but they were surrounded by slowly decaying tails. These tails contained a significant amount of power compared to the power within the focus, usually more than 10%. This raised questions about the ability of RPI to study samples with low contrast, because error due to these tails would not disappear with increased photon dose. Figure 6-4 shows the presence of these tails in an example focal spot with $r = 1\,\mu\text{m}$.

Ultimately, there are two distinct features in the optic that cause the tails. The primary effect is due to phase vortices, which appear as fork dislocations in the optic. Because the grating of the naïvely designed optic has a uniform efficiency everywhere, the vortex core has a sharp discontinuity at its center. This causes unintended scattering to high angles.

The secondary effect arises from the sharp inner and outer edges of the zone plate. This discontinuity causes tails that decay with a power law in r. The resulting tails are typically hidden by the effect of phase vortices in naïvely designed zone plates, but become visible when the effect of the vortices is mitigated. A comparison between synthesized focal spots with both kinds of tails is shown in Figure 6-4.

Both effects arise because the zone plate optics imperfectly synthesize the desired wavefront at the optic plane, due to their inability to locally vary the amplitude of the diffracted wavefield. Fixing this means introducing a method for varying the efficiency of the optic locally. Several methods have been proposed previously to do this, but the method most compatible with standard lithography-based fabrication is

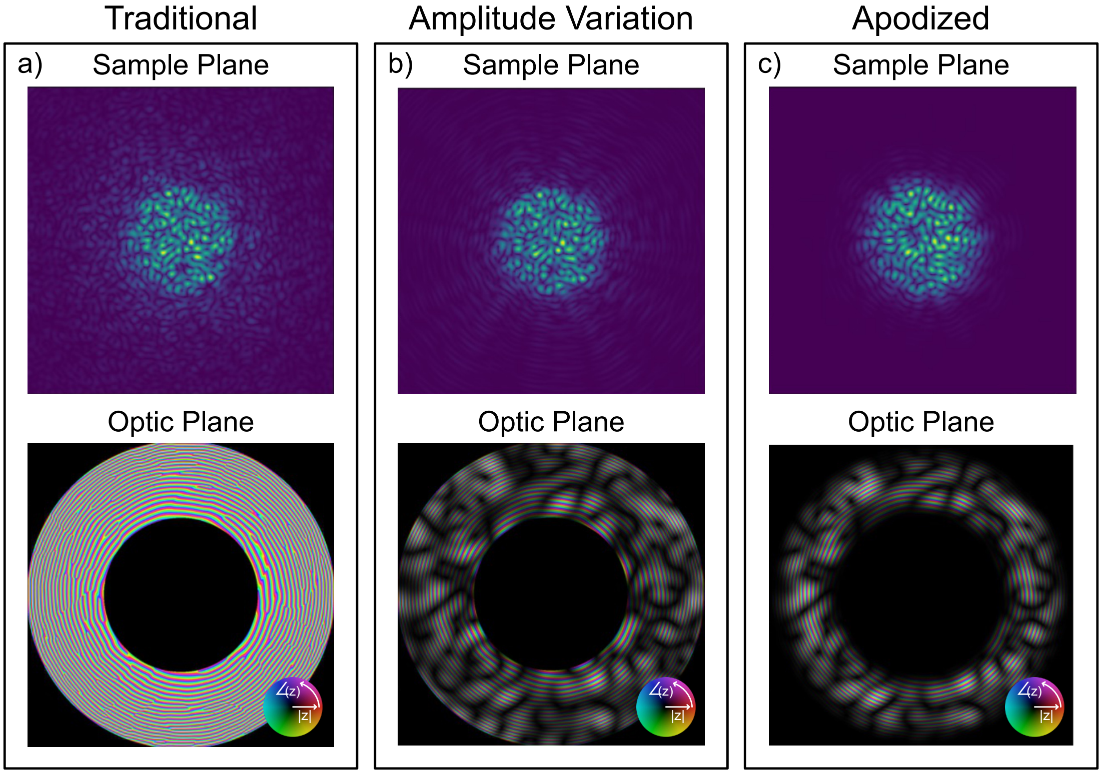

Figure 6-4: **Focal Spot Tails**. A comparison of demonstration focal spots simulated using three optic designs. (a) Upper, A focal spot simulated with a typical binary zone plate optic. Lower, the associated wavefield synthesized at the optic plane. (b) A focal spot simulated from a zone plate that also captures the amplitude variation in the target wavefield, removing discontinuities at vortex cores. (c) A focal spot from a zone plate that is apodized to remove the sharp discontinuity at the outer and inner edges of the optic

to introduce a secondary diffraction grating whose efficiency modulates the primary grating [75]—essentially, making the zones dashed lines instead of solid lines.

This method doesn't require modifying the thickness of the zones and it doesn't sensitively depend on difficult-to-control parameters like the width of the zones. Our method differs from this previously reported approach because we fixed the secondary diffraction grating to the zone plate's azimuthal direction. This keeps it always locally perpendicular to the zones, improving the ease of manufacture of the optic. We developed two methods to generate these gratings, discussed below.

Let us consider the problem of synthesizing an optic to match a target wavefield $\Psi(x, y)$. The method we originally used, defined as ZP_0 in [75], produces the binarized optic:

$$B(x,y) = \begin{cases} 1 & \text{if} \quad 0 \leq \angle\Psi(x,y) < \pi \\ 0 & \text{otherwise.} \end{cases} \tag{6.7}$$

For the first new approach, we first generate an azimuthal grating with a desired angular pitch P, in zones/radian:

$$G(x,y) = P\angle(x + iy) \mod 2\pi \tag{6.8}$$

Next, we define a diffraction grating with a locally varying efficiency, based on this azimuthal phase ramp:

$$A(x,y) = \begin{cases} 1 & \text{if} \quad 0 \leq G(x,y) < 2\pi\frac{|\Psi(x,y)|}{\max_{x,y}(|\Psi(x,y)|)} \\ 0 & \text{otherwise.} \end{cases} \tag{6.9}$$

This azimuthal diffraction grating has a duty cycle that is 100% when the target wavefield is at its maximum intensity and 0% when the target wavefield's intensity passes through zero. The final zone plate design is:

$$B2(x,y) = A(x,y)B(x,y) \tag{6.10}$$

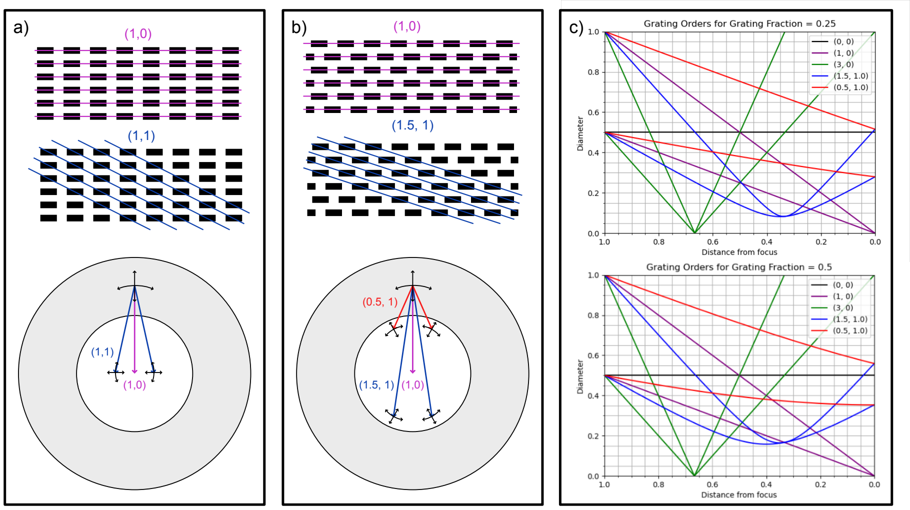

Figure 6-5: **Additional Diffraction Orders**. Introducing a second grating perpendicular to the main zone plate grating causes light to be diffracted into additional orders. (a) The relevant diffraction orders caused by a simple grating of style $B2$. Upper: Illustration of the naming convention. Lower: Diagram showing the location of the resulting diffraction orders at the focal plane of an optic. (b) The relevant diffraction orders caused by a centered rectangular grating of style $B3$. (c) caustics of the beams generated by each diffraction order, at two typical values of the grating fraction. These caustics show the constraints on order selecting aperture (OSA) placement and diameter.

This design method produces a zone plate that looks like the overlap between a set of radial spokes and a traditional zone plate. To keep a relatively uniform pitch in real space, the zone plate can be broken up into several radial regions with different azimuthal gratings.

However, this design has a potential issue: the new grating will create diffraction of its own which can pass through the order selecting aperture. Figure 6-5a diagrams the relevant diffraction orders. Comparison with Figure 6-5b highlights the fact that this issue could be alleviated slightly, given a fixed azimuthal pitch, by placing dashes in a centered-rectangular lattice rather than a simple rectangular lattice.

We can accomplish this by slightly modifying the azimuthal grating above using an "effective zone number." This is defined as the difference between the approximate unwrapped phase of an ideal zone plate [147]

$$ZP(x,y) = 2\pi \frac{x^2 + y^2}{f\lambda} \tag{6.11}$$

and the desired, wrapped phase of the wavefield to synthesize:

$$N(x,y) = \frac{1}{2\pi}(ZP(x,y) - \angle\Psi(x,y)). \tag{6.12}$$

Both wavefields are calculated for the same focal distance f and wavelength λ. This generates a value that locally increases by 1 with every zone, but globally varies continuously as the zone placement shifts due to the distortions present in the randomized zone plate. This is shown in Figure 6-6a. We can then generate the centered-rectangular azimuthal grating as:

$$C(x,y) = \begin{cases} 1 & \text{if} \quad 0 \le G(x,y) + \pi N(x,y) < 2\pi \frac{|\Psi(x,y)|}{\max_{x,y}(|\Psi(x,y)|)} \\ 0 & \text{otherwise,} \end{cases} \tag{6.13}$$

shown in Figure 6-6. The final zone plate design is then found as:

$$B3(x,y) = C(x,y)B(x,y). \tag{6.14}$$

The optics designated for work at beamline 7.0.1.2 of the ALS were designed using this $B3$ method, and comparison between the focal spots shown in Figure 6-7 and those reported in Chapter 4 shows the clear improvement possible as a result.

6.4 Deep-K Learning for RPI

One other direction we explored was the viability of using machine learning methods to reconstruct *RPI* data. I collaborated on this project with Zhen Guo and George Barbastathis, who are experts in the use of machine learning to solve inverse problems in imaging. We wrote a paper describing a hybrid machine learning and iterative system for reconstructions of simple objects [148].

The framework built on the concept of an approximant, the output of a small,

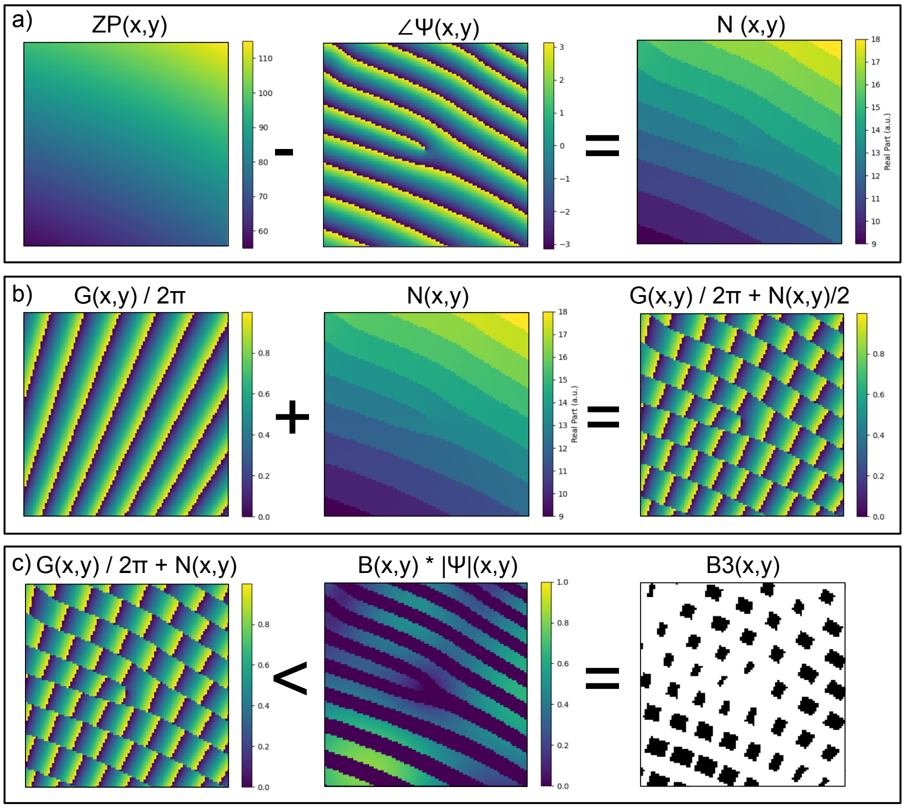

Figure 6-6: **Alternating Style Gratings**. A visualization of the zone plate design process, from a section of a demonstration zone plate design. (a) Calculation of the effective zone number. (b) Generation of the alternating azimuthal grating. (c) Use of the alternating grating to vary the amplitude of the design beam.

fixed number of iterations of an iterative algorithm [149, 150]. We then trained a neural network, in this case a U-net with attention layers, to map the noisy approximant to an estimate of the underlying image. This process turned out to work surprisingly well for simple phase objects whose structure could be well-described by a prior distribution over natural images, outperforming iterative algorithms in low-light conditions. However, we found that the utility of this approach was limited to weak phase objects, and it struggled to reconstruct mixed amplitude and phase objects, as well as objects containing phase singularities and other complicated phase textures.

This is an important direction for future work, and the outcome has clear value in many situations where RPI can be used.

6.5 Development of RPI at ALS beamline 7.0.1.2

The nanosurveyor instrument at beamline 7.0.1.2 [142] is one of the world's premier ptychography beamlines. It is often used for spectroscopic measurements of battery materials [9, 151] and magnetic nanostructures [152]. Because it was designed specifically with ptychography in mind, it has high stability, good coherence, and a large, fast detector (FastCCD, [153]). This makes it an ideal test-bed for RPI. Furthermore, its existing user community creates opportunities for RPI to be discovered by scientists who are not experts in computational imaging.

Consequently, I designed a set of RZP optics for this beamline and developed computational tools to enable straightforward analysis of RPI data at the beamline. I then performed several experiments to test the reliability of both ptychography and RPI with the newly designed optics and the computational systems.

6.5.1 Design and testing of optics

I produced a design for an optic with an outer zone width of 30 nm and a focal spot with a 6.75 µm diameter. This zone plate used the design strategy $B3$ outlined in section 6.3. The nominal design gives reconstructions at $R = 1$ (60 nm full pitch resolution) a space-bandwidth product of just over $30,000$, with 225 speckles across the field of view.

This design was sent to two organizations to be fabricated, Applied Nanotools (Optic 1) and the Center for X-ray Optics (Optic 2) at Lawrence Berkeley Lab. We then tested the optics on samples of gold nanoparticles with diameters of 100 nm and 40 nm dispersed on a Silicon Nitride membrane. Ptychography datasets were collected to characterize the probe functions of the optics at 707 eV. The reconstructions from data collected with Optic 1 are shown in Figure 6-7 and achieved a resolution of 21 nm over a field of view of nearly 10 µm. The quality and resolution of these results reinforces the following adage:

Adage 6.2

Randomized zone plate optics are great for ptychography, too.

The optics have a slight astigmatism, possibly due to their mounting and possibly due to the fabrication itself. This astigmatism would cause major blurring in scanning transmission x-ray microscope (STXM) from a traditional zone plate, but causes essentially no degradation in either ptychography or RPI imaging, as shown here.

Finally, I performed RPI on the same samples, with an example result shown in Figure 6-8. These reconstructions, as expected, are noisier and lower resolution, achieving only a full-pitch resolution of 190 nm. The results highlight the extent to which RPI is a photon-hungry technique, requiring a high coherent flux. This photon hungriness arises because RPI uses a single exposure to capture an entire full-field image. This naturally requires more photons than are needed for one exposure from a series of ptychography data, which can easily include hundreds or even thousands of exposures. In this case, the ptychography reconstruction shown in Figure 6-7 included data from $3,843$ exposures spaced out over a 5 µm × 5 µm field of view, compared with 1 exposure for the RPI data.

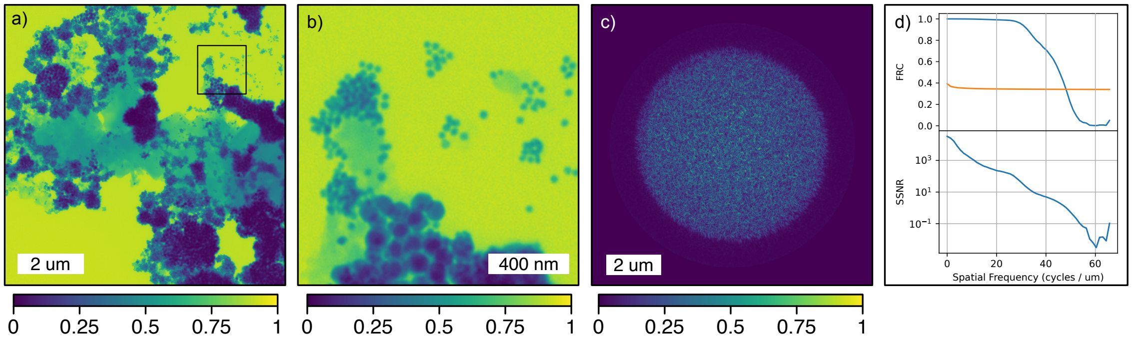

Figure 6-7: **Gold Nanoparticles**. Ptychography reconstruction of gold nanoparticles on a SiN membrane, using a zone plate designed as described in section 6.3. (a) The full reconstructed field of view, (b) A closeup of one section, showing the uniform high quality of the reconstruction, (c) The reconstructed top probe mode, (d) Fourier ring correlation (FRC) compared to a half-bit threshold and spectral signal-to-noise ratio (SSNR) of the reconstruction, calculated via a 50/50 split of the data. The final resolution at SSNR=0.5 was 21 nm full pitch.

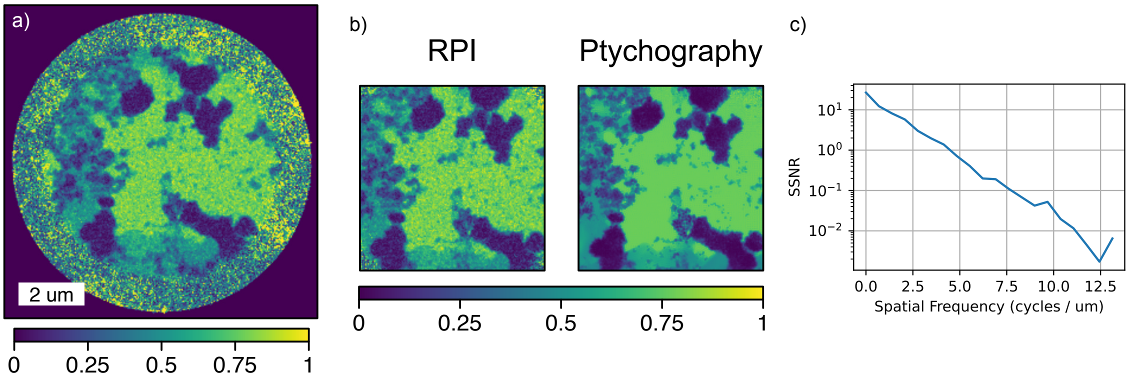

Figure 6-8: **Gold Nanoparticle RPI**. RPI reconstruction of gold nanoparticles on a SiN membrane. (a) A RPI reconstruction. (b) A comparison between the RPI reconstruction and the same region of a ptychography reconstruction downsampled to the RPI pixel size of $36nm$. (c) The SSNR calculated by comparing the RPI and ptychography reconstructions on the window shown. The final achieved resolution at SSNR= 0.5 was 190 nm full pitch.

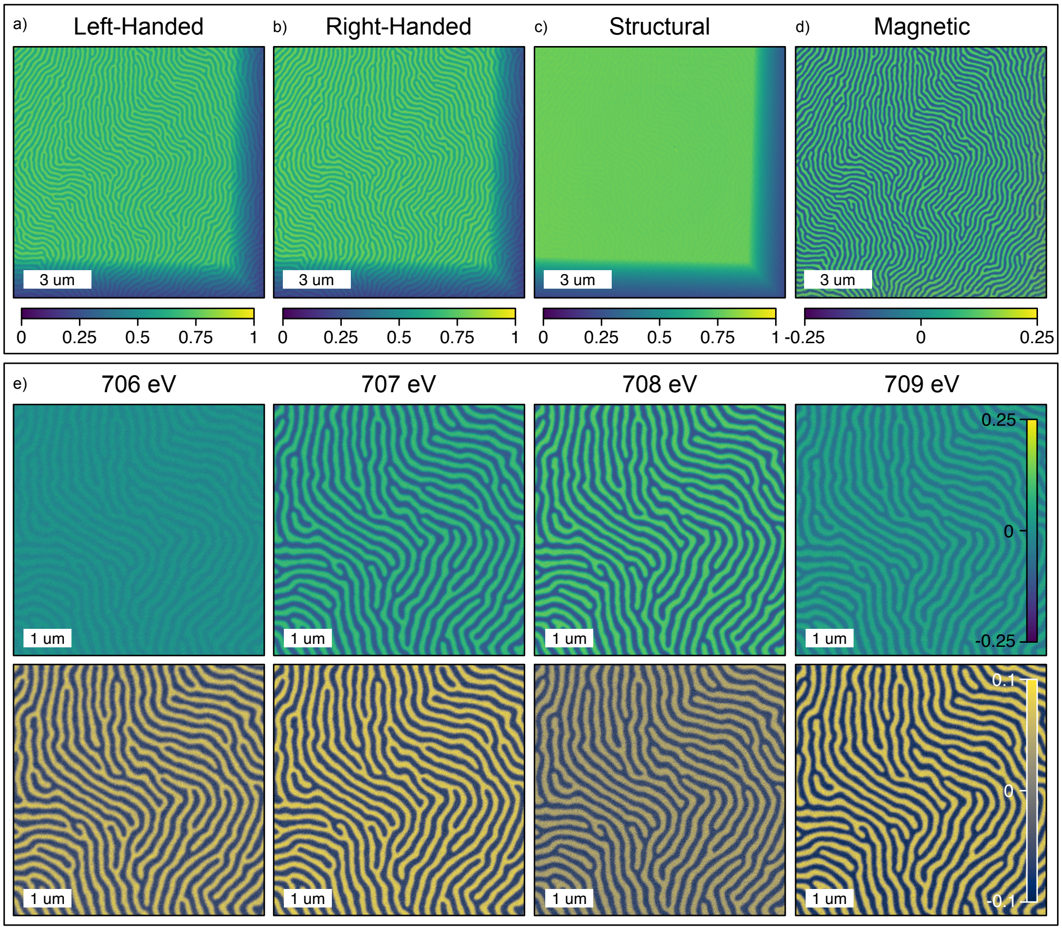

Figure 6-9: **GdFe Film**. Ptychography reconstructions of a thin film of GdFe with circular dichroism. (a,b) Reconstructions with left- and right-hand circularly polarized light. (c) The extracted amplitude of the structural component. (d) The real part of the magnetic scattering factor, extracted according to Equation (6.16). (e) Energy dependence of the magnetic signal.

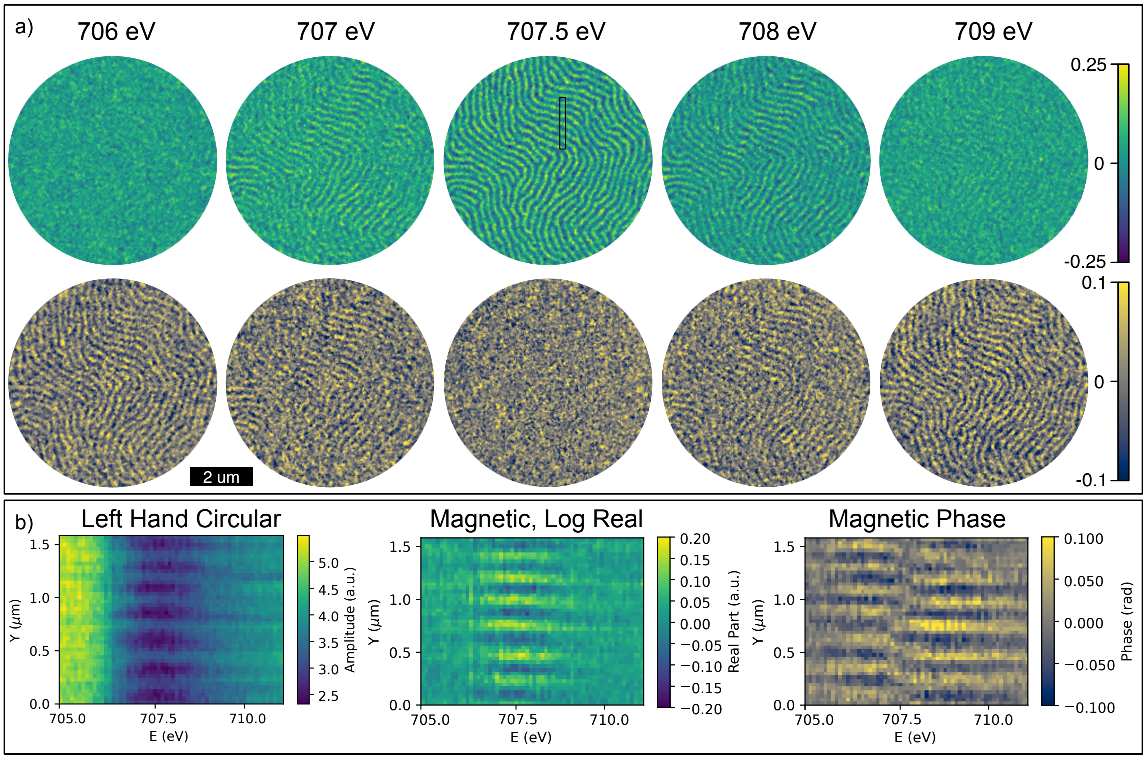

Figure 6-10: **GdFe Film RPI**. (a) Circular Dichroic RPI of a GdFe thin film at various energies. Upper: the real part of the magnetic scattering factor, extracted according to Equation (6.16). Lower: the imaginary part of the magnetic scattering factor. (b) Linecuts from the region boxed in (a), showing the full energy dependence in 0.1 eV steps.

6.5.2 Spectroscopic Imaging of Magnetic Domains in GdFe

Next, I tested the value of RZP based ptychography and RPI for magnetic imaging. To do this, I studied a GdFe(55nm)/Pt(3nm) film on Silicon Nitride, grown by Ales Hrabec. Ptychographic imaging at the Fe L3 edge revealed strong magnetic contrast, as seen in Figure 6-9. The magnetic contrast can be separated from the structural contrast using the following model for circularly polarized light impinging upon a sample with magnetization out of plane [154]:

$$\mathbf{O}_{\pm} = e^{\pm f_{\text{mag}}}\mathbf{O}_{\text{struct}} \tag{6.15}$$

Where $\mathbf{O}_+$ and $\mathbf{O}_-$ are the object when viewed under circular left and right hand polarized light. This model suggests a method of isolating the structural and magnetic contributions from a measured circular dichroic pair of images:

$$f_{\text{mag}} = \frac{1}{2}(\ln(\mathbf{O}_+) - \ln(\mathbf{O}_-)) \tag{6.16}$$

$$\mathbf{O}_{\text{struct}} = e^{\frac{1}{2}(\ln(\mathbf{O}_+)+\ln(\mathbf{O}_-))} \tag{6.17}$$

The structural object is calculated by exponentiating a summed logarithm. This preserves phase information that can be lost in direct multiplication followed by a square root. An example of this process is shown in Figure 6-9c and 6-9d. It is important to note that the magnetic scattering factor f_{mag} is in general complex-valued. It takes a form typical of any resonant process, where the phase shifts by π radians through the resonance. This can be seen in the sparse energy-dependent set of dichroic images in Figure 6-9e.

Naturally, these ptychography scans take a long time to run. Each image above derives from 30×30 images with 165 ms long exposures, a total of roughly 2 minutes 30 seconds of acquisition time. Including the overhead for motor motion, each scan took roughly 10 minutes of wall-clock time. Because of this, it is quite time-intensive to collect a full spectroscopic scan of ptychographic x-ray images. For this, I turned to RPI.

I collected a densely spaced RPI scan over the Fe L3 edge, from 705 eV to 711 eV in 0.1 eV steps, a total of 60 exposures and a total exposure time per polarization of 10 s. By wall-clock time, the scan took roughly 20 minutes, in this case because the monochromator is not set up for continuous energy sweeps and dominates the time required for a full spectroscopic scan.

6.6 Live analysis algorithm for RPI

Because RPI requires only one frame per reconstruction, it is possible to provide nearly instantaneous feedback if the proper analysis systems are in place. In an ideal situation, the reconstruction would proceed transparently, with a user seeing images of their sample live as though they are using a normal microscope. To do this, we need an algorithm that is significantly faster than the L-BFGS based algorithm described in section 6.1 but still producing robust, interpretable images.

To this end, we implemented a conjugate-gradient based method which uses the Fletcher-Reeves conjugate gradient algorithm [155] to determine update directions. The major speedup comes from an estimate of the optimal step size using a fast, direct calculation inspired by the procedure described in [53]. The algorithm is described below and explicitly defined in Algorithms 6-1 and 6-2.

The global structure of the algorithm is to minimize the mean squared error (MSE) of the simulated diffraction magnitudes when compared to the square root of the measured intensities. This error metric is ubiquitous in the ptychography literature, and can be derived as an approximation to the Poisson negative log-likelihood function [40]. The first step in the algorithm is to calculate the gradient of this error metric with respect to the object, in our case using automatic differentiation. This gradient is then transformed into a search direction with the Fletcher-Reeves [155] algorithm as laid out in Algorithm 6-1. This initial step takes roughly the same amount of time as a full iteration of the basic automatic differentiation based algorithm.

The speedup comes from quickly choosing a near-optimal step size in this direction.

Instead of doing an explicit line search, we expand the error metric explicitly to second order. The first step is to abstract out the forward model, because the wavefield at the detector plane is guaranteed to be linear in the object function:

$$\text{Forward}(\mathbf{P}, \mathbf{O} + \mathbf{S}) = \text{Forward}(\mathbf{P}, \mathbf{O}) + \text{Forward}(\mathbf{P}, \mathbf{S}). \tag{6.18}$$

This means that we can expand the error metric in the detector-plane wavefield. The second order expansion of the loss function in the direction $\Delta\tilde{\mathbf{E}}$ with step size α is:

$$\begin{aligned}
\mathcal{L} &= \sum(\sqrt{\mathbf{I}} - |\tilde{\mathbf{E}} + \alpha\Delta\tilde{\mathbf{E}}|)^2 \\
&= \sum(\sqrt{\mathbf{I}} - \sqrt{(\tilde{\mathbf{E}} + \alpha\Delta\tilde{\mathbf{E}})(\tilde{\mathbf{E}} + \alpha\Delta\tilde{\mathbf{E}})\dagger})^2 \\
&= \sum\left(\sqrt{\mathbf{I}} - |\tilde{\mathbf{E}}|\sqrt{1 + \alpha\frac{2\text{Real}[\tilde{\mathbf{E}}(\Delta\tilde{\mathbf{E}})^\dagger]}{|\tilde{\mathbf{E}}|^2} + \alpha^2\frac{|\Delta\tilde{\mathbf{E}}|^2}{|\tilde{\mathbf{E}}|^2}}\right)^2 \\
&\approx \sum\left(\mathbf{D} - \alpha\mathbf{A} - \alpha^2\frac{|\Delta\tilde{\mathbf{E}}|^2}{|\tilde{\mathbf{E}}|} + \alpha^2\frac{\mathbf{A}^2}{4|\tilde{\mathbf{E}}|}\right)^2 \\
&\approx \sum \mathbf{D}^2 - 2\alpha\mathbf{A}\mathbf{D} + \alpha^2\mathbf{A}^2 - \alpha^2\mathbf{D}\frac{4|\Delta\tilde{\mathbf{E}}|^2 - \mathbf{A}^2}{2|\tilde{\mathbf{E}}|}.
\end{aligned} \tag{6.19}$$

Above, the sums are defined to run over all pixels, and $\mathbf{D}$ and $\mathbf{A}$ are defined as in Algorithm 6-2. From this expansion, we can estimate the step size required to hit the minimum of the loss function:

$$\begin{aligned}
0 &= \frac{\partial\mathcal{L}}{\partial\alpha} \\
&= \sum -2\mathbf{A}\mathbf{D} + \alpha\left(2\mathbf{A}^2 - \mathbf{D}\frac{4|\Delta\tilde{\mathbf{E}}|^2 - \mathbf{A}^2}{|\tilde{\mathbf{E}}|}\right) \\
\alpha &= \frac{\sum \mathbf{A}\mathbf{D}}{\sum \mathbf{A}^2}
\end{aligned} \tag{6.20}$$

This calculated step size is therefore a near-optimal step calculated in constant time. In practice, however, the final term in the denominator is typically dominated by pixels where $|\tilde{\mathbf{E}}|$ is small, precisely the pixels where the Taylor series approximation will perform the worst. Therefore, we remove this term, resulting in the estimated step

$$\alpha = \frac{\sum \mathbf{A}\mathbf{D}}{\sum \mathbf{A}^2 - \mathbf{D}\frac{4|\Delta\tilde{\mathbf{E}}|^2 - \mathbf{A}^2}{2|\tilde{\mathbf{E}}|}}. \tag{6.21}$$

This step is very stable in practice, and tends to either accurately estimate the optimal step or slightly underestimate it. This bias toward small steps is preferable to a bias toward large steps, as it makes instabilities unlikely. The final algorithms are defined in Algorithms 6-1 and 6-2

Algorithm 6-1 The conjugate-gradient based algorithm for RPI

```
$\mathbf{S}_{-1} = 0$
for $i \in [0, n_{\text{iter}})$ do
    $\tilde{\mathbf{E}}_i \leftarrow \text{Forward}(\mathbf{P}, \mathbf{O}_i)$                          ▷ From Equation (4.1)
    $\mathbf{D}_i \leftarrow \sqrt{\mathbf{I}} - |\tilde{\mathbf{E}}_i|$
    $\Delta\mathbf{O}_i \leftarrow -(\sum \mathbf{D}_i^2).\text{backward}()$                  ▷ Calculate gradients with AD
    $\mathbf{S}_i \leftarrow \Delta\mathbf{O}_i + (\sum |\Delta\mathbf{O}_i|^2) \,/\, (\sum |\Delta\mathbf{O}_{i-1}|^2) \cdot \mathbf{S}_{i-1}$     ▷ Fletcher-Reeves CG Step
    $\mathbf{O}_{i+1} \leftarrow \mathbf{O}_i + \text{EstimateOptimalStep}(\mathbf{S}_i, \tilde{\mathbf{E}}_i, \mathbf{D}_i)$
end for
```

Algorithm 6-2 Estimate optimal step for an object function along a given direction

```
$\mathbf{S} \leftarrow$ Desired step direction
$\tilde{\mathbf{E}} \leftarrow$ Current simulated exit wavefield
$\mathbf{D} \leftarrow$ Difference between simulated and target diffraction pattern magnitudes
function EstimateOptimalStep($\mathbf{S}, \tilde{\mathbf{E}}, \mathbf{D}$)
    $\Delta\tilde{\mathbf{E}} \leftarrow \text{Forward}(\mathbf{P}, \mathbf{S})$
    $\mathbf{A} \leftarrow \left(\text{Real}\left[\tilde{\mathbf{E}}(\Delta\tilde{\mathbf{E}})^\dagger\right]\right) / \left(|\tilde{\mathbf{E}}|\right)$
    $\alpha \leftarrow (\sum \mathbf{D}\mathbf{A}) \,/\, (\sum \mathbf{A}^2)$
    return $\alpha_i \mathbf{S}$
end function
```

A numerical experiment comparing a reconstruction of optical RPI data with the fast algorithm to a reconstruction with the full algorithm is shown in Figure 6-11. Both objects were set to 512 × 512 pixels, using 1000 × 1000 pixel detector images. The full automatic differentiation algorithm used L-BFGS. Because no line search

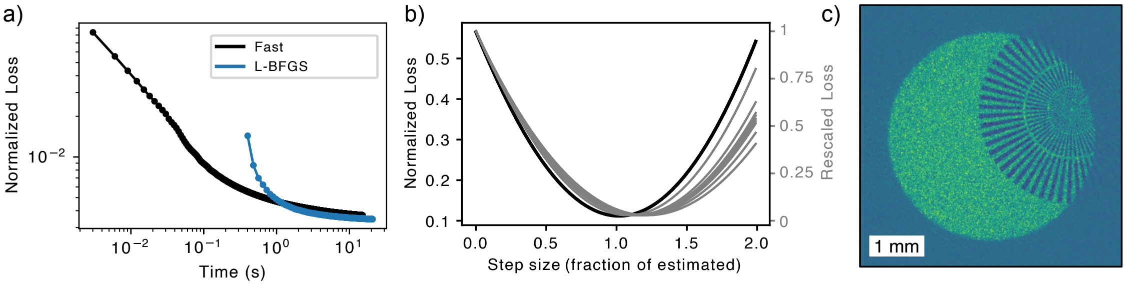

Figure 6-11: **Fast RPI algorithm performance**. (a) The normalized loss metric (euclidean distance of magnitudes) as a function of time for reconstructions run on optical data with the fast and full algorithms. (b) Dependence of the loss metric on step size for the first ten iterations of the fast reconstruction. The first iteration is highlighted in black. (c) The reconstructed object after 3 iterations, or 0.01 seconds, showing that very fast reconstructions can still be valuable for transparency into the experiment state.

is needed, the fast iterations can complete much faster (0.003 seconds per iteration, vs 0.125 seconds for the starting iterations of L-BFGS). The difference in speed also includes some specific optimizations to make each iteration compute faster.

At long times, the algorithm converges more slowly than L-BFGS, likely due to the less sophisticated method of selecting the update direction. Further work could address this, but the algorithm as it now exists provides sufficient quality for fast visibility into an experiment's state, and has proven useful at the beamline.

Chapter 7

Outlook

Coherent lensless imaging is a rapidly growing field, full of exciting opportunities. The work I explored in this thesis pulled at a few of these threads. In chapter 3, I discussed how automatic differentiation can be used to make custom, flexible reconstruction algorithms tailored to specific experimental needs. This program of research, aimed at developing easy and reliable algorithms for everyone's needs, has been underway for over a decade and will continue apace.

In chapter 4, I then discussed how several ideas that have been "in the air" can be melded into an imaging method which is greater than the sum of its parts. The method I define in that chapter, randomized probe imaging (RPI), has very different limitations from traditional methods, making it an ideal option in many scenarios. Its robustness and single-frame nature make it a natural match for experiments at free electron lasers (FELs).

In chapter 5, I discussed exactly that, demonstrating how RPI can be implemented in the unique environment of a FEL. This project showed how the large space-bandwidth product of RPI reconstructions is unusual for single-shot imaging, and how valuable this kind of imaging can be.

Finally, in chapter 6 I discussed a collection of small projects, primarily from the final year of my degree work, that were designed to push the technical side of RPI forward. This includes improved algorithms, a better understanding of the limits of RPI, and improved designs for randomized zone plates (RZPs). This culminated in a discussion of the implementation of RPI at a beamline of the Advanced Light Source, designed to make the method accessible to users.

7.1 A larger role for coherent imaging

One of the biggest drivers behind the growth of lensless imaging is the development of FELs and fourth-generation synchrotron light sources, which produce nearly fully coherent illumination. The first few fourth-generation synchrotrons have already come online, and many synchrotron light sources will be undergoing upgrades in the coming years.

These light sources typically don't make more intense beams of x-rays, but a much larger fraction of the light they produce is coherent. Making use of that coherence has therefore become a significant goal, and lensless imaging is one of the best ways to do that. This means that the balance is shifting ever faster in favor of coherent imaging, and consequently coherent imaging methods that are a bit esoteric today will become the gold standard for many kinds of imaging in the future.

Although the community hasn't yet settled on a replacement for full-field transmission x-ray microscopy, RPI is an ideal candidate. Developing full-field methods is especially important because the brighter coherent beams can quickly become brighter than they need to be for ptychography scans, which are limited by the mechanical motion of a probe. The large footprint of the focused beams from RZP spread out the flux and make it easier to image at low doses.

7.2 Bragg mode lensless imaging with soft x-rays

To date, very few experiments have found success using soft x-rays to study electronic structure in the Bragg and reflection geometries. Nonetheless, this experiment is particularly valuable because it can, in principle, measure the amplitude and phase of electronic order parameters in a symmetry-broken material. For example, it is possible to measure the antiferromagnetic order parameter in spin density wave material, or the local order parameter in a charge-density wave. Access to phase information means that structures such as vortices and phase slips can be visualized directly, and correlated with other physics that may be present in a sample.

These measurements have been hard to perform because very few sources of soft x-ray light are bright enough to get a measurable signal from a coherent scattering experiment. As the number of coherent soft x-ray synchrotron sources grows, more opportunities will arise to develop this imaging method and slowly the community will come to understand how to overcome the experimental difficulties facing their work. Some methods will involve error correction models like those discussed in Chapter 3, and others will involve changes to the beamlines and endstations. Continued work in this direction is critical, because the ability to routinely study the mesoscale structure of these subtle electronic correlations is critical to improving our understanding of many classes of materials.

7.3 Applications of RPI

I am especially hopeful for the role that RPI will play in this research and other fields. Following my year working at the Advanced Light Source, a setup now exists to perform RPI experiments at beamline 7.0.1.2. This includes a randomized zone plate that is well matched to the experimental geometry at transition metal l-edge wavelengths. This setup has already proven useful for quantitative imaging over large areas using ptychography, and with the soon-to-arrive upgrade to the Advanced Light Source, the flux will become sufficient for high quality single-frame imaging.

Work is also underway to make reflection and Bragg mode imaging possible with RPI at the Advanced Light Source, enabling fast imaging of the electronic order parameters discussed above. Already, a strong research program exists to study the direct scattered light emerging from these phases and extra information about the correlation timescales. Upgrading these experiments to visualize the dynamics directly in real space will unlock new kinds of insight. Achieving this goal involves overcoming challenging technical hurdles, but the potential value is high. RPI is unique among reference-free single frame lensless imaging methods because it is robust to complicated phase textures with phase vortices.

7.4 Variations on and extensions of RPI

The framework behind RPI can also be extended in a number of exciting ways beyond simply applying the existing method to studies of new materials. One proposal I have developed is a multi-frame RPI experiment. In this experiment, a single optic produces multiple time-delayed randomized probes on the sample. The diffraction from this time-delayed series of probes can be analyzed to generate a full movie of a single event.

In the world of optical light, a number of fun possibilities exist for extending RPI. One obvious direction is to add polarization dependence, where the speckled input wavefield could contain a spatially varying polarization. This opens up the possibility of extracting quantitative Jones matrix information from single exposures on moving subjects, such as biological cells.

Another interesting direction would be to implement a Fourier-transformed version of RPI, related to RPI in the same way that Fourier ptychography is related to ptychography. Such a method could enforce a finite support constraint in real space explicitly, using a pinhole, and perform the multiplication with a diffuser in the Fourier plane, with the diffuser's length scale chosen to be small enough to enable sufficient oversampling in the ultimate real-space intensity measurement.

Surely other possible extensions exist that I have not yet considered, and the process of exploring and developing this space is likely to lead to a rich and useful program of research.

7.5 Conclusion

As more and more x-ray sources come online with perfect or near-perfect coherence, the value of coherent lensless imaging will continue to grow. The work presented in this thesis constitutes a step forward for the practicality of these lensless imaging methods, and will form a part of the foundation of work in this field moving into the future. I sincerely hope that some of the concepts and methods explored in this dissertation contribute to the developments of the future.

Appendix A

Supplement to single-frame far-field diffractive imaging with randomized illumination

Abstract

In this supplement, we provide additional information that may be helpful for researchers seeking to reproduce the results of the aforementioned paper. This includes the full set of parameters used in our numerical experiments, an expanded explanation of our zone plate design process, a discussion of the method used to calculate the Fourier ring correlation curves reported in the main paper, and a description of our fabrication method for the x-ray zone plate optics. In addition, we discuss our calculation of the resolution limit for reconstructions, further discuss the effect of the illumination intensity's nonuniformity, and provide a set of design considerations for RIP setups.

A.1 Numerical Experiment Parameters

We describe here the setup of our numerical experiments. Those interested in further inspecting the details of the experiments are encouraged to contact the authors for the original code.

In all cases, ideal BLR illumination was defined on an array with size larger than $2k_p \times 2k_p$, where k_p is the intended maximum probe frequency measured in pixels.

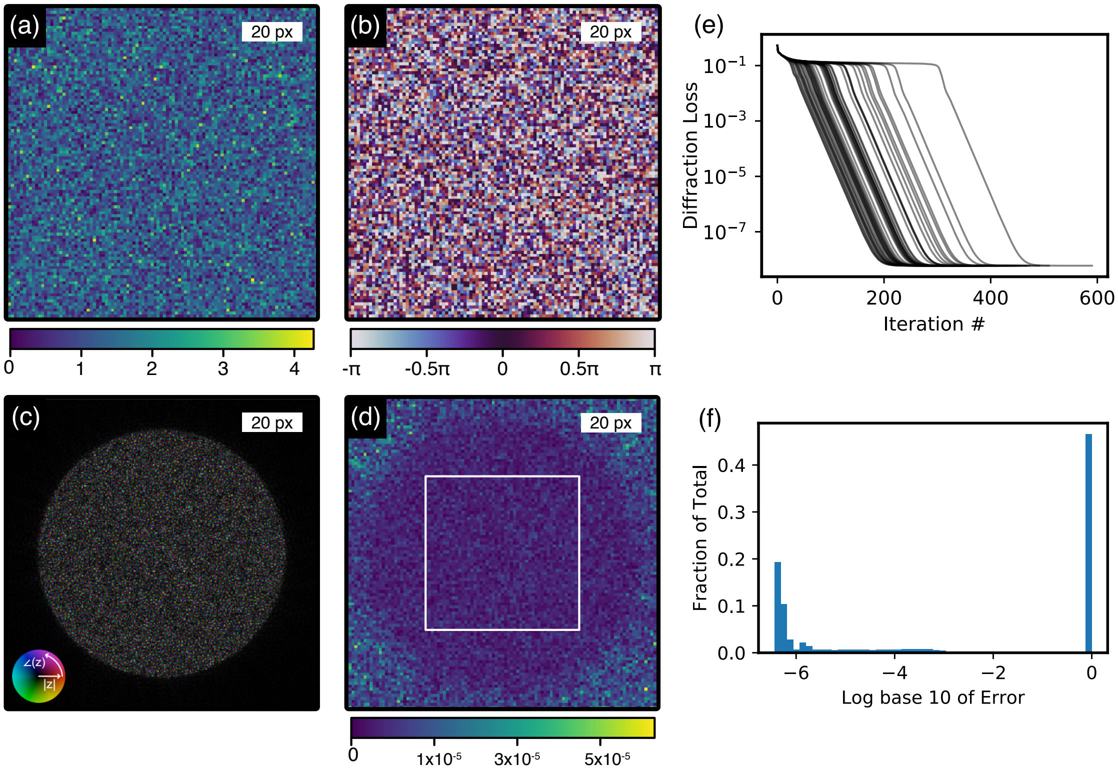

Figure A-1: **Typical Reconstruction**. (a), (b) the amplitude and phase of a typical random object. (c), a typical randomized probe used in the simulations. (d), The magnitude of the difference between a completed noise-free RPI reconstruction and the ground truth, with the central region from which error is reported marked in white. (e) The progression of diffraction loss over time for 50 independently initialized reconstructions on the same typical diffraction pattern. (f) The distribution of final RMS error for all reconstructions reported in Figure 3a.

To generate the illumination, a central circular region with a diameter $\frac{4}{5}$ of the array size is filled with uniform amplitude phase noise. This is then propagated into the far-field via a 2D Fourier transform. The region in Fourier space within $|k| = \frac{k_p}{2}$ is set to zero, as is the region outside $|k| = k_p$, leaving a ring in reciprocal space. This is then propagated back into the near field to form an ideal BLR probe for simulation.

Before simulating diffraction from the objects, the probes were upsampled by padding in Fourier space such that they are defined on an array of at least $(2k_p + 2k_o) \times (2k_p + 2k_o)$ pixels. This ensures that the multiplicative interaction between the probe and object doesn't lead to aliasing. We show the results of a typical ensemble of reconstructions from a randomly generated object in A-1.

For the numerical experiments which discuss the impact of noise sources, all reconstructions were performed on a probe with maximum frequency $k_p = 128$ and a band-limiting frequency of $R = 0.4$. In each case, we performed reconstructions on between 50 and 200 randomly generated images per noise level. The error is determined from a square central region of the reconstructions, with a side length half that of the overall array, as shown in Figure S1d.

A.2 Design of Randomized Zone Plates

Our technique relies on illuminating a large region with highly speckled light containing high frequency components. This illumination function was achieved with diffractive optics that fill a clearly defined field of view with light whose intensity is as uniform as possible. The field of view is chosen to have a sharp drop-off simply so that the eventual reconstructions achieve comparable noise levels across the entire field of view. Other approaches which, for example, lead to a Gaussian envelope on the probe are equally applicable in principle, but lead to a less uniform noise profile in the reconstruction.

To generate these diffractive optics, we start by simulating a discretized light field that consists of uniform amplitude pure phase noise within a circular aperture whose diameter matches the intended final field of view diameter. We then numerically

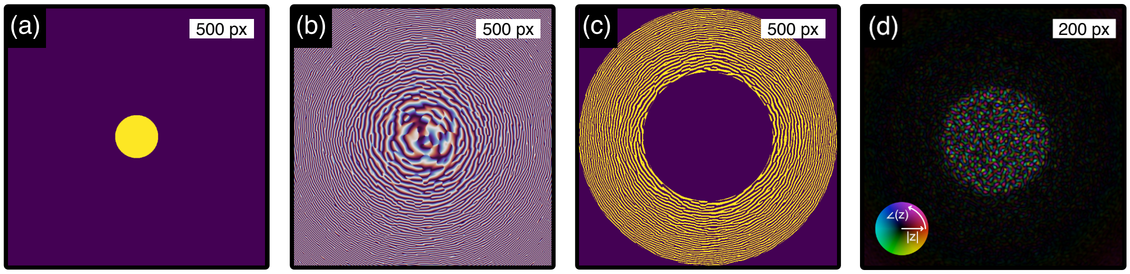

Figure A-2: **Optics Design**. (a), The amplitude of the initial focal spot, filled with randomly varying phase. (b), The phase of this focal spot, propagated to the plane of the optic. (c), The binarized optic defined from this phase distribution. (d), The focal spot simulated from this binarized optic.

propagate this light field to the plane of our diffractive optic, and record the phase of the light field at each pixel. We generate a binarized optic with a specified outer diameter and beamstop diameter from this phase field. The zones are defined by setting all pixels with phase below a certain threshold to 0 and all other pixels to 1. This process matches the simplest "ZP_0" approach described in detail in [75]. This design process is further outlined visually in Supplementary figure A-2. All zone plates used in this paper used a beamstop with a diameter of half the outer zone plate diameter.

A.3 Experimental Resolution Calculations

All reported resolutions were calculated via the Fourier Ring Correlation (FRC)[101] method, by comparison with ptychography results captured from the same sample region. Although this is potentially an underestimate of the true resolution if (as we saw in our reconstructions from FeGd samples) the ptychography itself exhibits pathologies, we believe that it is a conservative estimate which addresses the fundamental question of how RPI compares with ptychography.

In each case, we started by cutting out a region of the ptychography scan which overlapped with a region of the RPI reconstruction entirely within the RPI field of view. This ptychography cutout was then downsampled by extracting the central region in Fourier space corresponding to the band-limiting constraint in the RPI reconstruction. Any linear phase ramps in the RPI reconstructions (arising when the detector alignment shifts slightly between the calibration and RPI reconstructions) were manually removed and the cropped images were apodized with a Hann window. We then calculated a subpixel shift between the RPI and downsampled ptychography reconstructions and shifted the ptychography reconstruction to overlap with the RPI result. Finally, an FRC was calculated between the two images, excluding the outer 4 pixels to avoid artifacts from the circular shift of the ptychography reconstruction.

A.4 X-ray Optics Fabrication

The thin film stack was prepared by magnetron sputtering (Leica EM ACE 600, Germany) of a 100 nm thick Au film on a 100 nm thick SiN membrane with a window size of 500×500 μm on a 200-μm thick Si frame (Silson ltd, UK). The ion beam lithography was done using a dualbeam focused ion beam instrument (Nova Nanolab 600, FEI, Netherlands) with a pattern generator attachment (ELPHY MultiBeam, Raith GmbH, Germany). The design bitmap pattern was converted to a dot map GDSII stream file and used as input. The binarization, using a 0.15:0.85 line-to-space ratio, resulted in a pattern that can be machined using a process that resembles a single-pixel-single-pass process discussed in [156], and which resulted in the best pattern quality. The 60-μm wide computer-generated-hologram with 40-nm outermost width was milled using a 30-kV Ga+ ion beam and 50-pA current (19 nm nominal beam size) and 0.5-ms dwell time resulting in a dosage of 0.025 pC per dot.

A.5 Calculation of Resolution Limit

The theoretical limit on resolution is found by comparing the number of intensity measurements contained in the diffraction pattern M to the number of complex parameters in the object N. It is generally believed that, in order for the phase retrieval to be well posed, the number of intensity measurements must typically be greater than four times the number of complex parameters in the object [94, 95]. This belief is conditioned on relatively weak assumptions regarding the diversity of information contained within the individual measurements. We note that M is proportional to the area of the support of $\tilde{P} * \tilde{O}$ (the diffraction patttern). In contrast, N scales with the support of $\tilde{O}$, with the same constant of proportionality. In our test geometry, we use a circular zone plate which fills a ring in reciprocal space $\frac{k_p}{2} < |k| < k_p$ with random phase noise. The band-limiting constraint applied to the object, however, is a square such that $|k_x| < k_o$,$|k_y| < k_o$.

For $k_o > \frac{k_p}{2}$, the support of $\tilde{P} * \tilde{O}$ will be a rounded square covering an area of $A_I = \pi k_p^2 + 8k_pk_o + 4k_o^2$. The support of $\tilde{O}$ will simply be $A_O = 4k_o^2$. We can solve for the ratio, $R = \frac{k_o}{k_p}$, such that $A_I = 4A_O$, finding $R \approx 0.944$. In our paper, we chose to band-limit the object to a square region to allow for a simple computational approach and a well-defined interpretation of the resulting images. In a potentially more elegant reconstruction approach where the object is band-limited to a circular region in Fourier space, one finds the simpler result that the reconstructions are theoretically limited to $R = 1$.

This result should be considered as a limit on the potential reconstruction resolution of RPI rather than an estimate of the likely achievable resolution. In order for the reconstruction problem to be well posed at $M = 4N$, the measurement vectors must meet conditions which we do not guarantee. In addition, practical phase retrieval algorithms typically require more stringent constraints on the ratio of measurements to complex parameters to lead to convergence with high probability.

A.6 Illumination Uniformity

One issue which arises when using structured illumination is a variation in the intensity of the light which interacts with each pixel of the low-resolution object. In a shot-noise limited experiment, this naturally leads to higher uncertainty in the reconstruction of weakly illuminated pixels. Other noise sources which uniformly affect the diffraction pattern are also expected to lead to higher errors in the reconstruction of weakly illuminated pixels.

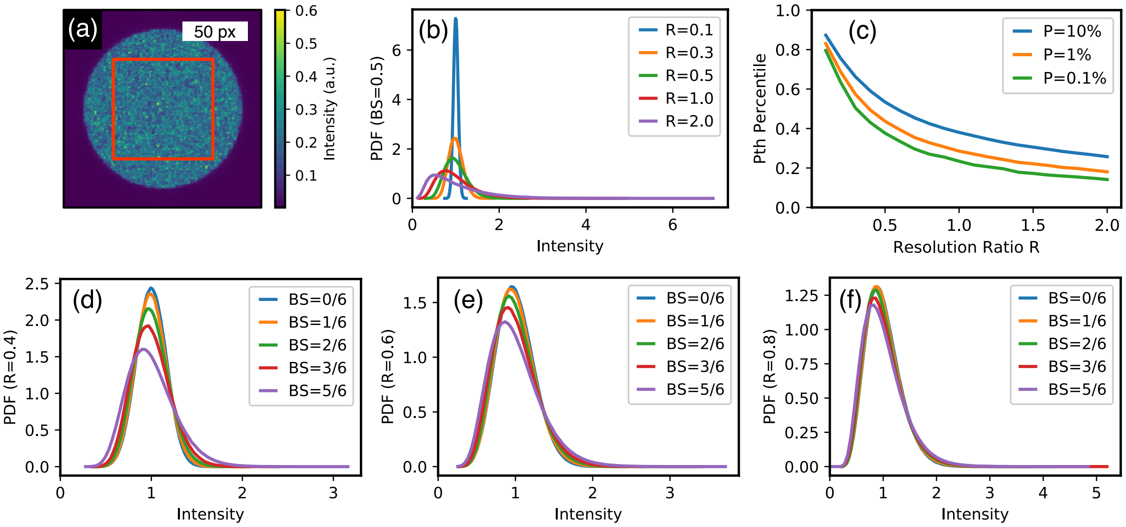

Figure A-3: **Illumination Uniformity**. (a), The distribution of intensities over the focal spot for $R = 0.6$, $BS = 0.5$, with the region used to extract the PDFs outlined in orange. (b) The intensity PDF arising from BLR light with $BS = 0.5$ at various resolution ratios. (c) The ratio of various percentiles to the median flux as a function of R for $BS = 0.5$. (d), (e), (f), The dependence of the illumination PDF on the beamstop diameter BS at $R = 0.4$,$R = 0.6$, and $R = 0.8$ respectively.

Conceptually, one can break the intensity variations into two distinct categories. Most important are the high-frequency variations that manifest themselves in the network of zeros interpenetrating the illumination. However, in addition to these nodes between the speckles, the speckles themselves have nonuniform intensities. This variation becomes important at high R as the number of speckles per pixel decreases. Additionally, at large beamstop diameters correlations between neighboring speckles emerge. This reduces our ability to improve the uniformity by lowering R, because the various speckles within a low-resolution pixel remain correlated with one another.

To inspect the extent of this issue in RPI, we generated many BLR probes with $k_p = 128$ at a variety of beamstop diameters using the framework outlined in Supplementary Section 1. The beamstop diameters are defined by BS, the ratio of the beamstop diameter to the optic diameter. Then, we ran the forward model on low-resolution objects at a collection of resolution ratios R. For each pixel in the object, we recorded the total diffracted intensity resulting from an object with that pixel set to 1 and all others set to 0. This measures how sensitive the final diffraction pattern is to the complex free parameter associated with that pixel. We then studied the distribution of resulting intensities arising from the pixels within the illuminated field of view.

The results are presented in Figure A-3. We find several important trends. The dominant trend is, as expected, an increase in the intensity variation with resolution ratio, R, at all beamstop diameters. In addition, illumination with a smaller diameter beamstop leads to smaller deviations at low R. This is a result of the beamstop-induced low-frequency intensity correlations. Nevertheless, we find that the distribution of intensities remains relatively small in the region $R < 0.6$ where reconstructions are experimentally feasible. This remains true even when using a beamstop with the extraordinarily large diameter of $\frac{5}{6}$.

In conclusion, the illumination non-uniformity is likely to become an increasingly relevant issue if algorithmic improvements allow for numerically stable reconstructions at values of R closer to 1. However, for currently feasible reconstruction ratios and standard beamstop diameters, the variability in illumination is tolerable. To be specific, at $R = 0.6$ and $BS = 0.5$, 99.9% of all pixels are illuminated by light which is at least 0.39 times as intense as the flux through the median pixel.

A.7 Application Recommendations

First, we discuss recommendations for ptychography beamlines interested in adding a new zone plate to enable RPI. The most important parameter to consider is the numerical aperture of the randomized zone plate. In cases where the numerical aperture of the detector is large enough to enable reconstructions below a full pitch resolution of 20 nm, the appropriate zone plate numerical aperture is likely to be “as high as possible”, because it may not be practical to design a zone plate with a numerical aperture large enough to fill the detector.

However, in a situation where it is possible to design such a zone plate with a numerical aperture matched to a detector, we recommend choosing a zone plate design which fills roughly $\frac{2}{3}$ of the detector at the lowest commonly used energy. This is because RPI reconstructions are likely to be reliable out to a resolution ratio of $R \approx 0.5$. For a zone plate which fills $\frac{2}{3}$ of a detector, the highest reconstructed object frequencies at $R = 0.5$ will be pushed exactly to the edge of the detector. One is tempted to increase the filling of the detector, however there is an advantage to leaving a portion of the detector unfilled. Because the high frequency components of an object are typically much weaker than the low-frequency ones, overfilling a detector with a high NA zone plate will swamp all the high frequency components with Poisson noise from the intense low frequency region. The portion of the diffraction pattern beyond the zone plate filling only includes high frequency components, and it is therefore desirable to capture it on the detector.

In all cases, we recommend designing the zone plate such that its focal spot fills roughly 80% of the detector conjugate coordinate space at the highest commonly used energy. It may seem ideal to simply fill this space to get the most out of each pixel. An obvious downside is that, as the real space oversampling decreases, the finite pixel size on the detector leads to a reduced speckle contrast. In addition, including a small region around the focal spot allows some room for the probe to be defocused and therefore reduces the complexity of probe alignment. We finally note that, if a large energy range is used at a beamline, it may be worthwhile to design a collection of zone plates optimized for use at different energies.

We next discuss the issues that are relevant when designing a system from the ground up with RPI in mind. First, and most obviously, one must have set of motors capable of reliably scanning a test sample through the beam with accuracy high enough for ptychography to succeed. Second, there is an incentive to use a detector with as many pixels as possible. This is, unsurprisingly, because the space bandwidth product of the final RPI reconstruction will be related to the number of pixels in the detector. Fortunately, randomized zone plates lead to ptychography and RPI reconstructions which are robust to missing data, so there are no major issues with using segmented detectors to increase the pixel count.

If one is interested in the highest resolution imaging, it is prudent to design the highest numerical aperture zone plate which is reasonable, and ensure that the detector can subtend a numerical aperture which is at least $\frac{3}{2}$ times that. If one is interested in capturing lower resolution images with a larger field of view, it may be more appropriate to start by designing a detector geometry which leads to the desired field of view before designing a zone plate to match that detector's numerical aperture. Finally, designing with beamline stability in mind is especially critical, because a well characterized beamline would (for example) enable energy or polarization sweeps without the need for repeat calibration ptychography scans.

Appendix B

Error Metrics for Partially Coherent Wavefields

The following appendix is based on a publication which appeared in Optics Letters [138], discussing the problem of comparing partially coherent wavefields with one another. It is included below with modifications to fit the thesis formatting, and followed by a discussion highlighting further progress in our understanding following the publication of this paper.

The paper was coauthored by myself and Riccardo Comin. I was responsible for initiating the line of inquiry and deriving the presented results. Copyright is held by Optica Publishing Group. See the published version [138] for full affiliation and funding info.

B.1 Abstract

Lensless imaging methods that account for partial coherence have become very common in the past decade. However, there are no metrics in use for comparing partially coherent light fields, despite the widespread use of such metrics to compare fully coherent objects and wavefields. Here, we show how reformulating the mean squared error and Fourier ring correlation in terms of quantum state fidelity naturally generalizes them to partially coherent wavefields. These results fill an important gap in the

lensless imaging literature and will enable quantitative assessments of the reliability and resolution of reconstructed partially coherent wavefields.

B.2 Body

When demonstrating a new imaging method or reconstruction algorithm, it is important to quantitatively study its accuracy and reliability. In the ptychography literature, this is usually done by calculating a variation on the Mean Squared Error (MSE) [157] between reconstructed and ground truth objects, as a function of noise level or other parameters. Because the signal quality in an image is typically length-scale-dependent, it is also common to consider the resolution of an imaging method, i.e. the smallest length scale at which the retrieved image is sufficiently accurate. This is often accomplished with the Fourier Ring Correlation (FRC), which [158] has the additional advantage that it can determine an empirical resolution from experimental data even when the ground truth is not known.

However, because of two recent trends these practices are no longer always sufficient. First, mixed state (or "multi-mode") ptychography methods have become popular since their introduction in 2013 [43]. These methods treat the illumination as an incoherent mixture of wavefields. In doing so, they account for the inevitable presence of partial coherence due to, for example, the intrinsic properties of the source [159–161] or physical motion of the sample [162–165].

Second, ptychography is now routinely used to characterize the probe itself, with applications ranging from understanding the statistical source properties of synchrotrons [166] and free electron lasers [55, 167, 168] to measuring the aberrations of x-ray optics [169, 170]. Therefore, it has become important to place the veracity of probes retrieved via mixed state ptychography on a firm footing. Because the the traditional metrics (MSE and FRC) only apply to fully coherent light, though, they can't be used for this analysis.

Through a literature search we have identified two methods which have been used to solve this problem and study the consistency of multi-mode probe reconstructions.

The first, most common method is to plot the orthogonalized modes of the light fields to be compared [43, 159–161, 164, 166–168]. This is a comprehensive overview, but it can't be used when quantitative comparison is required. It is often augmented by a breakdown showing the relative power in each orthogonalized mode [162, 163, 165]. The mode breakdown can be used to generate quantitative comparisons, e.g. comparing the global degree of coherence between reconstruction attempts [55], but these metrics are insensitive to variations in the spatial structure of the probe modes.

The second, less common method is to compare each pair of orthogonalized probes in series using a metric such as the normalized MSE [171]. This takes into account the spatial structure of the probe, but it has the problem that the ordering of the modes can be unstable. There is also no natural way to reduce this list of comparisons to a single error metric, and it can't compare light fields with different numbers of modes.

A metric which could avoid these pitfalls should satisfy a few basic requirements:

1. It should be independent of the representation used, e.g. the ordering or number of modes.
2. It should be minimized only when comparing formally indistinguishable fields.
3. It should reduce to a metric already in widespread use when applied to coherent wavefields.

The first condition is especially important because while the multi-mode expansion is the most widespread, it is not the only way to treat partial coherence [172]. This, together with the second condition, implies that such a metric should have a definition in terms of the density matrix $\boldsymbol{\rho} = \rho(\vec{r}, \vec{r}')$ [43], also known as the mutual coherence function [173, 174]. This is because $\boldsymbol{\rho}$ is the most general description of a monochromatic partially coherent wavefield, which all other representations can be rephrased in terms of.

However, to find a metric which meets condition 3 it will be helpful to link the density matrix representation to one defined explicitly in terms of coherent wavefields. We do this by exploring a common model of partial coherence, as the consequence of averaging over a time-varying coherent wavefield. Specifically, integrating the diffracted intensity from a time-varying wavefield $|\psi(t)\rangle$ over a period of time T is equivalent to simulating that diffraction using the density matrix [43]

$$\boldsymbol{\rho} = \int_0^T \boldsymbol{\rho}(t) = \int_0^T dt\, |\psi(t)\rangle \langle\psi(t)|\,. \tag{B.1}$$

Crucially, the time-dependent representation still describes a coherent wavefield, but with an extra dimension (time). This suggests that we might be able to generalize metrics from coherent to partially coherent wavefields by applying them in the time-dependent representation.

Initially this seems like a fool's errand because each density matrix $\boldsymbol{\rho}$ corresponds to infinitely many wavefields $|\psi(t)\rangle$. However, there will still be a unique minimum, corresponding to the best case which is consistent with the known information. In the case of the MSE calculated between two wavefields $|\psi_1(t)\rangle$ and $|\psi_2(t)\rangle$, we can set up the following minimization problem:

$$\begin{aligned} \min_{\psi_1,\psi_2} \quad & \langle\psi_1|\psi_1\rangle + \langle\psi_2|\psi_2\rangle - 2\,\mathrm{Re}\left[\langle\psi_1|\psi_2\rangle\right] \\ \text{s.t.} \quad & \boldsymbol{\rho_i} = \int dt\, |\psi_i(t)\rangle \langle\psi_i(t)|\,, \quad i \in \{1, 2\}, \end{aligned} \tag{B.2}$$

where the time-dependence is suppressed in the inner products to indicate that they integrate over time as well as the spatial/pixel dimensions. The first two terms are the traces (Tr) of $\boldsymbol{\rho_1}$ and $\boldsymbol{\rho_2}$ respectively, and do not depend on the choice of $|\psi_i(t)\rangle$. Further, because $|\psi_1(t)\rangle$ and $|\psi_2(t)\rangle$ have a global phase degree of freedom, their overlap can be chosen real and non-negative. Therefore, (B.2) simplifies to

$$\mathrm{Tr}(\boldsymbol{\rho_1}) + \mathrm{Tr}(\boldsymbol{\rho_2}) - 2\max_{\psi_1,\psi_2} |\langle\psi_1|\psi_2\rangle| \tag{B.3}$$

With the maximization problem operating under the same constraints as (B.2). Happily, this problem has a well known solution. In the language of quantum states, $|\psi_1(t)\rangle$ and $|\psi_2(t)\rangle$ are purifications of the density matrices $\boldsymbol{\rho_1}$ and $\boldsymbol{\rho_2}$. A classic result in quantum information is that the maximum overlap between the purifications of a pair of density matrices is equal to their square-root fidelity ($\mathbb{F}$) [175, 176]:

$$\mathbb{F}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2) = \operatorname{Tr}\left(\sqrt{\sqrt{\boldsymbol{\rho}_1}\boldsymbol{\rho}_2\sqrt{\boldsymbol{\rho}_1}}\right) = \max_{\psi_1,\psi_2} |\langle\psi_1|\psi_2\rangle| \tag{B.4}$$

Consequently, the solution to (B.2) (which we define as the partially coherent MSE, PCMSE) can be written directly in terms of the density matrices:

$$\operatorname{PCMSE}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2) = \operatorname{Tr}(\boldsymbol{\rho}_1) + \operatorname{Tr}(\boldsymbol{\rho}_2) - 2\mathbb{F}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2). \tag{B.5}$$

As required, this reduces to the MSE maximized over the phase degree of freedom [157] when $\boldsymbol{\rho}_1$ and $\boldsymbol{\rho}_2$ represent pure states, and it depends only on the information in the density matrix representation. But, does it satisfy condition 2? It is straightforward to show that

$$0 \leq \mathbb{F}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2) \leq \sqrt{\operatorname{Tr}(\boldsymbol{\rho}_1)\operatorname{Tr}(\boldsymbol{\rho}_2)} \tag{B.6}$$

with the upper equality achieved only when $\frac{\boldsymbol{\rho}_1}{\operatorname{Tr}(\boldsymbol{\rho}_1)} = \frac{\boldsymbol{\rho}_2}{\operatorname{Tr}(\boldsymbol{\rho}_2)}$ and the lower equality achieved only when $\boldsymbol{\rho}_1\boldsymbol{\rho}_2 = 0$ (see supplement 1.1). Consequently,

$$0 \leq \operatorname{PCMSE}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2) \leq \operatorname{Tr}(\boldsymbol{\rho}_1) + \operatorname{Tr}(\boldsymbol{\rho}_2) \tag{B.7}$$

with the lower equality achieved only when $\boldsymbol{\rho}_1 = \boldsymbol{\rho}_2$. Not only does this metric clearly satisfy condition 2, its bounds mirror those of the standard MSE. Consequently, we can generalize the normalized MSE by normalizing to $\operatorname{Tr}(\boldsymbol{\rho})$. It also is worth noting that the minimization of the normalized PCMSE over a global amplitude degree of freedom has an especially simple definition, which parallels that derived in [157]:

$$\min_a \left[\frac{\operatorname{PCMSE}(\boldsymbol{\rho}_1, a\boldsymbol{\rho}_2)}{\operatorname{Tr}(\boldsymbol{\rho}_1)}\right] = 1 - \frac{\mathbb{F}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2)^2}{\operatorname{Tr}(\boldsymbol{\rho}_1)\operatorname{Tr}(\boldsymbol{\rho}_2)}. \tag{B.8}$$

This form, which is just 1 minus the fidelity of the normalized density matrices, is very useful in ptychography where the intensity ratio of two probe reconstructions cannot always be determined.

The PCMSE therefore seems like an ideal error metric, but based on the discussion so far it remains completely impractical to calculate. This is because the density matrices are usually so large that it is not even possible to hold them in memory, much less calculate their square roots. To make it practical, we need a way to calculate the PCMSE directly from the multi-mode representation. In this representation, an $N \times M$ matrix $\boldsymbol{\psi}$ is stored, such that

$$\boldsymbol{\rho} = \boldsymbol{\psi}\boldsymbol{\psi}^{\dagger}. \tag{B.9}$$

N is the number of pixels in the image, and M is the number of modes, i.e. the assumed maximum rank of $\boldsymbol{\rho}$. There are, unsurprisingly, strong connections between this expression and the time-dependent breakdown in (B.1). To start, we note that $\mathrm{Tr}(\boldsymbol{\rho}) = \mathrm{Tr}\big(\boldsymbol{\psi}^{\dagger}\boldsymbol{\psi}\big)$, the sum of the integrated intensities of each mode. Perhaps more surprisingly, the square root fidelity is also cheap to calculate,

$$\mathbb{F}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2) = ||\boldsymbol{\psi}_1^{\dagger}\boldsymbol{\psi}_2||_* \tag{B.10}$$

where $||||_*$ is the nuclear norm, i.e. the sum of the singular values (see supplement 1.2). This form only requires a singular value decomposition of an $M_1 \times M_2$ matrix. As a result, calculating the PCMSE between two mixed state reconstructions can be done in the same amount of time it would take to calculate the MSE between $M_1 M_2$ pairs of coherent wavefields. Finally, because this form does not rely on any properties of $\boldsymbol{\psi}$ other than (B.9) it still satisfies condition 1 - the two expansions need not even have the same number of modes!

The derivations above constitute a complete framework for reducing the difference between partially coherent wavefields to a single number. However, there is also a need for metrics which can empirically assess the resolution of experimental results, i.e. an extension of the FRC. In practice, FRCs are rarely used to characterize retrieved probes, but it is our belief that as ptychography becomes a standard diagnostic of illumination sources the need for frequency-dependent analysis of the reliability of retrieved probe functions will grow.

The coherent FRC is calculated by splitting each of two images into a set of concentric rings in Fourier space, and calculating the correlation coefficient between each pair of rings. This curve, as a function of spatial frequency, is used to determine the length scale below which the images can no longer be considered reliable [101, 158]. We formalize this by defining a projection operator for each ring spanned by frequencies $k_1 < k_2$ in Fourier space:

$$\boldsymbol{P}_{k_1,k_2} = \int d^2k \, |\vec{k}\rangle \, \langle\vec{k}| \, \theta(|\vec{k}| - k_1)\theta(k_2 - |\vec{k}|). \tag{B.11}$$

Here $|\vec{k}\rangle$ is a complex exponential at frequency $\vec{k}$ and θ is the Heaviside step function. We can then define the FRC for coherent fields as

$$\mathrm{FRC}(|\psi_1\rangle, |\psi_2\rangle; k_1, k_2) = \frac{\langle\psi_1| \, \boldsymbol{P} \, |\psi_2\rangle}{\sqrt{\langle\psi_1| \, \boldsymbol{P} \, |\psi_1\rangle \, \langle\psi_2| \, \boldsymbol{P} \, |\psi_2\rangle}}, \tag{B.12}$$

where the subscript k_1, k_2 is dropped from the projection operators for compactness. We should quickly note that the FRC, as originally envisioned, is a real function of real-valued fields which can be positive or negative. However, for complex-valued fields, it is complex. Universally, what is reported in the ptychography literature is the magnitude of the FRC, i.e. the maximum over a phase degree of freedom.

Generalizing this definition to density matrices in the vein of our earlier approach means maximizing (B.12) over the space of all consistent purifications. This results in a simple expression for the Partially Coherent FRC (PCFRC) as the normalized Fidelity of the projected density matrices:

$$\mathrm{PCFRC}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2; k_1, k_2) = \frac{\mathbb{F}(\boldsymbol{P}\boldsymbol{\rho}_1\boldsymbol{P}, \boldsymbol{P}\boldsymbol{\rho}_2\boldsymbol{P})}{\sqrt{\mathrm{Tr}(\boldsymbol{P}\boldsymbol{\rho}_1\boldsymbol{P}) \, \mathrm{Tr}(\boldsymbol{P}\boldsymbol{\rho}_2\boldsymbol{P})}}. \tag{B.13}$$

This is not surprising because the FRC is essentially a normalized overlap, and the Fidelity is just the generalization of the overlap to mixed states. It is evident from (B.6) that this expression is constrained between 0 and 1, just like the standard FRC. In addition, because the projection operators are easy to implement in the multi-mode framework, the PCFRC remains cheap to calculate, equivalent to the

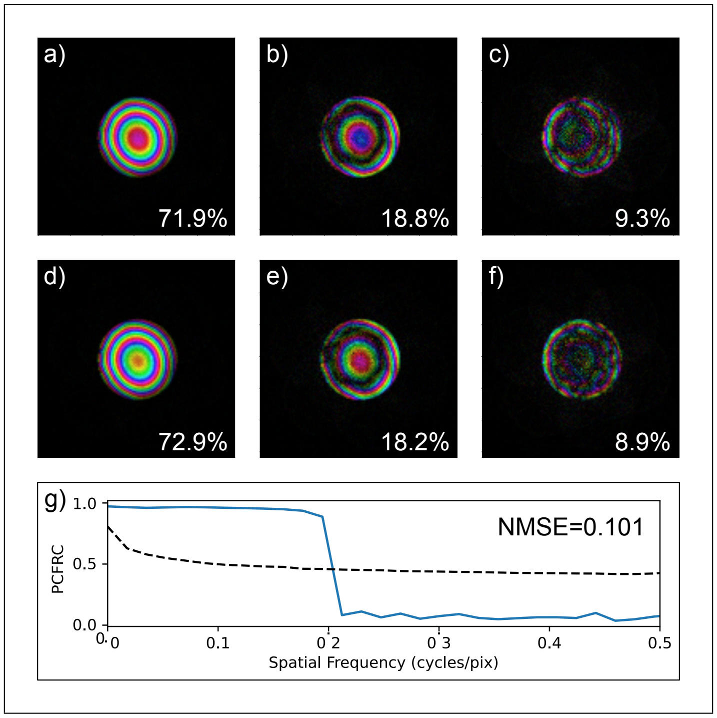

Figure B-1: Example use of the PCFRC and PCMSE to estimate the reliability of a probe reconstruction. (a-c) and (d-f) are two three-mode probe reconstructions from separate halves of a reference dataset from Chen et al. [160]. The probes are shown in Fourier space to emphasize the connection with the PCFRC. (g) contains the PCFRC between the two reconstructed partially coherent probes, compared to a half-bit threshold [101].

cost of calculating $M_1 M_2$ coherent FRCs.

Finally, to demonstrate how these metrics perform in a typical case, in Figure B-1 we show a pair of reconstructed probes from two halves of the electron ptychography dataset made publicly available by Chen et al. [160]. The reconstructions remain nonzero up to a maximum frequency defined by the condenser's aperture. This is reflected in the PCFRC, which falls sharply from a high value near 1 within the aperture to a low value near zero beyond it. The amplitude-minimized normalized PCMSE is 0.101, indicating that although the reconstructions are visually similar, a wavefield containing at least 10.1% of the power in the first probe would be needed to map it onto the second probe.

In sum, we have generalized the MSE and FRC from coherent images to the space of partially coherent wavefields by finding an analogy with the quantum state fidelity.

The results reduce to simple expressions which are computationally cheap when the density matrices are stored in the low rank multi-mode approximation. These metrics address a major need in the coherent imaging community for quantitative analysis of reconstructed partially coherent wavefields, and we hope they will find widespread use as the computational study of partial coherence continues to expand.

B.3 Supplement

B.3.1 Bounds on Fidelity and PCMSE

For any two matrices with trace 1, the bound

$$0 \leq \mathbb{F}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2) \leq 1$$

is well established [176], with equality only when $\boldsymbol{\rho}_1 = \boldsymbol{\rho}_2$. For any matrix $\boldsymbol{\rho}$ with arbitrary trace, the matrix $\frac{\boldsymbol{\rho}}{\mathrm{Tr}(\boldsymbol{\rho})}$ has trace 1, so

$$0 \leq \mathbb{F}\left(\frac{\boldsymbol{\rho}_1}{\mathrm{Tr}(\boldsymbol{\rho}_1)}, \frac{\boldsymbol{\rho}_2}{\mathrm{Tr}(\boldsymbol{\rho}_2)}\right) \leq 1.$$

From the definition of the square-root fidelity $\mathbb{F}$, it is clear that

$$\mathbb{F}(a\boldsymbol{\rho}_1, b\boldsymbol{\rho}_2) = \sqrt{ab}\, \mathbb{F}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2)$$

for scalar a, b. Therefore, we find

$$0 \leq \mathbb{F}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2) \leq \sqrt{\mathrm{Tr}(\boldsymbol{\rho}_1)\, \mathrm{Tr}(\boldsymbol{\rho}_2)}.$$

Plugging this into the expression for the PCMSE, we find

$$\mathrm{Tr}(\boldsymbol{\rho}_1) + \mathrm{Tr}(\boldsymbol{\rho}_2) \geq \mathrm{PCMSE}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2) \geq \mathrm{Tr}(\boldsymbol{\rho}_1) + \mathrm{Tr}(\boldsymbol{\rho}_2) - 2 * \sqrt{\mathrm{Tr}(\boldsymbol{\rho}_1)\, \mathrm{Tr}(\boldsymbol{\rho}_2)}.$$

Finally, as a consequence of the AM-GM inequality

$$\mathrm{Tr}(\boldsymbol{\rho}_1) + \mathrm{Tr}(\boldsymbol{\rho}_2) \geq \mathrm{PCMSE}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2) \geq 0,$$

with the lower bound holding only when $\mathrm{Tr}(\boldsymbol{\rho}_1) = \mathrm{Tr}(\boldsymbol{\rho}_2)$ and $\frac{\boldsymbol{\rho_1}}{\mathrm{Tr}(\boldsymbol{\rho_1})} = \frac{\boldsymbol{\rho_2}}{\mathrm{Tr}(\boldsymbol{\rho_2})}$, i.e. $\boldsymbol{\rho_1} = \boldsymbol{\rho_2}$

B.3.2 Multi-mode expression for Fidelity

We expand the definition of fidelity using the multi-mode representation of $\boldsymbol{\rho}_2$:

$$\mathbb{F}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2) = \mathrm{Tr}\left(\sqrt{\sqrt{\boldsymbol{\rho}_1}\boldsymbol{\psi}_2\boldsymbol{\psi}_2^\dagger\sqrt{\boldsymbol{\rho}_1}}\right)$$

Noting that $\sqrt{\boldsymbol{\rho}_1}$ is Hermetian, we can see that this expression is equivalent to

$$\mathbb{F}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2) = ||\sqrt{\boldsymbol{\rho}_1}\boldsymbol{\psi}_2||_*$$

by the definition of the nuclear norm. However, it is also the case that

$$\sqrt{\boldsymbol{\rho}_1}\,(\sqrt{\boldsymbol{\rho}_1})^\dagger = \boldsymbol{\rho}_1 = \boldsymbol{\psi}_1\boldsymbol{\psi}_1^\dagger$$

Therefore, there must exist a semi-unitary matrix $\boldsymbol{U}$ satisfying $\boldsymbol{U}^\dagger\boldsymbol{U} = \boldsymbol{I}$ such that

$$(\sqrt{\boldsymbol{\rho}_1}) = (\sqrt{\boldsymbol{\rho}_1})^\dagger = U^\dagger\boldsymbol{\psi}_1^\dagger,$$

from [177] Theorem (7.3.11). Therefore,

$$\mathbb{F}(\boldsymbol{\rho}_1, \boldsymbol{\rho}_2) = ||U^\dagger\boldsymbol{\psi}_1^\dagger\boldsymbol{\psi}_2||_* = ||\boldsymbol{\psi}_1^\dagger\boldsymbol{\psi}_2||_*,$$

due to the unitary invariance of the nuclear norm.

B.3.3 Comparing reconstructions with varying numbers of modes

Because the PCMSE and PCFRC allow for comparisons between reconstructions of probe functions with different numbers of modes, we confirmed that comparisons

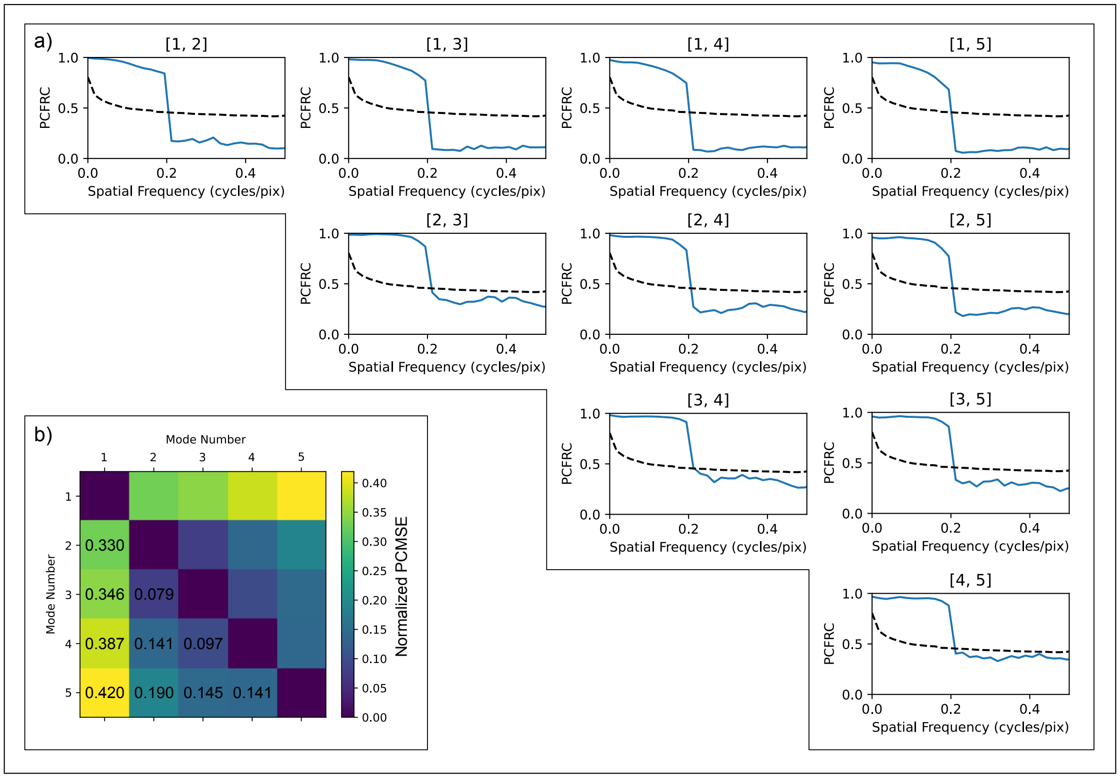

Mode Number	1	2	3	4	5
1					
2	0.330				
3	0.346	0.079			
4	0.387	0.141	0.097		
5	0.420	0.190	0.145	0.141	

Figure B-2: Comparisons between reconstructions with varying numbers of modes. (a) shows comparisons between each pair of reconstructions using the PCFRC. (b) compares the same reconstructions using the amplitude-minimized normalized PCMSE.

perform as expected on typical experimental data. Therefore, we performed 5 reconstructions on the full dataset from reference [160] which was used to generate Figure 1 of the main paper. The reconstructions were run with a varying number of modes from 1 to 5.

In Figure B-2, we compare the partially coherent probes resulting from each reconstruction using the amplitude-minimized normalized PCMSE and the PCFRC. The results conform to our expectations. First, the introduction of even a single additional mode causes a dramatic change in the reconstructed probe as measured by PCMSE. This is expected because, in a two-mode reconstruction, the subdominant mode contains significant power (23.5%). However, the mode-power breakdown hides the fact that the dominant mode in the two-mode reconstruction also has a different spatial structure from the single-mode result. This difference is captured by the normalized PCMSE of 0.330, which is higher than the power fraction in the subdominant mode due to the differing spatial structure of the two results.

Continuing on to higher modes, we find that the relative error introduced by each additional mode decreases until the fourth mode is added, at which point it begins to increase again. This is likely caused by higher order modes accumulating artifacts. This effect is again hidden by the orthogonalized mode power breakdown. For example, even though the fourth mode in a four-mode reconstruction only contains 5.1% of the power, the normalized PCMSE between the three- and four-mode reconstructions is nearly double at 0.97. This again indicates that the structure of the lower modes has been modified to accommodate the additional mode, an effect which the PCMSE.

Finally, looking at the PCFRCs reveals more information about the spatial structure of these changes. Notably, the PCFRCs calculated between the one-mode and remaining results show a dip near the edge of the probe's Fourier space representation. This region corresponds to the the edge of the aperture in the electron microscope's condenser. The dip is therefore quantifying our expectation that the effects of incoherence are magnified near the edges of the condenser aperture, a well known phenomenon. As the number of modes increases, the edge of the PCFRC curve becomes more defined. This indicates that although minor changes are occurring as each mode is added, the essential structure of the probe has been sufficiently captured.

Bibliography

1. Galilei, G. *Sidereus Nuncius* (Thomas Baglioni, Republic of Venice, Mar. 13, 1610).
2. Van Leeuwenhoek, A. Observations, Communicated to the Publisher by Mr. Antony van Leewenhoeck, in a Dutch Letter of the 9th Octob. 1676. Here English'd: Concerning Little Animals by Him Observed in Rain-Well-Sea- and Snow Water; as Also in Water Wherein Pepper Had Lain Infused. *Philosophical Transactions of the Royal Society of London* **12,** 821–831. doi:`10.1098/rstl.1677.0003` (Mar. 25, 1677).
3. Franklin, R. E. & Gosling, R. G. The Structure of Sodium Thymonucleate Fibres. I. The Influence of Water Content. *Acta Crystallographica* **6,** 673–677. doi:`10.1107/S0365110X53001939` (Sept. 10, 1953).
4. Callaway, E. The Revolution Will Not Be Crystallized: A New Method Sweeps through Structural Biology. *Nature* **525,** 172–174. doi:`10.1038/525172a` (Sept. 10, 2015).
5. Withers, P. J., Bouman, C., Carmignato, S., Cnudde, V., Grimaldi, D., Hagen, C. K., Maire, E., Manley, M., Du Plessis, A. & Stock, S. R. X-ray computed tomography. *Nature Reviews Methods Primers* **1,** 18. doi:`10.1038/s43586-021-00015-4` (Feb. 25, 2021).
6. Takahashi, Y., Nishino, Y., Tsutsumi, R., Zettsu, N., Matsubara, E., Yamauchi, K. & Ishikawa, T. High-Resolution Projection Image Reconstruction of Thick Objects by Hard x-Ray Diffraction Microscopy. *Physical Review B* **82,** 214102. doi:`10.1103/PhysRevB.82.214102` (Dec. 2, 2010).
7. Jiang, Y., Chen, Z., Han, Y., Deb, P., Gao, H., Xie, S., Purohit, P., Tate, M. W., Park, J., Gruner, S. M., Elser, V. & Muller, D. A. Electron Ptychography of 2D Materials to Deep Sub-Ångström Resolution. *Nature* **559,** 343–349. doi:`10.1038/s41586-018-0298-5` (July 2018).
8. Holler, M., Odstrcil, M., Guizar-Sicairos, M., Lebugle, M., Müller, E., Finizio, S., Tinti, G., David, C., Zusman, J., Unglaub, W., Bunk, O., Raabe, J., Levi, A. F. J. & Aeppli, G. Three-Dimensional Imaging of Integrated Circuits with Macro- to Nanoscale Zoom. *Nature Electronics* **2,** 464–470. doi:`10.1038/s41928-019-0309-z` (Oct. 2019).

9. Csernica, P. M., Kalirai, S. S., Gent, W. E., Lim, K., Yu, Y.-S., Liu, Y., Ahn, S.-J., Kaeli, E., Xu, X., Stone, K. H., Marshall, A. F., Sinclair, R., Shapiro, D. A., Toney, M. F. & Chueh, W. C. Persistent and Partially Mobile Oxygen Vacancies in Li-rich Layered Oxides. *Nature Energy* **6,** 642–652. doi:`10.1038/s41560-021-00832-7` (June 14, 2021).
10. Fink, J., Schierle, E., Weschke, E. & Geck, J. Resonant Elastic Soft X-Ray Scattering. *Reports on Progress in Physics.* doi:`10.1088/0034-4885/76/5/056502` (2013).
11. Fischer, P., Schütz, G., Schmahl, G., Guttmann, P. & Raasch, D. Imaging of Magnetic Domains with the X-ray Microscope at BESSY Using X-ray Magnetic Circular Dichroism. *Zeitschrift für Physik B Condensed Matter* **101,** 313–316. doi:`10.1007/s002570050214` (Dec. 1997).
12. Li, J., Pelliciari, J., Mazzoli, C., Catalano, S., Simmons, F., Sadowski, J. T., Levitan, A., Gibert, M., Carlson, E., Triscone, J.-M., Wilkins, S. & Comin, R. Scale-Invariant Magnetic Textures in the Strongly Correlated Oxide NdNiO3. *Nature Communications* **10,** 4568. doi:`10.1038/s41467-019-12502-0` (Dec. 2019).
13. Born, M., Wolf, E., Bhatia, A. B., Clemmow, P. C., Gabor, D., Stokes, A. R., Taylor, A. M., Wayman, P. A. & Wilcock, W. L. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light* 7th ed. doi:`10.1017/CBO9781139644181` (Cambridge University Press, Oct. 13, 1999).
14. Sayre, D. Some Implications of a Theorem Due to Shannon. *Acta Crystallographica* **5,** 843–843. doi:`10.1107/S0365110X52002276` (1952).
15. Shannon, C. Communication in the Presence of Noise. *Proceedings of the IRE* **37,** 10–21. doi:`10.1109/JRPROC.1949.232969` (Jan. 1949).
16. Bates, R. Fourier Phase Problems Are Uniquely Solvable in More than One Dimension I: Underlying Theory. *Optik* **61,** 247–262 (1982).
17. Candes, E., Romberg, J. & Tao, T. Robust Uncertainty Principles: Exact Signal Reconstruction from Highly Incomplete Frequency Information. *IEEE Transactions on Information Theory* **52,** 489–509. doi:`10.1109/TIT.2005.862083` (Feb. 2006).
18. Fienup, J. R. Phase Retrieval Algorithms: A Comparison. *Applied Optics* **21,** 2758–2769. doi:`10.1364/AO.21.002758` (1982).
19. Fienup, J. R. Reconstruction of a Complex-Valued Object from the Modulus of Its Fourier Transform Using a Support Constraint. *Journal of the Optical Society of America A* **4,** 118. doi:`10.1364/JOSAA.4.000118` (1987).
20. Miao, J., Charalambous, P., Kirz, J. & Sayre, D. Extending the Methodology of X-ray Crystallography to Allow Imaging of Micrometre-Sized Non-Crystalline Specimens. *Nature* **400,** 342–344. doi:`10.1038/22498` (July 1999).

21. Seibert, M. M., Ekeberg, T., Maia, F. R. N. C., Svenda, M., Andreasson, J., Joensson, O., Odic, D., Iwan, B., Rocker, A., Westphal, D., Hantke, M., DePonte, D. P., Barty, A., Schulz, J., Gumprecht, L., Coppola, N., Aquila, A., Liang, M., White, T. A., Martin, A., Caleman, C., Stern, S., Abergel, C., Seltzer, V., Claverie, J.-M., Bostedt, C., Bozek, J. D., Boutet, S., Miahnahri, A. A., Messerschmidt, M., Krzywinski, J., Williams, G., Hodgson, K. O., Bogan, M. J., Hampton, C. Y., Sierra, R. G., Starodub, D., Andersson, I., Bajt, S., Barthelmess, M., Spence, J. C. H., Fromme, P., Weierstall, U., Kirian, R., Hunter, M., Doak, R. B., Marchesini, S., Hau-Riege, S. P., Frank, M., Shoeman, R. L., Lomb, L., Epp, S. W., Hartmann, R., Rolles, D., Rudenko, A., Schmidt, C., Foucar, L., Kimmel, N., Holl, P., Rudek, B., Erk, B., Hoemke, A., Reich, C., Pietschner, D., Weidenspointner, G., Strueder, L., Hauser, G., Gorke, H., Ullrich, J., Schlichting, I., Herrmann, S., Schaller, G., Schopper, F., Soltau, H., Kuehnel, K.-U., Andritschke, R., Schroeter, C.-D., Krasniqi, F., Bott, M., Schorb, S., Rupp, D., Adolph, M., Gorkhover, T., Hirsemann, H., Potdevin, G., Graafsma, H., Nilsson, B., Chapman, H. N. & Hajdu, J. Single Mimivirus Particles Intercepted and Imaged with an X-ray Laser. *Nature* **470,** 78–U86. doi:`10.1038/nature09748` (2011).

22. Takahashi, Y., Suzuki, A., Zettsu, N., Oroguchi, T., Takayama, Y., Sekiguchi, Y., Kobayashi, A., Yamamoto, M. & Nakasako, M. Coherent Diffraction Imaging Analysis of Shape-Controlled Nanoparticles with Focused Hard X-ray Free-Electron Laser Pulses. *Nano Letters* **13,** 6028–6032. doi:`10.1021/nl403247x` (Dec. 11, 2013).

23. McNulty, I., Kirz, J., Jacobsen, C., Anderson, E. H., Howells, M. R. & Kern, D. P. High-Resolution Imaging by Fourier Transform X-ray Holography. *Science* **256,** 1009–1012. doi:`10.1126/science.256.5059.1009` (May 15, 1992).

24. Müller, L., Schleitzer, S., Gutt, C., Pfau, B., Schaffert, S., Geilhufe, J., von Korff Schmising, C., Schneider, M., Günther, C. M., Büttner, F., Capotondi, F., Pedersoli, E., Düsterer, S., Redlin, H., Al-Shemmary, A., Treusch, R., Bach, J., Frömter, R., Vodungbo, B., Gautier, J., Zeitoun, P., Popescu, H., Lopez-Flores, V., Beaulieu, N., Sirotti, F., Jaouen, N., Malinowski, G., Tudu, B., Li, K., Lüning, J., Oepen, H. P., Kiskinova, M., Eisebitt, S. & Grübel, G. Ultrafast Dynamics of Magnetic Domain Structures Probed by Coherent Free-Electron Laser Light. *Synchrotron Radiation News* **26,** 27–32. doi:`10.1080/08940886.2013.850384` (Nov. 1, 2013).

25. Willems, F., von Korff Schmising, C., Weder, D., Günther, C. M., Schneider, M., Pfau, B., Meise, S., Guehrs, E., Geilhufe, J., Merhe, A. E. D., Jal, E., Vodungbo, B., Lüning, J., Mahieu, B., Capotondi, F., Pedersoli, E., Gauthier, D., Manfredda, M. & Eisebitt, S. Multi-Color Imaging of Magnetic Co/Pt Heterostructures. *Structural Dynamics* **4,** 014301. doi:`10.1063/1.4976004` (Jan. 2017).

26. Hoppe, W. Beugung Im Inhomogenen Primärstrahlwellenfeld. I. Prinzip Einer Phasenmessung von Elektronenbeungungsinterferenzen. *Acta Crystallographica Section A* **25,** 495–501. doi:`10.1107/S0567739469001045` (1969).

27. Hoppe, W. & Strube, G. Beugung in Inhomogenen Primärstrahlenwellenfeld. II. Lichtoptische Analogieversuche Zur Phasenmessung von Gitterinterferenzen. *Acta Crystallographica Section A* **25,** 502–507. doi:`10.1107/S0567739469001057` (1969).

28. Hoppe, W. Beugung Im Inhomogenen Primärstrahlwellenfeld. III. Amplituden- Und Phasenbestimmung Bei Unperiodischen Objekten. *Acta Crystallographica Section A* **25,** 508–514. doi:`10.1107/S0567739469001069` (1969).

29. Hegerl, R. & Hoppe, W. Dynamische Theorie der Kristallstrukturanalyse durch Elektronenbeugung im inhomogenen Primärstrahlwellenfeld. *Berichte der Bunsengesellschaft für physikalische Chemie* **74,** 1148–1154. doi:`10.1002/bbpc.19700741112` (Nov. 1970).

30. Hawkes, P. W. Is the STEM a Ptychograph? *Ultramicroscopy* **9,** 27–30. doi:`10.1016/0304-3991(82)90225-X` (1982).

31. Hoppe, W. Trace Structure Analysis, Ptychography, Phase Tomography. *Ultramicroscopy* **10,** 187–198 (1982).

32. Rodenburg, J. & Bates, R. H. T. The Theory of Super-Resolution Electron Microscopy via Wigner-distribution Deconvolution. *Philosophical Transactions of the Royal Society of London. Series A: Physical and Engineering Sciences* **339,** 521–553. doi:`10.1098/rsta.1992.0050` (June 15, 1992).

33. McCallum, B. C. & Rodenburg, J. M. Two-Dimensional Demonstration of Wigner Phase-Retrieval Microscopy in the STEM Configuration. *Ultramicroscopy* **45,** 371–380. doi:`10.1016/0304-3991(92)90149-E` (1992).

34. Nellist, P. D., McCallum, B. C. & Rodenburg, J. M. Resolution beyond the 'information Limit' in Transmission Electron Microscopy. *Nature* **374,** 630–632. doi:`10.1038/374630a0` (Apr. 1995).

35. Faulkner, H. M. & Rodenburg, J. M. Movable Aperture Lensless Transmission Microscopy: A Novel Phase Retrieval Algorithm. *Physical Review Letters* **93,** 023903–1. doi:`10.1103/PhysRevLett.93.023903` (2004).

36. Guizar-Sicairos, M. & Fienup, J. R. Phase Retrieval with Transverse Translation Diversity: A Nonlinear Optimization Approach. *Optics Express* **16,** 7264–7278. doi:`10.1364/OE.16.007264` (2008).

37. Thibault, P., Dierolf, M., Menzel, A., Bunk, O., David, C. & Pfeiffer, F. High-Resolution Scanning X-ray Diffraction Microscopy. *Science* **321,** 379–382. doi:`10.1126/science.1158573` (July 18, 2008).

38. Maiden, A. M. & Rodenburg, J. M. An Improved Ptychographical Phase Retrieval Algorithm for Diffractive Imaging. *Ultramicroscopy* **109,** 1256–1262. doi:`10.1016/j.ultramic.2009.05.012` (2009).

39. Thibault, P., Dierolf, M., Bunk, O., Menzel, A. & Pfeiffer, F. Probe Retrieval in Ptychographic Coherent Diffractive Imaging. *Ultramicroscopy* **109,** 338–343. doi:`10.1016/j.ultramic.2008.12.011` (2009).

40. Thibault, P. & Guizar-Sicairos, M. Maximum-Likelihood Refinement for Coherent Diffractive Imaging. *New Journal of Physics* **14.** doi:`10.1088/1367-2630/14/6/063004` (2012).

41. Maiden, A. M., Humphry, M. J., Sarahan, M. C., Kraus, B. & Rodenburg, J. M. An Annealing Algorithm to Correct Positioning Errors in Ptychography. *Ultramicroscopy* **120,** 64–72. doi:`10.1016/j.ultramic.2012.06.001` (2012).

42. Tripathi, A., McNulty, I. & Shpyrko, O. G. Ptychographic Overlap Constraint Errors and the Limits of Their Numerical Recovery Using Conjugate Gradient Descent Methods. *Optics Express* **22,** 1452. doi:`10.1364/OE.22.001452` (2014).

43. Thibault, P. & Menzel, A. Reconstructing State Mixtures from Diffraction Measurements. *Nature* **494,** 68–71. doi:`10.1038/nature11806` (2013).

44. Griewank, A. & Walther, A. *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation, Second Edition* Second. doi:`10.1137/1.9780898717761` (Society for Industrial and Applied Mathematics, Jan. 2008).

45. Rall, L. B. *Automatic Differentiation: Techniques and Applications Lecture Notes in Computer Science* **120** (Springer-Verlag Springer e-books, Berlin, Heidelberg, 1981).

46. Ghosh, S., Nashed, Y. S. G., Cossairt, O. & Katsaggelos, A. *ADP: Automatic Differentiation Ptychography* in *2018 IEEE International Conference on Computational Photography (ICCP)* 2018 IEEE International Conference on Computational Photography (ICCP) (IEEE, Pittsburgh, PA, May 2018), 1–10. doi:`10.1109/ICCPHOT.2018.8368470`.

47. Kandel, S., Maddali, S., Allain, M., Hruszkewycz, S. O., Jacobsen, C. & Nashed, Y. S. G. Using Automatic Differentiation as a General Framework for Ptychographic Reconstruction. *Optics Express* **27,** 18653–18672. doi:`10.1364/OE.27.018653` (June 24, 2019).

48. Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization. Version 9. doi:`10.48550/ARXIV.1412.6980` (2014).

49. Liu, D. C. & Nocedal, J. On the Limited Memory BFGS Method for Large Scale Optimization. *Mathematical Programming* **45,** 503–528. doi:`10.1007/BF01589116` (Aug. 1989).

50. Kandel, S., Maddali, S., Nashed, Y. S. G., Hruszkewycz, S. O., Jacobsen, C. & Allain, M. *A Matrix-Free Levenberg-Marquardt Algorithm for Efficient Ptychographic Phase Retrieval* Feb. 26, 2021.

51. Rodenburg, J. M. & Faulkner, H. M. A Phase Retrieval Algorithm for Shifting Illumination. *Applied Physics Letters* **85,** 4795–4797. doi:`10.1063/1.1823034` (2004).

52. Enders, B., Dierolf, M., Cloetens, P., Stockmar, M., Pfeiffer, F. & Thibault, P. Ptychography with Broad-Bandwidth Radiation. *Applied Physics Letters* **104,** 171104. doi:`10.1063/1.4874304` (Apr. 28, 2014).

53. Odstrčil, M., Menzel, A. & Guizar-Sicairos, M. Iterative Least-Squares Solver for Generalized Maximum-Likelihood Ptychography. *Optics Express.* doi:`10.1364/OE.26.003108` (2018).

54. Odstrcil, M., Baksh, P., Boden, S. A., Card, R., Chad, J. E., Frey, J. G. & Brocklesby, W. S. Ptychographic Coherent Diffractive Imaging with Orthogonal Probe Relaxation. *Optics Express* **24,** 8360. doi:`10.1364/OE.24.008360` (2016).

55. Kharitonov, K., Mehrjoo, M., Ruiz-Lopez, M., Keitel, B., Kreis, S., Seyrich, M., Pop, M. & Plönjes, E. Flexible Ptychography Platform to Expand the Potential of Imaging at Free Electron Lasers. *Optics Express* **29,** 22345. doi:`10.1364/OE.426931` (July 5, 2021).

56. Fannjiang, A. Raster Grid Pathology and the Cure. *Multiscale Modeling & Simulation* **17,** 973–995. doi:`10.1137/18M1227354` (Jan. 2019).

57. Zhang, F., Peterson, I., Vila-Comamala, J., Diaz, A., Berenguer, F., Bean, R., Chen, B., Menzel, A., Robinson, I. K. & Rodenburg, J. M. Translation Position Determination in Ptychographic Coherent Diffraction Imaging. *Optics Express* **21,** 13592. doi:`10.1364/OE.21.013592` (June 3, 2013).

58. Maiden, A. M., Humphry, M. J., Zhang, F. & Rodenburg, J. M. Superresolution Imaging via Ptychography. *Journal of the Optical Society of America. A, Optics, image science, and vision* **28,** 604–12. doi:`10.1364/JOSAA.28.000604` (2011).

59. Seaberg, M. D., Zhang, B., Gardner, D. F., Shanblatt, E. R., Murnane, M. M., Kapteyn, H. C. & Adams, D. E. Tabletop Nanometer Extreme Ultraviolet Imaging in an Extended Reflection Mode Using Coherent Fresnel Ptychography. *Optica* **1,** 39. doi:`10.1364/OPTICA.1.000039` (July 22, 2014).

60. Porter, C. L., Tanksalvala, M., Gerrity, M., Miley, G., Zhang, X., Bevis, C., Shanblatt, E., Karl, R., Murnane, M. M., Adams, D. E. & Kapteyn, H. C. General-Purpose, Wide Field-of-View Reflection Imaging with a Tabletop 13 Nm Light Source. *Optica* **4,** 1552. doi:`10.1364/OPTICA.4.001552` (Dec. 20, 2017).

61. Hruszkewycz, S. O., Holt, M. V., Murray, C. E., Bruley, J., Holt, J., Tripathi, A., Shpyrko, O. G., McNulty, I., Highland, M. J. & Fuoss, P. H. Quantitative Nanoscale Imaging of Lattice Distortions in Epitaxial Semiconductor Heterostructures Using Nanofocused X-ray Bragg Projection Ptychography. *Nano Letters* **12,** 5148–5154. doi:`10.1021/nl303201w` (2012).

62. Fuoss, P. H., Kim, D., Hruszkewycz, S. O., Folkman, C. M., Stephenson, G. B., Hong, S., Tripathi, A., Highland, M. J., Holt, M. V. & Thompson, C. Imaging Local Polarization in Ferroelectric Thin Films by Coherent X-Ray Bragg Projection Ptychography. *Physical Review Letters* **110,** 1–5. doi:`10.1103/physrevlett.110.177601` (2013).

63. Wu, L., Shen, Y., Barbour, A. M., Wang, W., Prabhakaran, D., Boothroyd, A. T., Mazzoli, C., Tranquada, J. M., Dean, M. P. M. & Robinson, I. K. Real Space Imaging of Spin Stripe Domain Fluctuations in a Complex Oxide. *Physical Review Letters* **127,** 275301. doi:`10.1103/PhysRevLett.127.275301` (Dec. 28, 2021).

64. Levitan, A. L., Keskinbora, K., Sanli, U. T., Weigand, M. & Comin, R. Single-Frame Far-Field Diffractive Imaging with Randomized Illumination. *Optics Express* **28,** 37103. doi:`10.1364/OE.397421` (Dec. 7, 2020).

65. Chapman, H. N. & Nugent, K. A. Coherent Lensless X-ray Imaging. *Nature Photonics* **4,** 833–839. doi:`10.1038/nphoton.2010.240` (2010).

66. Miao, J., Ishikawa, T., Robinson, I. K. & Murnane, M. M. Beyond Crystallography: Diffractive Imaging Using Coherent x-Ray Light Sources. *Science* **348,** 530–535. doi:`10.1126/science.aaa1394` (2015).

67. Marchesini, S., He, H., Chapman, H. N., Hau-Riege, S. P., Noy, A., Howells, M. R., Weierstall, U. & Spence, J. C. H. X-Ray Image Reconstruction from a Diffraction Pattern Alone. *Physical Review B* **68,** 140101. doi:`10.1103/PhysRevB.68.140101` (Oct. 28, 2003).

68. Shapiro, D., Thibault, P., Beetz, T., Elser, V., Howells, M., Jacobsen, C., Kirz, J., Lima, E., Miao, H., Neiman, A. M. & Sayre, D. Biological Imaging by Soft X-Ray Diffraction Microscopy. *Proceedings of the National Academy of Sciences* **102,** 15343–15346. doi:`10.1073/pnas.0503305102` (Oct. 25, 2005).

69. Chapman, H. N., Barty, A., Bogan, M. J., Boutet, S., Frank, M., Hau-Riege, S. P., Marchesini, S., Woods, B. W., Bajt, S., Benner, W. H., London, R. A., Plönjes, E., Kuhlmann, M., Treusch, R., Düsterer, S., Tschentscher, T., Schneider, J. R., Spiller, E., Möller, T., Bostedt, C., Hoener, M., Shapiro, D. A., Hodgson, K. O., van der Spoel, D., Burmeister, F., Bergh, M., Caleman, C., Huldt, G., Seibert, M. M., Maia, F. R. N. C., Lee, R. W., Szöke, A., Timneanu, N. & Hajdu, J. Femtosecond Diffractive Imaging with a Soft-X-ray Free-Electron Laser. *Nature Physics* **2,** 839–843. doi:`10.1038/nphys461` (2006).

70. Stockmar, M., Cloetens, P., Zanette, I., Enders, B., Dierolf, M., Pfeiffer, F. & Thibault, P. Near-Field Ptychography: Phase Retrieval for Inline Holography Using a Structured Illumination. *Scientific Reports* **3,** 1927. doi:`10.1038/srep01927` (2013).

71. Pfeiffer, F. X-Ray Ptychography. *Nature Photonics* **12,** 9–17. doi:`10.1038/s41566-017-0072-5` (Jan. 22, 2018).

72. Maiden, A. M., Morrison, G. R., Kaulich, B., Gianoncelli, A. & Rodenburg, J. M. Soft X-ray Spectromicroscopy Using Ptychography with Randomly Phased Illumination. *Nature Communications* **4,** 1666–1669. doi:`10.1038/ncomms2640` (2013).

73. Marchesini, S., Tu, Y. C. & Wu, H. T. Alternating Projection, Ptychographic Imaging and Phase Synchronization. *Applied and Computational Harmonic Analysis* **41,** 815–851. doi:`10.1016/j.acha.2015.06.005` (2016).

74. Morrison, G. R., Zhang, F., Gianoncelli, A. & Robinson, I. K. X-Ray Ptychography Using Randomized Zone Plates. *Opt. Express* **26,** 14915–14927. doi:`10.1364/OE.26.014915` (2018).

75. Marchesini, S. & Sakdinawat, A. Shaping Coherent X-Rays with Binary Optics. *Optics Express* **27,** 907–917. doi:`10.1364/OE.27.000907` (Jan. 21, 2019).

76. Odstrčil, M., Lebugle, M., Guizar-Sicairos, M., David, C. & Holler, M. Towards Optimized Illumination for High-Resolution Ptychography. *Optics Express* **27,** 14981–14997. doi:`10.1364/oe.27.014981` (2019).

77. Guizar-Sicairos, M., Holler, M., Diaz, A., Vila-Comamala, J., Bunk, O. & Menzel, A. Role of the Illumination Spatial-Frequency Spectrum for Ptychography. *Physical Review B - Condensed Matter and Materials Physics* **86,** 100103(R). doi:`10.1103/PhysRevB.86.100103` (2012).

78. Fannjiang, A. & Liao, W. Phase Retrieval with Random Phase Illumination. *Journal of the Optical Society of America A* **29,** 1847–1859. doi:`10.1364/JOSAA.29.001847` (Aug. 15, 2012).

79. Fannjiang, A. Absolute Uniqueness of Phase Retrieval with Random Illumination. *Inverse Problems* **28,** 075008. doi:`10.1088/0266-5611/28/7/075008` (July 1, 2012).

80. Horisaki, R., Egami, R. & Tanida, J. Single-Shot Phase Imaging with Randomized Light (SPIRaL). *Optics Express* **24,** 3765–3773. doi:`10.1364/oe.24.003765` (2016).

81. Horisaki, R., Kojima, T., Matsushima, K. & Tanida, J. Subpixel Reconstruction for Single-Shot Phase Imaging with Coded Diffraction. *Applied Optics* **56,** 7642–7647. doi:`10.1364/ao.56.007642` (2017).

82. Williams, G. J., Quiney, H. M., Dhal, B. B., Tran, C. Q., Nugent, K. A., Peele, A. G., Paterson, D. & De Jonge, M. D. Fresnel Coherent Diffractive Imaging. *Physical Review Letters* **97,** 025506. doi:`10.1103/PhysRevLett.97.025506` (2006).

83. Abbey, B., Nugent, K. A., Williams, G. J., Clark, J. N., Peele, A. G., Pfeifer, M. A., De Jonge, M. & McNulty, I. Keyhole Coherent Diffractive Imaging. *Nature Physics* **4,** 394–398. doi:`10.1038/nphys896` (2008).

84. Williams, G. J., Quiney, H. M., Peele, A. G. & Nugent, K. A. Fresnel Coherent Diffractive Imaging: Treatment and Analysis of Data. *New Journal of Physics* **12,** 035020. doi:`10.1088/1367-2630/12/3/035020` (2010).

85. Zhang, F., Chen, B., Morrison, G. R., Vila-Comamala, J., Guizar-Sicairos, M. & Robinson, I. K. Phase Retrieval by Coherent Modulation Imaging. *Nature Communications* **7,** 13367. doi:`10.1038/ncomms13367` (2016).
86. Dong, X., Pan, X., Liu, C. & Zhu, J. Single Shot Multi-Wavelength Phase Retrieval with Coherent Modulation Imaging. *Optics Letters* **43,** 1762–1765. doi:`10.1364/OL.43.001762` (2018).
87. Ulvestad, A., Cha, W., Calvo-Almazan, I., Maddali, S., Wild, S. M., Maxey, E., Duparaz, M. & Hruszkewycz, S. O. Bragg Coherent Modulation Imaging: Strain- and Defect- Sensitive Single Views of Extended Samples. *Submitted* (July 31, 2018).
88. Tang, W., Yang, J., Yi, W., Nie, Q., Zhu, J., Zhu, M., Guo, Y., Li, M., Li, X. & Wang, W. Single-Shot Coherent Power-Spectrum Imaging of Objects Hidden by Opaque Scattering Media. *Applied Optics* **58,** 1033–1039. doi:`10.1364/ao.58.001033` (2019).
89. Pan, X., Liu, C. & Zhu, J. Single Shot Ptychographical Iterative Engine Based on Multi-Beam Illumination. *Applied Physics Letters* **103,** 171105. doi:`10.1063/1.4826273` (2013).
90. Sidorenko, P. & Cohen, O. Single-Shot Ptychography. *Optica* **3,** 9–14. doi:`10.1364/OPTICA.3.000009` (Jan. 20, 2016).
91. Zhou, Y., Wu, J., Suo, J., Han, X., Zheng, G. & Dai, Q. Single-Shot Lensless Imaging via Simultaneous Multi-Angle LED Illumination. *Optics Express* **26,** 21418–21432. doi:`10.1364/OE.26.021418` (Aug. 20, 2018).
92. He, X., Pan, X., Liu, C. & Zhu, J. Single-Shot Phase Retrieval Based on Beam Splitting. *Applied Optics* **57,** 4832–4838. doi:`10.1364/ao.57.004832` (2018).
93. Candès, E. J., Strohmer, T. & Voroninski, V. PhaseLift: Exact and Stable Signal Recovery from Magnitude Measurements via Convex Programming. *Communications on Pure and Applied Mathematics* **66,** 1241–1274. doi:`10.1002/cpa.21432` (Aug. 2013).
94. Balan, R., Casazza, P. & Edidin, D. On Signal Reconstruction without Phase. *Applied and Computational Harmonic Analysis* **20,** 345–356. doi:`10.1016/j.acha.2005.07.001` (2006).
95. Bandeira, A. S., Cahill, J., Mixon, D. G. & Nelson, A. A. Saving Phase: Injectivity and Stability for Phase Retrieval. *Applied and Computational Harmonic Analysis* **37,** 106–125. doi:`10.1016/j.acha.2013.10.002` (2014).
96. Candès, E. J., Li, X. & Soltanolkotabi, M. Phase Retrieval from Coded Diffraction Patterns. *Applied and Computational Harmonic Analysis* **39,** 277–299. doi:`10.1016/j.acha.2014.09.004` (Sept. 2015).
97. Yurtsever, A., Udell, M., Tropp, J. & Cevher, V. Sketchy Decisions: Convex Low-Rank Matrix Optimization with Optimal Storage. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics* **54,** 1188–1196 (2017).

98. Elser, V. Phase Retrieval by Iterated Projections. *Journal of the Optical Society of America A* **20,** 40–55. doi:`10.1364/JOSAA.20.000040` (Jan. 1, 2003).

99. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J. & Chintala, S. in *Advances in Neural Information Processing Systems 32* 8024–8035 (Curran Associates, Inc., 2019).

100. Thibault, P., Elser, V., Jacobsen, C., Shapiro, D. & Sayre, D. Reconstruction of a Yeast Cell from X-ray Diffraction Data. *Acta Crystallographica Section A Foundations of Crystallography* **62,** 248–261. doi:`10.1107/S0108767306016515` (July 1, 2006).

101. Van Heel, M. & Schatz, M. Fourier Shell Correlation Threshold Criteria. *Journal of Structural Biology* **151,** 250–262. doi:`10.1016/j.jsb.2005.05.009` (2005).

102. Ordavo, I., Ihle, S., Arkadiev, V., Scharf, O., Soltau, H., Bjeoumikhov, A., Bjeoumikhova, S., Buzanich, G., Gubzhokov, R., Günther, A., Hartmann, R., Holl, P., Kimmel, N., Kühbacher, M., Lang, M., Langhoff, N., Liebel, A., Radtke, M., Reinholz, U., Riesemeier, H., Schaller, G., Schopper, F., Strüder, L., Thamm, C. & Wedell, R. A New pnCCD-based Color X-ray Camera for Fast Spatial and Energy-Resolved Measurements. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **654,** 250–257. doi:`10.1016/j.nima.2011.05.080` (Oct. 2011).

103. Clark, J. N., Beitra, L., Xiong, G., Fritz, D. M., Lemke, H. T., Zhu, D., Chollet, M., Williams, G. J., Messerschmidt, M. M., Abbey, B., Harder, R. J., Korsunsky, A. M., Wark, J. S., Reis, D. A. & Robinson, I. K. Imaging Transient Melting of a Nanocrystal Using an X-ray Laser. *Proceedings of the National Academy of Sciences* **112,** 7444–7448. doi:`10.1073/pnas.1417678112` (June 16, 2015).

104. Ihm, Y., Cho, D. H., Sung, D., Nam, D., Jung, C., Sato, T., Kim, S., Park, J., Kim, S., Gallagher-Jones, M., Kim, Y., Xu, R., Owada, S., Shim, J. H., Tono, K., Yabashi, M., Ishikawa, T., Miao, J., Noh, D. Y. & Song, C. Direct Observation of Picosecond Melting and Disintegration of Metallic Nanoparticles. *Nature Communications* **10,** 2411. doi:`10.1038/s41467-019-10328-4` (June 3, 2019).

105. Bogan, M. J., Boutet, S., Chapman, H. N., Marchesini, S., Barty, A., Benner, W. H., Rohner, U., Frank, M., Hau-Riege, S. P., Bajt, S., Woods, B., Seibert, M. M., Iwan, B., Timneanu, N., Hajdu, J. & Schulz, J. Aerosol Imaging with a Soft X-Ray Free Electron Laser. *Aerosol Science and Technology* **44,** i–vi. doi:`10.1080/02786820903485800` (Feb. 10, 2010).

106. Loh, N. D., Hampton, C. Y., Martin, A. V., Starodub, D., Sierra, R. G., Barty, A., Aquila, A., Schulz, J., Lomb, L., Steinbrener, J., Shoeman, R. L., Kassemeyer, S., Bostedt, C., Bozek, J., Epp, S. W., Erk, B., Hartmann, R., Rolles, D., Rudenko, A., Rudek, B., Foucar, L., Kimmel, N., Weidenspointner, G., Hauser, G., Holl, P., Pedersoli, E., Liang, M., Hunter, M. S., Gumprecht, L., Coppola, N., Wunderer, C., Graafsma, H., Maia, F. R. N. C., Ekeberg, T., Hantke, M., Fleckenstein, H., Hirsemann, H., Nass, K., White, T. A., Tobias, H. J., Farquar, G. R., Benner, W. H., Hau-Riege, S. P., Reich, C., Hartmann, A., Soltau, H., Marchesini, S., Bajt, S., Barthelmess, M., Bucksbaum, P., Hodgson, K. O., Strüder, L., Ullrich, J., Frank, M., Schlichting, I., Chapman, H. N. & Bogan, M. J. Fractal Morphology, Imaging and Mass Spectrometry of Single Aerosol Particles in Flight. *Nature* **486,** 513–517. doi:`10.1038/nature11222` (June 28, 2012).

107. Barty, A., Boutet, S., Bogan, M. J., Hau-Riege, S., Marchesini, S., Sokolowski-Tinten, K., Stojanovic, N., Tobey, R., Ehrke, H., Cavalleri, A., Düsterer, S., Frank, M., Bajt, S., Woods, B. W., Seibert, M. M., Hajdu, J., Treusch, R. & Chapman, H. N. Ultrafast Single-Shot Diffraction Imaging of Nanoscale Dynamics. *Nature Photonics* **2,** 415–419. doi:`10.1038/nphoton.2008.128` (July 2008).

108. Jones, C. F., Bernando, C., Tanyag, R. M. P., Bacellar, C., Ferguson, K. R., Gomez, L. F., Anielski, D., Belkacem, A., Boll, R., Bozek, J., Carron, S., Cryan, J., Englert, L., Epp, S. W., Erk, B., Foucar, L., Hartmann, R., Neumark, D. M., Rolles, D., Rudenko, A., Siefermann, K. R., Weise, F., Rudek, B., Sturm, F. P., Ullrich, J., Bostedt, C., Gessner, O. & Vilesov, A. F. Coupled Motion of Xe Clusters and Quantum Vortices in He Nanodroplets. *Physical Review B* **93,** 180510. doi:`10.1103/PhysRevB.93.180510` (May 26, 2016).

109. Suzuki, A., Tanaka, H., Yamashige, H., Orikasa, Y., Niida, Y., Kimura, T., Tono, K., Yabashi, M., Ishikawa, T., Bessho, Y., Joti, Y. & Nishino, Y. Femtosecond X-ray Laser Reveals Intact Sea–Island Structures of Metastable Solid-State Electrolytes for Batteries. *Nano Letters* **22,** 4603–4607. doi:`10.1021/acs.nanolett.1c04392` (June 8, 2022).

110. Gorkhover, T., Ulmer, A., Ferguson, K., Bucher, M., Maia, F. R. N. C., Bielecki, J., Ekeberg, T., Hantke, M. F., Daurer, B. J., Nettelblad, C., Andreasson, J., Barty, A., Bruza, P., Carron, S., Hasse, D., Krzywinski, J., Larsson, D. S. D., Morgan, A., Mühlig, K., Müller, M., Okamoto, K., Pietrini, A., Rupp, D., Sauppe, M., van der Schot, G., Seibert, M., Sellberg, J. A., Svenda, M., Swiggers, M., Timneanu, N., Westphal, D., Williams, G., Zani, A., Chapman, H. N., Faigel, G., Möller, T., Hajdu, J. & Bostedt, C. Femtosecond X-ray Fourier Holography Imaging of Free-Flying Nanoparticles. *Nature Photonics* **12,** 150–153. doi:`10.1038/s41566-018-0110-y` (Mar. 2018).

111. Hantke, M. F., Hasse, D., Maia, F. R. N. C., Ekeberg, T., John, K., Svenda, M., Loh, N. D., Martin, A. V., Timneanu, N., Larsson, D. S. D., van der Schot, G., Carlsson, G. H., Ingelman, M., Andreasson, J., Westphal, D., Liang, M.,

Stellato, F., DePonte, D. P., Hartmann, R., Kimmel, N., Kirian, R. a., Seibert, M. M., Mühlig, K., Schorb, S., Ferguson, K., Bostedt, C., Carron, S., Bozek, J. D., Rolles, D., Rudenko, A., Epp, S., Chapman, H. N., Barty, A., Hajdu, J. & Andersson, I. High-Throughput Imaging of Heterogeneous Cell Organelles with an X-ray Laser. *Nature Photonics* **8,** 943–949. doi:`10.1038/nphoton.2014.270` (2014).

112. Takayama, Y., Inui, Y., Sekiguchi, Y., Kobayashi, A., Oroguchi, T., Yamamoto, M., Matsunaga, S. & Nakasako, M. Coherent X-Ray Diffraction Imaging of Chloroplasts from *Cyanidioschyzon Merolae* by Using X-Ray Free Electron Laser. *Plant and Cell Physiology* **56,** 1272–1286. doi:`10.1093/pcp/pcv032` (July 2015).

113. Pan, D., Fan, J., Nie, Z., Sun, Z., Zhang, J., Tong, Y., He, B., Song, C., Kohmura, Y., Yabashi, M., Ishikawa, T., Shen, Y. & Jiang, H. Quantitative Analysis of the Effect of Radiation on Mitochondria Structure Using Coherent Diffraction Imaging with a Clustering Algorithm. *IUCrJ* **9,** 223–230. doi:`10.1107/S2052252521012963` (Mar. 1, 2022).

114. Kimura, T., Joti, Y., Shibuya, A., Song, C., Kim, S., Tono, K., Yabashi, M., Tamakoshi, M., Moriya, T., Oshima, T., Ishikawa, T., Bessho, Y. & Nishino, Y. Imaging Live Cell in Micro-Liquid Enclosure by X-ray Laser Diffraction. *Nature Communications* **5,** 3052. doi:`10.1038/ncomms4052` (Jan. 7, 2014).

115. Van der Schot, G., Svenda, M., Maia, F. R. N. C., Hantke, M., DePonte, D. P., Seibert, M. M., Aquila, A., Schulz, J., Kirian, R., Liang, M., Stellato, F., Iwan, B., Andreasson, J., Timneanu, N., Westphal, D., Almeida, F. N., Odic, D., Hasse, D., Carlsson, G. H., Larsson, D. S. D., Barty, A., Martin, A. V., Schorb, S., Bostedt, C., Bozek, J. D., Rolles, D., Rudenko, A., Epp, S., Foucar, L., Rudek, B., Hartmann, R., Kimmel, N., Holl, P., Englert, L., Duane Loh, N.-T., Chapman, H. N., Andersson, I., Hajdu, J. & Ekeberg, T. Imaging Single Cells in a Beam of Live Cyanobacteria with an X-ray Laser. *Nature communications* **6,** 5704. doi:`10.1038/ncomms6704` (2015).

116. Kobayashi, A., Takayama, Y., Hirakawa, T., Okajima, K., Oide, M., Oroguchi, T., Inui, Y., Yamamoto, M., Matsunaga, S. & Nakasako, M. Common Architectures in Cyanobacteria Prochlorococcus Cells Visualized by X-ray Diffraction Imaging Using X-ray Free Electron Laser. *Scientific Reports* **11,** 3877. doi:`10.1038/s41598-021-83401-y` (Feb. 16, 2021).

117. Sandberg, R. L., Bolme, C., Ramos, K., McCulloch, Q., Barber, J. L., Martinez, R., Greenfield, M., McGrane, S. D., Abbey, B., Schropp, A., Sieboth, F., Heiman, P., Nagler, B., Galtier, E. C. & Granados, E. *Ultrafast Imaging of Shocked Material Dynamics with X-ray Fee Electron Laser Pulses* in *CLEO: 2014 Postdeadline Paper Digest* CLEO: Science and Innovations (OSA, San Jose, California, 2014), STh5C.8. doi:`10.1364/CLEO_SI.2014.STh5C.8`.

118. Schropp, A., Hoppe, R., Meier, V., Patommel, J., Seiboth, F., Ping, Y., Hicks, D. G., Beckwith, M. A., Collins, G. W., Higginbotham, A., Wark, J. S., Lee, H. J., Nagler, B., Galtier, E. C., Arnold, B., Zastrau, U., Hastings, J. B. & Schroer, C. G. Imaging Shock Waves in Diamond with Both High Temporal and Spatial Resolution at an XFEL. *Scientific reports* **5,** 11089. doi:`10.1038/srep11089` (2015).

119. Nagler, B., Schropp, A., Galtier, E. C., Arnold, B., Brown, S. B., Fry, A., Gleason, A., Granados, E., Hashim, A., Hastings, J. B., Samberg, D., Seiboth, F., Tavella, F., Xing, Z., Lee, H. J. & Schroer, C. G. The Phase-Contrast Imaging Instrument at the Matter in Extreme Conditions Endstation at LCLS. *Review of Scientific Instruments* **87,** 103701. doi:`10.1063/1.4963906` (Oct. 2016).

120. Seiboth, F., Fletcher, L. B., McGonegle, D., Anzellini, S., Dresselhaus-Cooper, L. E., Frost, M., Galtier, E., Goede, S., Harmand, M., Lee, H. J., Levitan, A. L., Miyanishi, K., Nagler, B., Nam, I., Ozaki, N., Rödel, M., Schropp, A., Spindloe, C., Sun, P., Wark, J. S., Hastings, J., Glenzer, S. H. & McBride, E. E. Simultaneous 8.2 keV Phase-Contrast Imaging and 24.6 keV X-ray Diffraction from Shock-Compressed Matter at the LCLS. *Applied Physics Letters* **112,** 221907. doi:`10.1063/1.5031907` (May 28, 2018).

121. Vassholz, M., Hoeppe, H. P., Hagemann, J., Rosselló, J. M., Osterhoff, M., Mettin, R., Kurz, T., Schropp, A., Seiboth, F., Schroer, C. G., Scholz, M., Möller, J., Hallmann, J., Boesenberg, U., Kim, C., Zozulya, A., Lu, W., Shayduk, R., Schaffer, R., Madsen, A. & Salditt, T. Pump-Probe X-ray Holographic Imaging of Laser-Induced Cavitation Bubbles with Femtosecond FEL Pulses. *Nature Communications* **12,** 3468. doi:`10.1038/s41467-021-23664-1` (Dec. 2021).

122. Hagemann, J., Vassholz, M., Hoeppe, H., Osterhoff, M., Rosselló, J. M., Mettin, R., Seiboth, F., Schropp, A., Möller, J., Hallmann, J., Kim, C., Scholz, M., Boesenberg, U., Schaffer, R., Zozulya, A., Lu, W., Shayduk, R., Madsen, A., Schroer, C. G. & Salditt, T. Single-Pulse Phase-Contrast Imaging at Free-Electron Lasers in the Hard X-ray Regime. *Journal of Synchrotron Radiation* **28,** 52–63. doi:`10.1107/S160057752001557X` (Jan. 1, 2021).

123. Vagovič, P., Sato, T., Mikeš, L., Mills, G., Graceffa, R., Mattsson, F., Villanueva-Perez, P., Ershov, A., Faragó, T., Uličný, J., Kirkwood, H., Letrun, R., Mokso, R., Zdora, M.-C., Olbinado, M. P., Rack, A., Baumbach, T., Schulz, J., Meents, A., Chapman, H. N. & Mancuso, A. P. Megahertz X-Ray Microscopy at x-Ray Free-Electron Laser and Synchrotron Sources. *Optica* **6,** 1106. doi:`10.1364/OPTICA.6.001106` (Sept. 20, 2019).

124. Von Korff Schmising, C., Pfau, B., Schneider, M., Günther, C. M., Giovannella, M., Perron, J., Vodungbo, B., Müller, L., Capotondi, F., Pedersoli, E., Mahne, N., Lüning, J. & Eisebitt, S. Imaging Ultrafast Demagnetization Dynamics after a Spatially Localized Optical Excitation. *Physical Review Letters* **112,** 217203. doi:`10.1103/PhysRevLett.112.217203` (May 29, 2014).

125. Wang, T., Zhu, D., Wu, B., Graves, C., Schaffert, S., Rander, T., Müller, L., Vodungbo, B., Baumier, C., Bernstein, D. P., Bräuer, B., Cros, V., de Jong, S., Delaunay, R., Fognini, A., Kukreja, R., Lee, S., López-Flores, V., Mohanty, J., Pfau, B., Popescu, H., Sacchi, M., Sardinha, A. B., Sirotti, F., Zeitoun, P., Messerschmidt, M., Turner, J. J., Schlotter, W. F., Hellwig, O., Mattana, R., Jaouen, N., Fortuna, F., Acremann, Y., Gutt, C., Dürr, H. A., Beaurepaire, E., Boeglin, C., Eisebitt, S., Grübel, G., Lüning, J., Stöhr, J. & Scherz, A. O. Femtosecond Single-Shot Imaging of Nanoscale Ferromagnetic Order in Co / Pd Multilayers Using Resonant X-Ray Holography. *Physical Review Letters* **108,** 267403. doi:`10.1103/PhysRevLett.108.267403` (June 27, 2012).

126. Eisebitt, S., Lüning, J., Schlotter, W. F., Lörgen, M., Hellwig, O., Eberhardt, W. & Stöhr, J. Lensless Imaging of Magnetic Nanostructures by X-ray Spectro-Holography. *Nature* **432,** 885–888. doi:`10.1038/nature03139` (Dec. 2004).

127. Marchesini, S., Boutet, S., Sakdinawat, A. E., Bogan, M. J., Bajt, S., Barty, A., Chapman, H. N., Frank, M., Hau-Riege, S. P., Szöke, A., Cui, C., Shapiro, D. A., Howells, M. R., Spence, J. C., Shaevitz, J. W., Lee, J. Y., Hajdu, J. & Seibert, M. M. Massively Parallel X-ray Holography. *Nature Photonics* **2,** 560–563. doi:`10.1038/nphoton.2008.154` (2008).

128. Guizar-Sicairos, M. & Fienup, J. R. Direct Image Reconstruction from a Fourier Intensity Pattern Using HERALDO. *Optics Letters* **33,** 2668. doi:`10.1364/OL.33.002668` (Nov. 15, 2008).

129. Martin, A. V., D'Alfonso, A. J., Wang, F., Bean, R., Capotondi, F., Kirian, R. A., Pedersoli, E., Raimondi, L., Stellato, F., Yoon, C. H. & Chapman, H. N. X-Ray Holography with a Customizable Reference. *Nature Communications* **5,** 4661. doi:`10.1038/ncomms5661` (Dec. 17, 2014).

130. Garcia-Sucerquia, J., Xu, W., Jericho, S. K., Klages, P., Jericho, M. H. & Kreuzer, H. J. Digital In-Line Holographic Microscopy. *Applied Optics* **45,** 836. doi:`10.1364/AO.45.000836` (Feb. 10, 2006).

131. Gorniak, T., Heine, R., Mancuso, A. P., Staier, F., Christophis, C., Pettitt, M. E., Sakdinawat, A., Treusch, R., Guerassimova, N., Feldhaus, J., Gutt, C., Grübel, G., Eisebitt, S., Beyer, A., Gölzhäuser, A., Weckert, E., Grunze, M., Vartanyants, I. A. & Rosenhahn, A. X-Ray Holographic Microscopy with Zone Plates Applied to Biological Samples in the Water Window Using 3rd Harmonic Radiation from the Free-Electron Laser FLASH. *Optics Express* **19,** 11059. doi:`10.1364/OE.19.011059` (June 6, 2011).

132. Kharitonov, K., Mehrjoo, M., Ruiz-Lopez, M., Keitel, B., Kreis, S., Gang, S.-g., Pan, R., Marras, A., Correa, J., Wunderer, C. & Plonjes, E. *Single-Shot Ptychography at a Soft X-ray Free-Electron Laser* preprint (In Review, May 13, 2022). doi:`10.21203/rs.3.rs-1629463/v1`.

133. Osterhoff, M., Vassholz, M., Hoeppe, H. P., Rosselló, J. M., Mettin, R., Hagemann, J., Möller, J., Hallmann, J., Scholz, M., Schaffer, R., Boesenberg, U., Kim, C., Zozulya, A., Lu, W., Shayduk, R., Madsen, A. & Salditt, T. Nanosecond Timing and Synchronization Scheme for Holographic Pump–Probe Studies at the MID Instrument at European XFEL. *Journal of Synchrotron Radiation* **28,** 987–994. doi:`10.1107/S1600577521003052` (May 1, 2021).

134. Schropp, A., Patommel, J., Seiboth, F., Arnold, B., Galtier, E. C., Lee, H. J., Nagler, B., Hastings, J. B. & Schroer, C. G. *Developing a Platform for High-Resolution Phase Contrast Imaging of High Pressure Shock Waves in Matter* in. SPIE Optical Engineering + Applications (eds Moeller, S. P., Yabashi, M. & Hau-Riege, S. P.) (San Diego, California, USA, Oct. 17, 2012), 85040F. doi:`10.1117/12.929882`.

135. Faenov, A. Y., Pikuz, T. A., Mabey, P., Albertazzi, B., Michel, Th., Rigon, G., Pikuz, S. A., Buzmakov, A., Makarov, S., Ozaki, N., Matsuoka, T., Katagiri, K., Miyanishi, K., Takahashi, K., Tanaka, K. A., Inubushi, Y., Togashi, T., Yabuuchi, T., Yabashi, M., Casner, A., Kodama, R. & Koenig, M. Advanced High Resolution X-Ray Diagnostic for HEDP Experiments. *Scientific Reports* **8,** 16407. doi:`10.1038/s41598-018-34717-9` (Nov. 6, 2018).

136. Beckwith, M. A., Jiang, S., Schropp, A., Fernandez-Pañella, A., Rinderknecht, H. G., Wilks, S. C., Fournier, K. B., Galtier, E. C., Xing, Z., Granados, E., Gamboa, E., Glenzer, S. H., Heimann, P., Zastrau, U., Cho, B. I., Eggert, J. H., Collins, G. W. & Ping, Y. Imaging at an X-Ray Absorption Edge Using Free Electron Laser Pulses for Interface Dynamics in High Energy Density Systems. *Review of Scientific Instruments* **88,** 053501. doi:`10.1063/1.4982166` (May 2017).

137. Capotondi, F., Pedersoli, E., Mahne, N., Menk, R. H., Passos, G., Raimondi, L., Svetina, C., Sandrin, G., Zangrando, M., Kiskinova, M., Bajt, S., Barthelmess, M., Fleckenstein, H., Chapman, H. N., Schulz, J., Bach, J., Frömter, R., Schleitzer, S., Müller, L., Gutt, C. & Grübel, G. Invited Article: Coherent Imaging Using Seeded Free-Electron Laser Pulses with Variable Polarization: First Results and Research Opportunities. *Review of Scientific Instruments* **84,** 051301. doi:`10.1063/1.4807157` (May 2013).

138. Levitan, A. & Comin, R. Error Metrics for Partially Coherent Wave Fields. *Optics Letters* **47,** 2322. doi:`10.1364/OL.455955` (May 1, 2022).

139. Penczek, P. A. in *Methods in Enzymology* 73–100 (Elsevier, 2010). doi:`10.1016/S0076-6879(10)82003-8`.

140. Vila-Comamala, J., Diaz, A., Guizar-Sicairos, M., Mantion, A., Kewish, C. M., Menzel, A., Bunk, O. & David, C. Characterization of High-Resolution Diffractive X-ray Optics by Ptychographic Coherent Diffractive Imaging. *Optics Express* **19,** 21333. doi:`10.1364/OE.19.021333` (Oct. 24, 2011).

141. Donnelly, C., Scagnoli, V., Guizar-Sicairos, M., Holler, M., Wilhelm, F., Guillou, F., Rogalev, A., Detlefs, C., Menzel, A., Raabe, J. & Heyderman, L. J. *High Resolution Hard X-ray Magnetic Imaging with Dichroic Ptychography* 2016. doi:`10.1103/PhysRevB.94.064421`.

142. Shapiro, D. A., Babin, S., Celestre, R. S., Chao, W., Conley, R. P., Denes, P., Enders, B., Enfedaque, P., James, S., Joseph, J. M., Krishnan, H., Marchesini, S., Muriki, K., Nowrouzi, K., Oh, S. R., Padmore, H., Warwick, T., Yang, L., Yashchuk, V. V., Yu, Y.-S. & Zhao, J. An Ultrahigh-Resolution Soft x-Ray Microscope for Quantitative Analysis of Chemically Heterogeneous Nanomaterials. *Science Advances* **6,** eabc4904. doi:`10.1126/sciadv.abc4904` (Dec. 2020).

143. Pfau, B., Günther, C. M., Schaffert, S., Mitzner, R., Siemer, B., Roling, S., Zacharias, H., Kutz, O., Rudolph, I., Treusch, R. & Eisebitt, S. Femtosecond Pulse X-Ray Imaging with a Large Field of View. *New Journal of Physics* **12,** 095006. doi:`10.1088/1367-2630/12/9/095006` (Sept. 10, 2010).

144. Barnett, A. H., Epstein, C. L., Greengard, L. & Magland, J. *Geometry of the Phase Retrieval Problem: Graveyard of Algorithms* 1st ed. doi:`10.1017/9781009003919` (Cambridge University Press, Apr. 30, 2022).

145. Cox, I. J. & Sheppard, C. J. R. Information Capacity and Resolution in an Optical System. *Journal of the Optical Society of America A* **3,** 1152. doi:`10.1364/JOSAA.3.001152` (Aug. 1, 1986).

146. Neifeld, M. A. Information, Resolution, and Space–Bandwidth Product. *Optics Letters* **23,** 1477. doi:`10.1364/OL.23.001477` (Sept. 15, 1998).

147. Thompson, A. C., Vaughan, D., *et al. X-Ray Data Booklet* **4** (Lawrence Berkeley National Laboratory, University of California Berkeley, CA, 2001).

148. Guo, Z., Levitan, A., Barbastathis, G. & Comin, R. Randomized Probe Imaging through Deep K-Learning. *Optics Express* **30,** 2247. doi:`10.1364/OE.445498` (Jan. 17, 2022).

149. Goy, A., Arthur, K., Li, S. & Barbastathis, G. Low Photon Count Phase Retrieval Using Deep Learning. *Physical Review Letters* **121,** 243902. doi:`10.1103/PhysRevLett.121.243902` (2018).

150. Kang, I., Zhang, F. & Barbastathis, G. Phase Extraction Neural Network (PhENN) with Coherent Modulation Imaging (CMI) for Phase Retrieval at Low Photon Counts. *Optics Express* **28,** 21578. doi:`10.1364/OE.397430` (July 20, 2020).

151. Yu, Y.-S., Farmand, M., Kim, C., Liu, Y., Grey, C. P., Strobridge, F. C., Tyliszczak, T., Celestre, R., Denes, P., Joseph, J., Krishnan, H., Maia, F. R. N. C., Kilcoyne, A. L. D., Marchesini, S., Leite, T. P. C., Warwick, T., Padmore, H., Cabana, J. & Shapiro, D. A. Three-Dimensional Localization of Nanoscale Battery Reactions Using Soft X-ray Tomography. *Nature Communications* **9,** 921. doi:`10.1038/s41467-018-03401-x` (Mar. 2, 2018).

152. Rana, A., Liao, C.-T., Iacocca, E., Zou, J., Pham, M., Lu, X., Subramanian, E.-E. C., Lo, Y. H., Ryan, S. A., Bevis, C. S., Karl, R. M., Glaid, A. J., Rable, J., Mahale, P., Hirst, J., Ostler, T., Liu, W., O'Leary, C. M., Yu, Y.-S., Bustillo, K., Ohldag, H., Shapiro, D. A., Yazdi, S., Mallouk, T. E., Osher, S. J., Kapteyn, H. C., Crespi, V. H., Badding, J. V., Tserkovnyak, Y., Murnane, M. M. & Miao, J. Three-Dimensional Topological Magnetic Monopoles and Their Interactions in a Ferromagnetic Meta-Lattice. *Nature Nanotechnology* **18,** 227–232. doi:`10.1038/s41565-022-01311-0` (Mar. 2023).

153. Denes, P., Doering, D., Padmore, H. A., Walder, J.-P. & Weizeorick, J. A Fast, Direct x-Ray Detection Charge-Coupled Device. *Review of Scientific Instruments* **80,** 083302. doi:`10.1063/1.3187222` (Aug. 2009).

154. Donnelly, C., Scagnoli, V., Guizar-Sicairos, M., Holler, M., Wilhelm, F., Guillou, F., Rogalev, A., Detlefs, C., Menzel, A., Raabe, J. & Heyderman, L. J. High-Resolution Hard x-Ray Magnetic Imaging with Dichroic Ptychography. *Physical Review B* **94,** 1–9. doi:`10.1103/PhysRevB.94.064421` (2016).

155. Fletcher, R. Function Minimization by Conjugate Gradients. *The Computer Journal* **7,** 149–154. doi:`10.1093/comjnl/7.2.149` (Feb. 1, 1964).

156. Keskinbora, K., Sanli, U. T., Baluktsian, M., Grévent, C., Weigand, M. & Schütz, G. High-Throughput Synthesis of Modified Fresnel Zone Plate Arrays via Ion Beam Lithography. *Beilstein Journal of Nanotechnology* **9,** 2049–2056. doi:`10.3762/bjnano.9.194` (2018).

157. Fienup, J. R. Invariant Error Metrics for Image Reconstruction. *Applied Optics* **36,** 8352. doi:`10.1364/AO.36.008352` (Nov. 10, 1997).

158. Van Heel, M. Similarity Measures between Images. *Ultramicroscopy* **21,** 95–100. doi:`10.1016/0304-3991(87)90010-6` (Jan. 1987).

159. Cao, S., Kok, P., Li, P., Maiden, A. M. & Rodenburg, J. M. Modal Decomposition of a Propagating Matter Wave via Electron Ptychography. *Physical Review A* **94,** 063621. doi:`10.1103/PhysRevA.94.063621` (Dec. 19, 2016).

160. Chen, Z., Odstrcil, M., Jiang, Y., Han, Y., Chiu, M.-H., Li, L.-J. & Muller, D. A. Mixed-State Electron Ptychography Enables Sub-Angstrom Resolution Imaging with Picometer Precision at Low Dose. *Nature Communications* **11,** 2994. doi:`10.1038/s41467-020-16688-6` (Dec. 2020).

161. Pound, B. A., Mertes, K. M., Carr, A. V., Seaberg, M. H., Hunter, M. S., Ward, W. C., Hunter, J. F., Sweeney, C. M., Sewell, C. M., Weisse-Bernstein, N. R., Baldwin, J. K. S. & Sandberg, R. L. Ptychography at the Linac Coherent Light Source in a Parasitic Geometry. *Journal of Applied Crystallography* **53.** doi:`10.1107/S1600576720010778` (Oct. 1, 2020).

162. Clark, J. N., Huang, X., Harder, R. J. & Robinson, I. K. Continuous Scanning Mode for Ptychography. *Optics Letters* **39,** 6066. doi:`10.1364/OL.39.006066` (2014).

163. Pelz, P. M., Guizar-Sicairos, M., Thibault, P., Johnson, I., Holler, M. & Menzel, A. On-the-Fly Scans for X-ray Ptychography. *Applied Physics Letters* **105,** 3–8. doi:`10.1063/1.4904943` (2014).

164. Deng, J., Nashed, Y. S. G., Chen, S., Phillips, N. W., Peterka, T., Ross, R., Vogt, S., Jacobsen, C. & Vine, D. J. Continuous Motion Scan Ptychography: Characterization for Increased Speed in Coherent x-Ray Imaging. *Optics Express* **23,** 5438. doi:`10.1364/OE.23.005438` (2015).

165. Huang, X., Lauer, K., Clark, J. N., Xu, W., Nazaretski, E., Harder, R., Robinson, I. K. & Chu, Y. S. Fly-Scan Ptychography. *Scientific Reports* **5,** 1–5. doi:`10.1038/srep09074` (2015).

166. Tsai, Y.-W., Lin, J.-M., Chen, C.-Y., Chen, Y., Lin, B.-H., Yin, G.-C., Tang, M.-T. & Huang, Y.-S. Hard X-ray Ptychography at Taiwan Photon Source at 11–20 Nm Spatial Resolution. *Journal of Synchrotron Radiation* **28.** doi:`10.1107/S1600577521008638` (Nov. 1, 2021).

167. Sala, S., Daurer, B. J., Odstrcil, M., Capotondi, F., Pedersoli, E., Hantke, M. F., Manfredda, M., Loh, N. D., Thibault, P. & Maia, F. R. N. C. Pulse-to-Pulse Wavefront Sensing at Free-Electron Lasers Using Ptychography. *Journal of Applied Crystallography* **53,** 949–956. doi:`10.1107/S1600576720006913` (Aug. 1, 2020).

168. Daurer, B. J., Sala, S., Hantke, M. F., Reddy, H. K. N., Bielecki, J., Shen, Z., Nettleblad, C., Svenda, M., Ekeberg, T., Carini, G. A., Hart, P., Osipov, T., Aquila, A., Loh, N. D., Maia, F. R. N. C. & Thibault, P. *Ptychographic Wavefront Characterisation for Single-Particle Imaging at X-ray Lasers* Dec. 26, 2020.

169. Huang, X., Nazaretski, E., Xu, W., Hidas, D., Cordier, M., Stripe, B., Yun, W. & Chu, Y. S. Metrology of a Focusing Capillary Using Optical Ptychography. *Sensors* **20,** 6462. doi:`10.3390/s20226462` (Nov. 12, 2020).

170. Takeo, Y., Suzuki, A., Motoyama, H., Takei, Y., Kume, T., Matsuzawa, Y., Senba, Y., Kishimoto, H., Ohashi, H. & Mimura, H. Soft X-Ray Nanobeam Formed by an Ellipsoidal Mirror. *Applied Physics Letters* **116,** 121102. doi:`10.1063/1.5144932` (Mar. 23, 2020).

171. Li, P., Edo, T., Batey, D., Rodenburg, J. & Maiden, A. Breaking Ambiguities in Mixed State Ptychography. *Optics Express* **24,** 9038. doi:`10.1364/OE.24.009038` (Apr. 18, 2016).

172. Chang, H., Enfedaque, P., Lou, Y. & Marchesini, S. Partially Coherent Ptychography by Gradient Decomposition of the Probe. *Acta Crystallographica Section A: Foundations and Advances* **74,** 157–169. doi:`10.1107/S2053273318001924` (2018).

173. Wolf, E. New Theory of Partial Coherence in the Space–Frequency Domain Part I: Spectra and Cross Spectra of Steady-State Sources. *Journal of the Optical Society of America* **72,** 343. doi:`10.1364/JOSA.72.000343` (Mar. 1, 1982).

174. Wolf, E. New Theory of Partial Coherence in the Space-Frequency Domain Part II: Steady-state Fields and Higher-Order Correlations. *Journal of the Optical Society of America A* **3,** 76. doi:`10.1364/JOSAA.3.000076` (Jan. 1, 1986).

175. Uhlmann, A. The "Transition Probability" in the State Space of a *-Algebra. *Reports on Mathematical Physics* **9,** 273–279. doi:`10.1016/0034-4877(76)90060-4` (Apr. 1976).

176. Jozsa, R. Fidelity for Mixed Quantum States. *Journal of Modern Optics* **41,** 2315–2323. doi:`10.1080/09500349414552171` (Dec. 1994).

177. Horn, R. A. & Johnson, C. R. *Matrix Analysis* 1st ed. doi:`10.1017/CBO9780511810817` (Cambridge University Press, Dec. 27, 1985).