

# UC Berkeley

## UC Berkeley Electronic Theses and Dissertations

### Title

Dynamical Diffraction: Friend or Foe? 4D-STEM Measurements Robust To Multiple Scattering

### Permalink

<https://escholarship.org/uc/item/0cj8x6fv>

### Author

Zeltmann, Steven

### Publication Date

2022

Peer reviewed|Thesis/dissertation

Dynamical Diffraction: Friend or Foe? 4D-STEM Measurements Robust To Multiple  
Scattering

by

Steven Zeltmann

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering- Materials Science & Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Andrew M Minor, Chair

Professor Rebecca Abergel

Professor Mary Scott

Fall 2022

Dynamical Diffraction: Friend or Foe? 4D-STEM Measurements Robust To Multiple  
Scattering

Copyright 2022  
by  
Steven Zeltmann

## Abstract

## Dynamical Diffraction: Friend or Foe? 4D-STEM Measurements Robust To Multiple Scattering

by

Steven Zeltmann

Doctor of Philosophy in Engineering- Materials Science &amp; Engineering

University of California, Berkeley

Professor Andrew M Minor, Chair

Four-dimensional scanning transmission electron microscopy (4D-STEM) is a modern operating mode of a transmission electron microscope in which a focused electron probe is rastered across the sample and the diffraction pattern is recorded at each position. The resulting diffraction patterns can be analyzed to obtain a wealth of local structural information, such as deformation or strain, changes in symmetry or lattice distortions, orientation of a crystal lattice, as well as to measure electric and magnetic fields. More advanced analyses, i.e. ptychography, can also extract structural information at a spatial resolution finer than the size of the electron probe.

Several challenges arise in realizing these measurements: First, the sheer number of diffraction patterns recorded in a 4D-STEM experiment leads to computational challenges and puts demands on the complexity of the algorithms used to recover the structural information. Second, experimental considerations often strictly limit the number of electrons in each of the diffraction patterns, which can be mitigated through robust analysis approaches or by de-noising that takes advantage of the high dimensionality of the data. Most critically, all of the structural measurements described above are effectively trivial in the limit of thin and weakly scattering materials but become rather challenging when analyzing diffraction from a thick sample where multiple scattering effects are present.

In this work, we will explore several means to mitigate these challenges. First, to handle the large quantities of data and the low number of electrons recorded by modern detectors operated at their full speed, we will show a hyperspectral denoising method based on total variation denoising and show its application to 4D-STEM datasets.

The bulk of this work, however, will focus on the latter challenge: dynamical scattering. In 4D-STEM measurements of local strain or deformation, dynamical scattering causes unwanted contrast inside of diffraction disks which hinders accurate determination of the

lattice. To mitigate this, we demonstrate a method for imprinting known contrast into the diffraction disks to improve the precision of the measured lattice. In measurements of the local orientation of the crystal, multiple scattering causes the diffraction disk intensities to vary in a highly nonlinear way as the crystal tilts, and as a function of the thickness of the crystal. We present a hybrid pattern-matching and simulation-matching algorithm for precisely determining both the orientation and thickness of a crystalline sample from 4D-STEM measurements. Finally, many polar structures of technological interest exist only under exacting electrical and mechanical boundary conditions and so can only be studied as a thick and heterogeneous sample. To measure polarization structures in such materials, we construct a dynamical scattering model for the system and demonstrate an optimization procedure which recovers local polar order from large-area scans of a thick multilayer sample.

To WZP

My constant companion for the past decade, and many more to come.

# Contents

<b>Contents</b>	<b>ii</b>
<b>List of Figures</b>	<b>iv</b>
<b>1 Introduction</b>	<b>1</b>
1.1 4D-STEM . . . . .	1
1.2 Multiple Scattering . . . . .	2
1.3 Overview . . . . .	3
<b>2 Theory</b>	<b>5</b>
2.1 Electron Scattering . . . . .	5
2.2 Kinematic Diffraction . . . . .	9
2.3 Dynamical Diffraction . . . . .	10
2.4 Computation . . . . .	12
<b>3 Hyperspectral Denoising of 4D-STEM Datasets Using a Total Variation Method</b>	<b>18</b>
3.1 Introduction . . . . .	18
3.2 TV Denoising Algorithm . . . . .	20
3.3 Methods . . . . .	24
3.4 Results and Discussion . . . . .	24
3.5 Conclusions . . . . .	28
<b>I Dynamical Diffraction as a Foe</b>	<b>29</b>
Robust Measurements <i>Despite</i> Dynamical Diffraction . . . . .	30
<b>4 Patterned Probes for High Precision 4D-STEM Bragg Measurements</b>	<b>33</b>
4.1 Introduction . . . . .	33
4.2 Theory . . . . .	36
4.3 Methods . . . . .	41
4.4 Results and Discussion . . . . .	44
4.5 Conclusion . . . . .	50

<b>5</b>	<b>Automated Crystal Orientation Mapping in py4DSTEM using Sparse Correlation Matching</b>	<b>52</b>
5.1	Introduction . . . . .	53
5.2	Methods . . . . .	54
5.3	Results and Discussion . . . . .	69
5.4	Conclusion . . . . .	75
5.5	Source Code and Data Availability . . . . .	75
<b>II</b>	<b>Dynamical Diffraction as a Friend</b>	<b>76</b>
	Robust Measurements by <i>Inverting</i> Dynamical Diffraction . . . . .	77
<b>6</b>	<b>A Dynamical Diffraction Refinement Method for Precise 4D-STEM Measurements of Crystal Orientation and Thickness</b>	<b>79</b>
6.1	Introduction . . . . .	79
6.2	Orientation and Thickness Refinement Procedure . . . . .	80
6.3	Experimental Methods . . . . .	82
6.4	Results and Discussion . . . . .	82
6.5	Conclusion . . . . .	89
<b>7</b>	<b>Uncovering Polar Vortex Structures by Inversion of Multiple Scattering with a Stacked Bloch Wave Model</b>	<b>90</b>
7.1	Introduction . . . . .	91
7.2	Theory . . . . .	92
7.3	Methods . . . . .	99
7.4	Results & Discussion . . . . .	99
7.5	Conclusions . . . . .	105
<b>III</b>	<b>Conclusion</b>	<b>106</b>
<b>8</b>	<b>Summary and Future Work</b>	<b>107</b>
8.1	Conclusions . . . . .	107
8.2	Future Work . . . . .	108
	<b>Bibliography</b>	<b>110</b>



# List of Figures

2.1	Relativistic corrected (300 kV) single atom scattering factors for Si and Au using two different parameterizations. The Lobato and Van Dyck [98] method contains only an elastic part and is purely real, while the Weickenmeier and Kohl method simulated absorption due to phonon and core loss interactions as well as damping by the Debye-Waller factor [183]. . . . .	6
2.2	Comparison of IAM and DFT potentials for a unit cell of $\text{PbTiO}_3$ , computed using abTEM and GPAW. . . . .	8
3.1	Atomic model used for multislice simulations of a 4D-STEM dataset. Octahedral Au nanoparticles of approximately 5 nm size are distributed on an approximately 3 nm thick amorphous carbon support. . . . .	25
3.2	Denoising of the simulated Au nanoparticle dataset. The first two columns show diffraction patterns selected from a region of amorphous carbon and from the nanoparticle at the center of the scan. The third column shows the maximal diffraction pattern across the field of view. Virtual images are shown in the final two columns. Denoising recovers diffracted disks in the patterns that were rendered undetectable at low dose, at the cost of some blurring of the virtual images. . . . .	26
3.3	Effect of the data weighting parameters, $\mu_R$ for real space and $\mu_Q$ for diffraction space, on the SNR increase for the simulated Au nanoparticle dataset after denoising, using FISTA accelerated convergence. The left panel shows the SNR after running for 50 iterations, while the right panel shows the highest SNR achieved during the iterative reconstruction process. For some choices of the smoothing parameters, the algorithm reaches an optimal solution but proceeds to diverge as iteration continues except for a small range of parameter space. If iteration is stopped at the appropriate step, it is possible to achieve optimal performance of the reconstruction over a wide range of parameters. . . . .	27
3.4	Converged beam electron diffraction patterns from an 80 nm thick Si crystal near the (110) zone axis at 300 kV. The crystal is tilted by a small amount about the vertical axis, causing substantial intensity redistribution within the diffraction disks (complicating measurement of their positions, and thus determination of the local strain). Simulations were performed using py4DSTEM. . . . .	31

3.5	Diffraction intensities for several reflections from an 80 nm thick Si crystal near the (110) zone axis at 300 kV. The crystal is tilted amount about the vertical axis, causing nonlinear intensity redistribution between the diffraction disks (complicating orientation matching). Simulations were performed using py4DSTEM. . . . .	31
4.1	(a) Schematic of experimental setup for 4D-STEM strain mapping. A converged electron probe is rastered across the sample and a diffraction pattern is acquired at each probe position. Thick regions of the sample have complicated dynamical contrast inside the CBED disks that make accurate position determination difficult. In (b), a grating is inserted in the condenser system of the microscope to pattern the probe in momentum space. This pattern is imprinted on the diffracted disks, providing sharp edges in registry with the probe pattern that makes computational determination of their position more robust. . . . .	36
4.2	Numerical tests of image registration of an ideal STEM probe with a noisy measurement. Position error was measured for 1000 randomly generated probes along one dimension, for (a) circular disks with different radii, (b) varying numbers of concentric rings, and (c) varying numbers of intersecting rays. These measurements are compared to the theoretical precision given by Eq. 4.3. Inset images show examples of noisy measurements. . . . .	37
4.3	Numerical tests of the impact of sampling on relative position error variance. Concentric ring probes show the scaling expected from the theory at greater than Nyquist sampling, but have substantial position error due to aliasing at lower sampling. Dashed lines are the prediction of Equation 4.6. Inset: example images of under-, Nyquist, and over-sampled disk images for 3, 7, and 15 ring disks. . . . .	39
4.4	Numerical tests of the STEM probe size using different ring patterns at different convergence angles, for an aberration-free microscope at 300 kV. Adding the ring pattern to the probe causes the real-space probe size to grow by a factor determined by the number of rings. Inset: cumulative radial intensity profiles of probe intensity for a 2 mrad convergence semiangle. . . . .	40
4.5	SEM micrograph of the fabricated bullseye aperture plate. . . . .	44
4.6	Strain mapping precision of simulated silicon diffraction data at different samples thickness and electron dose. (a,e) Representative simulated diffraction patterns at different electron doses per pattern. (b,f) Comparison of the $u$ and $v$ reciprocal lattice vectors measured at each scan position in the simulation of a strain-free sample. The center of each histogram represents the average $u$ , $v$ positions obtained from the noise-free simulation data. (c,g) Cross-validation error and (d,h) RMS fit error, relative to the reciprocal lattice vector length (equivalent to the strain error in the small strain limit). (a-d) are obtained from a 5 nm model with largely kinematical scattering, while (e-h) are from a 20 nm model with dynamical contrast inside the CBED disks. The reciprocal lattice vectors are drawn in the bottom left panel of (a), and have length $\approx 70$ pixels. . . . .	45

- 4.7 **Cross-validation error determination from 4D-STEM experiments on a Si  $\langle 110 \rangle$  wedge.** Diffraction patterns from a thin region of the wedge with (a) the standard circular aperture and (b) with the bullseye amplitude grating. The cross validation error, computed by fitting a lattice to half of the identified diffraction disks and measuring the error of the remaining half, using (c) cross-correlation and (d) hybrid fitting. Diffraction patterns from a thick region of the wedge (e) without and (f) with the bullseye aperture. Cross-validation strain error for (g) cross-correlation and (h) hybrid correlation disk detection. Strain maps from each region of the Si wedge sample are shown in (i)–(l). The label on each strain map indicates the standard deviation of that strain component over the field of view. 47
- 5.1 **ACOM using correlation matching in py4DSTEM.** (a) Structure of fcc Au. (b) Atomic scattering factor of Au. (c) Structure factors for fcc Au. (d) Zone axes included in orientation plan. (e) Diffraction patterns for various orientations, and (f) corresponding orientation plan slices. (g) Correlogram maxima for each pattern in (e) as a function of zone axis, and (h) corresponding in-plane rotation correlation. Highest correlation scores are shown in (g) and (h) using red circles. 56
- 5.2 **Examples of alternative orientation plan types in py4DSTEM.** Fiber texture examples where (a) orientations fully orbit around a single zone axis (the fiber axis), or (b) contain only a symmetry-reduced wedge of zone axes which orbit around a the fiber axis. (c) Examples of orientation plans generated directly from Materials Project entries [70], using pymatgen symmetries [121]. . . . . 64
- 5.3 **Zone axis misorientation as a function of sampling and maximum scattering angle for kinematical simulations.** The mean tilt error and number of patterns matched per second are shown inset for each panel. . . . . 67
- 5.4 **ACOM of overlapping diffraction patterns.** (a) Three overlapping diffraction patterns with randomly chosen in-plane rotations. (b) First match, (c) second match, and (d) third match returned by ACOM code. The fitted zone axes and correlation scores are inset into fits. . . . . 68
- 5.5 **Dynamical simulated diffraction patterns.** (a) Example diffraction patterns for Au oriented to the [011] zone axis for 10-80 nm thick slices. (b) Plots showing the mean zone axis misorientation in degrees as a function of thickness for Cu, Ag, and Au. Each plot shows the errors for correlation prefactors of  $q_s|V_g|$  (red) and  $q_s$  (blue). . . . . 69
- 5.6 **4D-STEM scan of twisted polycrystalline AuAgPd nanowires.** (a) Diffraction image of probe over vacuum, showing bullseye pattern. (b) Maximum of each pixel in diffraction space over all probe positions. (c) Histogram of all peak locations detected by correlation in py4DSTEM of (a) with each pattern included in (b). (d) HAADF-STEM image of the sample. (e) 1D histogram of scattering vectors, with fcc AuAg inverse plane spacings overlaid. . . . . 72

5.7	<b>Orientation mapping of polycrystalline AuAgPd nanowires.</b> (a) Total of measured correlation signal for each probe position. (b) Estimated number of patterns indexed for each probe position. (c) Example of 2 orientations indexed from a single diffraction pattern, collected at the position indicated by the arrow shown in (b), with correlation scores inset. (d) Orientation maps of the 3 highest correlation signals for each probe position. A legend for the crystallographic orientation is shown above, and arrows indicate the direction of the x and y axes, while the zone axis direction is out of the page. . . . .	73
5.8	<b>Orientation analysis of grains in AuAgPd nanowires.</b> (a) Crystal grains, with in-plane (111) planes colored by orientation. (b) (111) planes shared by two overlapping grains. . . . .	74
6.1	Virtual images obtained from the 4D-STEM dataset for the direct beam (bright-field) and each of the first order (darkfield) reflections. . . . .	84
6.2	Crystal orientation measured using the kinematic dictionary matching approach of Ophus et al. [124], showing (left) the crystal orientation pointing downward in the plane of the scan and (right) the orientation of the crystal facing the electron beam. The kinematic matching library shows the sample to be single crystalline, and reveals a slight misorientation boundary near the bottom of the scan. . . .	85
6.3	Projections of the 3D cost function onto thickness and orientation dimensions during progression of the iterative refinement procedure for a single diffraction pattern. Each row corresponds to one iteration, and the orientation search range is progressively narrowed around the best match from the previous search (indicated with a blue circle on the tilt projection). The first iteration searches a neighborhood centered on the orientation obtained from the kinematic library matching routine. . . . .	86
6.4	(left) Crystal orientations obtained from the dynamical refinement, shown as mistilt from the (111) crystal orientation (the maximum mistilt shown on the left panel is $1^\circ$ ). (right) Histogram of misorientation of each pixel in the scan. The majority of the crystal is oriented within about $0.2^\circ$ of the perfect zone axis orientation, and the defect area on the bottom left of the scan region is misoriented by approximately $0.5^\circ$ . . . . .	87
6.5	Maps of local sample thickness obtained from the dynamical refinement procedure and compared with low-loss EELS log ratio estimates of the thickness. . . . .	88
7.1	Nanobeam electron diffraction signals from lattices with varying in-plane polarization, indicated by the arrows. Diffraction pattern simulations of (a) thin $\text{PbTiO}_3$ , (b) thick $\text{PbTiO}_3$ , and (c) multilayer with 16:16:16 unit cells of $\text{SrTiO}_3$ : $\text{PbTiO}_3$ : $\text{SrTiO}_3$ . Left to right, the in-plane $\text{PbTiO}_3$ polarization varies smoothly from zero, full left-facing, zero full right-facing, and zero polarization. . . . .	93

- 7.2 Confounding of tilt and polarization signals for a STO:PTO:STO trilayer. (a) Sequences of diffraction patterns simulated at relative  $x$ -direction polarization varying from 0–1 for various mistilts from a perfect [001] orientation. The two diffraction disks conventionally used to measure polarization, (200) and  $(\bar{2}00)$  are highlighted, and the inset text indicates the relative polarization measured as  $\frac{I_{200}-I_{\bar{2}00}}{I_{200}+I_{\bar{2}00}}$ . (b) Line traces of selected diffraction intensities for different mistilts, with the (200) and  $(\bar{2}00)$  reflections highlighted. (c) Map of apparent polarization signal for a trilayer with  $P_x = 1$  as measured from the asymmetry of the (200) and  $(\bar{2}00)$  disks for different mistilts. Contrast reversals in the polarization signal occur with at little as 5 mrad mistilt. The overlaid numbers indicate the tilt values corresponding to the rows of (a). . . . . 100
- 7.3 Derivatives of the diffracted intensities with respect to  $x$ -polarization  $P_x$  and tilt about the  $x$ -axis  $\theta_x$ . The inset numbers show the value of  $\frac{\partial I_{200}}{\partial \square} - \frac{\partial I_{\bar{2}00}}{\partial \square}$  in arbitrary units, where  $\square$  represents polarization for the left 3 columns, and tilt for the right 3 columns. . . . . 102
- 7.4 (a),(b) Polarization and (c),(d) tilt of the STO:PTO:STO sample recovered from the optimization procedure. Approximate polarization directions are labeled above. 103

## Acknowledgments

Since arriving at UC Berkeley and Lawrence Berkeley National Lab in 2017, not a day has gone by that I did not learn and grow as a scientist. This is entirely thanks to the amazing crew of people I have had the delight to work with over the past five years, especially the staff and fellow users at the National Center for Electron Microscopy. Exceptional thanks go to Colin Ophus, Karen Bustillo, Rohan Dhall, and Jim Ciston, for their constant mentorship on all aspects of electron microscopy. Peter Ercius mentored me in the dual arts of scripting control of a TEM and small batch coffee roasting, both of which have been essential to my success here. Help with getting up to speed from Tom Pekin in the early days of my PhD were also critical in making this work possible. A steady stream of customers have passed through my office for a shot of espresso and I am grateful to all of them for giving color to my day and liveliness to my work space. Numerous colleagues and collaborators also made this work possible: Stephanie Ribet, Ben Savitzky, Alex Müller, Lauren Hughes, John Turner, Rebecca Abergel and Mary Scott at LBL; Ruopeng Zhang, Jenna Tan, Paul Corbae, Rebecca Wai, Sahar Saremi, Lei Jin, Sandhya Susarla, Lane Martin, Ramamoorthy Ramesh, and Junqiao Wu at UC Berkeley; Shelly Conroy, Scott Findlay, Hamish Brown, Sven Donges, and Markus Raschke. I will finally extend my greatest thanks to Andy Minor, my advisor and mentor, who took the risk of hiring an unknown and mysterious student and then expertly guided me to where I am today. Everyone mentioned here, and the many I am sure to have forgotten, have my utmost gratitude.

Reaching this point would also not have been possible without the continuous support of my family. I had to move across the country in order to undertake this journey, keeping me away from them for much of the past few years. But with their support I have been able to keep going strong. Especially during the height of the pandemic this separation was difficult, but their support and assurance was constant. Thank you, Mom and Dad.

My graduate work was supported by the National Science Foundation under STROBE Grant No. DMR 1548924. Work at the Molecular Foundry was supported by the Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. Part of this research was performed while visiting the Institute for Pure and Applied Mathematics (IPAM), which is supported by the National Science Foundation (Grant No. DMS-1925919). This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231.

# Chapter 1

## Introduction

Transmission electron microscopy (TEM) is one of the most powerful techniques for studying structures at the nanoscale, both because of its unmatched spatial resolution and the flexibility of the instrument to perform a wide variety of imaging, diffraction, and spectroscopy measurements. In a TEM, a beam of electrons is accelerated to relativistic speeds and focused by a sequence of magnetic lenses onto the sample. Following interaction of the beam with the sample, lenses carry the beam on towards a detector where it is measured. A modern TEM instrument is capable of a wide variety of operating modes: the sample can be illuminated with a plane wave or a focused probe, the scattered wave can be projected from the plane of the sample or into the far field, and detectors of different shapes and sizes can record the scattered electrons. In scanning transmission electron microscopy (STEM), we utilize the focused probe illumination mode and place the detector in the far field, where a converged beam electron diffraction (CBED) pattern is present. In this work, we focus on one particular operating mode of the microscope: 4D-STEM. In this mode, a focused probe is scanned across the sample, and a pixelated detector is used to capture an image of the diffraction pattern at each position. This mode relies on advanced pixelated detector hardware, and so is a modern innovation in the field of TEM. By collecting a detailed image of the diffraction by the sample at every point it has enabled a wide range of structural measurements not possible before.

### 1.1 4D-STEM

In a 4D-STEM experiment, a diffraction pattern is acquired at each scan position of a focused electron probe. This operating mode of the TEM instrument was not practical until the advent of fast digital pixelated detectors, which have frame rates fast enough acquire a useful number of patterns before drift or instability degrade the experiment. The first reported experiment that could be termed 4D-STEM was identified in Ophus' review [123] to be that of Zaluzec [196] in 2002, who recorded diffraction patterns from a 30x30 pixel grid in real space and measured the deflection caused by the Lorentz force in a magnetic

sample. Today, further advances in detector technology have enabled much larger scans to be acquired in less time, and various analyses have been subsequently developed to make use of these large quantities of diffraction data. The simplest analysis, and one often conducted routinely from datasets that are also used for more sophisticated analyses, is virtual imaging. In this analysis, intensities from regions of diffraction space are summed to produce images corresponding to various arbitrary bright- or dark-field images [49, 195, 76]. Structural measurements dominate the materials science applications of 4D-STEM. By matching the diffracted spots in each pattern to the crystal structure of a material, orientation maps can be produced [124, 160, 82, 37, 149]. By precisely determining the position of each diffracted disk the local lattice parameters can be computed, and their change can be mapped over a field of view to image the strain in a sample [37, 147, 149, 148]. Disordered samples can also be analyzed, for example using fluctuation microscopy to detect short- and medium-range ordering [16, 177, 93, 51]

## Nanobeam 4D-STEM

The experiments and analysis in this work will all utilize a subset of 4D-STEM, which we refer to as “nanobeam” diffraction. This term refers to the configuration where the convergence angle of the electron beam is set such that the Bragg disks in a CBED pattern do not overlap. This configuration is compatible with many of the structural measurements of interest to materials scientists, such as strain and orientation mapping and electric field measurements. By avoiding overlap between diffracted disks, we obtain a diffraction pattern where different scattering vectors do not interfere coherently. This coherent interference transfers information at sub-unit-cell length scales, and is the basis of differential phase contrast and ptychography at atomic resolution. However, this interference greatly complicates the simulation of the diffraction pattern as each plane wave in the converged probe can no longer be considered completely independently. Where the diffraction disks overlap their intensity is not simply summed, and so finding the position and intensity of each diffraction disk independently is no longer feasible. 4D-STEM experiments where diffraction disks overlap and interfere coherently are often termed phase-contrast measurements, since the intensities in the overlap regions depend on the relative phase between the interfering beams.

## 1.2 Multiple Scattering

Thanks to the strong interaction between fast electrons and the sample, the diffraction patterns recorded in a 4D-STEM experiment contain a wealth of information about the material structure. But the downside of this strong interaction is that in most samples of materials science interest, the electron interacts with the sample multiple times. The multiple scattering of the probe greatly complicates the analysis of the diffraction patterns, as the simple kinematic model of scattering that only depends on the geometry of the problem must be replaced with the dynamical model which is fully quantum mechanical. The dynamical



model is challenging to solve, but numerous strategies for computing diffraction patterns with it are available and will be discussed in this work. Inverting the dynamical scattering problem is an open area of research, particularly for atomic resolution data [30, 47, 133]. In the work presented here we will explore dynamical inversion of nanobeam patterns, a problem which is far simpler and that we will show can be performed from large scale 4D-STEM data and thus be applied to practical materials science problems.

In electron microscopy, the term “multiple scattering” is usually taken to encompass the phenomenon where the electron probe undergoes any two or more scattering events. Thus this term also encompasses inelastic scattering and other phenomena. In this work, we are specifically concerned with the situation where the electron scatters in multiple elastic events: this is termed as dynamical diffraction, and we will maintain this distinction throughout the work. Many instruments now place their 4D-STEM detectors behind energy filters, which make it possible to exclude inelastically scattered electrons from the recorded patterns, so that the only multiple scattering observed is dynamical diffraction.

### 1.3 Overview

The focus of this work is on advancements in the analysis of 4D-STEM data for materials science, particularly on developments in the analysis of diffraction data from thick and complex samples to extract robust information that is useful to materials science. 4D-STEM diffraction data contains a wealth of structural information about a material, but extracting this information in a reliable manner is a challenge due to the scale of the data obtained in each experiment. Unlike conventional diffraction experiments, in 4D-STEM we typically analyze patterns of relatively low dose, since long integration times mean that acquiring a scan from a large region is prohibitively slow. Further, while a conventional TEM diffraction analysis might examine a handful of converged beam diffraction patterns, in 4D-STEM we typically find ourselves analyzing tens to hundreds of thousands of diffraction patterns.

In this work, we will explore several developments that aid in overcoming these challenges and making 4D-STEM a useful and robust technique for materials science. The work is organized as follows: In Chapter 2, we will review the physics of electron scattering, with a particular focus on computational aspects related to dynamical scattering, as much of this work will rely on using large-scale computation to interpret diffraction data.

Chapter 3 addresses the challenge of low dose per diffraction pattern in 4D-STEM through the development of a high-performance implementation of a hyperspectral total variational denoising algorithm. This approach to denoising exploits the high dimensionality of the data to allow neighboring pixels to share information, but without blurring of sharp edges. Our implementation is highly optimized to quickly denoise scanning diffraction data, and can also be used for even the largest datasets currently being collected by using distributed computing.

One of the greatest challenges to 4D-STEM analysis in materials science is that it is often impossible to produce samples of a given material that are thin enough that the electron

beam will only scatter one time as it traverses the sample. This means that the scattering in most of the 4D-STEM datasets in materials science is dynamical, and this greatly hinders interpretation of the diffraction patterns. Artifacts like phantom strain and inversion of electric field contrast can lead to spurious interpretations of the physics of a material, and must be addressed to make 4D-STEM robust. The bulk of this dissertation will focus on this challenge.

In Part I we will first meet dynamical diffraction as a foe, and explore two measurements that can be made robust *despite* dynamical effects. In Chapter 4, we address the problem of strain mapping from thick samples, where dynamical diffraction causes unwanted contrast inside of diffraction disks that makes it impossible to precisely find their locations. Using special patterned apertures, we are able to impart known, strong contrast onto the diffraction disks, which allows them to be located with much higher precision. This approach allows us to obtain much more precise strain maps from thick samples and using a lower electron dose. In Chapter 5, we explore the problem of determining the orientation of a crystalline sample from 4D-STEM diffraction patterns. We devise an efficient algorithm for matching orientations against a simulated library, but we find that dynamical effects limit the maximum attainable precision.

In Part II, we will instead meet dynamical diffraction as a friend, and explore two measurements that can be made robust by *inverting* dynamical effects. In this part, we utilize large scale computational modeling of the dynamical diffraction process, and develop efficient algorithms for recovering structural properties from 4D-STEM data that contains significant dynamical scattering. In both cases we find that we can achieve better precision in our measurements by utilizing, rather than suppressing, dynamical diffraction effects. In Chapter 6, we extend the orientation mapping algorithm developed in Chapter 5 to include dynamical scattering using an adaptive grid search and Bloch wave simulations. We find that not only can we match orientations from thick samples and low-dose measurements with higher precision, but we can also obtain a measurement of the thickness of the sample simultaneously with the refined orientation. In Chapter 7, we take this approach further by analyzing a sample consisting of multiple distinct crystal layers along the beam direction. We develop a multiple scattering model based on the multilayer structure and parameterize it over a number of physically relevant structural distortions. We then develop an optimization procedure for matching these parameters to 4D-STEM experimental data, and use it to recover the structure of polar vortices in the material without interference from tilting of the foil.

# Chapter 2

## Theory

The scattering of fast electrons by crystalline materials has been studied since the early days of the electron microscope, with Bethe providing a description of dynamical scattering in 1928 [14], even before Ruska’s electron microscope was constructed [81] in 1932. In this section, we will give an overview of the process of electron diffraction in the context of 4D-STEM, focusing on effects that are important for nanobeam diffraction and giving less emphasis to effects that dominate in plane wave or highly converged electron probes. Numerous modern textbooks give a full account of the material we will discuss here. Fultz and Howe provide a good foundation beginning from x-ray and neutron scattering and building up to electron diffraction [48]. Zuo and Spence<sup>1</sup> provide a comprehensive account of modern electron microscopy, with detailed descriptions of scattering by partially ordered materials, thermal scattering, and defect contrast [205]. Kirkland provides a particular focus on the multislice formalism [79], while De Graef focuses on the Bloch wave description [40]. These two books are particularly invaluable for practical computations as they are accompanied by full-featured computer programs for performing simulations using their chosen methods.

### 2.1 Electron Scattering

#### Atomic Form Factors

The fundamental interaction that gives rise to all of the crystalline diffraction phenomena we will discuss is the single-atom scattering of the incoming electron. Electrons are charged particles, and so scatter off of the electrostatic potential around each atom. The electrostatic potential  $V^a(\mathbf{r})$  is determined by both the electron density  $\rho_e(\mathbf{r})$  and the nuclear charge density  $\rho_n(\mathbf{r})$  as

$$\Delta V^a(\mathbf{r}) = -\frac{|e|}{\epsilon_0} [\rho_n(\mathbf{r}) - \rho_e(\mathbf{r})] \quad (2.1)$$

---

<sup>1</sup>I would like to thank John Spence who, shortly before his passing in 2021, gifted me a copy of this textbook after we were introduced at the Microscopy and Microanalysis virtual conference.

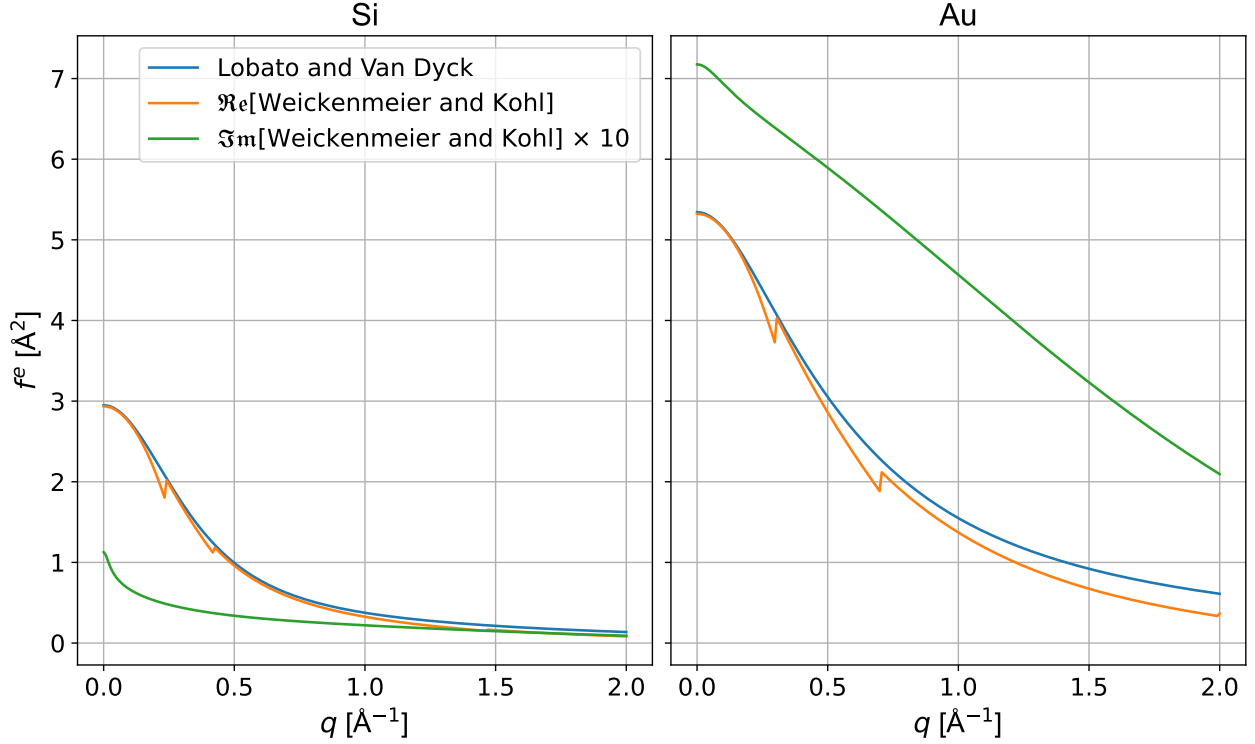


Figure 2.1: Relativistic corrected (300 kV) single atom scattering factors for Si and Au using two different parameterizations. The Lobato and Van Dyck [98] method contains only an elastic part and is purely real, while the Weickenmeier and Kohl method simulated absorption due to phonon and core loss interactions as well as damping by the Debye-Waller factor [183].

where  $\Delta$  is the Laplacian operator,  $e$  is the electron charge, and  $\epsilon_0$  is the permittivity of free space. The single atom electron scattering factor is defined as the probability for an electron plane wave with direction  $\mathbf{k}$  to be scattered by the electrostatic potential into the direction  $\mathbf{k}'$ , and is equal to the Fourier transform of the atomic potential

$$f^e(\Delta\mathbf{k}) = \int_{\mathbb{R}^3} V^a(\mathbf{r}) \exp[-2\pi i \Delta\mathbf{k} \cdot \mathbf{r}] d^3\mathbf{r} \quad (2.2)$$

where  $\Delta\mathbf{k} = \mathbf{k}' - \mathbf{k}$  [40]. The form factors are usually assumed to be spherically symmetric, and so are typically written to depend simply on  $k$ , the magnitude of the change in direction.

The electron scattering factors can be related to the widely tabulated x-ray scattering factors by the inclusion of the nuclear potential (as x-rays do not interact with the nuclear charge). However, for electron microscopy there are a number of parameterizations of the scattering factors which are commonly used in calculations. The most modern parameterization is the one published by Lobato and Van Dyck [98], which has the correct asymptotic

behaviors and is highly smooth. Kirkland also has a modern and widely-used parameterization, which is often the one chosen for multislice calculations [79].

It is possible to define an effective single atom scattering factor that takes into account inelastic phenomena, by adding an imaginary component  $f'(\mathbf{k})$  to the elastic factor.

$$f^e(\mathbf{k}) = f_{\text{elastic}}^e(\mathbf{k}) + i f'(\mathbf{k}) \quad (2.3)$$

These are referred to as absorptive potentials, and their computation and parameterization are more challenging as the physics of inelastic scattering and absorption are far more complex than the elastic scattering we have described here. Early on, Hashimoto and coworkers proposed a simple model where the inelastic component was assumed to be simply proportional to the elastic factor, and they found that a proportionality constant of 0.1 was generally in agreement with experiments [64]. A commonly used model of the absorptive component was derived by Hall and Hirsch [61], who used the Einstein model to derive the following absorptive form factor due to thermal diffuse scattering

$$f'(k) = \frac{1}{k} \int f(\mathbf{q}) f(\mathbf{q} - \mathbf{g}) (\exp[-Mg^2] - \exp[-M[q^2 - (\mathbf{q} - \mathbf{g})^2]]) d^2 \mathbf{q} \quad (2.4)$$

While it is possible to evaluate this integral numerically, Weickenmeier and Kohl give a useful parameterization of the model [183]. A FORTRAN code which computed the factors from the parameterization was released along with their paper, which was freely available at the time of publication but is no longer directly accessible as the BITNET network for contacting the authors has not existed for over a decade [184] and  $5\frac{1}{4}$  in floppy disks are no longer routinely distributed. Luckily this code lives on as part of the open-source EMSOFT distribution [41], and has been translated to Python as part of `py4DSTEM` [158]. Both of these implementations also include code to compute a core-loss contribution, supplied to De Graef by Weickenmeier and Kohl and following the theory by Rose [151].

A comparison of the single atom scattering factors for Si and Au is shown in Figure 2.1, using both the Lobato-Van Dyck (L-VD) and Weickenmeier-Kohl (W-K) parameterizations. The L-VD factors and the real part of the W-K factors are in close agreement, though the L-VD parameterization is smoother because the W-K potentials are defined over multiple domains of the argument and are not necessarily continuous. The W-K factors are lower at high scattering angle because they are computed with Debye-Waller damping while the L-VD values are not (due to implementation differences in `py4DSTEM`). The imaginary part of the W-K factors are also shown in the plots, scaled up for clarity, and have a distinct dependence on scattering magnitude that differs from the elastic part.

In addition to these parameterizations of the atomic form factors, it is also possible to derive factors from more comprehensive models of the electron density in an atom. Wu et al used multipole density functions to allow the electron density to have anisotropy due to bonding effects [188], and was able to fit the parameters of this model to converged beam electron diffraction patterns. Their model also included the effect of charge transfer on the atomic potentials, which allows the form factor for atoms which gain electrons to become




Figure 2.2: Comparison of IAM and DFT potentials for a unit cell of  $\text{PbTiO}_3$ , computed using abTEM and GPAW.

negative. They found the effect of the more accurate form factors to be most prominent at large scattering angles, and that the different models produce a detectable difference in the diffraction patterns.

## Crystal Potential

In a crystalline sample, the potential that interacts with the incoming electron wave is comprised of contributions from each of the atoms. Using the parameterized potentials described above, we typically proceed using the independent atom model (IAM), which assumes that the potential of the crystal is simply the sum of the potentials at each atomic site. For a unit cell containing  $N$  atoms at positions  $\mathbf{r}_j$ , the potential of the unit cell  $V_{\text{cell}}(\mathbf{r})$  is thus given by

$$V_{\text{cell}}(\mathbf{r}) = \sum_{j=1}^N V_j^a(\mathbf{r} - \mathbf{r}_j) \quad (2.5)$$

This potential has the periodicity of the lattice and so can be expanded as a Fourier series over the reciprocal lattice points  $\mathbf{g}$

$$V(\mathbf{r}) = \sum_{\mathbf{g}} V_{\mathbf{g}} \exp[2\pi i \mathbf{g} \cdot \mathbf{r}] \quad (2.6)$$

which holds for an infinite crystal, where the potential has discrete peaks at the reciprocal lattice points and is zero elsewhere in reciprocal space.

In any other case than an infinite perfect crystal, the potential is a function of a continuous reciprocal coordinate  $\mathbf{q}$ :

$$V(\mathbf{q}) = \frac{1}{\Omega} \int_{\mathbb{R}^3} V(\mathbf{r}) \exp[-2\pi i \mathbf{q} \cdot \mathbf{r}] d^3\mathbf{r} \quad (2.7)$$

It is common to separate the potential as a convolution between a unit-cell contribution corresponding to an infinite crystal and a shape function  $\mathcal{T}(\mathbf{r})$  that modifies the extents of the crystal, so that

$$V(\mathbf{q}) = \frac{1}{\Omega} \left[ \sum_{j=1}^N f_j^e \exp(-2\pi i \mathbf{g} \cdot \mathbf{r}) \right] [\mathcal{T}(\mathbf{r})] \quad (2.8)$$

In addition to the above expression using the single atom scattering factors, it is also possible to derive the crystal potential from *ab initio* calculations. This method is able to fully capture the effects of charge transfer and bonding, which is particularly important for thin samples where it may be possible to image these effects directly [103, 171]. *ab initio* potential calculations are implemented as part of the abTEM multislice simulation code [102], which uses GPAW to perform its density functional theory (DFT) calculations [113, 86]. A comparison of DFT and IAM potentials for a unit cell of PbTiO<sub>3</sub> with distortions induced by a built-in polarization is shown in Figure 2.2. In this example, there are large differences between the two potential calculations but they are largely confined to the regions between the atoms where the density is low. The electrostatic potential is dominated by the nuclear contribution and core electrons, which are not changed drastically due to bonding, and so at the atomic sites the relative difference is negligible. Diffraction simulations performed using these potentials resulted in a difference in intensities of at most 0.1%. However, for light elements the valence contributions can become important.

## 2.2 Kinematic Diffraction

Diffraction is a phenomenon that arises from coherent interference of wave scattered by a regular arrangement of scatterers. By assuming that the scattering interaction is weak, we can develop a simple description of the diffraction process which is referred to as the “kinematic” theory, as it relies only on the geometry of the scatterers (as opposed to the “dynamical” theory, which also accounts for propagation). Derivation of the kinematic theory uses the first Born approximation, which assumes that the wave does not lose intensity upon scattering and is scattered only once. A complete derivation of the equation for the scattered wave in the first Born approximation is beyond the scope of this work, and can be found in the textbook of Fultz and Howe [48]. The wavefunction  $\psi$  of an electron originally entering as a plane wave with direction  $\mathbf{k}$  and scattered by a potential is written as

$$\psi_{\text{scattered}}(\Delta\mathbf{k}, \mathbf{r}) = \frac{-m}{2\pi\hbar^2} \frac{e^{i\mathbf{k}\cdot\mathbf{r}}}{|\mathbf{r}|} \int_{\mathbb{R}^3} V(\mathbf{r}') \exp[-2\pi i \Delta\mathbf{k} \cdot \mathbf{r}'] d^3\mathbf{r}' \quad (2.9)$$

where  $m$  is the electron mass and  $\hbar$  is the reduced Planck constant. An intuitive explanation for this relationship is that the potential is comprised of an ensemble of individual scatterers, and upon illumination by a plane wave each of these scatterers coherently emits a spherical wave, and these waves interfere at some location  $\mathbf{r}$  away from the sample. Since we are generally not interested in the near-field wave, we can use a coordinate transformation (again,

with the details to be found in [48]) to obtain an expression that only depends on  $\Delta\mathbf{k}$ . If we assume that the scatterers are individual atoms in a perfect crystal, scattered wave simplifies to

$$\psi(\Delta\mathbf{k}) = \sum_{j=1}^N f_j^e(\Delta\mathbf{k}) e^{-2\pi i \Delta\mathbf{k} \cdot \mathbf{r}_j} \quad (2.10)$$

Finally, the scattered intensity is simply

$$I_{\text{scattered}} = \psi_{\text{scattered}}^* \psi_{\text{scattered}} \quad (2.11)$$

The kinematic theory brings us to a simple conclusion: the diffracted wave is proportional to the Fourier transform of the distribution of scatterers. The diffracted wave arises from coherent superposition of waves scattered by each atom (or, in a continuous sense, from each voxel of potential). In a perfect crystal, these waves interfere coherently only along directions corresponding to reciprocal lattice points, and relativistic effects cause the intensity scattered by more than a few tens of milliradians to be negligible.

## 2.3 Dynamical Diffraction

Dynamical diffraction is derived not from the geometry of the scattering medium but from the Schrödinger equation that describes the evolution of a quantum mechanical wave through a potential [40]. There are many conventions for writing the Schrödinger equation for a relativistic electron moving along the  $z$  direction through a potential and scattering by small angles, all of which are equivalent. Here we will use the notation of Kirkland [79], who gives the equation as

$$\frac{\partial \psi(\mathbf{r})}{\partial z} = \left[ \frac{i\lambda}{4\pi} \nabla_{xy}^2 + i\sigma V(\mathbf{r}) \right] \psi(\mathbf{r}) \quad (2.12)$$

where  $\lambda$  is the relativistically corrected wavelength,  $\nabla_{xy}^2$  is the Laplace operator in 2D, and  $\sigma$  is the relativistically corrected interaction constant.

Let us examine the three components of Equation 2.12 in turn. The first term, on the left side of the equality, describes the evolution of the wave as it progresses along the  $z$  axis. This description of the scattering is a differential equation along a spatial dimension, and so it will include the effect of propagation of the wave. The two terms on the right then describe what happens to the wave as it propagates. The Laplace operator  $\nabla_{xy}^2$  in the second term is a double spatial derivative—its magnitude is large where the curvature of the wave is large, and vice versa. This term thus describes the spreading out of the wave as it propagates, causing the wavefront to tend to become smoother. The third term encapsulates interaction with the sample potential, and causes the wave to gain a phase factor proportional to the potential  $V(\mathbf{r})$ . The constant of proportionality for this interaction is  $\sigma$ , defined as

$$\sigma = \frac{2\pi\gamma m_0 e\lambda}{h^2} \quad (2.13)$$



where  $\gamma$  is the relativistic mass ratio,  $m_0$  and  $e$  are the rest mass and charge of the electron, and  $h$  is Planck's constant.

Evaluation of this governing equation for dynamical diffraction is nontrivial. Examining this equation, we observe that it takes the form of  $\frac{\partial \psi}{\partial z} = \mathcal{A}\psi$ , where for the moment we will understand  $\mathcal{A}$  to be an operator<sup>2</sup> consisting of the terms in the square brackets in 2.12. On this basis, the solution should have the form of

$$\psi(z) = e^{i\mathcal{A}z}\psi(0) \quad (2.14)$$

If we write  $\mathcal{A}$  as a combination of two operators,  $\mathcal{A} = \mathcal{T} + \mathcal{V}$ , where  $\mathcal{T}$  corresponds to the transmission term and  $\mathcal{V}$  to the potential interaction term inside the square brackets, the exponential solution becomes

$$\psi(z) = e^{i\mathcal{A}z}\psi(0) = e^{i(\mathcal{T}+\mathcal{V})z}\psi(0) = e^{i\mathcal{T}z}e^{i\mathcal{V}z}\psi(0) \quad \leftarrow \text{this is wrong!}$$

Unfortunately, the two terms encapsulated by  $\mathcal{A}$  do not commute, and we cannot evaluate this solution directly. In other words, dynamical diffraction is not simply kinematic diffraction plus a spreading term—it is the processes of spreading and interacting occurring simultaneously.

There are two common methods for solving Equation 2.12 and computing diffraction intensities, and we will explain them here as representing different approaches to performing the exponentiation in Equation 2.14. The first, and most common in practice today, is referred to as the multislice method. The multislice method takes advantage of the fact that over a very small distance  $\epsilon$ , it is indeed true that

$$e^{i(\mathcal{T}+\mathcal{V})\epsilon} = e^{i\mathcal{T}\epsilon} \times e^{i\mathcal{V}\epsilon} \quad (2.15)$$

Thus it is valid to write

$$\psi(z) = e^{i\mathcal{A}z}\psi(0) = e^{i\mathcal{T}\epsilon}e^{i\mathcal{V}\epsilon}e^{i\mathcal{T}\epsilon}e^{i\mathcal{V}\epsilon} \dots e^{i\mathcal{T}\epsilon}e^{i\mathcal{V}\epsilon}e^{i\mathcal{T}\epsilon}e^{i\mathcal{V}\epsilon}\psi(0) \quad (2.16)$$

where the products are evaluated enough times such that the total of the small distances  $\epsilon$  equals the sample thickness. The second method is the Bloch wave approach, where we write the wavefunction as the sum of Bloch waves, which are plane waves modulated by a function that has the periodicity of the lattice. This choice of a representation is guaranteed, by Bloch's theorem, to satisfy the Schrödinger equation in a periodic potential. This representation is beneficial because it leads to a matrix representation of  $\mathcal{A}$  which is diagonalizable. The exponential of a matrix that has been diagonalized as  $\mathbf{U}\mathbf{\Lambda}\mathbf{U}^{-1}$  is easy to compute:

$$e^{\mathbf{U}\mathbf{\Lambda}\mathbf{U}^{-1}} = \mathbf{U}e^{\mathbf{\Lambda}}\mathbf{U}^{-1} \quad (2.17)$$

and since  $\mathbf{\Lambda}$  is diagonal its exponential can be evaluated elementwise. In the next section, we will explore these two methods in further detail.

---

<sup>2</sup>It is possible, using either the Darwin-Howie-Whelan form or the Bloch wave form (as will be seen below), to write  $\mathcal{A}$  as a matrix, in which context the equations in this paragraph are more rigorously true. Here I am calling them “operators” since they represent the transformations the matrices impart to the wave on multiplication, which is an abuse of notation.

## 2.4 Computation

As described above, computation of diffraction intensities under dynamical diffraction conditions are nontrivial. In this section, we will provide an overview of the two primary methods for solving Equation 2.12 for a given sample and experimental condition. We will also discuss the PRISM algorithm, which is a variant of the multislice algorithm which borrows concepts from the Bloch wave approach and is widely used in 4D-STEM simulations.

### Multislice Method

As described above, the multislice method relies on evaluating the two terms inside the brackets of Equation 2.12 over small distances in sequence. In this section, we will describe the procedure for evaluating these operators. For a complete description of the method, with careful attention to its implementation, we refer to the textbook by Kirkland [79].

The first operator in Equation 2.12, which represents propagation of the electron wave, is challenging to solve in real space because of the presence of the double differential operator. However, there is a simple expression for evaluating it in reciprocal space, where derivatives become a multiplication by the frequency coordinate [40]. Evaluation of the propagation operator is thus given by

$$\psi(z + \epsilon) = \mathcal{F}^{-1} \left\{ e^{i\lambda\epsilon k^2} \mathcal{F} \{ \psi(z) \} \right\} \quad (2.18)$$

where  $\mathcal{F}$  is the Fourier transform and  $\mathcal{F}^{-1}$  is the inverse Fourier transform.

The second operator, which represents interaction with the sample potential, is simple to evaluate. First, we project the potential within the slice (of thickness  $\epsilon$ ) onto the  $xy$ -plane:

$$V_\epsilon(z) = \int_{z-\epsilon/2}^{z+\epsilon/2} V(\mathbf{r}) dz \quad (2.19)$$

The action of the operator is to add a phase to the wave proportional to this slice of the potential

$$\psi(z + \epsilon) = e^{i\sigma V_\epsilon(z)} \psi(z) \quad (2.20)$$

If we denote the  $j$ -th potential slice along the thickness of the sample as  $V_\epsilon^{(j)}$ , the full multislice computation is written as

$$\psi(z) = \prod_{j=1}^N \left\{ \mathcal{F}^{-1} \left[ e^{i\lambda\epsilon k^2} \left\{ \mathcal{F} \left[ e^{i\sigma V_\epsilon^{(j)}} \right] \right\} \right] \right\} \psi(0) \quad (2.21)$$

### Bloch Wave Method

Here we will summarize the Bloch wave method, following the notation given in the textbook by DeGraef [40]. Each Bloch wave  $\Psi^{(j)}(\mathbf{r})$ , indexed by the superscript ( $j$ ), takes the form of

a plane wave (with wavevector  $\mathbf{k}^{(j)}$ ) modulated by a function  $C^{(j)}(\mathbf{r})$  with the periodicity of the lattice:

$$\Psi^{(j)}(\mathbf{r}) = C^{(j)}(\mathbf{r})e^{2\pi\mathbf{k}^{(j)}\cdot\mathbf{r}} = \sum_{\mathbf{g}} C_{\mathbf{g}}^{(j)} e^{2\pi i(\mathbf{k}^{(j)}+\mathbf{g})\cdot\mathbf{r}} \quad (2.22)$$

In the rightmost equality, the Bloch wave is written as a Fourier series with coefficients  $C_{\mathbf{g}}$ . The restriction that  $C^{(j)}$  has the periodicity of the lattice allows us to write the Fourier series on the right hand side as a sum over the reciprocal lattice of the crystal, i.e. the  $\mathbf{g}$  vectors of the Fourier decomposition of the Bloch wave are the same as the diffraction vectors of the crystal. The total wavefunction in the crystal is a combination of these Bloch waves, weighted by their excitation amplitudes  $\alpha^{(j)}$

$$\begin{aligned} \psi(\mathbf{r}) &= \sum_j \alpha^{(j)} C^{(j)}(\mathbf{r}) e^{2\pi\mathbf{k}^{(j)}\cdot\mathbf{r}} \\ &= \sum_j \alpha^{(j)} \sum_{\mathbf{g}} C_{\mathbf{g}}^{(j)} e^{2\pi i(\mathbf{k}^{(j)}+\mathbf{g})\cdot\mathbf{r}} \end{aligned} \quad (2.23)$$

Substituting the Bloch wave description into the Schrödinger equation yields a set of relationships, one for each Bloch wave, between the wavevector  $\mathbf{k}$  of the Bloch wave and the incident electron wavevector  $k_0$ :

$$[k_0^2 - (\mathbf{k} + \mathbf{g})]C_{\mathbf{g}} + \sum_{\mathbf{h} \neq \mathbf{g}} U_{\mathbf{g}-\mathbf{h}} C_{\mathbf{h}} = 0. \quad (2.24)$$

$U_{\mathbf{g}-\mathbf{h}}$  is the Fourier component of the sample electrostatic potential corresponding to the scattering vector  $\mathbf{g} - \mathbf{h}$ . By applying continuity of the wavefunction and its first derivative at the entrance and exit planes of the crystal, ignoring backscattered electrons, and assuming high incident electron energy, it is possible to write the Bloch wave equations as [40]:

$$2k_0 s_{\mathbf{g}} C_{\mathbf{g}}^{(j)} + \sum_{\mathbf{h} \neq \mathbf{g}} U_{\mathbf{g}-\mathbf{h}} C_{\mathbf{h}}^{(j)} = 2k_n \gamma^{(j)} C_{\mathbf{g}}^{(j)} \quad (2.25)$$

where  $k_n$  is the normal component of the incident electron wave and  $\gamma^{(j)}$  is the normal component of the Bloch wave:  $\mathbf{k}^{(j)} = \mathbf{k}_0 + \gamma^{(j)}\mathbf{n}$ .  $s_{\mathbf{g}}$  is the excitation error for the beam  $\mathbf{g}$ , given by

$$s_{\mathbf{g}} = \frac{-\mathbf{g} \cdot (2\mathbf{k} + \mathbf{g})}{2|\mathbf{k} + \mathbf{g}| \cos \alpha} \quad (2.26)$$

where  $\mathbf{k}$  is the wavevector of the incident electron beam, and  $\alpha$  is the angle between the sample normal and the incident beam direction. Solution of this system of equations is possible by casting it as an eigenvalue/eigenvector decomposition. In this form the equations are written as

$$\bar{\mathcal{A}}\mathbf{C}^{(j)} = 2k_n \gamma^{(j)} \mathbf{C}^{(j)} \quad (2.27)$$

where  $\mathbf{C}^{(j)}$  is interpreted as a matrix whose column vectors contain the Bloch wave coefficients  $C_{\mathbf{g}}^{(j)}$ . The matrix  $\bar{\mathcal{A}}$  contains the term  $2k_0s_{\mathbf{g}}$  on its diagonal and  $U_{\mathbf{g}-\mathbf{h}}$  on the off-diagonal.

$$\bar{\mathcal{A}} = \begin{bmatrix} 0 & U_{-\mathbf{g}} & \cdots & U_{-\mathbf{h}} \\ U_{\mathbf{g}} & 2k_0s_{\mathbf{g}} & \cdots & U_{\mathbf{g}-\mathbf{h}} \\ \vdots & \vdots & \ddots & \vdots \\ U_{\mathbf{h}} & U_{\mathbf{h}-\mathbf{g}} & \cdots & 2k_0s_{\mathbf{h}} \end{bmatrix} \quad (2.28)$$

When absorption is included, the imaginary component of the potential  $U'_{\mathbf{g}}$  is included in the matrix  $\bar{\mathcal{A}}$  by adding  $iU'_{\mathbf{0}}$  to the diagonal and  $iU'_{\mathbf{g}-\mathbf{h}}$  to the off-diagonal.

In this form, the matrix  $\bar{\mathcal{A}}$  is fully determined by the structure of the material and the crystal orientation—by computing the eigenvalue/eigenvector decomposition of this matrix, we obtain the Bloch wave coefficients  $C_{\mathbf{g}}^{(j)}$  and the normal components  $\gamma^{(j)}$ . The final step is to obtain the scattered wave amplitudes for a given crystal thickness. The electron wave  $\psi$  at a depth  $z$  in the crystal is

$$\psi(z) = \mathcal{C}\mathcal{E}(z)\mathcal{C}^{-1}\psi(0) = \mathcal{S}\psi(0) \quad (2.29)$$

where  $\mathcal{C}$  is the matrix containing the eigenvectors,  $\psi(0)$  is a vector with the value of 1 at the index corresponding to the incident beam direction and zero elsewhere, and  $\mathcal{E}(z) = e^{2\pi i\gamma^{(j)}z}\delta_{ij}$  is a diagonal matrix which depends on the thickness and the Bloch wave normal components. We have also introduced the scattering matrix  $\mathcal{S}$  which maps the vector representing the incident electron wave to the scattered wave.

### Implementation in **py4DSTEM**

The approach above allows us to compute the dynamical diffraction intensities for a perfect crystal for a given orientation and thickness. We have implemented this method in **py4DSTEM** as part of the Crystal module, extending the kinematic diffraction calculation routines that were developed there by Ophus et al [124].

First, the crystal structure is input by specifying the lattice and atomic basis, importing from a CIF file, or by searching the Materials Project database. From the crystal structure, all reciprocal lattice points are computed up to a maximum scattering vector  $k_{\max}$  specified by the user. For calculation of kinematic diffraction patterns we also compute the Fourier coefficients of the sample potential at each reciprocal lattice point, using the atomic form factors as parameterized by Lobato and Van Dyck [96].

In principle, one could include every reciprocal lattice point (up to  $k_{\max}$ ) in a Bloch wave calculation, and obtain the diffraction intensities for every reflection. However, the computational complexity of the eigendecomposition used in the Bloch method scales as the third power of the number of beams included, so the computation time can be massively reduced by first evaluating which beams may be excited for a given sample orientation and only performing the Bloch computation on these beams. In **py4DSTEM** this is performed by first

computing a kinematic diffraction pattern: we input the orientation of the crystal and choose thresholds for the maximum excitation error  $s_g$ , and compute a kinematic pattern containing a list of the reflections that meet the thresholds along with their kinematic intensities. The datastructure for this return data contains a list of reflections, with the  $hkl$  indices, kinematic diffraction intensity (computed as the magnitude of the Fourier coefficient of the reflection damped by a function depending on the excitation error), and projected locations in the diffraction plane  $(q_x, q_y)$  for each. Further reduction of computational complexity can be achieved by sorting the reflections into weak and strong beams and eliminating the couplings between weak beams, as described in [178], but this is not currently implemented in `py4DSTEM`.

Several other Bloch wave implementations have been described in the literature: EMSOFT by DeGraef [41] is open source, under continued development, and served as a reference for our own implementation; Zuo et al maintain a cloud-based, paid version of their closed-source implementation; and several accelerated schemes have been reported such as the GPU-based one in [138].

## Modeling Absorption

Accurately reproducing the dynamical diffraction intensities for thick samples requires treatment of diffuse scattering, which causes electrons to scatter out of the Bragg beams and into the pattern background. In multislice simulations, this effect is captured by the “frozen phonon” method, which involves simulating multiple configurations of the atoms using a purely elastic (i.e. non-absorbing) potential but with each configuration having the atoms slightly displaced from their ideal positions by a thermal vibration factor. This approach reproduces diffuse scattering contrast outside the Bragg spots such as Kikuchi lines. The potential is generally computed using the elastic form factors for isolated atoms. Several parameterizations of these form factors are available, with the most common ones in current use being those by Lobato et al [96], Kirkland [79], and Weickenmeier and Kohl (which we will refer to as WK) [183]. These elastic form factors are very well supported by experiments using both electrons and X-rays, and when used with the frozen phonon method they yield generally similar diffraction patterns that agree well with experiments [98].

By contrast, in the Bloch wave approach off-Bragg scattering is not explicitly included and absorption is reproduced by instead adding an imaginary component to the atomic form factors, and thus to the sample potential. These absorptive form factors are generally derived from assumptions about the absorption mechanisms and not from experiments. Early attempts to determine absorptive form factors simply modeled the imaginary part as being proportional to the real part, and a factor of 0.1 was found to give the best match to experiments [64]. Later work found that the absorptive part of the potential is much more tightly localized to the atomic sites than the elastic part, leading to substantial experimental errors from the “proportional” approach, and that applying the Einstein model of atomic vibrations allows computation of more accurate absorptive form factors [152]. The Weickenmeier-Kohl parameterization of the elastic form factors was designed to admit an analytical evaluation

of the Einstein absorption model, leading to the WK form factors being predominantly used in absorptive Bloch calculations. The original Fortran code for the WK model is available open-source as a component of the EMSOFT library [41] that accompanies the excellent textbook by DeGraef [40]; we have re-implemented the WK subroutines in Python as part of the Bloch wave computation functionality of py4DSTEM [158]. These routines also have the option of including a “core” contribution to the absorption due to electron energy loss, using the model of [151].

## PRISM

In the above description of the computational methods for computing diffraction patterns, we considered the problem of computing a single pattern. In the case of multislice, we made no assumptions on the content of the initial electron wave, while for the Bloch method our derivation assumed plane wave illumination. When simulating 4D-STEM experiments, we aim to generate a sequence of diffraction patterns for each scan position of a converged beam electron probe. A natural strategy for accomplishing this in the multislice method is to generate a separate incident wave representing the shifted, converged probe, for each scan position. This particular multislice approach to 4D-STEM has computation time that scales with *the number of probe positions*. In contrast, since the Bloch wave method can only produce plane wave patterns<sup>3</sup>, the strategy is different. In this case, we instead perform a separate Bloch wave computation for each plane wave inside the probe-forming aperture. These plane waves are then combined to form the position-dependent diffraction patterns by summing them coherently, weighted by phase factors that shift the resulting wave to the desired scan position. This particular Bloch wave approach to 4D-STEM has computation time that scales with *the number of pixels inside the probe-forming aperture*.

Because the multislice method is able to simulate samples other than perfect crystals, it is generally preferred in modern practice to produce simulations that faithfully represent complex experiments. However, to simulate 4D-STEM data with large scan dimensions (and often with coarse reciprocal space sampling), the strategy described above that scales with number of scan positions is inefficient. The PRISM, or plane wave interpolated scattering matrix, method developed by Ophus [122] and later extended by DaCosta et al [146] is a strategy that combines the flexibility of multislice with the strategy described for Bloch wave calculations above. Instead of using a separate converged beam probe for each scan position, PRISM performs a multislice simulation for each plane wave pixel in the probe-forming aperture and stores the resulting complex exit waves in a “compact” scattering matrix. The position-dependent patterns are then produced as in the Bloch wave strategy, by coherently summing the plane wave exit waves with the appropriate phase factors. Ophus also intro-

---

<sup>3</sup>It is possible to combine multiple plane waves into a single Bloch wave computation, but the subproblems corresponding to each plane wave do not interact. This implies that  $\mathcal{A}$  will be a block diagonal matrix, so each of the submatrices can be diagonalized independently and combined afterwards. It is more computationally efficient to diagonalize a sequence of small, dense matrices than one large, sparse block matrix.

duced an interpolation scheme that allows for a tradeoff between simulation accuracy and computation time, which we will not describe here.

## Software

A wide range of software packages are available for computing electron diffraction. Here, we will briefly summarize a few packages which are under current development and which are capable of simulating 4D-STEM experiments. The majority of the current packages use the multislice algorithm, often with the option to use PRISM for 4D-STEM scans. The first implementation of PRISM was in the package `Prismatic` [142], and the software has undergone a large upgrade in version 2.0 [146]. `Prismatic` is one of the most highly optimized codes available, implemented in C++ and has a streaming multi-GPU algorithm for use on clusters and supercomputers. Thus it is often considered the best choice for the largest simulations, but its interface is primarily through special plaintext configuration files that can present a steep learning curve. The `abTEM` code implements multislice and PRISM, and has a single-GPU algorithm for both (though multi-GPU support is promised to be available soon) [102]. `abTEM` is written in Python and uses the Atomic Simulation Environment for handling structures [86]. `abTEM` also naturally interfaces with *ab initio* simulation codes, and can perform diffraction simulations directly from electron densities computed with DFT. `EMSoft` contains both a GPU-accelerated multislice code and a CPU-based Bloch wave code, as well as the ability to simulate Kikuchi diffraction in a scanning electron microscope; while written in Fortran, it is still under active development [41]. `py4DSTEM` contains a Bloch wave simulation code that is largely based on `EMSoft`, and all of the Bloch wave simulations in this work were produced using `py4DSTEM`. Finally, a number of packages can also perform simulations that include diffraction of electrons that have undergone inelastic scattering. These include `muSTEM` [2], `MULTEM` [97], and `py_multislice` [20].

## Chapter 3

# Hyperspectral Denoising of 4D–STEM Datasets Using a Total Variation Method

### 3.1 Introduction

The rise of direct electron detectors for scanning transmission electron microscopy has brought data collection into a new speed regime, opening up the possibility of new techniques that capture complex scattering information at every position of the STEM probe without requiring prohibitively long dwell times. These techniques include both high speed electron energy loss spectroscopy (EELS) mapping, where an EEL spectrum is acquired at every scan position and which is now routinely possible at atomic resolution [84, 55], as well as four dimensional STEM (4D-STEM), where a 2D diffraction image of the probe is acquired at every probe position [123].

The fastest detectors available currently operate at speeds approaching 100,000 frames per second [134], for a pixel rate on the order of  $10^{10}$  /s. To realize these high speeds, these detectors operate in electron counting mode where readout noise is mitigated by operating at very low fluence. By ensuring that only one out of every 40–100 pixels is illuminated by an electron in each frame, it is possible to set a threshold that yields a binary image that rejects electronic noise and that is not affected by the large variance of the energy deposited per incident electron [58]. The resulting data, when acquired at full speed, is thus incredibly low dose; for example, the 4D Camera at the National Center for Electron Microscopy is  $576 \times 576$  pixels large and so could handle 3,000–8,000 electrons per frame at the fill factors cited above. However it must be noted that this maximum fill factor applies *locally*, and so in a nanobeam STEM experiment where the Bragg disks are small in comparison to the distance between them, only a fraction of this dose can be utilized.

Some STEM imaging modalities can produce useful results at these low doses, in particu-



lar ptychographic reconstructions<sup>1</sup>. [120, 75]. O’Leary et al reported using electron counted data containing fewer than 300 electrons per pattern to image a zeolite structure at atomic resolution [120]. However, many nanobeam 4D–STEM measurements break down when applied to sparse, low dose data. Strain mapping by cross-correlation of a template image with a known reference requires approximately  $10^6$  electrons per pattern to achieve good accuracy and is often not possible below  $10^3$ – $10^4$  electrons per pattern (see Chapter 4 or [199]). Other nanobeam measurements that also rely on Bragg disk detection, such as automated crystal orientation mapping (Chapter 5 or [124]) or phase mapping [71], share similar requirements on dose per pattern.

For these types of nanobeam diffraction experiments, when acquired using fast counting mode detectors, it is therefore necessary to use some form of denoising in order to make useful measurements. Several examples of denoising for electron microscopy data, including volumetric and hyperspectral data, are available in the literature. A powerful method for image denoising is block matching and 3D-filtering, or BM3D, which is a form of nonlocal means that searches for similar patches across an image [36]. The same block-matching concept has also been used as the basis for machine learning denoising models [8, 179]. Block matching has been extended for time series of images, where it is known as BM4D [105]. BM4D could in principle be extended to 4D data as well (which would presumably be known as BM5D), but likely at very high computational cost as the set of possible patches becomes considerably larger than the already-large original data as the dimensionality increases. For 4D–STEM data, tensor singular value decomposition has been explored as a straightforward and efficient means of denoising by finding a low-rank approximation to the measurement, similar to principal component analysis [201].

Total variational (TV) denoising is a popular method for reducing noise in images which is attributed to Rudin, Osher, and Fatemi [156]. It has since been extended to higher dimensional data such as 3D volume imaging [5] and hyperspectral data[29]. It is particularly known for its edge-preserving properties, which is advantageous for nanobeam 4D–STEM experiments where diffraction disks are expected to have sharp boundaries. TV is also useful as a regularizer in the solution of inverse problems, for example when reconstructing tomograms from a limited number of tilts [159], where the TV algorithm is often preferred for its ability to remove noise while preserving sharp edges at the surface of the sample.

In this work, we first extend a computationally efficient TV algorithm by Jia and Zhao originally devised for 2D images to work on 3D and 4D hyperspectral datasets with varying weights along each dimension. We demonstrate a high performance implementation of this modified algorithm using a FISTA accelerated convergence scheme, and MPI-based

---

<sup>1</sup>It is important to distinguish between dose per diffraction pattern and dose per area. By simply stepping the beam very finely, ptychography can realize a substantial dose per area even with limited dose per pattern. A good reconstruction algorithm is ideally capable of stitching this dose together effectively, and so a high signal-to-noise can be obtained from a large number of low signal-to-noise measurements, limited by the dose per area. In the case of other nanobeam measurements, where the individual diffraction patterns are analyzed individually, this accumulation of signal does not occur and the measurement is limited by dose per pattern

distributed implementation that allows the processing of terabyte-scale datasets on a super-computer. We investigate the capability of the algorithm to reduce the effects of Poisson noise on simulated data and also show its use on a number of experimental datasets.

## 3.2 TV Denoising Algorithm

The total variational denoising method recovers a denoised signal  $u$  from a noisy measurement  $f$  by solving an optimization problem, written here for a 2D image:

$$\min_u \left[ \|\nabla_x u\|_1 + \|\nabla_y u\|_1 + \frac{\mu}{2} \|u - f\|_2^2 \right], \quad (3.1)$$

written here for a two dimensional image, and where  $\mu$  is an appropriately chosen positive parameter. The operators  $\nabla_x$  and  $\nabla_y$  are backwards difference operators along the  $x$  and  $y$  directions, respectively. For example:

$$\nabla_x u(i, j) = u(i, j) - u(i - 1, j), \quad i = 2, \dots, N, j = 1, \dots, N \quad (3.2)$$

with the boundary value  $\nabla_x u(1, j) = 0$ . The algorithm gets the name ‘‘total variation’’ from the use of the 1-norm of the difference operator applied to the signal inside the optimization:

$$\|u\|_1 := \sum_{i,j} |u(i, j)|, \quad (3.3)$$

which represents the total magnitude of variation between neighboring pixels of the image.

Numerous approaches for solving this optimization problem are available in the literature. The original paper by Rudin et al cast the problem to the solution of a differential equation and showed an approximate, iterative numerical scheme for its evaluation [156]. Later Osher and coworkers introduced the split Bregman iteration scheme that remains popular [54, 126]. In this work, we adapt the algorithm of Jia and Zhao [74], originally shown for 2D images, for TV denoising of 3D and 4D datasets. This algorithm (which we will refer to here and in the accompanying code as J-Z) is advantageous for large and high-dimensional datasets because each iteration involves only evaluating local gradients of the function, and is easily adaptable to accelerated convergence schemes. The J-Z algorithm solves the minimization problem in Eq. 3.1 using the following iteration scheme:

$$b_x^k := \text{cut}(\nabla_x u^k + b_x^{k-1}, 1/\lambda) \quad (3.4)$$

$$b_y^k := \text{cut}(\nabla_y u^k + b_y^{k-1}, 1/\lambda) \quad (3.5)$$

$$u^{k+1} := f - \frac{\lambda}{\mu} (\nabla_x^T b_x^k + \nabla_y^T b_y^k) \quad (3.6)$$

where  $b_x$  and  $b_y$  are auxiliary arrays and the cut operator is defined as

$$\text{cut}(c, 1/\lambda) := \begin{cases} 1/\lambda & \text{for } c > 1/\lambda \\ c & \text{for } -1/\lambda \leq c \leq 1/\lambda \\ -1/\lambda & \text{for } c < -1/\lambda \end{cases} \quad (3.7)$$

and  $\nabla_x$  and  $\nabla_x^T$  are backward and forward difference operators along the dimension  $x$ :

$$\nabla_x u(i, j) = u(i, j) - u(i - 1, j) \quad (3.8)$$

$$\nabla_x^T u(i, j) = u(i, j) - u(i + 1, j) \quad (3.9)$$

The parameter  $\mu$  controls the balance between TV minimization and fidelity to the measured data. The  $\lambda$  parameter is effectively a step size of the algorithm, which Jia and Zhao showed to be optimally chosen such that  $\lambda/\mu = 1/8$ .

A simplified listing of a Python program to evaluate this algorithm for 2D images is given below, where `raw_data` is the noisy measurement, `denoised` is the array that holds the denoised data and which is initialized with zeros, and `lam` and `mu` are the parameters  $\lambda$  and  $\mu$ , respectively.

Listing 3.1: A simplified Python implementation of the J-Z total variation denoising algorithm in two dimensions.

---

```

import numpy as np

denoised = np.zeros_like(raw_data)
bx = np.zeros_like(raw_data)
by = np.zeros_like(raw_data)

for i in range(number_iterations):
    bx = np.clip(
        denoised - np.roll(denoised, 1, axis=0) + bx,
        -1/lam, 1/lam
    )
    by = np.clip(
        denoised - np.roll(denoised, 1, axis=1) + by,
        -1/lam, 1/lam
    )

denoised = raw_data - \
    ((bx - np.roll(bx, -1, axis=0)) * (lam/mu) +
     (by - np.roll(by, -1, axis=1)) * (lam/mu))

```

---

Convergence of the algorithm is proven by Jia and Zhao [74], but iterative thresholding algorithms of this type are known to have a slow rate of convergence. Beck and Teboulle demonstrated a general method for accelerating the convergence of this type of algorithm, which they refer to as Fast Iterative Shrinkage-Thresholding (FISTA) [11]. Their method stores both the current iterate and the previous one, and combines them using a specially chosen and variable weighting. In the remainder of this work, all results are presented using a FISTA-accelerated version of the J-Z algorithm. For each auxiliary array  $b_{x,y,\dots}$  we also store a history array  $d_{x,y,\dots}$ . At each iteration, the update is modified to become (presented

in one dimension here for compactness):

$$d^k := \text{cut}(\nabla u^k + b^{k-1}, 1/\lambda) \quad (3.10)$$

$$b^k = d^k + \left(\frac{t_k - 1}{t_{k+1}}\right)(d^k - d^{k-1}) \quad (3.11)$$

$$u^{k+1} := f - \frac{\lambda}{\mu}(\nabla^T b^k) \quad (3.12)$$

The FISTA acceleration parameter  $t_k$  is initialized to be 1 at the first iteration, and updated according to the following rule:

$$t_{k+1} = \frac{1 + \sqrt{1 + 4t_k^2}}{2} \quad (3.13)$$

FISTA acceleration has the benefit that a much faster rate of convergence is achieved without additional computation at each step. The price of this speedup is that we need to store an additional history array, with the same size as the raw data, for each dimension of the problem. When dealing with large scale 4D-STEM datasets, this can become a substantial challenge as the memory demand for TV denoising becomes 10 times the original data size (one needs to store the raw array, four auxiliary arrays, four history arrays, and the reconstruction).

In this work, we extend the J-Z algorithm to operate in higher dimensions and reformulate the smoothing parameter  $\mu$  so that it can be specified separately along each axis of the reconstruction. For example, in the case of a 4D-STEM dataset (where we label the dimensions as  $r_x$ ,  $r_y$ ,  $q_x$ , and  $q_y$ , corresponding to real and reciprocal space axes), the hyperspectral TV problem becomes:

$$u^* = \min_u \left[ \frac{1}{\mu_{r_x}} \|\nabla_{r_x} u\|_1 + \frac{1}{\mu_{r_y}} \|\nabla_{r_y} u\|_1 + \frac{1}{\mu_{q_x}} \|\nabla_{q_x} u\|_1 + \frac{1}{\mu_{q_y}} \|\nabla_{q_y} u\|_1 + \frac{1}{2} \|u - f\|_2^2 \right] \quad (3.14)$$

Note that we have now allowed  $\mu$  to vary for each dimension of the problem. This allows for the relative smoothing to be different between real space and reciprocal space, to account for different expected variation of the signal in each space. The new dimensions are added to the J-Z algorithm by adding extra auxiliary arrays, which are associated with the difference operator along each dimension, and using the same iterative procedure. We have found that the ideal step size  $\lambda$  is dependent on the number of dimensions  $D$  as  $\lambda(D)/\mu = \frac{1}{4D}$

### **cyTVDN: A Cython Implementation of 3D/4D TV Denoising**

The J-Z algorithm for TV denoising is simple to implement, with our nearly complete 2D example (Listing 3.1) fitting in ten lines of Python. Unfortunately, this naïve implementation proves to be extremely inefficient when applied to large datasets. This is because each of the individual operations performed on the arrays by numpy [63] will silently produce a

temporary array. While only the final one is retained on each line, this causes a large number of allocations and deallocations of memory, and for large datasets puts a lot of pressure on the memory management. To get around this limitation, we have developed a highly optimized total variation denoising code using Cython [13], which is a library for writing optimized extensions to Python using a C-like syntax that gives access to low-level features and high performance. Using Cython we are able to concatenate several steps of the algorithm in order to reduce memory traffic, and utilize OpenMP directives to enable parallel computation on multi-core processors [28]. The Cython code is automatically compiled into a library that is then accessible from within an ordinary Python interpreter. Our implementation is known as `cyTVDN` and is freely distributed as open-source code [197]. This implementation focuses specifically on 3D and 4D hyperspectral data, and includes the original J-Z algorithm and the FISTA acceleration scheme described above. Other optimized TV denoising programs are also available, such as `proxTV` which generalizes to any number of dimensions with variable weighting, is accessible through Python and MATLAB, and is distributed open-source [7, 6].

### Distributed Algorithm

While the algorithm we have chosen is relatively memory-efficient, as we are specifically interested in denoising of very large (several to hundreds of gigabytes) datasets acquired by fast electron detectors, the memory required by this approach will frequently exceed the level available even in substantial workstation computers. For the unaccelerated algorithm (*i.e.* without FISTA) the total memory required during evaluation is 5 times the dataset size for 3D data, and 6 times the dataset size for 4D data; adding FISTA acceleration increases these to 8 and 10 times the original dataset size, respectively. In order to make this method feasible for larger datasets, we have devised a distributed parallel algorithm for use on computing clusters or supercomputers, using MPI for communication between workers residing on separate computers.

To divide the problem between multiple workers, the dataset is divided into a 2D grid of sub-arrays, along the first two axes (the “slow” storage axes of the C-ordered array), and each worker loads and operates on one sub-array. If each worker were to operate independently on its own sub-array, the final solution would not respect the global boundary conditions of the algorithm, but rather would enforce those constraints at each of the boundaries between sub-arrays. Thus, distributed evaluation of the algorithm requires that the sub-arrays overlap one another, and synchronize data at each iteration so that the final solution respects the overall boundary conditions. Each worker’s local data thus also contains one pixel in each direction of overlap with its neighboring tiles, except at the edges of the array where the ordinary boundary conditions should apply. At each step of the iteration, inter-worker communication is used to synchronize the overlap regions between each tile, shifting data from the tile which has correctly evaluated the updated result in the overlap region to its neighboring tile which has inappropriately applied the boundary conditions. On the accumulator update step, the overlap regions on the low index overlap regions are invalid, while the overlap regions on the high index side are valid; each worker with a neighbor towards the low index side copies

data from that neighbor into its local accumulator. On the reconstruction update step, the difference operator works in the opposite direction and the situation is flipped: tiles with neighbors to the low index side will contain invalid data in the overlap region, and so copies data from that neighbor into the local reconstruction.

### 3.3 Methods

Simulated 4D-STEM data was produced using the multislice method as implemented in abTEM [102]. The simulation cell, shown in Figure 3.1 consists of a substrate of amorphous carbon with several fcc Au nanoparticles, with octahedral shape and approximately 5 nm in size, in random orientations on the substrate. Simulations were performed using a 2.5 mrad semiconvergence angle to produce well-separated nanobeam diffraction disks, at an accelerating voltage of 300 keV. We summed 20 frozen phonon configurations to ensure that each pattern contained a diffuse background resulting from thermal distortions of the lattice. The 4D-STEM scan was produced using a step size of 4 Å and diffraction patterns were binned to achieve a pixel size of approximately  $0.8 \text{ \AA}^{-1}$ . The resulting dataset had a size of  $50 \times 50$  scan positions and  $100 \times 100$  pixels in reciprocal space, with the real space extents spanning the full  $(200 \text{ \AA})^2$  area of the simulation cell and the diffraction patterns cropped to a maximum scattering angle of 40 mrad. Diffraction patterns from the amorphous carbon and the central Au nanoparticle and virtual images are shown in the first row of Figure 3.2. The noisy dataset was produced by applying Poisson noise statistics at a signal level of 500 electrons per pattern, as shown in the second row of Figure 3.2.

Denoising was performed using the `cyTVDN` package with the FISTA acceleration scheme enabled. The binned dataset is approximately 100 MB in size, so FISTA denoising required approximately 1 GB of RAM to evaluate. Each denoising run was performed for 50 iterations and the mean square error between the simulated data and the reconstruction was tracked at each step. Each denoising run took approximately 2.5 seconds to complete on a 32-core workstation with two AMD EPYC 7252 CPUs running at 3.1 GHz.

### 3.4 Results and Discussion

The results of denoising of the simulated 4D-STEM dataset are shown in the bottom row of Figure 3.2. Denoising was performed using  $\mu_R = 1.6 \times 10^{-2}$  and  $\mu_Q = 2.0 \times 10^{-1}$ , and iteration was halted when the relative magnitude of the update step was 1% (this choice of parameters is discussed further below). With these selection of parameters, an improvement in SNR of 9.1 dB is achieved after 10 iterations. The maximal diffraction pattern, computed by taking the maximum pixel value across the real space scan for each pixel in diffraction space, is a useful way of summarizing the information in a 4D-STEM scan, as it highlights the diffraction signals and suppresses the background. In the maximal diffraction pattern from the noisy dataset, we observe that many of the higher order reflections that were




Figure 3.1: Atomic model used for multislice simulations of a 4D-STEM dataset. Octahedral Au nanoparticles of approximately 5 nm size are distributed on an approximately 3 nm thick amorphous carbon support.

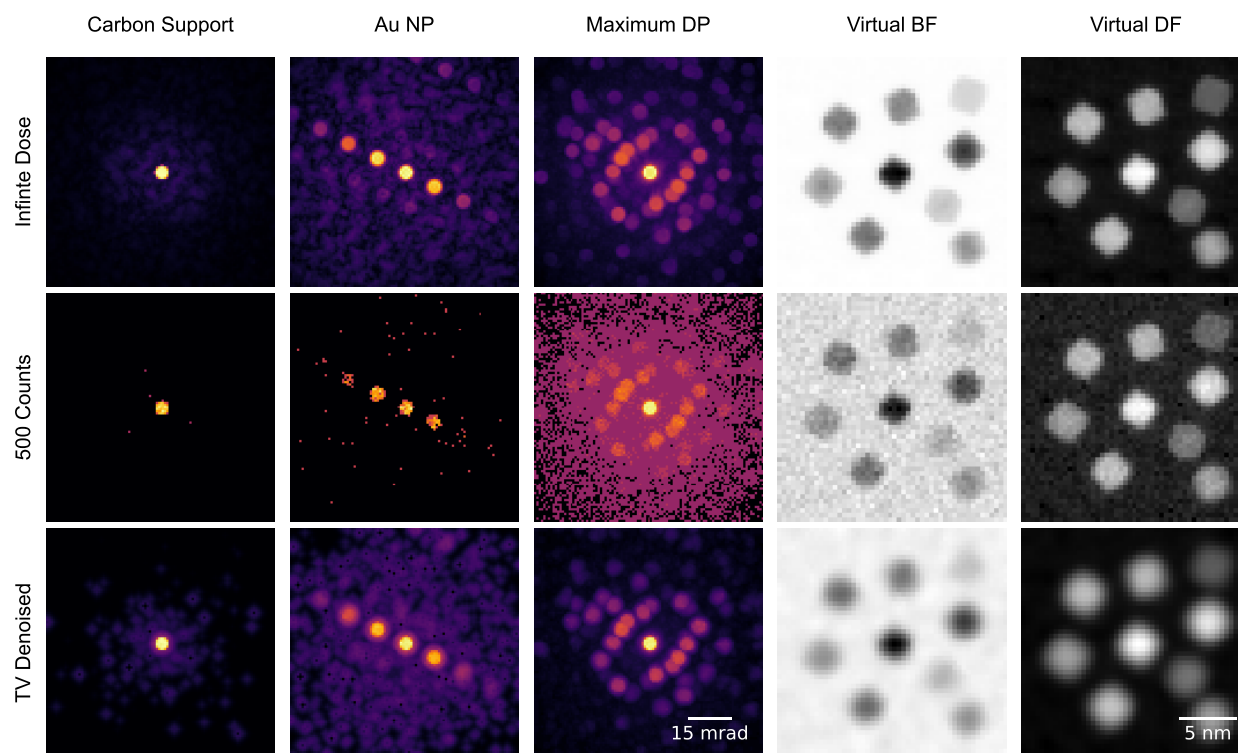


Figure 3.2: Denoising of the simulated Au nanoparticle dataset. The first two columns show diffraction patterns selected from a region of amorphous carbon and from the nanoparticle at the center of the scan. The third column shows the maximal diffraction pattern across the field of view. Virtual images are shown in the final two columns. Denoising recovers diffracted disks in the patterns that were rendered undetectable at low dose, at the cost of some blurring of the virtual images.

easily observed in the infinite dose simulation are now buried beneath the noise. In the denoised data, these higher order reflections are once again visible. This recovery is possible because while the pixel values within these reflections are on the same order as the noise in the noisy data, the intensity is correlated between neighboring scan positions and between neighboring diffraction pixels. Total variation denoising is able to use these correlations, without excessive blurring between neighboring pixels, to recover the original signal. In the individual diffraction pattern chosen from over the Au NP, diffraction disks that were not clearly observable in the noisy data are recovered with intensity substantially above the background, but weak reflections in the original data are not recovered. Some blurring between pixels is still observed, as each of the diffraction disks in the denoised data has a small “halo,” and the virtual brightfield and darkfield images show some loss in the sharpness of the edges of the particles, though the contrast is increased.



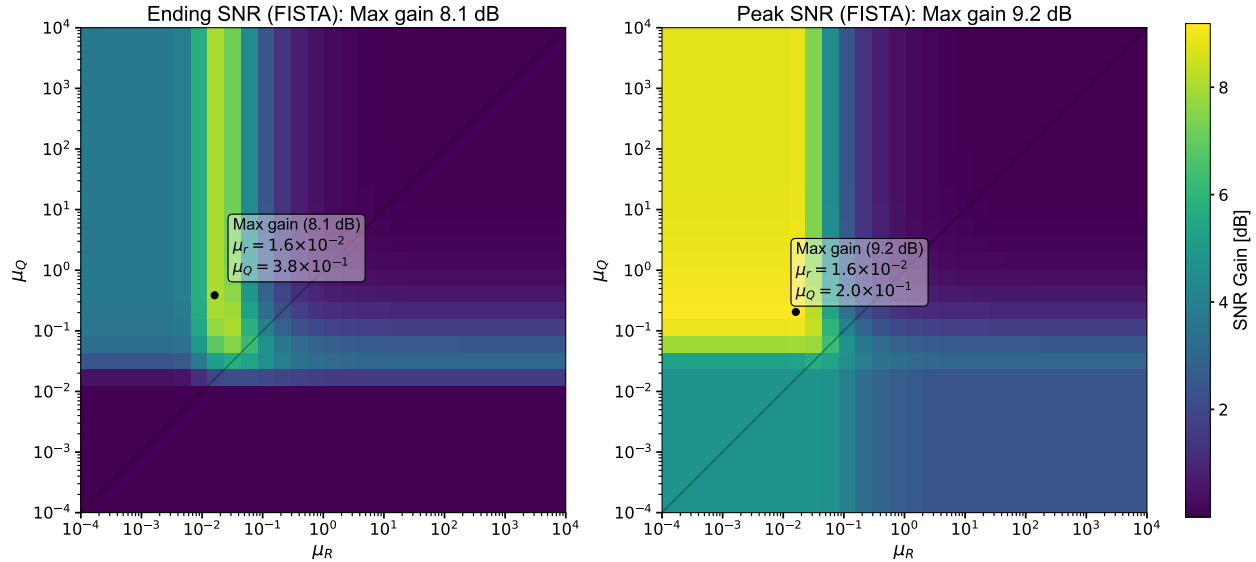


Figure 3.3: Effect of the data weighting parameters,  $\mu_R$  for real space and  $\mu_Q$  for diffraction space, on the SNR increase for the simulated Au nanoparticle dataset after denoising, using FISTA accelerated convergence. The left panel shows the SNR after running for 50 iterations, while the right panel shows the highest SNR achieved during the iterative reconstruction process. For some choices of the smoothing parameters, the algorithm reaches an optimal solution but proceeds to diverge as iteration continues except for a small range of parameter space. If iteration is stopped at the appropriate step, it is possible to achieve optimal performance of the reconstruction over a wide range of parameters.

## Parameter Selection and Convergence

As noted in [95], the optimal solution of the TV problem for a given choice of the smoothing parameters may in fact have a lower total variation than that of the underlying data, and so the reconstructed dataset may diverge from the underlying data (which has a nonzero total variation) despite continuing to minimize the objective function. Thus the physically optimal solution is not the same as the the global optimum solution of the TV minimization. While the algorithms used for TV minimization have been shown to be guaranteed to converge to the minimum of the objective function, for many choices of  $\mu$  the reconstruction will deviate strongly from the underlying signal.

In order to explore the dependence on the real space and reciprocal space smoothing parameters, we performed TV denoising for 50 iterations using FISTA acceleration for a wide range of the parameter space, as shown in Figure 3.3. We found that a signal to noise

ratio (SNR) improvement of approximately 8 dB could be achieved after 50 iterations, using the optimal set of parameters. However, the range of parameters that yield this optimal reconstruction is relatively small, and is particularly sensitive to the choice of the real space parameter. We also tracked the SNR gain at each iteration step for each parameter choice, and found that over a wide range of parameter space an even higher signal improvement was realized (up to approximately 9 dB), but continued iteration caused the result to diverge from the underlying signal. By tracking the magnitude of each update step between iterations, we found that the peak SNR was typically observed when the magnitude of the update step dropped to about 1% the magnitude of the reconstruction. In other numerical experiments with multislice simulations and arbitrary test objects, we found this 1% cutoff to be generally optimal. This stopping condition gives the user greater flexibility in choosing the smoothing parameters that prevents oversmoothing of the data.

### 3.5 Conclusions

In this work, we demonstrated an efficient method for total variational denoising of hyperspectral datasets based on the algorithm of Jia and Zhao [74]. Our accelerated implementation of this algorithm, `cyTVDN`, is implemented in Cython, using OpenMP for local parallel computation and MPI for distributed parallelism. We demonstrated the effectiveness of the method for recovery of diffraction signals from sparse, low dose 4D-STEM data using simulated data to compare the reconstructions against the ground truth signal. Since the total variation method requires a weighting parameter (or multiple weighting parameters, in the case where real space and reciprocal space are weighted differently), we explored the performance and convergence of the algorithm for different choices of the parameter. In general this smoothing parameter is dependent on the level and character of the noise in the measurement, and so predicting the optimal value *a priori* is difficult. We found that the algorithm performs poorly except for a narrow range of parameters when evaluated for a fixed number of iterations, but that when an appropriate stopping criterion is utilized there is a broad range of the parameter space that yields nearly optimal reconstructions. This finding should greatly improve the relevance of the method to experimental datasets, collected under conditions where the noise model has not been painstakingly calibrated beforehand.

## Part I

# Dynamical Diffraction as a Foe

## Robust Measurements *Despite* Dynamical Diffraction

In the following two Parts of this work, we will meet dynamical scattering in opposite ways. Here in Part I, we meet dynamical diffraction as our nemesis, a confounding effect that degrades or even inverts our measurements, and we will explore the possibility of making robust measurements despite its effects. The measurements we'll consider in this part, strain and orientation mapping, are founded on simple models of how the diffraction patterns change due to the structural features of interest. For strain mapping, the simple model tells us that diffraction disks rigidly shift due to changes in the lattice vectors of the crystal—dynamical scattering violates this by causing intensity redistributions within disks that make it difficult to measure their positions as if they were simple flat disks. For orientation mapping, the simple model for relating diffraction positions and intensities to the crystal orientation fails because it does not account for interaction between scattered beams or any thickness effects—dynamical scattering causes a model mismatch, though one that may be suppressed through adjusting the algorithm.

In Part II we will meet dynamical diffraction from the other direction, as something to utilize in our measurements, and find that we can not only find vindication in orientation mapping but also fully take advantage of multiple scattering and recover even greater detailed information from our samples.

### Bragg Disk Position Detection

In order to measure the local strain from a crystalline sample in nanobeam diffraction, the positions of each of the Bragg disks must be located to high precision. Strain is then obtained by fitting a lattice to these positions and measuring the change in these lattice vectors over a region of interest, or by calibrating the lattice distances and comparing to known atomic spacings. In the limit of a perfectly thin sample, the Bragg disks in a nanobeam pattern are perfect replicas of the direct beam in reciprocal space, and they can be located efficiently by cross-correlation of an image of the unscattered beam with the diffraction pattern. When sample thickness increases, the disks no longer appear as perfect copies of the unscattered beam but instead develop complex internal contrast. The converged beam is composed of a combination of plane waves with a cone of incident angles, and under dynamical diffraction conditions the diffracted intensities vary strongly for even these small tilts. This contrast within the disk causes it to not match exactly, and so the peak

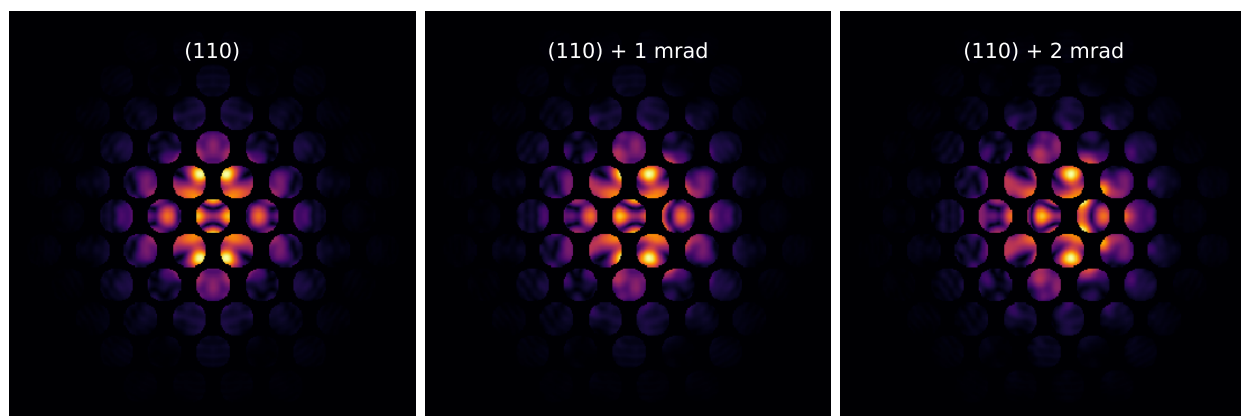


Figure 3.4: Converged beam electron diffraction patterns from an 80 nm thick Si crystal near the (110) zone axis at 300 kV. The crystal is tilted by a small amount about the vertical axis, causing substantial intensity redistribution within the diffraction disks (complicating measurement of their positions, and thus determination of the local strain). Simulations were performed using `py4DSTEM`.

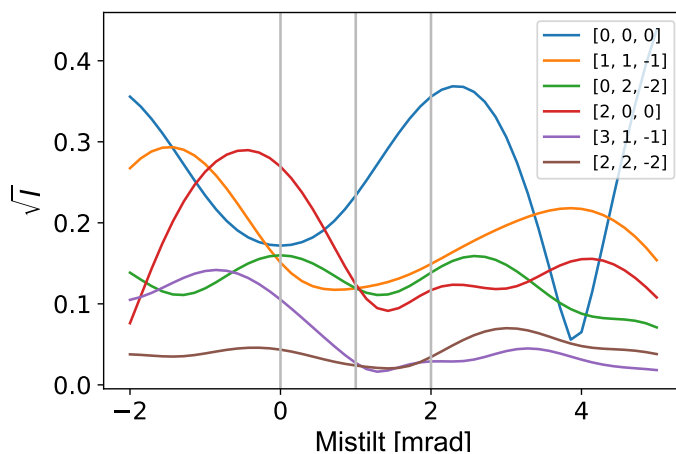


Figure 3.5: Diffraction intensities for several reflections from an 80 nm thick Si crystal near the (110) zone axis at 300 kV. The crystal is tilted amount about the vertical axis, causing nonlinear intensity redistribution between the diffraction disks (complicating orientation matching). Simulations were performed using `py4DSTEM`.

correlation with the unscattered beam is no longer a suitable way to measure the diffraction vector. An example of these phenomena is shown in Figure 3.4, where the center of mass of each diffracted disk is often far from its geometric center, and tilt of just 1 mrad (less than  $0.1^\circ$ ) causes complex redistribution of the intensities within each disk. In Chapter 4, we investigate a practical approach to overcoming this limitation, by use of a patterned aperture. The “bullseye” apertures we develop in this Chapter add numerous sharp edge features into the diffraction disk, and when cross correlation is performed these edges simply overwhelm whatever contrast fluctuations exist due to dynamical diffraction.

## Orientation Determination by Dictionary Matching

In order to measure the local orientation of a crystalline sample from nanobeam diffraction, the intensities of each of the Bragg disks must be determined and matched to a suitable model. Figure 3.5 highlights the difficulty of high-precision orientation mapping from thick samples. Using the same model as above, we track the intensity of various diffracted beams as a function of small tilts away from the zone axis (the tilts corresponding to the CBEDs are marked with gray lines). We observe that the diffraction intensities vary in a complicated and nonlinear means with respect to even small tilts of the crystal. Further, each diffracted disk appears to have a different functional dependence on the tilt, so a simple model for including thickness in a kinematic simulation like using a shape factor will not capture these physics. In Chapter 5, we will find that despite these fluctuations, we are still able to produce useful maps of the crystal orientation, and that these effects can be suppressed somewhat by our library matching algorithm. However, we will find that our precision is indeed limited by this effect—this will be addressed in Part II.

## Chapter 4

# Patterned Probes for High Precision 4D-STEM Bragg Measurements

Nanoscale strain mapping by four-dimensional scanning transmission electron microscopy (4D-STEM) relies on determining the precise locations of Bragg-scattered electrons in a sequence of diffraction patterns, a task which is complicated by dynamical scattering, inelastic scattering, and shot noise. These features hinder accurate automated computational detection and position measurement of the diffracted disks, limiting the precision of measurements of local deformation. Here, we investigate the use of patterned probes to improve the precision of strain mapping. We imprint a “bullseye” pattern onto the probe, by using a binary mask in the probe-forming aperture, to improve the robustness of the peak finding algorithm to intensity modulations inside the diffracted disks. We show that this imprinting leads to substantially improved strain-mapping precision at the expense of a slight decrease in spatial resolution. In experiments on an unstrained silicon reference sample, we observe an improvement in strain measurement precision from 2.7% of the reciprocal lattice vectors with standard probes to 0.3% using bullseye probes for a thin sample, and an improvement from 4.7% to 0.8% for a thick sample. We also use multislice simulations to explore how sample thickness and electron dose limit the attainable accuracy and precision for 4D-STEM strain measurements.

### 4.1 Introduction

Strain at the nanoscale is important in understanding deformation mechanisms of structural materials [181], as well as for engineering of transport properties in semiconductor devices

---

The results presented in this chapter have been published as a journal article titled “Patterned probes for high precision 4D-STEM bragg measurements” in *Ultramicroscopy* 209, 112890, by Steven E Zeltmann, Alexander Müller, Karen C Bustillo, Benjamin Savitzky, Lauren Hughes, Andrew M Minor, and Colin Ophus [199]. The material is presented here with permission of the co-authors.

[12]. Nanostructures can support strains of up to  $\approx 10\%$  without relaxation, providing great opportunities to engineer properties in ways that are not available in bulk materials [90]. A variety of techniques exist for measuring deformation with nanometer-scale resolution, including X-ray ptychography [65] or coherent diffraction [150], though at present the highest spatial resolution is achieved in the transmission electron microscope (TEM). TEM strain measurements have been accomplished by dark-field holography [83, 34], atomic resolution imaging [15, 50, 68], and converged-beam techniques [77, 202, 33, 69].

In scanning transmission electron microscopy (STEM), a converged electron probe is rastered across the sample, and some of the scattered electrons (usually those scattered incoherently by thermal diffuse scattering) are measured to assign a value to each pixel [140]. Modern electron detector technology allows the full scattering pattern at each STEM probe position to be recorded, an experiment referred to as four-dimensional scanning transmission electron microscopy (4D-STEM) [123]. This method, also referred to as scanning electron nanodiffraction (SEND) or nanobeam electron diffraction (NBED), has been used in analyses of crystal orientation [148, 23, 128], local ordering of glassy states [94], sample thickness [87, 203], and other analyses as described in a recent review [123].

4D-STEM is used for mapping strain at the nano-scale by locating the Bragg scattered electrons in each pattern, whose position on the detector is related to the local lattice spacing. This approach has been used to map strain in electronic devices [176], structural materials [132], including *in situ* deformed samples [51, 131], two-dimensional materials [62], and other systems where nanoscale deformation is of interest. 4D-STEM allows a large field of view and flexibility with regards to sample type and orientation [9, 154]. Figure 4.1a shows a schematic view of the experimental setup for 4D-STEM strain mapping—the convergence angle is chosen so that non-overlapping convergent beam electron diffraction (CBED) disks are obtained in each pattern.

In investigations of mechanical deformation and strain-engineered semiconductor devices the strains of interest are generally on the order of  $\approx 1\%$ , which is much larger than the currently achievable precision, reported to be  $6 \times 10^{-4}$  [59] using the standard microprobe-STEM mode (i.e. without precession or patterned probes). This precision is not sufficient for several potential applications of 4D-STEM strain mapping, such as temperature mapping by thermal expansion measurement or mapping certain structural transformations via the lattice parameters, where strains may be on the order of  $10^{-4}$ . We note that direct comparison between precision limits reported in the literature is difficult because the precision limit depends on the sample properties, microscope image distortions, and the electron dose [106, 132, 56].

The precision of the strain maps obtained by 4D-STEM is governed by the precision with which the Bragg scattered electrons can be located in each diffraction pattern. Non-uniform intensity of the diffracted disks, which can be caused by sample bending or dynamical diffraction in thick samples [35], makes accurate detection of the positions of the diffraction disks difficult. Reducing the convergence angle of the electron probe shrinks the diffraction disks, hiding some of the dynamical effects at the expense of a larger real-space probe size. For this reason, much of the existing literature on 4D-STEM strain mapping uses convergence



angles 0.2–1 mrad. When operating at larger convergence angles, the centers of mass of the diffracted disks are not necessarily at the reciprocal lattice points, thereby requiring methods sensitive to the locations of the edges of the disk [153, 132]. Disk position detection is often accomplished by cross- or phase-correlation of the diffraction pattern with a template image. These methods are still not ideal, as simulations performed by Mahr *et al* [106] found that the inner structure of the CBED disks is the limiting factor for precision of 4D-STEM strain measurements.

Post-processing of the 4D-STEM data and sophisticated data analysis methods have been shown to improve the precision of strain measurements. Pekin *et al* [132] investigated the optimal image filtering and correlation algorithms for diffraction disk detection, as well the robustness to non-uniform diffracted disks and signal-to-noise level. They found that the precision of disk location measurements can be degraded by an order of magnitude due to uneven illumination of the CBED disks. More computationally intensive disk-finding algorithms have also been implemented [106, 114].

Changes to the experimental setup provide another route to improve precision. Precession of the incident electron beam with incoherent summation of the diffraction patterns at each beam tilt “averages out” dynamical contrast and illuminates higher-order diffraction disks, which can yield a substantial improvement in strain precision to  $2 \times 10^{-4}$  [153]. However, this procedure requires specialized hardware in order to precess the beam in combination with scanning, and longer acquisition times. Mahr *et al* [106] showed simulations of the precision of 4D-STEM strain measurements for different experimental conditions, and suggested the use of patterned probes, but found no substantial improvement over standard circular apertures when imprinting a single cross on the probe. “Hollow-cone” or Bessel structured probes, produced using an annular condenser aperture, are akin to precession diffraction, but with all tilts illuminated simultaneously (and thus added coherently). Such probes were simulated and realized experimentally by Guzzinati *et al* [59], yielding strain precision of  $2.5 \times 10^{-4}$ , rivaling precession diffraction. This approach also allows for higher convergence angles, as the sparsity of the patterned probe reduces the interference between the scattered beams. Diffraction patterns through thick samples also contain a large background intensity due to inelastic scattering, which can be effectively eliminated by zero-loss energy filtering [60, 182].

In this paper, we investigate the use of probes with patterning in momentum space to improve the robustness of cross-correlation disk detection. Using an amplitude grating in the probe-forming aperture of the condenser system imprints known patterning on the diffraction pattern that allows accurate position location even in the presence of highly non-uniform illumination of the diffracted disks, as shown schematically in Figure 4.1b. Such patterned apertures are easily fabricated by physical vapor deposition and focused ion beam (FIB) machining; are mechanically stable; and, due to high conductivity, do not suffer from charging artifacts. We used mutlislice simulations to optimize the design, and estimate the improvement in accuracy and precision of disk detection for patterned apertures relative to typical circular probes. We also carried out 4D-STEM strain measurement experiments on unstrained silicon samples and characterized the improvement of precision when using patterned probes.




Figure 4.1: (a) Schematic of experimental setup for 4D-STEM strain mapping. A converged electron probe is rastered across the sample and a diffraction pattern is acquired at each probe position. Thick regions of the sample have complicated dynamical contrast inside the CBED disks that make accurate position determination difficult. In (b), a grating is inserted in the condenser system of the microscope to pattern the probe in momentum space. This pattern is imprinted on the diffracted disks, providing sharp edges in registry with the probe pattern that makes computational determination of their position more robust.

## 4.2 Theory

### Measuring Disk Positions

We determine the position of both scattered and unscattered Bragg disks by measuring the relative translation between a template image  $I_{\text{ref}}(\mathbf{r})$  and a disk image  $I(\mathbf{r})$  using digital image correlation. This correlation image  $I_{\text{corr}}(\mathbf{r})$  can be determined efficiently by taking the Fourier transforms  $\mathcal{F}\{\}$  of each image,

$$\begin{aligned} G(\mathbf{q}) &\equiv \mathcal{F}\{I(\mathbf{r})\} \\ G_{\text{ref}}(\mathbf{q}) &\equiv \mathcal{F}\{I_{\text{ref}}(\mathbf{r})\} \end{aligned}$$

and then using the expression,

$$I_{\text{corr}}(\mathbf{r}) = I(\mathbf{r}) \otimes I_{\text{ref}}(\mathbf{r}) \quad (4.1)$$

$$= \mathcal{F}^{-1} \left\{ \frac{G(\mathbf{q})G_{\text{ref}}^*(\mathbf{q})}{|G(\mathbf{q})G_{\text{ref}}^*(\mathbf{q})|^p} \right\}, \quad (4.2)$$

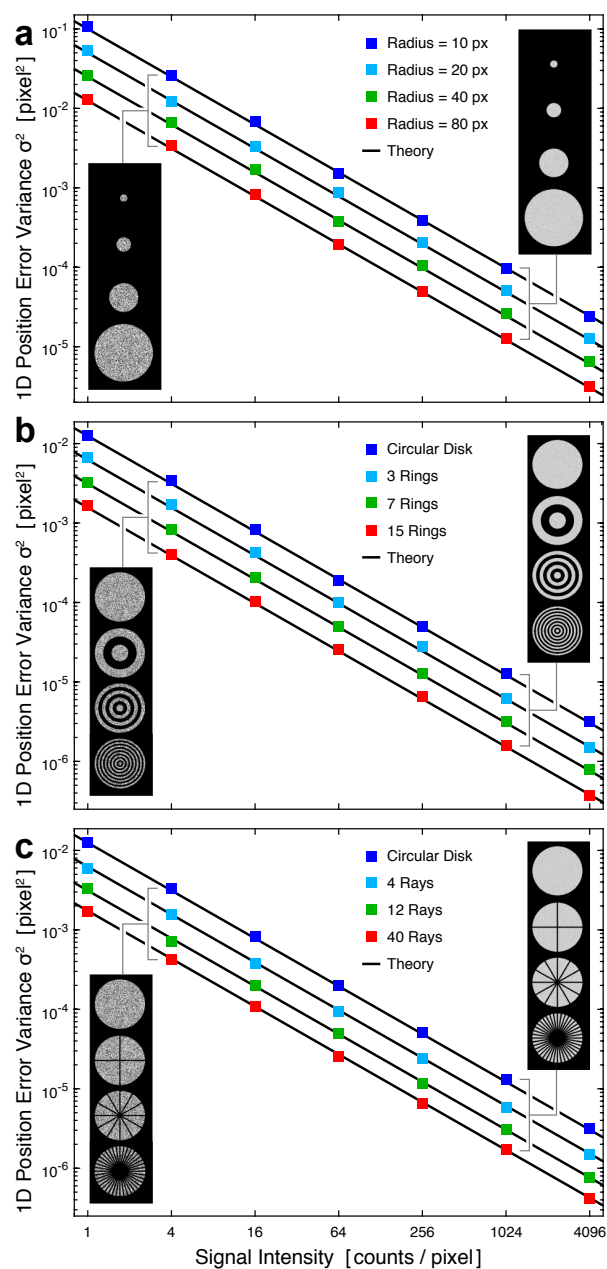


Figure 4.2: Numerical tests of image registration of an ideal STEM probe with a noisy measurement. Position error was measured for 1000 randomly generated probes along one dimension, for (a) circular disks with different radii, (b) varying numbers of concentric rings, and (c) varying numbers of intersecting rays. These measurements are compared to the theoretical precision given by Eq. 4.3. Inset images show examples of noisy measurements.

where  $\mathbf{r} = (x, y)$  and  $\mathbf{q} = (q_x, q_y)$  represent the real space and reciprocal space coordinates respectively,  $\otimes$  is the correlation operator,  $*$  indicates the complex conjugate, and  $p$  is the correlation power law coefficient. The cross-correlation is given when  $p = 0$ , and phase correlation is defined by  $p = 1$ . Values of  $p$  between 0 and 1 define a hybrid image correlation [132]. In this work, we use cross-correlation with  $p = 0$  for all simulations, and both cross and hybrid ( $p = 0.25$ ) for the experimental data.

To estimate the error of a measured disk position, we follow the methods of Clement *et al* [32]. We first assume an ideal, noise-free measurement of the template probe image  $I_{\text{ref}}(\mathbf{r})$  is available, from careful measurements of the vacuum probe image. Next, we assume the measured image of a disk  $I(\mathbf{r})$  has a signal given by a Poisson distribution with a mean of  $n$  counts per pixel, and therefore also a variance of  $n$ . The variance  $\sigma_x^2$  of a cross-correlation measurement of the image translation error along the  $x$  direction is given by

$$\sigma_x^2 = \frac{1}{nD_x}, \quad (4.3)$$

where  $D_x$  is the normalized ‘‘image roughness’’ [32] along the  $x$  direction, given by

$$D_x = \frac{1}{L_x L_y} \sum_{q_x, q_y} (2\pi q_x)^2 |G(\mathbf{q})|^2 \quad (4.4)$$

where  $L_x$  and  $L_y$  are the image dimensions. This expresses that the addition of more edges to the image template will lead to greater precision, as the presence of more edges will weight the higher Fourier coefficients more heavily. In addition, upsampling a band-limited image will increase the image dimensions  $L_x, L_y$  without increasing the higher Fourier components and lead to decreased precision.

If all units are in pixels, the image roughness for a circular disk with radius  $R$  is given by  $D_x \approx R$ . Using this expression in Eq. 4.3 gives a variance of

$$\sigma_x^2 = \frac{1}{nR}. \quad (4.5)$$

Note that this expression will often have a small numerical prefactor  $\approx 1$  due to image details such as the maximum bandwidth and sharpness of the edges. The 2D variance will be given by  $\sigma_x^2 + \sigma_y^2$ . To verify the above analysis, we performed numerical measurements of the disk position error for circular disks with various radii and counts per pixel. These measurements are shown in Figure 4.2a, and are in excellent agreement with Eq. 4.5.

To lower the disk position error, we must increase the image roughness  $D_x$  and  $D_y$ . One possibility is to add a series of concentric rings, as in Figure 4.2b. For  $M$  total concentric rings that are linearly spaced, the image roughness is given by

$$\sigma_x^2 = \frac{2}{nR(M+1)}. \quad (4.6)$$

Increasing the number of concentric rings to 3, 7, or 15 will decrease the disk position error variance by factors of  $1/2$ ,  $1/4$ , and  $1/8$ , respectively.




Figure 4.3: Numerical tests of the impact of sampling on relative position error variance. Concentric ring probes show the scaling expected from the theory at greater than Nyquist sampling, but have substantial position error due to aliasing at lower sampling. Dashed lines are the prediction of Equation 4.6. Inset: example images of under-, Nyquist, and over-sampled disk images for 3, 7, and 15 ring disks.

An alternative method of increasing the image roughness  $D_x$  and  $D_y$  is to add linear ray features, radiating out from the center of the disk as shown in Figure 4.2c. An increasing number of rays lowers the position error variance, by increasing the image roughness  $D_x$ . Interestingly, combining concentric rings with linear rays does not further decrease the position error, though it can, in some circumstances, reduce the total number of counts while maintaining the same position error variance.

Finally, we note that the position error of a circular disk in terms of the total electron dose  $N = \pi R^2 n$  is

$$\sigma_x^2 = \frac{\pi R}{N}. \quad (4.7)$$

Thus we see that for a constant disk radius  $R$ , the variance has the expected scaling of  $1/N$ . For a constant electron dose  $N$ , the variance scales linearly with radius  $R$ . This represents the fundamental trade-off between real space and reciprocal space error for Bragg disk position measurements. Increasing the probe’s outer angle will generate a smaller probe in real space and thus improve real space resolution, but will worsen the measurement precision in reciprocal space.

In order to realize the benefits of the patterned probes on the disk detection precision, the imprinted features inside the disks must be sufficiently resolved by the detector. The effect of the detector resolution is shown in Figure 4.3 for probes with varying number of concentric rings and a constant dose of 1024 *total* counts. Note that the lines on the plot for the theory drop the factor of 2 in Eq. 4.6, which arose from the “missing” pixels cut off by the pattern,




Figure 4.4: Numerical tests of the STEM probe size using different ring patterns at different convergence angles, for an aberration-free microscope at 300 kV. Adding the ring pattern to the probe causes the real-space probe size to grow by a factor determined by the number of rings. Inset: cumulative radial intensity profiles of probe intensity for a 2 mrad convergence semiangle.

whereas here we fix the total dose such that the intensity per pixel roughly doubles inside the illuminated portion of the pattern. Nyquist sampling of the patterned probes requires one pixel per ring (as marked on the figure). We observe that at slightly below Nyquist sampling the patterned probes show substantially worse performance as compared to even an unpatterned probe, and as the pattern collapses into only a few pixels it shows the same performance as an unpatterned probe. Thus while sampling of just above one pixel per ring is sufficient, it is preferable to oversample the pattern to avoid the catastrophic drop-off in precision at just below Nyquist sampling.

The above analysis for ideal disk position measurement will often underestimate the potential gains of using patterned probes because real experiments often contain a significant amount of background signal and fine structure imparted to the disks by dynamical diffraction. In the following sections, we will show how adding various amplitude features to the STEM probe can reduce the disk position error for both multislice STEM image simulations and STEM experiments.

## Probe Size

In STEM imaging experiments, especially at atomic resolution, it is usually advantageous to form as small a converged probe as possible. This is achieved when an aberration-free (flat phase) plane wave illuminates the probe-forming aperture, with as large a semi-convergence angle as possible. However, in 4D-STEM strain measurements, the minimum probe size

should be the dimensions of the crystalline unit cell being measured. Increasing the probe size can be achieved by reducing the semi-convergence angle, or by adding amplitude patterns to the probe as described in the previous section. The dependence of the real space probe size on the number of patterned rings added to an aberration-free STEM probe is shown in Figure 4.4. In order to include the effects of both increasing the size of the central lobe and increasing the intensity of the probe tails, we have defined the STEM probe size as the radius containing 80% of the total probe intensity. The inset of Figure 4.4 shows the cumulative radial intensity of different patterned STEM probes with a semiconvergence angle of 2 mrad at 300 kV. The patterned probes have long tails that extend out from the center, decreasing the realspace resolution. Compared to the typically reported full width at half maximum, this probe size metric will overestimate the size of the probe, but better capture the effect of the long tails of the structured probes.

All STEM probe patterns lead to the same scaling law for probe size, where the probe size varies inversely with the semi-convergence angle. Adding additional amplitude rings to the probe will increase the prefactor of these power laws. For example, Figure 4.4 shows that forming a 1 nm radius probe using 1, 3, 7, and 15 rings would require semi-convergence angles of 0.9, 3, 7, and 14 mrad respectively. Thus, when using patterned STEM probes, we will generally need to use somewhat larger semi-convergence angles to produce probes of the same size. When estimating the probe size in an experiment, the best practice is always to record real space images of the STEM probe in order to obtain an accurate estimate of the probe size and thus the spatial resolution.

## 4.3 Methods

### Multislice Simulations

In 4D-STEM experiments, the strain mapping precision is not only dependent on single-disk matching precision, but rather on the precision of the lattice fit to several diffracted disks in a whole (near) zone axis pattern. To investigate the strain mapping precision taking account of the whole pattern fitting, we performed multislice simulations on an unstrained Si  $\langle 110 \rangle$  model using a custom MATLAB code and potentials from Kirkland's parametrization [78], with some implementation details give in [122]. Poisson random noise was applied to the diffraction patterns to simulate shot noise for different numbers of electrons per diffraction pattern. We simulated a 5 nm thick model to obtain diffraction patterns with largely kinematical scattering, and a 20 nm thick model to obtain patterns with dynamical contrast in the CBED disks. The convergence angle was chosen to be 2.7 mrad at 300 kV to provide nearly-touching CBED disks, which maximizes the real-space resolution while avoiding interference between the diffracted beams, and gives the worst-case scenario for disk location; the simulations are aberration-free, which also maximizes the local variations in the CBED disks. The bullseye pattern is rotated by an arbitrary amount to prevent aliasing artifacts that may arise if the bars with the simulation grid.

## Numerical Measurement of Disk Positions

The procedure for obtaining strain maps from a 4D-STEM dataset involves (a) precisely locating the diffracted disks in each diffraction pattern, (b) obtaining a guess of the reciprocal lattice vectors  $u_0$  and  $v_0$ , (c) using the approximate reciprocal lattice vectors to index each diffracted disk, and finally (d) solving an (overdetermined) linear least squares problem to obtain the best-fit  $u$  and  $v$  vectors for each diffraction pattern, from which we calculate the strain. All of the analysis in this work was performed using the open-source py4DSTEM Python module [157] available at the [py4DSTEM github repository](#).

In each diffraction pattern we locate the diffracted disks by taking the Fourier correlation (Eq. 4.2) of a convolution kernel, or template image, with each diffraction pattern. For experimental data, the convolution kernel can be obtained either by imaging the probe in diffraction through vacuum or by averaging the direct beam from many diffraction patterns. For simulated data, we use the initial wavefunction as the convolution kernel. The peaks in the correlation image between the kernel and the diffraction pattern correspond to the locations of the diffracted disks. The positions of the diffracted disks are further refined by subpixel registration using the matrix-multiplication discrete Fourier transform upsampling approach [57, 166] and a final local parabolic fitting [53]. This subpixel refinement method locally upsamples the correlation image in a 1.5 px wide window around each correlation peak by a given factor (16 in this work), without computing the entire upsampled correlation image. Each identified peak is indexed based on an initial guess of the lattice vectors, and linear least squares fitting is used to determine the reciprocal lattice vectors in each diffraction pattern. Each indexed peak is weighted by the correlation intensity in the least squares fit. Strain maps are then obtained by mapping the change in the lattice vectors. There are several thresholds and filters applied in this procedure—while we slightly tune these parameters for the different simulated models and experimental samples, in all cases the normal probe and bullseye probe at each condition are processed with identical parameters.

## Cross-validation

In measuring the strain mapping precision from the simulated data, we make use of ground truth knowledge of the sample, i.e. that the model was completely strain-free and there were no projection distortions. For real experimental data, there are artifacts that complicate this analysis: the sample may be bent or strained due to fabrication artifacts or beam heating, and the microscope projection system introduces astigmatism that distorts the pattern. Since strain information in 4D-STEM is calculated from the lattice fitted to the diffracted disks in each diffraction pattern, we can estimate the precision of the strain measurement by evaluating the agreement between the fitted lattice and the individual disk position measurements.

While the residual error from the linear least squares fit of the lattice vectors is one such metric, because of the limited number of diffraction patterns in a dataset and the effects of the artifacts described above, for the experimental diffraction patterns we calculate a “cross-



validation” error. Cross-validation is often used to evaluate the quality of high-dimensional models [4]. In each diffraction pattern, half of the identified disks are chosen at random and a best-fit lattice is obtained from only these disks. The expected positions of the other half of the disks in the same pattern are computed from this lattice, and we define the error as the root mean square (RMS) difference between these predicted positions and the actual measured disk positions. For each diffraction pattern, we repeat this procedure of training on a random subset and testing against the other measurements 200 times per diffraction pattern to ensure statistical relevance.

## Bullseye Aperture Fabrication

We fabricated a set of bullseye apertures by FIB milling a gold-coated silicon nitride TEM window. Although the theory indicates adding linear rays to the concentric ring pattern does not improve strain precision, we included four rays for structural support in the fabricated apertures. An approximately 1  $\mu\text{m}$  thick layer of gold was thermally evaporated onto the flat side of a 200 nm thick silicon nitride TEM window (Norcada, Canada) with a single 250  $\mu\text{m}$  square window. Approximately one gram of gold was evaporated at a pressure better than  $2 \times 10^{-6}$  torr, with the substrate kept at room temperature.

The bullseye apertures were milled into the gold-coated window using a FEI Helios G4 UX dual beam SEM/FIB at 30 kV. The milled aperture plate is shown in Figure 4.5. We milled bullseye patterns with 2, 3, and 4 rings and with 70, 40, 20, and 10  $\mu\text{m}$  diameters. The 70, 40, and 10  $\mu\text{m}$  bullseyes match the sizes of the standard circular apertures installed in our microscope, which simplifies beam alignments. In addition, we milled a set of circular apertures of 20, 10, 5, 2, and 1  $\mu\text{m}$  diameter, which can be used to produce STEM probes with very small convergence angles or low beam current for imaging very dose-sensitive materials. Since the apertures are more closely spaced than is typical, electrons pass through all of them and a third condenser beam-forming aperture was therefore used to isolate a single probe for nanodiffraction experiments.

## Strain Map Acquisition

The bullseye aperture plate was installed in the second condenser aperture holder of a FEI TitanX operated at 300 kV. A silicon  $\langle 110 \rangle$  sample was prepared by wedge polishing followed by Ar ion milling. 4D-STEM datasets were acquired with a scan size of  $25 \times 25$  pixels, diffraction pattern image size of  $512 \times 512$  pixels, and a probe semi-convergence angle of approximately 3 mrad. Diffraction patterns were acquired using a Gatan Orius 830 CCD. We obtained scanning diffraction datasets from two regions of the wedge sample: a “thin” region with relatively even illumination of the diffracted disks, and a “thick” region with substantial dynamical contrast.




Figure 4.5: SEM micrograph of the fabricated bullseye aperture plate.

## 4.4 Results and Discussion

### Multislice Simulations

Multislice simulations of 5 and 20 nm thick unstrained silicon along the  $\langle 110 \rangle$  zone axis are shown in Fig. 4.6. The diffraction patterns in Fig. 4.6a from the 5 nm model show even illumination of the CBED disks and the (002) forbidden reflection is not excited. The diffraction patterns in Fig. 4.6e from the 20 nm model show uneven illumination of the disks and the (002) reflection is partially illuminated due to double diffraction.

Figure 4.6b and f show the locations of the  $u$  and  $v$  reciprocal lattice vectors identified in each diffraction pattern of the simulated 4D-STEM scans, illustrating the variation in the measured lattice vectors as the probe scans across a totally strain-free sample. In the limit of small strains, the uncertainty in the reciprocal lattice vectors relative to the reciprocal lattice vector length is equal to the uncertainty in the measured strain. The center of each histogram corresponds to the lattice vectors measured from the 5 nm model with infinite dose, which we take as the ground truth. The  $u$  and  $v$  vectors correspond to the  $(1\bar{1}1)$  and  $(\bar{1}11)$  reflections (drawn in the bottom left panel of Fig. 4.6a) and each have a length of  $\approx 70$  pixels. Both the normal and bullseye probes converge to the same lattice vectors at high dose, though in all cases the spread of values is substantially larger for the normal probes. These wide variations in the lattice vectors from an unstrained sample lead to correspondingly large fluctuations in the calculated strain values. The asymmetric error in

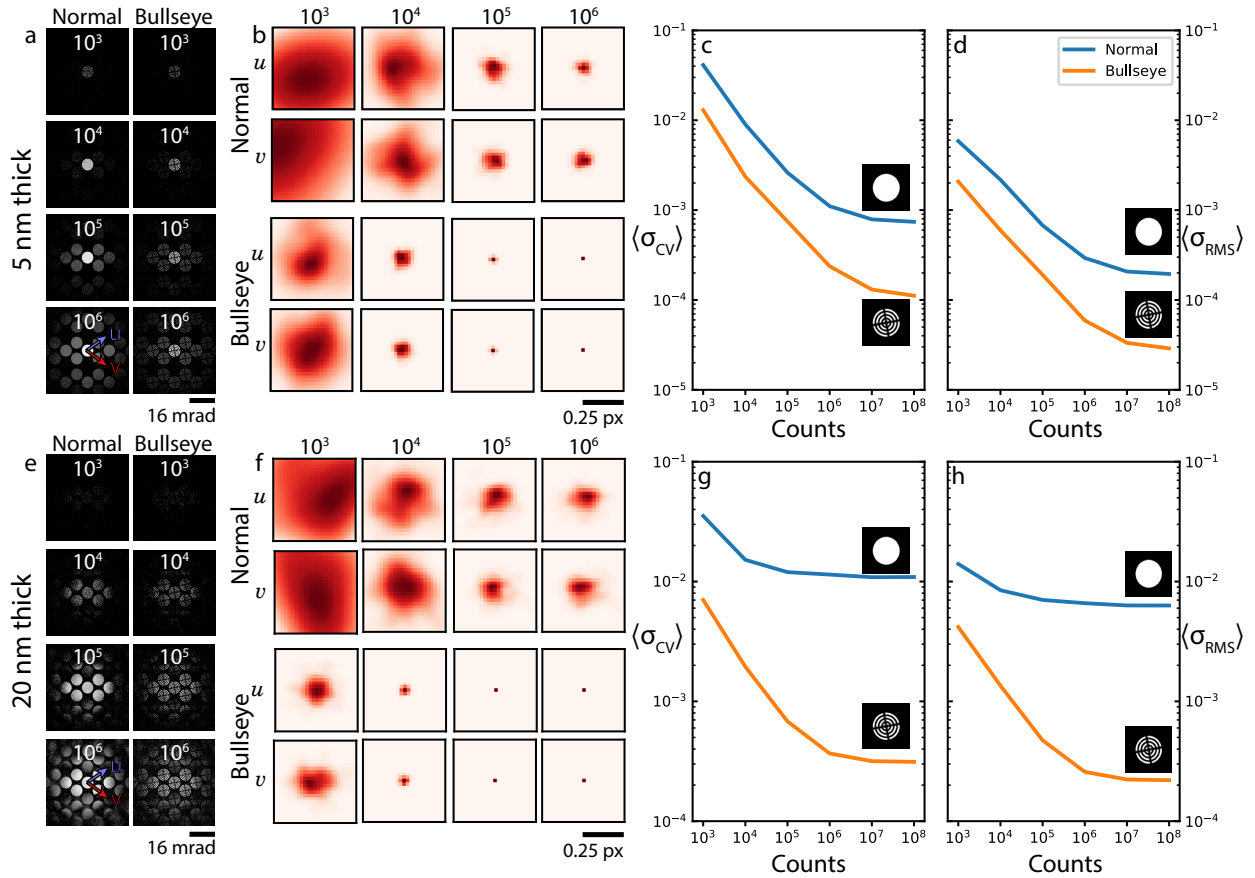


Figure 4.6: Strain mapping precision of simulated silicon diffraction data at different samples thickness and electron dose. (a,e) Representative simulated diffraction patterns at different electron doses per pattern. (b,f) Comparison of the  $u$  and  $v$  reciprocal lattice vectors measured at each scan position in the simulation of a strain-free sample. The center of each histogram represents the average  $u$ ,  $v$  positions obtained from the noise-free simulation data. (c,g) Cross-validation error and (d,h) RMS fit error, relative to the reciprocal lattice vector length (equivalent to the strain error in the small strain limit). (a-d) are obtained from a 5 nm model with largely kinematical scattering, while (e-h) are from a 20 nm model with dynamical contrast inside the CBED disks. The reciprocal lattice vectors are drawn in the bottom left panel of (a), and have length  $\approx 70$  pixels.

the histograms is likely due to the presence of the partially-illuminated forbidden reflections, which causes the cross-correlation peak uncertainty to be larger in one direction [132].

Figure 4.6d and h show the RMS residual error of the linear least squares lattice vector fitting, relative to the length of the  $(\bar{1}11)$  reciprocal lattice vector. This error is one metric for the precision of the strain measurement as it reflects the uncertainty in the fitted lattice vectors. The error decreases with increasing electron dose and the bullseye apertures have  $\approx 3.5$  times lower error at up to  $\approx 10^5$  counts. For this number of rings, the image roughness metric (Eq. 4.6) predicts a 4-fold increase in the precision of locating a single diffraction disk, without accounting for the presence of the incoherent background counts found in the multislice results. As the illumination of each diffraction disk varies across the pattern, the location precision of each diffraction disk also varies, complicating comparison with the single-disk location precision theory. At higher signal levels, the error stops decreasing as we approach the limits of the subpixel fitting algorithm. At  $10^8$  counts and above, the bullseye apertures show  $\approx 7$  times improved precision.

When dynamical contrast causes intensity variations inside the diffracted disks, the normal probes show substantially worse performance. When locating the disks by cross-correlation, as in this calculation, the location assigned to each disk is biased towards the center of mass of the disk. In the simulated diffraction patterns, many disks are seen to be half-illuminated, which leads to substantial position errors regardless of the number of counts. The patterned probes are less sensitive to this type of error, as the cross-correlation intensity should peak when the rings are in registry even if the rings are not fully illuminated. Thus we observe in Fig 4.6g and h that the precision of the normal probe saturates by  $10^5$  counts while the bullseye probe precision improves with increasing dose until  $10^7$  counts. This robustness to uneven disk illumination gives the bullseye probes an even larger precision advantage compared to the kinematical case for thin specimens, with the minimum error decreasing by  $\approx 30$  times at high electron counts.

## Experimental Measurements of Strain in Silicon

Representative diffraction patterns from the scans are shown in the top row of Figure 4.7. In the thick scan region, the (002) forbidden reflection is fully illuminated and there is substantial dynamical contrast inside the CBED disks.

As we cannot guarantee that the silicon specimen is strain-free, we cannot use the spread in the measured lattice vectors as an indicator of the precision of the measurement, and instead report only the cross-validation error for the experimental scans. The cross-validation (CV) error relative to the length of the (111) reciprocal lattice vector for the experimental scans is shown in Figure 4.7c and g. When finding the disk locations in the thin region by cross-correlation, use of the bullseye patterned probe causes the mean CV error score to decrease to 0.3% from 3.6%, an improvement of 12 times. In the thick region, the CV score decreases to 0.8% from 10.3%, an improvement of  $\approx 13$  times. The improvement in strain precision we observe in experiments is larger than predicted by the image roughness theory and observed in the multislice results. The inelastic component present in the experimental

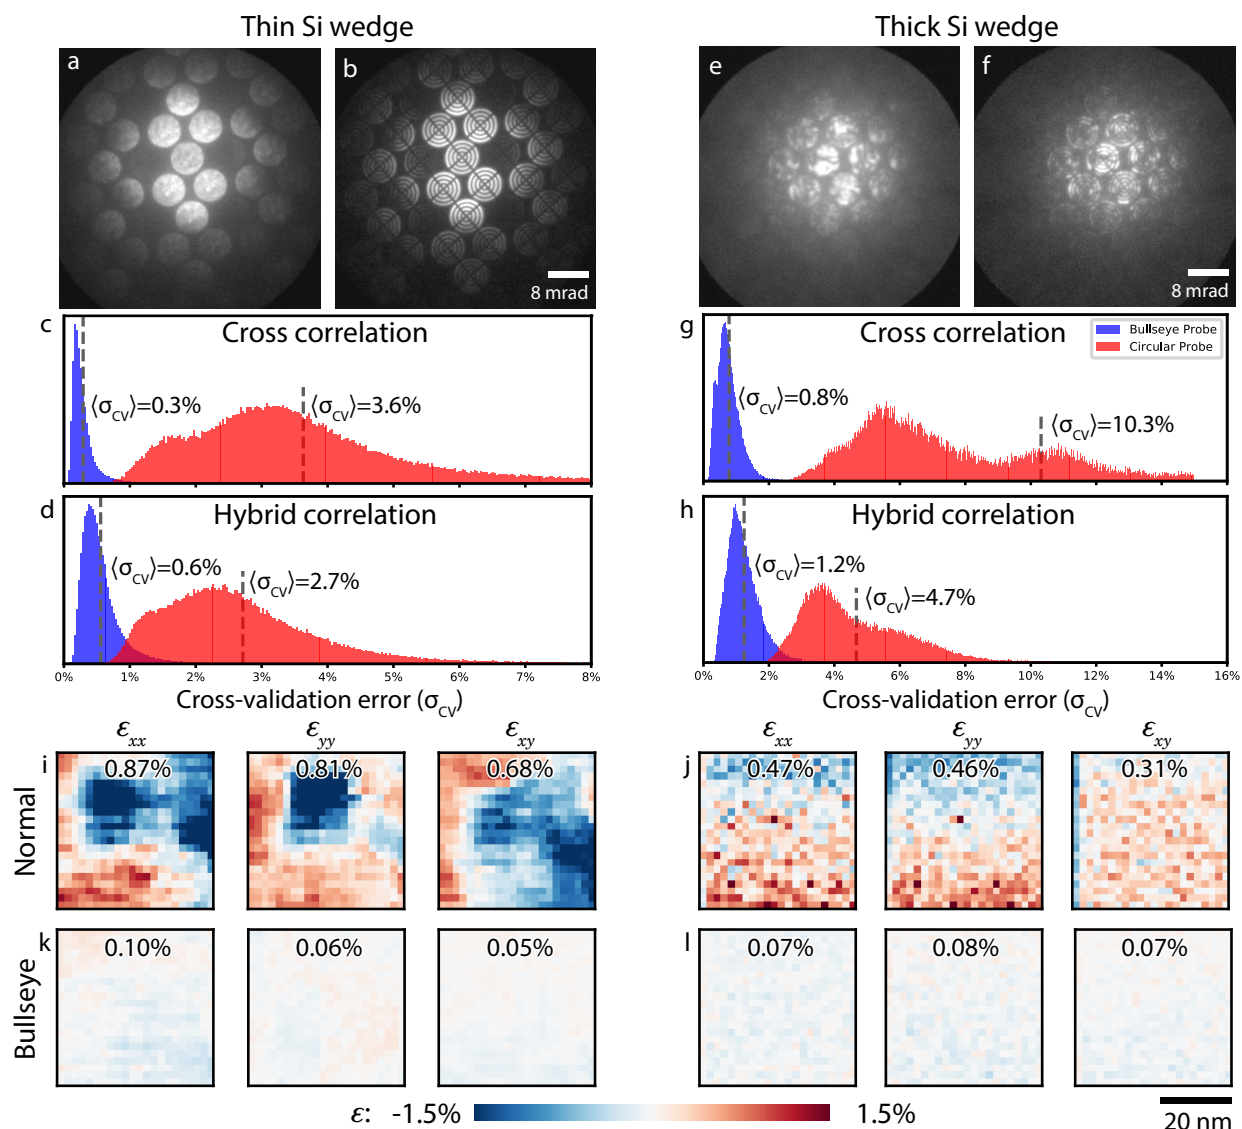


Figure 4.7: Cross-validation error determination from 4D-STEM experiments on a Si  $\langle 110 \rangle$  wedge. Diffraction patterns from a thin region of the wedge with (a) the standard circular aperture and (b) with the bullseye amplitude grating. The cross validation error, computed by fitting a lattice to half of the identified diffraction disks and measuring the error of the remaining half, using (c) cross-correlation and (d) hybrid fitting. Diffraction patterns from a thick region of the wedge (e) without and (f) with the bullseye aperture. Cross-validation strain error for (g) cross-correlation and (h) hybrid correlation disk detection. Strain maps from each region of the Si wedge sample are shown in (i)–(l). The label on each strain map indicates the standard deviation of that strain component over the field of view.

data likely plays a large part in this discrepancy, as the presence of substantial intensity between the Bragg disks reduces the contrast between the disks and the background, compounding the center-of-mass bias in the unpatterned probes and increasing the impact of the patterned bars in the disks on matching precision.

While cross-correlation performs well when the pattern background is low, hybrid correlation has been shown to better handle the ‘plasmonic blur’ in real samples [132]. To test if hybrid correlation can also improve the strain precision when using patterned probes we repeated the disk finding procedure with  $p = 0.25$  (Eq. 4.2). We observe that the CV error decreased substantially compared to cross-correlation for the normal probes, as shown in Figure 4.7d and h. However, in both the thick and thin Si wedge regions, the CV error of the bullseye apertures was 50–100% worse when using hybrid correlation. The image roughness theory discussed above does not generalize easily to hybrid correlation and does not account for additive background noise, and so cannot be used to explain the change in precision.

Using an identical procedure, we also computed the CV error for the multislice simulations, shown in Figure 4.6c and g. Because the lattice fitting in the CV approach uses only half the identified Bragg spots, the CV error is always higher than the RMS least squares residual. The trend is the same as for the RMS error in both thicknesses tested. For the 5 nm model the CV error is 3–7 times higher for the normal probes compared to the bullseyes, while for the 20 nm model the CV error is up to 29 times higher for the normal probe compared to the bullseye probe.

Strain maps produced from each region of the Si wedge using cross-correlation to locate the diffraction disks are shown in Figure 4.7(i)–(l), and the standard deviation of the measured strain in each region is marked on the maps. Strain values are referenced to the median measured lattice in each scan region. In the thin region of the Si wedge, the normal probe registers strains of approximately  $\pm 1.5\%$  across the scanned area. Bending in the thin region of the wedge leads to mistilt of a few milliradians across the scan region, which shifts the center of illumination of the pattern and the centers of mass of the diffraction disks. While sample mistilt does change the projected lattice spacing and thus the expected diffraction disk positions very slightly (on the order of 0.1% for the magnitude of tilting we observed), the strong position bias towards the center of mass of the disks when using normal probes leads to large variation in the measured strain. When using bullseye probes on the same sample region, the strain is measured as only  $\pm 0.1\%$ . In the thick region of the wedge the sample appears flat across the field of view—here, the normal probes show smoothly varying strain from top-to-bottom of the scan, likely due to variation in sample thickness altering the fine structure inside the diffraction disks. The bullseye probe again reveals a flat strain distribution with standard deviation of about 0.1%.

Recently, Gizzinati *et al* [59] demonstrated a Bessel beam structured probe for improved 4D-STEM strain mapping precision. By inserting an annulus in the second condenser aperture of an aberration corrected microscope, they produce a hollow cone probe with a semi-convergence angle of 6 mrad, giving spatial resolution of 1.3 nm. In the present work (without aberration correction), we measured the full width at half maximum of the probe to be 2.7 nm at 3 mrad. This follows the expected scaling, where using a semi-convergence

angle half as large leads to a doubling of the real-space probe size. Probes with strong amplitude structuring will necessarily sacrifice real space resolution for momentum resolution because of missing frequencies in the wavefunction. In the Bessel beam approach this broadening of the probe is partially mitigated because the beam is so sparse in momentum space that diffracted beams can overlap without substantial interference—this permits high convergence angles, leading to smaller realspace probes. Our approach is practically limited to  $\alpha < \theta_B/4$ , where  $\alpha$  is the semi-convergence angle and  $\theta_B$  is the Bragg scattering angle. By observing the variation in strain in a flat region of the sample, they estimated their strain precision as  $2.5 \times 10^{-4}$ . In this work, by comparison, we observed strain precision as good as  $2.2 \times 10^{-4} - 2.9 \times 10^{-5}$  for different simulated models (by the RMS residual metric), and  $8 \times 10^{-3} - 3 \times 10^{-3}$  in experiments on an Si wedge (by the CV metric).

The patterned probe approach has parallels to Sobel (edge-enhancing) filtering of the diffraction images. In using the Sobel filter, we assume that the CBED disks should have sharp edges and uniform intensity, and so once filtered the disks become rings. Naturally, dynamical structure in the CBED disks will also create edges that are exaggerated by the filter, and indeed Pekin *et al.* found that Sobel filtering improves precision for flat disks but causes artifacts when dynamical structure is present [132]. By applying the patterning to the probe before the sample, we avoid this drawback by adding many edges that are defined by the template.

Compared to other TEM strain mapping techniques, 4D-STEM has generally been reported to have lower precision and lower resolution than other TEM strain mapping techniques, such as atomic resolution imaging and darkfield holography. In particular, 4D-STEM strain mapping is not possible at atomic resolution as phase interference between scattered beams complicates measurement of the Bragg scattering. However, 4D-STEM offers the greatest flexibility with regards to sample type and orientation, allowing analyses of partially or completely amorphous samples, polycrystals, highly defective materials, and low-symmetry oriented crystals. Simultaneous measurement of other signals is also possible from 4D-STEM data, such as differential phase contrast (DPC) for electric field mapping. With patterned probes, the precision of 4D-STEM strain measurements can rival that of other techniques, though still with the trade-off between resolution and flexibility. Detailed comparisons of the various TEM strain mapping techniques are available in the literature [34, 10].

Early studies on 4D-STEM strain mapping were limited in acquisition speed by CCD detectors, so most of these works used well-exposed diffraction patterns with high electron doses (qualitatively, these patterns match those in our simulations where we find precision saturates and becomes dose-insensitive). With the latest generation of fast detectors operated at full speed, the dose per pattern is limited by the brightness and coherence of the source, limiting the attainable precision. A potential drawback of the patterned probe approach is the reduction in probe intensity. Our bullseye grating reduces the beam intensity by roughly half, leading to either a twofold increase in exposure time for the same dose (with accompanying increase in sample drift during a measurement) or the use of less coherent illumination to increase the probe current (which will degrade the probe size and quality of

the diffraction patterns). However, as the bullseye probes tested here give a roughly four-fold improvement in precision at equal dose while the precision scales as the square root of the dose, higher precision can be realized without compensating for the lost current. For thin samples, where the scaling laws (Eqs. 4.6 & 4.5) hold, using the bullseye aperture with identical microscope settings would give roughly  $4/\sqrt{2} \approx 2.8$  times improvement, while for thick samples the improvement can be larger.

Using the bullseye patterned probes also requires more pixels per CBED spot in order to resolve the fine pattern features with high fidelity. In some cases, particularly for thick samples and CCD detectors, this requirement necessitates “spreading out” the diffracted beam intensity over more pixels, lowering the signal-to-noise ratio. Conversely, when using direct electron detectors with limited dynamic range, the ability to operate at a higher convergence angle and distribute intensity over more pixels can be advantageous.

## 4.5 Conclusion

We have demonstrated how electron probes with patterning in momentum space can improve the precision of the CBED disk detection procedure used for calculating strain from scanning diffraction data. This approach greatly improves the precision of strain measurements from thicker samples by reducing the systematic errors that arise when locating the Bragg disks in diffraction patterns through thick samples, and potentially enabling more reliable temperature and subtle deformation measurements. In strain maps from a nominally unstrained silicon sample we observe that the anomalous strain measurements caused by dynamical effects are reduced from  $\epsilon = \pm 1.5\%$  to about  $\pm 0.1\%$ . The specific findings can be summarized as follows:

- Imprinting structure on the STEM probe in momentum space adds known, constant contrast to CBED disks which improves the precision of cross-correlation of a known template to the experimental data. For an evenly illuminated CBED disk the position measurement precision increases by a factor determined by the “roughness” of the pattern, independent of dose. For the “bullseye” pattern we used, a 4-fold improvement is expected.
- In multislice simulations of a thin sample with largely kinematic scattering, the strain mapping precision improved by a factor of  $\approx 4$  times at all doses, in agreement with theory. At high doses, the precision reaches a plateau, limited by the subpixel fitting. In simulations of thick samples, where dynamical scattering causes uneven illumination of the CBED disks, the precision improvement is even greater, up to a factor of 29 times.
- In experiments on an unstrained Si sample, we observe an improvement in precision of about 12 times for both thick and thin regions of the wedge sample. Due to the inelastic background scattering, the hybrid correlation algorithm performs better than cross-correlation when using a normal circular probe. Using the bullseye patterned



probe, the cross-correlation algorithm performs best. Strain maps produced from thick and thin regions of the silicon sample show substantially flatter strain across the same sample regions.

## Chapter 5

# Automated Crystal Orientation Mapping in py4DSTEM using Sparse Correlation Matching

Crystalline materials used in technological applications are often complex assemblies composed of multiple phases and differently oriented grains. Robust identification of the phases and orientation relationships from these samples is crucial, but the information extracted from the diffraction condition probed by an electron beam is often incomplete. We therefore have developed an automated crystal orientation mapping (ACOM) procedure which uses a converged electron probe to collect diffraction patterns from multiple locations across a complex sample. We provide an algorithm to determine the orientation of each diffraction pattern based on a fast sparse correlation method. We test the speed and accuracy of our method by indexing diffraction patterns generated using both kinematical and dynamical simulations. We have also measured orientation maps from an experimental dataset consisting of a complex polycrystalline twisted helical AuAgPd nanowire. From these maps we identify twin planes between adjacent grains, which may be responsible for the twisted helical structure. All of our methods are made freely available as open source code, including tutorials which can be easily adapted to perform ACOM measurements on diffraction pattern datasets.

---

The results presented in this chapter have been published as a journal article titled “Automated Crystal Orientation Mapping in py4DSTEM using Sparse Correlation Matching” in *Microscopy and Microanalysis* 28(2), 390-403, by Colin Ophus, Steven E Zeltmann, Alexandra Breufach, Alexander Rakowski, Benjamin H Savitzky, Andrew M Minor, and Mary C Scott [124]. The material is presented here with permission of the co-authors.

## 5.1 Introduction

Polycrystalline materials are ubiquitous in technological applications. An ideal crystal structure can be fully defined with a small number of parameters: the 3 vectors defining its unit cell, and the position and species of each atom inside the unit cell [17]. To fully describe crystalline materials in the real world however, we require a description of both the crystal lattice, and all defects present in a given material. These include point defects such as dopants, vacancies, or interstitials [43], line defects such as dislocations [89], planar defects including internal boundaries and surfaces [172], and volume defects such as precipitates [80]. Strain fields in the surrounding material can be induced by each of these defects, or generated by the boundary or growth conditions of the material such as in thin film stresses [72]. One large subset of crystalline materials are polycrystalline phases, which consist of many small crystalline grains, arranged in either a random or organized fashion. Many material properties such as mechanical strength [174], optical response [129, 100], or thermal or electrical conductivity [26] are strongly modulated by the density and orientation of the boundaries between crystalline grains [175]. Thus characterizing the orientation of polycrystalline grains is essential to understanding these materials.

The two primary tools used to study the orientation of polycrystalline materials are electron backscatter diffraction (EBSD) in scanning electron microscopy (SEM), and transmission electron microscopy (TEM). EBSD can measure the orientation of crystalline grains with very high accuracy, but has limited resolution and is primarily sensitive to the surface of materials [67, 185, 186]. Alternatively, we can directly measure the atomic-scale structure and therefore the orientation of polycrystalline grains, either by using plane wave imaging in TEM [91], or by focusing the probe down to sub-atomic dimensions and scanning over the sample surface in scanning TEM (STEM) [141]. This is possible due to the widespread deployment of aberration correction for both TEM and STEM instruments [92, 145]. Atomic resolution imaging, however, strictly limits the achievable field-of-view, and requires relatively thin samples, and thus is primarily suited for measuring polycrystalline grain orientations of 2D materials [125, 143].

Another approach to orientation mapping in TEM is to use diffraction space measurements. For crystalline materials, diffraction patterns will contain Bragg spots with spacing inversely proportional to the spacing of atomic planes which are approximately perpendicular to the beam direction (described by both the Laue condition and Bragg equations [48]). To generate a spatially-resolved orientation map, we can focus a STEM probe down to dimensions of 0.5 to 50 nm, scan it over the sample surface, and record the diffraction pattern for each probe position, a technique referred to as nanobeam electron diffraction (NBED) [127], scanning electron nanobeam diffraction (SEND) [173], or four dimensional-scanning transmission electron microscopy (4D-STEM) (we choose this nomenclature for this text) due to the 4D shape of the collected data [24]. 4D-STEM experiments are increasingly enabled by fast direct electron detectors, as these cameras allow for much faster recording and much larger fields of view [123, 119, 130].

By performing template matching of diffraction pattern libraries on 4D-STEM datasets,

we can map the orientation of all crystalline grains with sufficient diffraction signal. This method is usually named automated crystal orientation mapping (ACOM), and has been used by many authors in materials science studies [194, 148, 187, 82, 101, 99, 73, 206]. ACOM experiments in 4D-STEM are highly flexible; two recent examples include [85] implementing ACOM measurements in liquid cell experiments, and [189] adapting the ACOM method to a scanning confocal electron diffraction (SCED) experimental configuration. ACOM is also routinely combined with precession electron diffraction, where the STEM beam is continually rotated around a cone incident onto the sample, in order to excite more diffraction spots and thus produce more interpretable diffraction patterns [147, 23, 112, 45]. Recently, [109] have combined simulations with machine learning segmentation to map orientations of 2D materials, and [192] have used machine learning methods to improve the resolution and sensitivity of orientation maps by training on simulated data. For more information, [193] has provided a review of ACOM methods in SEM and TEM.

In this study, we introduce a new sparse correlation framework for fast calculation of orientation maps from 4D-STEM datasets. Our method is based on template matching of diffraction patterns along only the populated radial bands of a reference crystal’s reciprocal lattice, and uses direct sampling of the first two Euler angles (which, in the convention we have adopted, correspond to the zone axis), and a fast Fourier transform correlation step to solve for the final Euler angle (in our convention, the in-plane rotation of the pattern). We test our method on both kinematical calculations, and simulated diffraction experiments incorporating dynamical diffraction. Finally, we generate orientation maps of polycrystalline AuAgPd helically twisted nanowires, and use clustering to segment the polycrystalline structure, and map the shared (111) twin planes of adjacent grains.

## 5.2 Methods

### Overview

The problem we are solving is to identify the relative orientation between a given diffraction pattern measurement and a parent reference crystal. We solve this problem with three steps:

1. First, we generate a diffraction pattern library which covers all unique crystal orientations using kinematical simulation. This library, stored in a sparse polar coordinate representation  $P$ , is called an “orientation plan.”
2. We find all diffracted spots/disks in each diffraction pattern, and convert them into the same sparse polar coordinate representation  $X$ .
3. We determine the best fit orientation(s) by finding the maximum value(s) of the correlation  $C$  between the diffraction patterns and the orientation plan.

All of the previously discussed ACOM implementations work in essentially the same way, i.e. by precomputing the diffraction library in some form, and then comparing each

diffraction pattern to this library using a cost function based on some form of correlation. Performing template matching directly on diffraction patterns, which may contain millions of pixels, against a library of similarly sized patterns, is computationally expensive. However, the underlying information we are interested in, i.e. the projected lattice in the pattern, is typically composed of at most a few dozen non-zero points. Our sparse correlation method involves reducing the diffraction patterns to a simpler representation where the correlation can be evaluated rapidly, by first detecting the positions of the Bragg disks in the pattern, then segmenting the data into radial bands, and only evaluating the correlation in the populated bands.

The primary advances of this paper are: (1) We use Fourier transforms along the annular direction in polar coordinates for both the diffraction library and diffraction patterns to efficiently solve for the in-plane image rotation. For a full polar coordinate transform, only a small number of radial bins will contain reciprocal lattice points, and thus the output is sparse along the radial direction. We utilize this sparsity by only evaluating the polar coordinate correlations on radial shells that contain reciprocal lattice points of the reference structure, making the calculations much faster. (2) We give users fine-grained control over the relative weighting of diffraction peak radii and intensities in the correlation calculation, as well defining a kernel size which can be increased to allow more pattern distortion, or decreased to reduce the chance of false positive signals from grains with close orientations. (3) We automatically determine the symmetry-reduced range of allowed zone axes from the input crystal. (4) We provide all methods and codes as an open source implementation for the community to freely use and modify. Below we detail each of the steps for our orientation matching algorithm, and their required input calculations.

## Structure Factor Calculations

The *structure factors* of a given crystalline material are defined as the complex coefficients of the Fourier transform of an infinite crystal [167]. We require these coefficients in order to simulate kinematical diffraction patterns, and thus we briefly outline their calculation procedure here.

First, we define the reference crystal structure. This structure consists of two components, the first being its unit cell defined by its lattice vectors  $\mathbf{a}$ ,  $\mathbf{b}$ , and  $\mathbf{c}$  composed of positions in  $\mathbf{r} = (x, y, z)$ , the 3D real space coordinate system. The second component of a crystal structure is an array with dimensions  $[N, 4]$  containing the fractional atomic positions  $\mathbf{p}_n = (p_a, p_b, p_c)_n$  and atomic number  $Z_n$ , for the  $n$ th index of  $N$  total atoms in the unit cell. Together these positions and atomic numbers are referred to as the atomic basis. Because  $\mathbf{p}_n$  is given in terms of the lattice vectors, all fractional positions have values inside the range  $[0, 1)$ . The unit cell and real space Cartesian coordinates of the fcc Au structure are plotted in Fig. 5.1a.

All subsequent calculations are performed in reciprocal space (also known as Fourier space or diffraction space). Thus the next step is to compute the reciprocal lattice vectors,

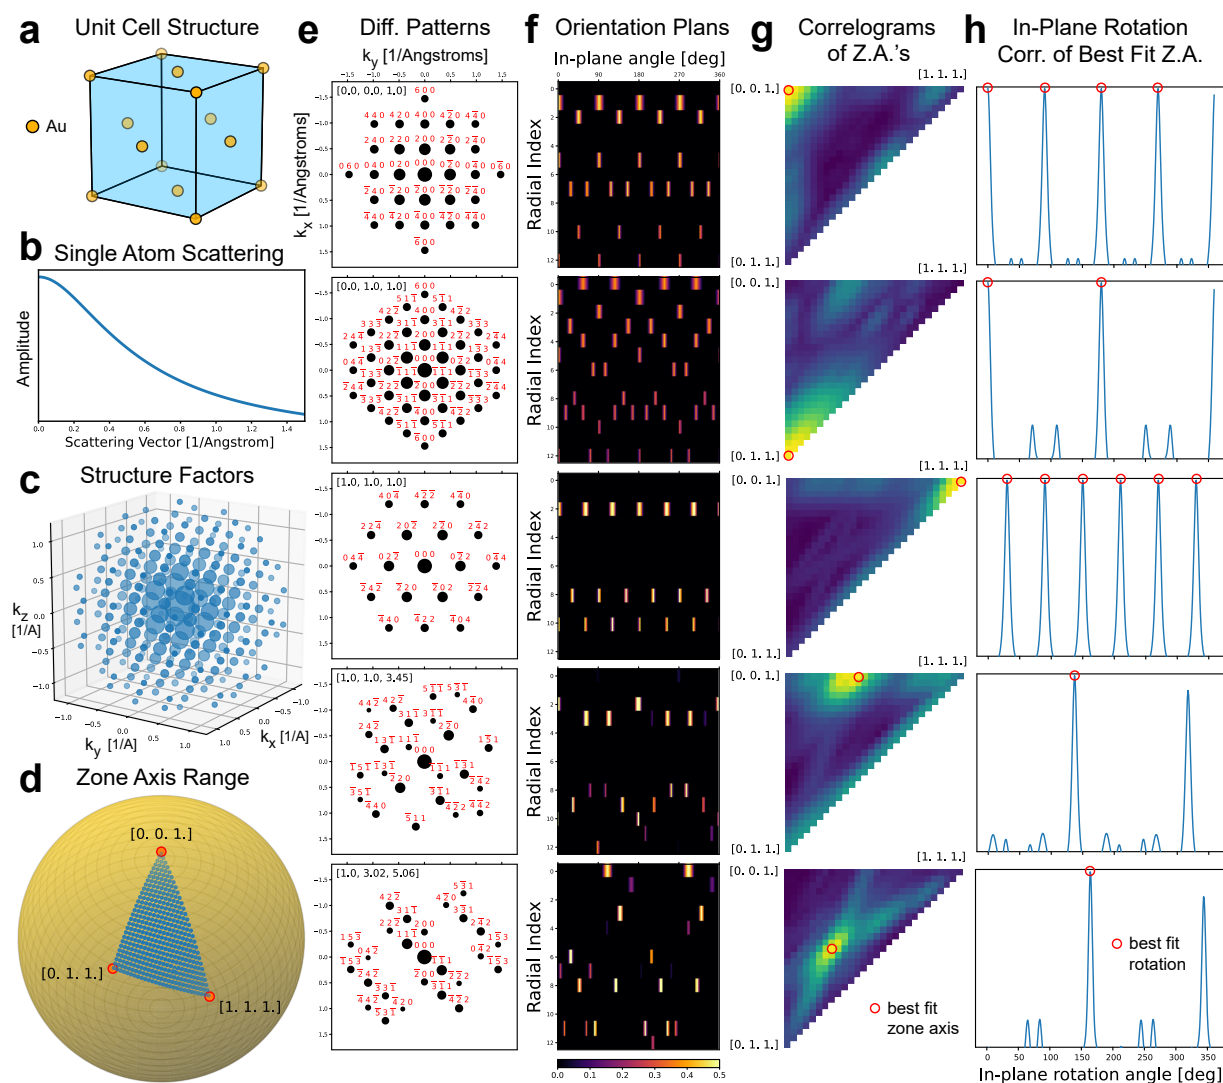


Figure 5.1: ACOM using correlation matching in `py4DSTEM`. (a) Structure of fcc Au. (b) Atomic scattering factor of Au. (c) Structure factors for fcc Au. (d) Zone axes included in orientation plan. (e) Diffraction patterns for various orientations, and (f) corresponding orientation plan slices. (g) Correlogram maxima for each pattern in (e) as a function of zone axis, and (h) corresponding in-plane rotation correlation. Highest correlation scores are shown in (g) and (h) using red circles.

defined by [52]

$$\begin{aligned} \mathbf{a}^* &= \frac{\mathbf{b} \times \mathbf{c}}{\mathbf{a} \cdot [\mathbf{b} \times \mathbf{c}]} = \frac{\mathbf{b} \times \mathbf{c}}{\Omega} \\ \mathbf{b}^* &= \frac{\mathbf{c} \times \mathbf{a}}{\mathbf{b} \cdot [\mathbf{c} \times \mathbf{a}]} = \frac{\mathbf{c} \times \mathbf{a}}{\Omega} \\ \mathbf{c}^* &= \frac{\mathbf{a} \times \mathbf{b}}{\mathbf{c} \cdot [\mathbf{a} \times \mathbf{b}]} = \frac{\mathbf{a} \times \mathbf{b}}{\Omega}, \end{aligned} \quad (5.1)$$

where  $\times$  represents the vector cross product and  $\Omega$  is the unit cell volume in real space. Note that this definition does not include factors of  $2\pi$ , and therefore all reciprocal coordinates have spatial frequency units.

Next, we calculate the position of all reciprocal lattice points required for our kinematical diffraction calculation, given by

$$\mathbf{g}_{hkl} = h \mathbf{a}^* + k \mathbf{b}^* + l \mathbf{c}^*, \quad (5.2)$$

where  $h$ ,  $k$ , and  $l$  are integers representing the reciprocal lattice index points corresponding to the Miller indices  $(h, k, l)$ . We include only points where  $|\mathbf{q}_{hkl}| < k_{\max}$ , where  $\mathbf{q} = (q_x, q_y, q_z)$  are the 3D coordinates in reciprocal space, i.e. those which fall inside a sphere given by the maximum scattering vector  $k_{\max}$ . To find all reciprocal lattice coordinates, we first determine the shortest vector given by linear combinations of  $(\mathbf{a}^*, \mathbf{b}^*, \mathbf{c}^*)$ , and divide  $k_{\max}$  by this vector length to give the range for  $(h, k, l)$ . We then tile  $(h, k, l)$  in both the positive and negative directions up to this value, and then remove all points with vector lengths larger than  $k_{\max}$ .

The reciprocal lattice defined above represents all possible coordinates where the structure factor coefficients  $V_g(\mathbf{q})$  could be non-zero. The structure factor coefficients depend only the atomic basis and are given by

$$F_{hkl} = \frac{1}{\Omega} \sum_{n=1}^N f_n(|\mathbf{g}_{hkl}|) \exp[-2\pi i(h, k, l) \cdot \mathbf{p}_n], \quad (5.3)$$

where  $f_n$  are the the single-atom scattering factors for the  $n$ th atom, which describe the scattering amplitude for a single atom isolated in space. There are multiple ways to parameterize  $f_n$ , but here we have chosen to use the factors defined by [96] which are implemented in `py4DSTEM`. Fig. 5.1b shows the atomic scattering factor for an Au atom.

We have now defined all structure factor coefficients for a perfect infinite crystal as

$$V_g(\mathbf{q}) = \begin{cases} F_{hkl} & \text{if } \mathbf{q} = \mathbf{g}_{hkl} \\ 0 & \text{otherwise.} \end{cases} \quad (5.4)$$

Fig. 5.1c shows the structure factors of fcc Au, where the marker size denotes the intensity (magnitude squared) of the  $F_{hkl}$  values.

## Calculation of Kinematical Diffraction Patterns

Here we briefly review the theory of kinematical diffraction of finite crystals, following [40]. We can fully describe an electron plane wave by its wavevector  $\mathbf{k}$ , which points in the direction of the electron beam and has a length given by  $|\mathbf{k}| = 1/\lambda$ , where  $\lambda$  is the (relativistically-corrected) electron wavelength. Bragg diffraction of the electron wave along a direction  $\mathbf{k}'$  occurs when electrons scatter from equally spaced planes in the crystal, described in reciprocal space as

$$\mathbf{k}' = \mathbf{k} + \mathbf{g}_{hkl}. \quad (5.5)$$

For elastic scattering,  $\mathbf{k}'$  has the same length as  $\mathbf{k}$ , and so scattering can only occur along the spherical surface known as the *Ewald sphere construction* [46]. For a perfect infinite crystal, scattering will seemingly almost never happen since it requires intersection of the Ewald sphere with the infinitesimally small points of the reciprocal lattice vectors. However, real samples have finite dimensions, and thus in reciprocal space their lattice points will be convolved by a *shape factor*  $D(\mathbf{q})$ . Therefore diffraction can still occur, as long as Eq. 5.5 is approximately satisfied.

If the sample foil is tilted an angle  $\alpha$  away from the beam direction, the vector between a reciprocal lattice point  $\mathbf{g}$  and its closest point on the Ewald sphere has a length equal to

$$s_{\mathbf{g}} = \frac{-\mathbf{g} \cdot (2\mathbf{k} + \mathbf{g})}{2|\mathbf{k} + \mathbf{g}| \cos(\alpha)}. \quad (5.6)$$

The  $s_{\mathbf{g}}$  term is known as the *excitation error* of a given reciprocal lattice point  $\mathbf{g}$ . When the excitation error  $s_{\mathbf{g}} = 0$ , the Bragg condition is exactly satisfied. When the length of  $s_{\mathbf{g}}$  is on the same scale as the extent of the shape factor, the Bragg condition is approximately satisfied.

A typical TEM sample can be approximately described as a slab or foil which is infinite in two dimensions, and with some thickness  $t$  along the normal direction. The shape function of such a sample is equal to

$$D(q_z) = \frac{\sin(\pi q_z t)}{\pi q_z}. \quad (5.7)$$

Because this expression is convolved with each reciprocal lattice point, we can replace  $q_z$  with the distance between the Ewald sphere and the reciprocal lattice point. For the orientation mapping application considered in this paper, we assume that  $\alpha = 0$ , and that the sample thickness  $t$  is unknown. Instead, we replace Eq. 5.7 with the approximation

$$D(q_z) = \exp\left(-\frac{q_z^2}{2\sigma^2}\right), \quad (5.8)$$

where  $\sigma$  represents the excitation error tolerance for a given diffraction spot to be included. We chose this expression for the shape function because it decreases monotonically with increasing distance between the diffraction spot and the Ewald sphere  $q_z$ , and produces smooth output correlograms.



To calculate a kinematic diffraction pattern for a given orientation  $\mathbf{w}$ , we loop through all reciprocal lattice points and use Eq. 5.6 to calculate the excitation errors. The intensity of each diffraction spot is given by the intensity of the structure factor  $|F_{hkl}|^2$ , reduced by a factor defined by either Eq. 5.7 or Eq. 5.8. We define the position of the diffraction spots in the imaging plane by finding two vectors perpendicular to the beam direction, and projecting the diffraction vectors  $q$  into this plane. The result is the intensity of each spot  $I_m$ , and its two spatial coordinates  $(q_{m_x}, q_{m_y})$ , or alternatively their polar coordinates  $q_m = \sqrt{q_{m_x}^2 + q_{m_y}^2}$  and  $\gamma_m = \arctan2(q_{m_y}, q_{m_x})$ . Note that the in-plane rotation angle is arbitrarily defined for kinematical calculations in the forward direction. The resulting diffraction patterns are defined by the list of  $M$  Bragg peaks, each defined by a triplet  $(q_{m_x}, q_{m_y}, I_m)$  in Cartesian or  $(q_m, \gamma_m, I_m)$  in polar coordinates.

Fig. 5.1e shows diffraction patterns for fcc Au, along five different zone axes (orientation directions). Each pattern includes Bragg spots out to a maximum scattering angle of  $k_{\max} = 1.5 \text{ \AA}^{-1}$ , and each spot is labeled by the  $(hkl)$  indices. The marker size shown for each spot scales with the amplitude of each spot's structure factor, decreased by Eq. 5.8 using  $\sigma = 0.02 \text{ \AA}^{-1}$ .

## Generation of an Orientation Plan

The orientation of a crystal can be uniquely defined by a  $[3 \times 3]$ -size matrix  $\overleftrightarrow{\mathbf{m}}$ , which rotates vectors  $\mathbf{d}_0$  in the sample coordinate system to vectors  $\mathbf{d}$  in the parent crystal coordinate system

$$\begin{aligned} \begin{bmatrix} d_x \\ d_y \\ d_z \end{bmatrix} &= \begin{bmatrix} u_x & v_x & w_x \\ u_y & v_y & w_y \\ u_z & v_z & w_z \end{bmatrix} \begin{bmatrix} d_{0_x} \\ d_{0_y} \\ d_{0_z} \end{bmatrix} \\ \mathbf{d} &= \overleftrightarrow{\mathbf{m}} \mathbf{d}_0, \end{aligned} \quad (5.9)$$

where the first two columns of  $\overleftrightarrow{\mathbf{m}}$  given by  $\mathbf{u}$  and  $\mathbf{v}$  represent the orientation of the in-plane  $x$  and  $y$  axis directions of the parent crystal coordinate system, respectively, and the third column  $\mathbf{w}$  defines the zone axis or out-plane-direction. The orientation matrix can be defined in many different ways, but we have chosen to use a  $Z - X - Z$  Euler angle scheme [155], defined as

$$\overleftrightarrow{\mathbf{m}} = \begin{bmatrix} C_1 & -S_1 & 0 \\ S_1 & C_1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & C_2 & S_2 \\ 0 & -S_2 & C_2 \end{bmatrix} \begin{bmatrix} C_3 & -S_3 & 0 \\ S_3 & C_3 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (5.10)$$

where  $C_1 = \cos(\phi_1)$ ,  $S_1 = \sin(\phi_1)$ ,  $C_2 = \cos(\theta_2)$ ,  $S_2 = \sin(\theta_2)$ ,  $C_3 = \cos(\phi_3)$ , and  $S_3 = \sin(\phi_3)$ . The Euler angles  $(\phi_1, \theta_2, \phi_3)$  chosen are fairly arbitrarily, as are the signs of rotation matrices given above.

In order to determine the orientation  $\overleftrightarrow{\mathbf{m}}$  of a given diffraction pattern, we use a two-step procedure. The first step is to calculate an *orientation plan*  $P((\phi_1, \theta_2), \phi_3, q_s)$  for a given reference crystal. The second step, which is defined in the following section, is to generate

a *correlogram* from each reference crystal, from which we directly determine the correct orientation.

The first two Euler angles  $\phi_1$  and  $\theta_2$  represent points on the unit sphere which will become the zone axis of a given orientation. The first step in generating an orientation plan is to select 3 vectors delimiting the extrema of the unique, symmetry-reduced zone axes possible for a given crystal. Fig. 5.1d shows these boundary vectors for fcc Au, which are given by the directions [001], [011], and [111]. We next choose a sampling rate or angular step size, and generate a grid of zone axes to test. We define a 2D grid of vectors on the unit sphere which span the boundary vectors by using spherical linear interpolation (SLERP) formula defined by [164]. These points with a step size of  $2^\circ$  are shown in Fig. 5.1d. The rotation matrices which transform the zone axis vector (along the z axis) are given by the matrix inverse of the first two terms in Eq. 5.10.

We then examine the vector lengths of all non-zero reciprocal lattice points  $\mathbf{g}_{hkl}$  and find all unique spherical shell radii  $q_s$ . These radii will become the first dimension of our orientation correlogram, where each radius is assigned one index  $s$ . We loop through all included zone axes, and calculate a polar coordinate representation of the kinematical diffraction patterns. [187] pointed out that a polar transformation can make the in-plane rotation matching step more efficient, as the pattern rotation becomes a simple translation. We will further speed up the in-plane matching by using Fourier correlation along the angular dimension after the polar transformation [39].

For each zone axis, the first step to compute the plan is to rotate all structure factor coordinates by the matrix inverse of the first two terms in Eq. 5.10. Next, we compute the excitation errors  $s_g$  for all peaks assuming a [0, 0, 1] projection direction, and the in-plane rotation angle of all peaks  $\gamma_q$ . The intensity values of the orientation plan for all  $q_s$  shells and in-plane rotation values  $\phi_3$  are defined using the expression

$$P_0((\phi_1, \theta_2), \phi_3, q_s) = \sum_{\{\mathbf{g}: |\mathbf{g}|=q_s\}} q_s^\gamma |V_{\mathbf{g}}|^\omega \times \max \left\{ 1 - \frac{1}{\delta} \sqrt{s_{\mathbf{g}}^2 + [\text{mod}(\phi_3 - \gamma_{\mathbf{g}} + \pi, 2\pi) - \pi]^2 q_s^2}, 0 \right\}, \quad (5.11)$$

where  $\delta$  is the correlation kernel size,  $\gamma$  and  $\omega$  represent the power law scaling for the radial and peak amplitude terms respectively,  $\max(\dots)$  is the maximum function, which returns the maximum of its two arguments,  $\text{mod}(\dots)$  is the modulo operator, and the summation includes only those peaks  $\mathbf{g}$  which belong to a given radial value  $q_s$ . We have used the combined indexing notation for  $(\phi_1, \theta_2)$  to indicate that in practice, this dimension of the correlation plan contains all zone axes, and thus the entire array has only 3 dimensions. The correlation kernel size  $\delta$  defines the azimuthal extent of the correlation signal for each reciprocal lattice point. Note that Eqs. 5.8 and 5.7 are not used for the calculation of orientation plans.

We normalize each zone axis projection using the function

$$A(\phi_1, \theta_2) = \frac{1}{\sqrt{\sum_{\phi_3} \sum_{q_s} P_0((\phi_1, \theta_2), \phi_3, q_s)^2}},$$

yielding the final normalized orientation plan

$$P((\phi_1, \theta_2), \phi_3, q_s) = A(\phi_1, \theta_2) P_0((\phi_1, \theta_2), \phi_3, q_s) \quad (5.12)$$

By default, we have weighted each term in the orientation plan with the prefactor  $q_s |V_g|$ , i.e. setting  $\gamma = \omega = 1$ . The  $q_s$  term gives slightly more weight to higher scattering angles, while the  $|V_g|$  term is used to weight the correlation in favour of peaks with higher structure factor amplitudes, which was found to be more reliable than weighting the orientation plan by  $|V_g|^2$ , which weights each peak by its structure factor intensity.

Fig. 5.1f shows 2D slices of the 3D orientation plan, for the 5 diffraction patterns shown in Fig. 5.1e. The in-plane rotational symmetry of each radial band is obvious for the low index zone axes, e.g. for the [001] orientated crystal, the first row of the corresponding orientation plan consists of four spots which maintains the 4-fold symmetry of the diffraction pattern and can be indexed as [020], [200], [0 $\bar{2}$ 0] and [ $\bar{2}$ 00]. The final step is to take the 1D Fourier transform along the  $\phi_3$  axis in preparation for the Fourier correlation step defined in the next section.

## Correlation Pattern Matching

For each diffraction pattern measurement, we first measure the location and intensity of each Bragg disk by using the template matching procedure outlined by [158]. The result is a set of  $M$  experimental diffraction peaks defined by the triplets  $(q_m, \gamma_m, I_m)$  in polar coordinates. Note that while all ACOM approaches we are aware of store the diffraction libraries in vector format [148], here we also reduce the experimental data to a list of peak position and intensity vectors. This has the effect of deconvolving the probe shape from the diffracted disk, and thus improving the resolution. From the experimental peaks, we calculate the sparse polar diffraction image  $X(\phi_3, q_s)$  using the expression

$$X(\phi_3, q_s) = \sum_{\{q_m : |q_m - q_s| < \delta\}} q_m^\gamma I_m^{\omega/2} \times \max \left\{ 1 - \frac{1}{\delta} \sqrt{(q_m - q_s)^2 + [\text{mod}(\phi_3 - \gamma_m + \pi, 2\pi) - \pi]^2 q_s^2}, 0 \right\}. \quad (5.13)$$

Note that the polar coordinates  $q_s$  and  $\phi_3$  used in this expression are identical to those used in the orientation plan calculation. The measured diffraction intensity is not normalized, as realistic sample thicknesses we expect the intensity to vary significantly from the kinematically predicted values.

By default, we again use prefactors weighted by the peak radius and estimated peak amplitude given by the square root of the measured disk intensities. However, if the dataset being analyzed contains a large number of different sample thicknesses, multiple scattering can cause strong oscillations in the peak amplitude values. The intensity weighting factor  $\omega$  provides a similar effect as the “gamma correction” used in many diffraction template matching routines [27], but acts on the measured disk intensities rather than the original diffraction pattern. As we will see in the simulations below, in these situations the best results may be achieved by setting  $\omega = 0$ , i.e. ignoring peak intensity and weighting only by the peak radii. Note that in the diffraction image, the correlation kernel size  $\delta$  again gives the azimuthal extent of the correlation signal. However, in Eq. 5.13 it also sets the range over which peaks are included in a given radial bin, and the fraction of the intensity assigned to each radial bin. To prevent experimental disk position errors from causing peaks to be assigned erroneously when the radial bins are near to one another (such as due to different reflections with nearly-similar spacing), experimental peaks can be included in multiple radial bins if they fall within the correlation kernel size of multiple bins. The kernel size  $\delta$  can be optimized for each type of sample: if the sample contains crystals with large lattice distortions, a larger kernel size can be used to increase the tolerance. Alternatively, if a sample consists of many overlapping grains then the kernel size can be decreased to lower the probability of false positive matches for nearby orientations.

Finally, we calculate the correlation  $C((\phi_1, \theta_2), \phi_3)$  of this image with the orientation plan using the expression

$$C((\phi_1, \theta_2), \phi_3) = \sum_{q_s} \mathcal{F}^{-1} \{ \mathcal{F} \{ P((\phi_1, \theta_2), \phi_3, q_s) \}^* \mathcal{F} \{ X(\phi_3, q_s) \} \}, \quad (5.14)$$

where  $\mathcal{F}$  and  $\mathcal{F}^{-1}$  are 1D forward and inverse fast Fourier transforms (FFTs) respectively along the  $\phi_3$  direction, and the  $*$  operator represents taking the complex conjugate. We use this correlation over  $\phi_3$  to efficiently calculate the in-plane rotation of the diffraction patterns. The maximum value in the correlogram will ideally correspond to the most probable orientation of the crystal. In order to account for mirror symmetry of the 2D diffraction patterns, we can also compute the correlation

$$C_{\text{mirror}}((\phi_1, \theta_2), \phi_3) = \sum_{q_s} \mathcal{F}^{-1} \{ \mathcal{F} \{ P((\phi_1, \theta_2), \phi_3, q_s) \}^* \mathcal{F} \{ X(\phi_3, q_s) \}^* \}, \quad (5.15)$$

where the mirror operation is accomplished by taking the complex conjugate of  $\mathcal{F} \{ X(\phi_3, q_s) \}$ . For each zone axis  $(\phi_1, \theta_2)$ , we take the maximum value of  $C$  and  $C_{\text{mirror}}$  in order to account for this symmetry. Figs. 5.1g and h show 5 output correlograms, for the 5 diffraction patterns shown in Fig. 5.1e. For each zone axis  $(\phi_1, \theta_2)$ , we have computed the maximum correlation value, which are plotted as a 2D array in Fig. 5.1g. In each case, the highest value corresponds to the correct orientation.

Note that to calculate the correlation values, we have re-binned the vector peak data from both the orientation plan and experimental peaks into a polar coordinate image with sparse radial bins. It is also possible to perform the correlations of Eqs. 5.14 and 5.15 directly on the inputs into Eq. 5.11 and experimental peaks  $(q_m, \gamma_m, I_m)$ . However, in our numerical tests, correlations computed from vector inputs were slower than the image correlation approach for all ranges of parameters tested. We attribute this to two factors: first, the polar coordinate images we use have a very small number of radial bins since we only operate on shells which contain reciprocal lattice vectors. Second, calculating the correlation of all in-plane rotations using Fourier transforms is highly efficient due to the high speed of the fast Fourier transform. This is why we have elected to compute the orientation correlations using radially-sparse polar coordinate images.

Fig. 5.1h shows the correlation values along the  $\phi_3$  axis, for the  $(\phi_1, \theta_2)$  bins with the highest correlation value in Fig. 5.1g. The symmetry of the correlation values in Fig. 5.1h reflect the symmetry of the underlying patterns. For the  $[0, 0, 1]$ ,  $[0, 1, 1]$ , and  $[1, 1, 1]$ , diffraction patterns, the in-plane angle  $\phi_3$  correlation signals have 4-fold, 2-fold, and 6-fold rotational symmetry respectively. By contrast, the asymmetric diffraction patterns with zone axes  $[1, 1, 3]$ , and  $[1, 3, 5]$  have only a single best in-plane orientation match.

The above default values are designed for matching of kinematical diffraction patterns. However, thermal excitation and multiple scattering can lead to non-zero intensities of the “kinematically forbidden peaks,” i.e. diffraction signals at reciprocal lattice points where the structure factor is zero. In order to include forbidden peaks, we can include all points where  $V = 0$  in Eq. 5.4 by setting the structure factor threshold to zero. In a future update of the code, we will use dynamical (i.e. including multiple scattering) structure factor calculations to include peaks which are likely to be excited by multiple scattering. Additionally, we can set  $\omega = 0$  in Eqs. 5.11 and 5.13, which removes the dependence of the correlation function on the peak intensities entirely, and uses only the peak positions. These steps will calculate the orientation correlation score using only the position of all scattering vectors, including the forbidden peaks.

## Matching of Overlapping Diffraction Patterns

In order to match multiple overlapping crystal signals, we have implemented an iterative detection process. First, we use the above algorithm to determine the best fit orientation for a given pattern. Next, the forward diffraction pattern is calculated for this orientation. We then loop through all experimental peaks, and any within a user-specified deletion radius are removed from the pattern. By default, this deletion radius is set to half of the correlation kernel size, i.e.  $0.5 \delta$ . This value can be modified by the user depending on how close together the diffraction peaks are for a given experiment. Peaks which are outside of the deletion radius, but within the correlation kernel size, have their intensities reduced by a factor defined by the linear distance between the experimental and simulated peaks divided by the distance between the correlation kernel size and the deletion radius. Then, the ACOM correlation matching procedure is repeated until the desired number of matches have been




Figure 5.2: **Examples of alternative orientation plan types in `py4DSTEM`.** Fiber texture examples where (a) orientations fully orbit around a single zone axis (the fiber axis), or (b) contain only a symmetry-reduced wedge of zone axes which orbit around a the fiber axis. (c) Examples of orientation plans generated directly from Materials Project entries [70], using pymatgen symmetries [121].

found, or no further orientations are found. Note that while we could update the correlation score after peak deletion, we output the original magnitude of the full pattern correlogram in order to accurately calculate the probability of multiple matches.

## ACOM Integration into `py4DSTEM`

The ACOM pattern matching described has been implemented into the `py4DSTEM` python toolkit written by [158]. A typical ACOM workflow starts with using `py4DSTEM` to import the 4D dataset and one or more images of the vacuum probe. We then use a correlation template matching procedure to find the positions of all diffracted disks at each probe position [132]. We use the correlation intensity of each detected peak as an estimate of the peak’s intensity. The resulting set of  $M$  peaks defined by the values  $(q_m, \gamma_m, I_m)$  are stored as a `PointList` object in `py4DSTEM`. Because the number of peaks detected at each probe position can vary, we store the full set of all detected peaks in a `PointListArray` object in `py4DSTEM`, which provides an interface to the ragged structured numpy data.

Most experimental datasets contain some degree of ellipticity, and the absolute pixel size must be calibrated. We perform these corrections on the set of measured diffraction disks using the `py4DSTEM` calibration routines defined by [158]. We know that the correlation approach is relatively robust against both ellipticity and small errors in the reciprocal space pixel size. However, precise phase mapping may require us to distinguish between crystals with similar lattice parameters; these experiments will require accurate calibration.

We perform ACOM in `py4DSTEM` by first creating a `Crystal` object, either by specifying the atomic basis directly, or by using the `pymatgen` package [121] to import structural data from crystallographic information files (CIF), or the Materials Project database [70]. The `Crystal` object is used to calculate the structure factors, and generate an orientation plan. The final step is to use the orientation plan to determine the best match (or matches) for each probe position, from the list of calibrated diffraction peaks. If the sample contains multiple phases, we perform the orientation plan calculation and correlation matching for each unique crystal structure.

In addition to specifying the orientation plan spanning 3 vectors as in Fig. 5.1, we define additional methods to describe the space of possible orientations. One such example is *fiber texture*, where we assume the crystals are all orientated near a single zone axis known as the fiber axis, shown in Figs. 5.2a and b. We can vary the angular range of zone axes included away from the fiber axis as in Fig. 5.2a, as well as choose the azimuthal range around this axis as in Fig. 5.2b to account for symmetry around the fiber axis. Alternatively, an “automatic” option is provided, which uses `pymatgen` to determine the symmetry of the structure and automatically choose the span of symmetrically unique zone axes which should be included in the orientation plan, based on the point group symmetry [40]. This is shown for a selection of different Materials Project database entries in Fig. 5.2c.

## Simulations of Diffraction Patterns from Thick Samples

One important metric for the performance of an orientation mapping algorithm is how well it performs when the diffraction patterns contain significant amounts of multiple scattering. We have therefore used our ACOM algorithm to measure the orientation of simulated diffraction patterns from samples tilted along many directions, over a wide range of thicknesses. We

performed these simulations using the multislice algorithm [35], and methods defined by [79] and [122]. These methods are implemented in the `Prismatic` simulation code by [146]. The diffraction patterns were generated using a acceleration potential of 300 keV, a 0.5 mrad convergence angle, with real space and reciprocal pixel sizes of 0.05 Å and 0.01 Å<sup>-1</sup> respectively, with 4 frozen phonons. In total we have simulated 3750 diffraction patterns from Cu, Ag, and Au fcc crystals, over 25 zone axes ([0, 0, 1] to [3, 4, 4] excluding symmetrically redundant reflections) and thicknesses up to 100 nm with a 2 nm step size. These diffraction patterns were generated using the simulation pipeline and database defined by [144].

## Chemical Synthesis of Twisted AuAgPd Nanowires

The performed synthesis was reproduced with minor modifications from a known method given by [180]. All reagents were purchased from Sigma Aldrich. We prepared the following solutions: 500 mM of Polyvinylpyrrolidone (PVP, MW 40,000) in Dimethylformamide (DMF), 50 mM Gold(III) chloride trihydrate (HAuCl<sub>4</sub> · 3H<sub>2</sub>O, >49.0% Au Basis) in DMF, 50 mM Silver nitrate (AgNO<sub>3</sub>) in deionized (DI) water (resistivity > 18 MΩ/cm), and 400 mM L-ascorbic acid (>99.0%, crystalline) in DI water. We created the reaction solution in a 4 mL glass vial (washed 3x with DI water and acetone) by mixing 800 μL DMF, 100 μL PVP, 20 μL HAuCl<sub>4</sub>, and 20 μL AgNO<sub>3</sub>. We mixed the solution for approximately 2 seconds using a Vortex-Genie 2 Mixer set to a value of 10, which spins the reaction solution at a speed of approximately 3200 rpm, then added 100 μL of L-ascorbic acid solution drop-wise to the mixture while gently swirling by hand. At this point, the color changed from pale yellow to clear. We left the solution at room temperature for 7 days, at which point the solution was light brown/purple. The primary product of this reaction was straight, ultrathin Au-Ag nanowires (2 nm in diameter).

To twist the underlying ultrathin Au-Ag nanowires, we prepared solutions of 1.875 mM L-ascorbic acid and 2 mM H<sub>2</sub>PdCl<sub>2</sub> in DI water. In a 4 mL glass vial (3x washed with DI water/acetone), we added 50 μL of the Au-Ag reacted solution to 640 μL of the L-ascorbic acid solution. Finally, we added 60 μL of the H<sub>2</sub>PdCl<sub>4</sub> solution and allowed the sample to incubate for at least 30 minutes. We purified the reaction solution by centrifuging the product down at 7500 rpm for 4 minutes. We decanted the supernatant, and then rinsed the reaction with DI water 3 times and re-dispersed in DI water. We prepared TEM samples of this material by depositing 10 μL of purified nanowire solution onto 400 mesh formvar/ultrathin carbon grids.

## 4D-STEM Experiments with Patterned Apertures

We collected the experimental data using a double aberration-corrected modified FEI Titan 80-300 microscope (the TEAM I instrument at the National Center for Electron Microscopy within Lawrence Berkeley National Laboratory). This microscope is equipped with a Gatan K3 detector and Continuum spectrometer, and was set to collect diffraction patterns integrated over 0.05 seconds, with 4x binning giving a calibrated pixel size of 0.00424 Å<sup>-1</sup>. We



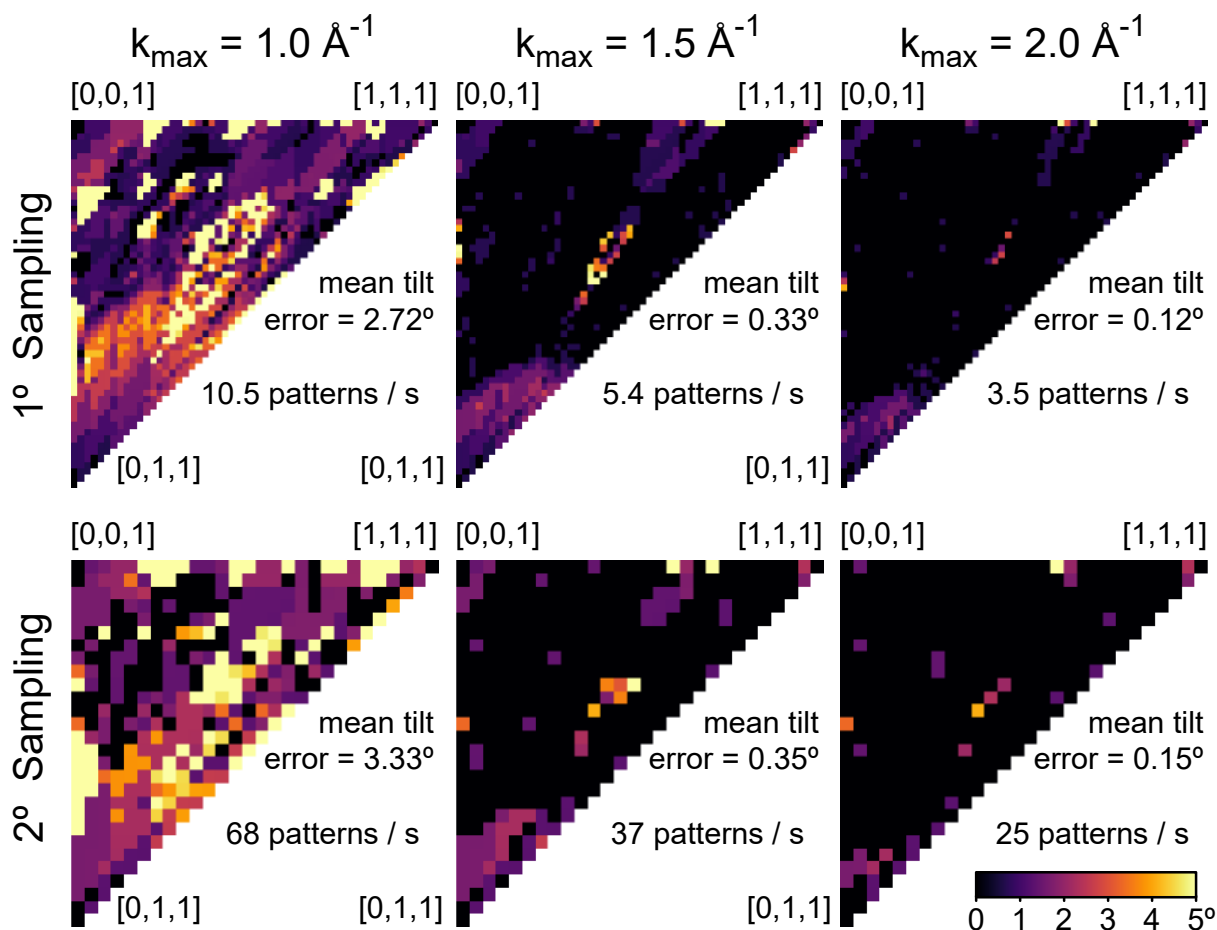


Figure 5.3: **Zone axis misorientation as a function of sampling and maximum scattering angle for kinematical simulations.** The mean tilt error and number of patterns matched per second are shown inset for each panel.

used an accelerating voltage of 300 keV, an energy slit of 20 eV, and a spot size of 6. The beam current was 6 pA. We used a 10  $\mu\text{m}$  bullseye aperture (probe size of approximately 1 nm) to form the STEM probe in order to improve detection precision of the Bragg disks [199]. We used a convergence semiangle of 2 mrad, with a camera length of 1.05 m. We recorded the experimental dataset using a probe step size of 5  $\text{\AA}$ , with a total of 286 and 124 steps in the x and y directions.

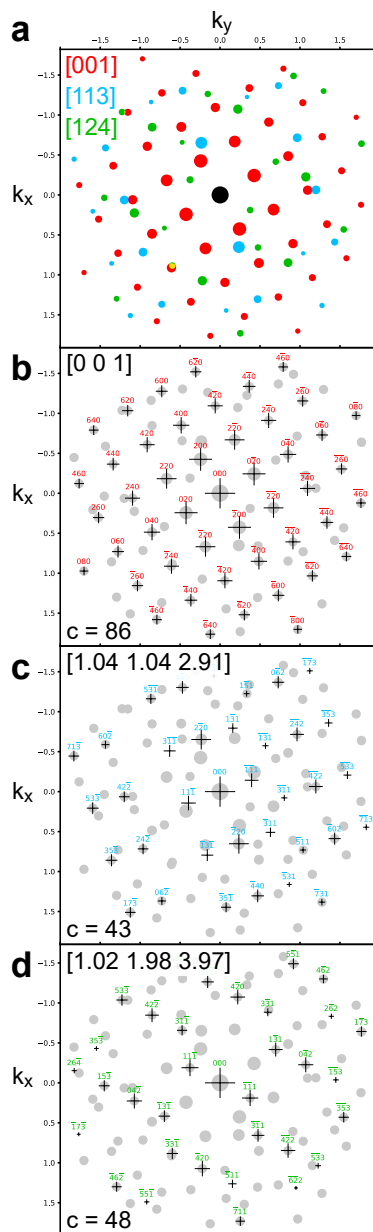


Figure 5.4: **ACOM of overlapping diffraction patterns.** (a) Three overlapping diffraction patterns with randomly chosen in-plane rotations. (b) First match, (c) second match, and (d) third match returned by ACOM code. The fitted zone axes and correlation scores are inset into fits.

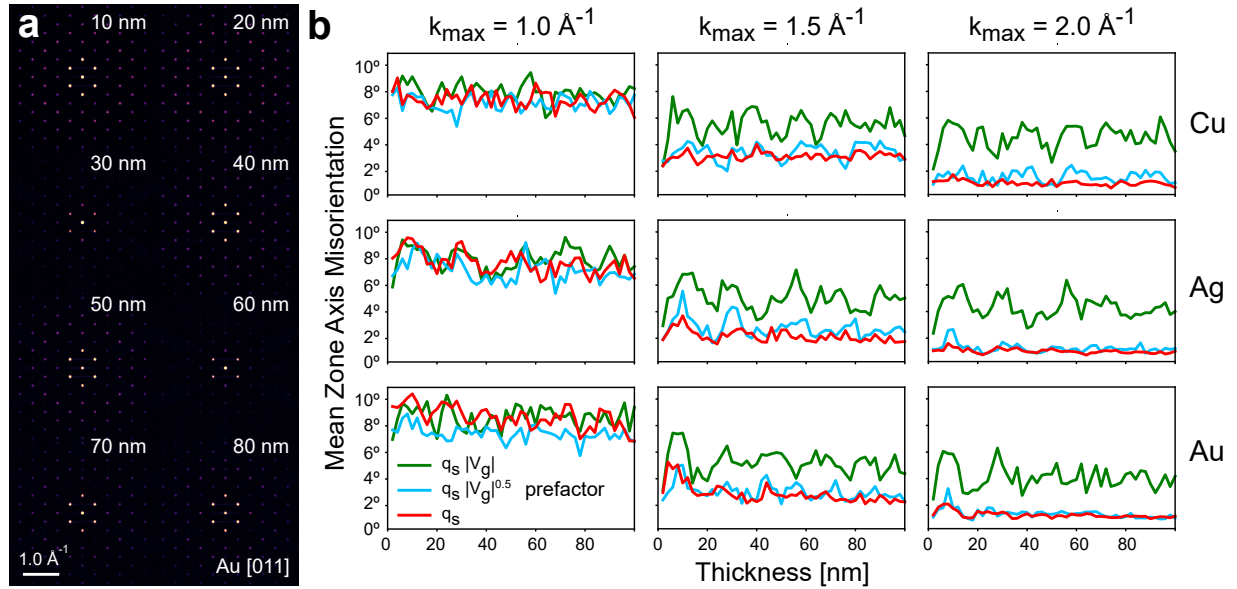


Figure 5.5: **Dynamical simulated diffraction patterns.** (a) Example diffraction patterns for Au oriented to the [011] zone axis for 10-80 nm thick slices. (b) Plots showing the mean zone axis misorientation in degrees as a function of thickness for Cu, Ag, and Au. Each plot shows the errors for correlation prefactors of  $q_s|V_g|$  (red) and  $q_s$  (blue).

### 5.3 Results and Discussion

#### ACOM of Kinematical Calculated Diffraction Patterns

For the first test of our correlation method, we applied it to the same patterns calculated to generate an orientation plan for fcc Au. Next, we measured the calculation time and angular error between the measured and ground truth zone axes for each pattern. The results are plotted in Fig. 5.3 for different maximum scattering angles  $k_{\text{max}}$ , and angular sampling of  $1^\circ$  and  $2^\circ$ .

The results in Fig. 5.3 show that the angular error in zone axis orientation is relatively insensitive to the angular sampling. However, the angular error drops by a factor of 10 from approximately  $\approx 3^\circ$  to  $\approx 0.3^\circ$  when increasing the maximum scattering angle included from  $k_{\text{max}} = 1 \text{ \AA}^{-1}$  to  $1.5 \text{ \AA}^{-1}$ , and by another factor of 2-3 when increasing  $k_{\text{max}}$  to  $2 \text{ \AA}^{-1}$ . This is unsurprising, as examining Fig. 5.1e shows that there is a large number of visible Bragg spots outside of  $k_{\text{max}} = 1 \text{ \AA}^{-1}$ , and because Bragg disks at higher scattering angles provide better angular precision relative to low  $k$  disks. This result emphasizes the importance of recording as many diffraction orders as possible when performing orientation matching of 4D-STEM data. More spots can be included by collecting data out to higher scattering angles, or by

reducing the convergence semiangle to bring weakly diffracting peaks above the noise floor. Setting  $k_{\max}$  beyond the highest angle detected disks will not yield any additional precision but will make the orientation plan larger, so  $k_{\max}$  should be chosen to correspond to the highest scattering angle peaks detected in an experiment.

The inset calculation times reported are for the single-threaded ACOM implementation in `py4DSTEM`, running in Anaconda [3] on a laptop with an Intel Core i7-10875H processor, running at 2.30 GHz. The calculation times can be increased by an order of magnitude or more when running in parallel, or by using a GPU to perform the matrix multiplication and Fourier transform steps.

## ACOM of Overlapping Diffraction Patterns

A common feature of polycrystalline samples is overlapping grains along the beam direction, leading to overlapping diffraction patterns. To demonstrate the ability of our method to work with overlapping grains, we have generated a combined set of diffraction patterns with three low index zone axes and random in-plane rotations, plotted in Figs. 5.4a. Figs. 5.4b-d shows the first three matches returned by our ACOM code using a kernel size of  $\delta = 0.08$  Å, a zone axis step size of  $1^\circ$ , and a prefactor of  $q_s|V_g|$ . The multi-pattern peak deletion radius was slightly decreased from the default value of  $0.04$  Å to  $0.02$  Å to prevent removal of adjacent peaks as matches are assigned. Our ACOM code has correctly returned 3 zone axes which match the ground truth values. This example also demonstrates a procedure which could be used to map the location and orientation of multiple phases, even when the diffraction patterns overlap.

## ACOM of Dynamical Simulated Diffraction Patterns

In diffraction experiments using thick specimens, strong dynamical diffraction effects such as multiple scattering can occur. This effect is especially pronounced in diffraction experiments along low index zone axes, where the diffracted peak intensities oscillate as a function of thickness. In order to test the effect of oscillating peak intensities on our ACOM method, we have simulated diffraction patterns for Cu, Ag, and Au fcc crystals, along multiple zone axes. Some example diffraction patterns for the [011] zone axis of Au are plotted in Fig. 5.5a. We see that all diffraction spots have intensities which oscillate multiple times as a function of thickness.

We performed ACOM by generating orientation plans with an angular sampling of  $2^\circ$ , a correlation kernel size of  $0.08$  Å<sup>-1</sup>, and maximum scattering angles of  $k_{\max} = 1.0, 1.5,$  and  $2.0$  Å<sup>-1</sup>. We kept the radial prefactor of weighting set to  $\gamma = 1$ , and tested peak amplitude prefactors of  $\omega = 1.0, 0.5$  and  $0.0$ . The average zone axis angular misorientation as a function of thickness is plotted in Fig. 5.5b. In total, we performed orientation matching on 3750 diffraction patterns, and a total of 33750 correlation matches on a workstation with an AMD Ryzen Threadripper 3960X CPU (2.2 GHz, baseclock). The typical number of

patterns matched per second were of 80-90, 45-55 and 25-30 patterns/s for  $k_{\max}$  values of 1.0, 1.5, and  $2.0 \text{ \AA}^{-1}$  respectively.

As expected, the errors are higher than those achieved under kinematic conditions, and the trend for smaller errors with larger  $k_{\max}$  is also preserved (mean errors of  $7.25^\circ$ ,  $3.09^\circ$  and  $1.39^\circ$  for  $k_{\max}$  values 1.0, 1.5 and  $2.0 \text{ \AA}^{-1}$  respectively,  $\gamma = 1$ ,  $\omega = 0.25$ ). We did not observe any dependence of the orientation accuracy on the simulation thickness. Despite the correlation prefactor  $|Vg|$  performing well for the examples shown in Fig. 5.1, for the dynamical diffraction simulations along zone axes it was out-performed by prefactors of both  $\sqrt{|Vg|}$  ( $\omega = 0.5$ ) and omitting the peak amplitude prefactor altogether ( $\omega = 0$ ). We therefore suggest that when mapping samples with a large range of thicknesses, or many crystals aligned to low index zone axes, the position of the diffracted peaks is significantly more important than their amplitudes or intensities. One possible method to increase the average accuracy for a randomly orientated sample while using higher amplitude prefactors is to perform an experiment which recovers more kinematical values for the diffracted peak intensities, for example by precessing the electron beam when recording diffraction patterns [111, 73]. A precession experiment could however make the diffraction patterns of some grains more dynamical, and thus worsen the orientation accuracy for some probe positions. We note that there is likely no global optimal choice of orientation mapping hyperparameters for all materials and thicknesses, and this may be a worthwhile topic for future investigations.

## 4D-STEM ACOM of Twisted AuAgPd Nanowires

We have tested our ACOM algorithm with a 4D-STEM dataset collected for an AuAgPd nanowires. These nanowires are morphologically twisted into double helices via a colloidal growth process as previously reported by [180]. An image of the vacuum bullseye STEM probe is shown in Fig. 5.6a. For each detector pixel, we have calculated the maximum value across all STEM probe positions to generate a *maximum diffraction pattern*, shown in Fig. 5.6b. The beamstop used to block the center beam is visible, as well as various crystalline diffraction rings out to approximately  $1.4 \text{ \AA}^{-1}$ .

After performing the correlation peak finding algorithm in `py4DSTEM`, we have an estimated position and intensity of all detected Bragg peaks. A 2D histogram of these peaks, known as a Bragg vector map, is plotted in Fig. 5.6c. Sharp polycrystalline diffraction rings are clearly visible, as well as false positives generated by the beamstop edge. These false positives were manually removed by using a mask generated from an image of the beamstop. A high angle annular dark field (HAADF) image was simultaneously recorded during the 4D-STEM data collection, which is shown in Fig. 5.6d.

The final experimental pre-processing steps are to calibrate the diffraction pattern center, the elliptical distortions, and the absolute pixel size. We performed these steps by fitting an ellipse to the (022) diffraction ring, and by assuming a lattice constant of  $4.08 \text{ \AA}$ , corresponding to the fcc Au structure [104]. This process is explained in more detail by [158]. We assumed that the Ag lattice constant is similar to that of Au. Despite the presence of

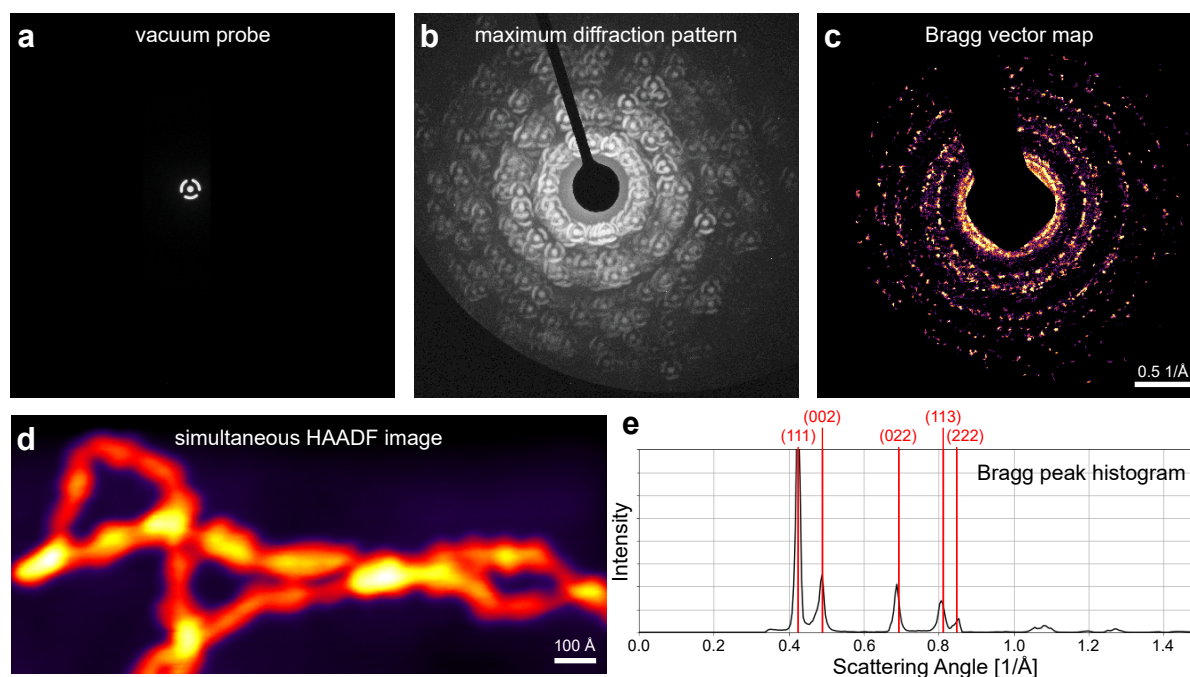


Figure 5.6: **4D-STEM scan of twisted polycrystalline AuAgPd nanowires.** (a) Diffraction image of probe over vacuum, showing bullseye pattern. (b) Maximum of each pixel in diffraction space over all probe positions. (c) Histogram of all peak locations detected by correlation in `py4DSTEM` of (a) with each pattern included in (b). (d) HAADF-STEM image of the sample. (e) 1D histogram of scattering vectors, with fcc AuAg inverse plane spacings overlaid.

Pd in the nanowires, there was no significant presence of secondary grains corresponding to the smaller lattice of fcc Pd grains. An intensity histogram of the corrected Bragg peak scattering angles are shown in Fig. 5.6e. We have overlaid the 5 smallest scattering angles of Au on Fig. 5.6e to show the accuracy of the correction.

We have performed ACOM on the AuAgPd nanowire sample, with the results shown in Fig. 5.7 shown for up to 3 matches for each diffraction pattern. For each probe position, the sum of the maximum detected correlation signals for up to three matches are shown in Fig. 5.7a. The structure is in good agreement with Fig. 5.6, though with additional modulations due to some grains generating more diffraction signal than others. Using a correlation intensity threshold of 0.5, we have plotted the number of matching patterns in Fig. 5.7b. The threshold of 0.5 was arbitrary chosen as a lower bound for a potential match, as the correlation values are scaled by the experimental intensity. Examples of 2 matches to a single diffraction pattern are plotted in Fig. 5.7c. In this figure the correlation score for the first matched pattern was higher than the second. The second match found shows some




Figure 5.7: **Orientation mapping of polycrystalline AuAgPd nanowires.** (a) Total of measured correlation signal for each probe position. (b) Estimated number of patterns indexed for each probe position. (c) Example of 2 orientations indexed from a single diffraction pattern, collected at the position indicated by the arrow shown in (b), with correlation scores inset. (d) Orientation maps of the 3 highest correlation signals for each probe position. A legend for the crystallographic orientation is shown above, and arrows indicate the direction of the x and y axes, while the zone axis direction is out of the page.




Figure 5.8: **Orientation analysis of grains in AuAgPd nanowires.** (a) Crystal grains, with in-plane (111) planes colored by orientation. (b) (111) planes shared by two overlapping grains.

deformation between the measured and simulated Bragg peak positions, and matches fewer peaks. It therefore produces a lower correlation score, which can be used to threshold the results as in Fig. 5.7d. Note that the threshold values for inclusion of any given match into the orientation maps is always user-defined.

Fig. 5.7d shows the 3D orientations plotted as inverse pole figures for all probe positions, with the 3 best matches shown. Each image is masked by the total correlation signal, so that low correlation values are colored black. Almost every diffraction pattern with Bragg disks detected was indexed for at least one orientation with high confidence. Additionally, the patterns are very consistent, with a large number of adjacent probe positions recording the same orientation. Some secondary grains are also clearly visible in the second-best match, while very few patterns have been assigned a third match with high confidence.

In order to investigate the grain organization of the AuAgPd nanowires, we have performed clustering analysis on the orientation maps. Grains with similar orientations have been clustered together by looping through each probe position and comparing its orientation to its neighbors. Grains with at least 10 contiguous probe positions are shown in Fig. 5.8a. (111) planes which lie in the image plane are overlaid onto the grain structure, colored by their orientation. Confirming our observations in Fig. 5.7d, only a few grains with substantial overlap were reliably identified. This might be due to the low thickness of the sample (only a single grain along the beam direction), some grains not being oriented close enough to a zone axis to be detected, or multiple scattering deviations in the diffracted signal. There is a noticeable bias in the orientation of the (111) planes, which tend to be



oriented horizontally near the growth direction of the nanowires.

One hypothesis for the growth mode of these twisted nanowires is that adjacent grains are connected by (111) twin planes, forming local helical structures to give the observed twisted structures. To test this hypothesis, we determined the position of (111) planes from Fig. 5.8a which are shared by two overlapping grains. Fig. 5.8b shows the location of these shared (111) planes (with plane normal differences below  $8^\circ$ ), colored by the normal vector of the plane. Many shared (111) planes were detected, most with normal vectors aligned to the wire growth direction. These observations support the hypothesis that these nanowires are composed of grains connected by (111) twin planes.

These experimental observations demonstrate the efficacy of our ACOM method. In order to improve these results, we will need to collect diffraction data with a wider angular range. This can be achieved by using precession electron diffraction [153], multibeam electron diffraction [66], or by tilting the sample or beam and recording multiple 4D-STEM datasets [110].

## 5.4 Conclusion

We have introduced an efficient and accurate method to perform automated crystal orientation mapping, using a sparse correlation matching procedure. We have implemented our methods into the open source `py4DSTEM` toolkit, and demonstrated the accuracy of our method using simulated diffraction patterns, where we show that lowering or removing the peak-intensity weighting can improve the accuracy for thick samples with substantial dynamical diffraction. We also applied ACOM to an experimental scan of a complex helical polycrystalline nanowire, where we were able to identify shared twin planes between adjacent grains which may be responsible for the twisted helical geometry. All of our methods have been made freely available to the microscopy community as open source codes. We believe that our implementation of ACOM is efficient and accurate enough to be incorporated into automated online TEM software [168]. In the future, we will improve our ACOM method using machine learning methods [116], and we will extend our ACOM methods to include multibeam electron diffraction experiments [66].

## 5.5 Source Code and Data Availability

All code used in this manuscript is available on the [py4DSTEM GitHub repository](#), and the tutorial notebooks are available on the [py4DSTEM tutorial repository](#). All simulated and experimental 4D-STEM datasets are available at [links will be added after publication].

## Part II

# Dynamical Diffraction as a Friend

## Robust Measurements by *Inverting* Dynamical Diffraction

In the previous Part, we explored two situations where dynamical scattering from thick samples could poison 4D-STEM measurements of sample properties. This occurred because for these measurements, experimental reality deviates from the simple assumptions about the scattering processes that the method utilized. In the case of strain mapping, intensity variations in converged beam diffraction patterns interfered with our algorithms for locating the diffraction disks. The strong variation of the diffracted intensity to the varying tilts in the converged beam could not be accounted for but were ultimately able to be overpowered by modifying the experiment. In the case of orientation mapping, diffraction intensities in experimental patterns from thick samples simply do not match the single scattering model that was used to generate the library for matching, and experimental modifications that could mitigate this (such as precession electron diffraction [147, 73]) were not available.

The challenge in the previous Part was that dynamical scattering causes the diffraction patterns to be *extremely* sensitive to thickness, tilt, lattice distortions, etc. *In this Part this sensitivity is not an artefact but the basis of our measurements.* By including a suitable model of the dynamical scattering process in the measurement pipeline, we can use this sensitivity to our advantage and obtain precise measurements of structural features of the sample from nanobeam 4D-STEM data. The challenge in this case becomes the scale of computation required to include the dynamical scattering model in our measurements. And the scale of this is exacerbated by the fact that there is no way to write down an analytical solution to the inverse diffraction problem when we simulate full patterns<sup>1</sup>.

In this section, we will discuss two methods for precise structural measurements that utilize Bloch wave calculations to match the parameters of a scattering model with experimental intensities. In both cases, we use strong prior knowledge about the crystal to simplify the space for refining the model. Further, we make a critical simplifying assumption: in a nanobeam diffraction pattern at a suitable shallow convergence angle, the summed intensity of each diffraction disk is approximately proportional to the intensities of the spots in the

---

<sup>1</sup>For the extremely simple case of a perfect two-beam condition we can write an expression for the forward model and potentially invert it; this expression is found in any TEM textbook, e.g. [48, 40]. There has been some research on analytic solutions to very particular three-beam conditions [19, 162], but since there is no general solution to polynomials of degree 5 and above an exact analytic solution to the N-beam case is unlikely.

plane wave pattern. This assumption separates our methods from those that fit each pixel in a CBED disk, such as the PRIMES family of techniques [137, 136, 135, 139], and allows our method to be orders of magnitude faster to evaluate, asuch that it is applicable to large area nanobeam diffraction scans.

## Crystal Orientation and Thickness

In Chapter 6, we extend the orientation mapping approach of Chapter 5 to include dynamical scattering. This extension is nontrivial: using the measurement from the kinematic library as a starting point, we perform thousands of Bloch wave diffraction calculations for a range of orientations in the neighborhood of the initial guess and for a dense sampling of thicknesses. The kinematic library matching provides us with a robust initial guess: in general, most of the diffraction vectors that can be excited in a certain orientation do get excited, and most of the ones that can't be kinematically excited don't get excited. Thus we only need to look at the fine variation in intensities in order to get precise measurements of the orientation and thickness thanks to the extreme sensitivity of the intensities to these quantities.

## Tilt and Polarization in a Stacked Heterostructure

In Chapter 7 we tackle a substantially harder problem, where two structural deviations that we are trying to recover from experimental data have substantial crosstalk, and cannot be separated using simple models of the scattering. In this chapter, we seek to disentangle the effects of polarization and tilting of the lattice in a trilayer superlattice structure. Using kinematic models of the scattering one would conventionally derive a metric for these variables based on symmetry breaking of the diffracted intensities. The challenge here is that in a sample with both varying polarization and tilt of the lattice, these two metrics will be nearly identical and the two measurements will be inextricably mixed, and are usually separated using Fourier filtering by assuming that one varies faster than the other. In our method, we derive a full dynamical model of the scattering that takes into account the fact that the sample is made of heterogenous materials and that only one of the layers can be polarized. Unlike in Chapter 6, where we were able to use effectively a brute-force search for the best match, in this section the problem has many more dimensions and none of them are inexpensive to evaluate. Thus, we derive a semi-analytic expression for computing the gradients of the diffracted intensities with respect to each structural distortion parameter, and perform gradient descent to arrive at the best fitting description of the local structure. Using this procedure we are able to effectively separate two signals which are normally totally confounded due to the complex scattering process and the similarities between them.

## Chapter 6

# A Dynamical Diffraction Refinement Method for Precise 4D–STEM Measurements of Crystal Orientation and Thickness

Diffraction intensities are extremely sensitive to the tilt and thickness of the crystal, in particular in the vicinity of a zone axis orientation. This sensitivity is highly nonlinear, and thus makes orientation mapping unreliable when the pattern library used for comparison does not include the same thickness effects as the experiment. Here, we demonstrate a refinement routine which compares the diffraction intensities measured in a 4D-STEM experiment to those from dynamical diffraction simulations, in order to precisely measure the local tilt of a crystalline sample. This refinement also matches the thickness of the crystal as a part of the refinement procedure. Using a wedge polished Si sample with a thickness gradient and a low-angle defect, we demonstrate the effectiveness of this method for improving the orientation mapping precision and for measuring sample thickness from single nanobeam electron diffraction patterns.

### 6.1 Introduction

Determining the orientation of a crystalline sample with respect to the electron beam in a (scanning) transmission electron microscope, (S)TEM, is a fundamental skill in electron microscopy. From a plane wave or nanobeam diffraction pattern and knowledge of the crystal structure, one can relate the diffraction spots to crystal directions and compute the orientation [48]. For a nanobeam four-dimensional STEM experiment (4D–STEM), where many thousands of diffraction patterns are acquired as the probe is rastered over the sample, automated methods have been developed to perform the same procedure. These methods are generally referred to as automated crystal orientation mapping (ACOM) after several papers

by Rauch and co-workers [147, 148, 149]. More recently, Cautaerts et al. have released a GPU-accelerated implementation of the same procedure as part of `pyxem` [27], and Ophus et al. have released a more computationally efficient algorithm as part of `py4DSTEM` [124], both of which are open-source packages. These automated methods work by computing a library of simulated diffraction patterns for a range of crystal orientations and comparing each experimental pattern to the library, and assigning each pattern an orientation based on which simulation it best matches.

Both the manual technique and the automated approaches utilize the geometry of the crystal lattice to compute the locations of diffracted spots for a given crystal orientation; if the intensities of the diffracted disks are also utilized, they are computed using the kinematic, single-scattering approximation. This model for the diffracted intensities is inaccurate for thick samples and for orientations nearby to high-symmetry crystal directions, where multiple scattering effects are important and a dynamical model is needed. This mismatch limits the accuracy of ACOM orientation mapping—Ophus et al found that under the best conditions a minimum orientation error of  $2^\circ$  was present at all examined thicknesses when using dynamical diffraction simulations as the input [124]. This level of precision is often sufficient to establish orientation relationships between grains in polycrystalline samples, but it is insufficient to detect bending or ripples in a thin film and is below the precision of electron backscatter diffraction orientation mapping in a scanning electron microscope, which can achieve absolute orientation precision of about  $0.5^\circ$  [67]. Previous work on matching experimental electron diffraction patterns to dynamical simulations have mostly focused on comparison with converged beam electron diffraction (CBED) patterns, since the fine details of these patterns are exquisitely sensitive to the structure, orientation [138], and even temperature of the a crystal [88, 204].

Here, we demonstrate the use of dynamical diffraction simulations for improving the precision of orientation measurements from nanobeam 4D-STEM experiments. This procedure operates by comparing the integrated diffraction disk intensities to dynamical diffraction calculations computed over a range of orientations in the neighborhood of the initial guess provided by matching against a pattern library using `py4DSTEM` [124]. These dynamical diffraction calculations are performed using the Bloch wave method, which allows us to perform all computations on lists of reflections and their intensities using sparse representations of the diffraction data.

## 6.2 Orientation and Thickness Refinement Procedure

### Initialization with Kinematic Library Matching

The dynamical refinement procedure builds upon the kinematic library matching approach of Ophus et al [124], which forms the first step in our method. The kinematic library matching is implemented in the ACOM module of `py4DSTEM`, and involves first detecting the positions and intensities of each of the Bragg disks in the 4D-STEM dataset by cross-correlation

with a template image. Then a library of kinematic diffraction patterns is computed based on the assumed known crystal structure of the material. This library contains a sparse representation of the diffraction pattern simulated for each orientation within a uniform sampling of the space of symmetrically unique orientations of the crystal. Each list of Bragg disks detected in the experiment is subsequently transformed into the same sparse representation and its correlation with each library pattern is computed. The orientation of the best matching pattern in the library is assigned to that pixel in the scan. The in-plane rotation of the pattern is determined to subpixel precision, while the rotations that bring the crystal orientation vector into alignment with the beam direction are only determined to the precision set by the sampling of the library.

## Dynamical Diffraction Calculations

Our method for refining the orientation using diffraction intensities involves comparison with dynamical diffraction calculations evaluated over a grid of orientations and thicknesses. We use the Bloch wave method (see Chapter 2 or the textbook by De Graef [40]) to perform these calculations. The Bloch wave method is ideal for a number of reasons: First, it is easy to include very small tilts of the crystal with high accuracy, since we do not need to generate an orthogonal simulation cell and our tilt resolution is not tied to the pixel size of the calculation, as is the case for multislice methods. In addition, with the Bloch wave method the primary computational effort comes in diagonalization of the structure matrix, but once it is obtained we can compute diffraction intensities for any thickness with minimal additional computation. This is very advantageous as we find that a dense sampling of thicknesses is beneficial in ensuring convergence of the refinement to match the experimental data. Finally, the output of the Bloch wave calculations are a list of reflections and their associated intensity, which naturally matches with the data structure that we use for the kinematic library sparse matching algorithm and with the Bragg disk detection routine. All of the dynamical diffraction calculations in this work were conducted using the `Crystal` module of `py4DSTEM` (version 0.13.7), which is available as open-source software. The dynamical refinement procedure shown in this work is in the process of being integrated into the main distribution of `py4DSTEM`.

## Adaptive Grid Refinement

With the approximate orientations of the crystal obtained from matching against the kinematic library in hand, we can proceed to perform the dynamical refinement. First, we use the knowledge of the approximate orientation at each pixel to assign Miller indices to each of the experimental Bragg peaks. We also generate a kinematic diffraction pattern for the same orientation, which we use as the list of beams to include in the Bloch wave simulation. The use of the kinematic pattern instead of the experimental one for the list of beams ensures that the Bloch wave simulation includes beams that may have been missed in the experiment because they fell below the detection threshold or had position errors that prevented proper

assignment of the Miller indices. Beams that were not detected must be included because they contribute to the scattering behavior, but are ignored in computing the loss function for refinement[138].

For each of the indexed experimental patterns, we use the Bloch wave simulation method to refine the orientation guess and determine the local thickness of the crystal using an iterated refinement procedure. At each refinement step, we simulate plane wave diffraction patterns for a cone of orientations around the previous guess, and for a range of thickness values. At each step, we choose the orientation that minimizes the cost function

$$C(\mathbf{R}, t) = \sum_{\mathbf{g}} (I_{\text{sim}(\mathbf{R}, t)}(\mathbf{g}) - I_{\text{exp}}(\mathbf{g})) \quad (6.1)$$

to be the refined guess, where  $\mathbf{R}$  is the rotation matrix for each test orientation,  $t$  is the thickness,  $I_{\text{exp}}$  is the list of experimentally measured disk intensities, and  $I_{\text{sim}(\mathbf{R}, t)}$  are the simulated intensities.

Following this step, we proceed by choosing a finer sampling of a smaller range of orientations for each subsequent iteration. This adaptive sampling approach allows us to hop over the many local minima that appear in the cost function while still finely exploring the region of the best match, without performing prohibitively many calculations for each pattern. As it is computationally inexpensive to compute the intensities at many thicknesses, we use the same dense sampling of thickness at each step.

### 6.3 Experimental Methods

Experimental data was acquired using the TEAM1 instrument at the National Center for Electron Microscopy, an FEI Titan equipped with probe and image correctors, a Gatan Continuum spectrometer/energy filter, and a Gatan K3 direct electron detector. A silicon wedge sample was produced from a commercial 111 oriented wafer by mechanical polishing followed by ion milling. This produced a sample with smoothly decreasing thickness which varied from the bulk down to approximately 15 nm at the edge, with likely a few nanometers of amorphous material on the surface. 4D-STEM data were acquired from this sample at an accelerating voltage of 300 kV, a semi-convergence angle of 1.4 mrad, and with an energy selecting slit with a width of 10 eV. We also acquired low-loss EELS maps from the same area of the sample at a semi-convergence angle of 17 mrad in order to obtain an independent measurement of the thickness of the sample by the log-ratio method.

### 6.4 Results and Discussion

#### Bragg Disk Detection

From the 4D-STEM dataset, we begin by using template matching to obtain the locations of all of the diffraction disks [132]. We then use the kinematic library matching routine



from [124] to identify the approximate orientation in each diffraction pattern. Results of the kinematic orientation matching are shown in Figure 6.2. As expected, the sample is identified as being single crystal and oriented near the (111) zone axis. Using the identified orientation at each scan position, we assign crystallographic indices to each of the detected Bragg disks.

While the cross-correlation intensities measure between each disk and the template in the first step are approximately equal to the disk intensity, modifications to the probe kernel and the presence of background intensity make this relationship imperfect. Thus the disk intensities are updated by integrating a circular region centered on the disk position and subtracting a local background value fitted to the neighborhood of the disk. In the present case, as the electron dose is relatively low at approximately  $10^5$  per pattern and the data was acquired without a patterned aperture [199], the disk positions contain substantial errors which result in large apparent variations in the lattice across the field of view. These errors in disk positions will translate into large errors in the integrated disk intensities, as we center our integration window on the measured disk positions. To mitigate this, we use the average lattice vector from the entire scan to determine the locations of the integration windows, correcting only for translation of the entire pattern across the detector caused by STEM descanner misalignment.

Virtual images obtained from intensity integration procedure described above for the direct beam and each of the first order reflections in the 4D-STEM dataset are shown in Figure 6.1, and show prominent fringes corresponding to variation of the thickness of the wedge polished sample. The fringes are most prominent in the brightfield image and the  $(0\bar{2}2)$  and  $(\bar{2}02)$  images but are present in all of them. From inspection of the diffraction patterns the sample is tilted so that the center of the Laue circle lies primarily between these three diffracted disks. The spacing of the fringes in these images does not necessarily correspond to those that would be observed in a two-beam bright/dark-field image as the sample is nearly perfectly on a zone axis and thus many diffraction conditions are operating. At the bottom left of the scan region, there is a defect which likely arose during the thinning process and which caused a large tilt of the lattice such that the thickness fringes disappear.

## Kinematic Orientation Mapping

Using the kinematic library matching code of Ophus et al [124], we generated a library of diffraction patterns which cover the span of unique orientation directions for a diamond cubic crystal at spacings of  $2^\circ$ . The lists of Bragg peaks from each scan position in the experiment are matched to the library in order to determine the best match, and these orientations are shown in Figure 6.2. The out-of-plane orientation corresponds to the crystal direction pointing along the electron beam direction, and the matching results show that the sample is oriented along the (111) direction as expected. Using these estimated orientations, we are able to assign crystallographic indices to each of the Bragg peaks identified in the experiment by simulating a kinematic pattern at the estimated orientation and pairing up experimental peaks with fall within a given cutoff distance. We also retain the simulated pattern, which




Figure 6.1: Virtual images obtained from the 4D-STEM dataset for the direct beam (bright-field) and each of the first order (darkfield) reflections.

contains all the possibly excited Bragg vectors out to a scattering vector of  $1.5 \text{ \AA}^{-1}$  to use as the basis for Bloch wave calculations. The in-plane orientations show the crystal direction pointing in the downward direction of the scan. As shown in Figure 6.2 these orientations are “symmetry-reduced,” and show a consistent (022) type orientation. However, because the orientation matches from the kinematic matching are close to, but not exactly on, a zone axis orientation there is an instability in the heuristic used to compute this in-plane rotation of the lattice. Thus pixels with different color on the out-of-plane map may give a different crystallographic index to the downward direction. For the dynamical refinement procedure this indexing should be consistent in order for the reference frames of each simulation to match one another, and so we apply a different heuristic that ensures that the final Euler angle for all of the patterns is reduced by the 6-fold symmetry of the (111) zone axis. For polycrystalline scans care must be taken that within each cluster of pixels corresponding to a grain, the indexing heuristic is similarly stable and globally consistent.




Figure 6.2: Crystal orientation measured using the kinematic dictionary matching approach of Ophus et al. [124], showing (left) the crystal orientation pointing downward in the plane of the scan and (right) the orientation of the crystal facing the electron beam. The kinematic matching library shows the sample to be single crystalline, and reveals a slight misorientation boundary near the bottom of the scan.

## Dynamical Refinement

At this stage we have, for each scan position in the dataset, an approximate orientation (obtained from kinematic library matching), a list of Bragg beams which may potentially be excited in this orientation (derived from the approximate orientation), and a list of each experimentally measured diffraction disk containing its crystallographic index and integrated intensity. With this information in hand we can proceed to the Bloch wave dynamical refinement procedure. This refinement procedure as applied to a single representative diffraction pattern chosen from the dataset is shown in Figure 6.3. The cost function from Eq. 6.1 is three dimensional, so for visualization we project it along the thickness and mistilt dimensions by taking the minimum across the other dimension(s). At each iteration, we sample a region of tilt vectors around the best match from the previous step, with the first step seeded with the estimate from the kinematic library matching. We used three refinement steps which progressed from exploring  $3.5^\circ$  of misorientation in  $0.25^\circ$  steps to exploring  $0.1^\circ$  of misorientation in  $0.01^\circ$  steps.

The cost function is highly oscillatory with respect to thickness but is relatively cheap to compute for many thickness values, so we evaluate it across the full set of candidate thicknesses at every iteration. It is also clear that there are many local minima as a function of tilt, so it is necessary to sample the entire grid around each guess (i.e. gradient descent

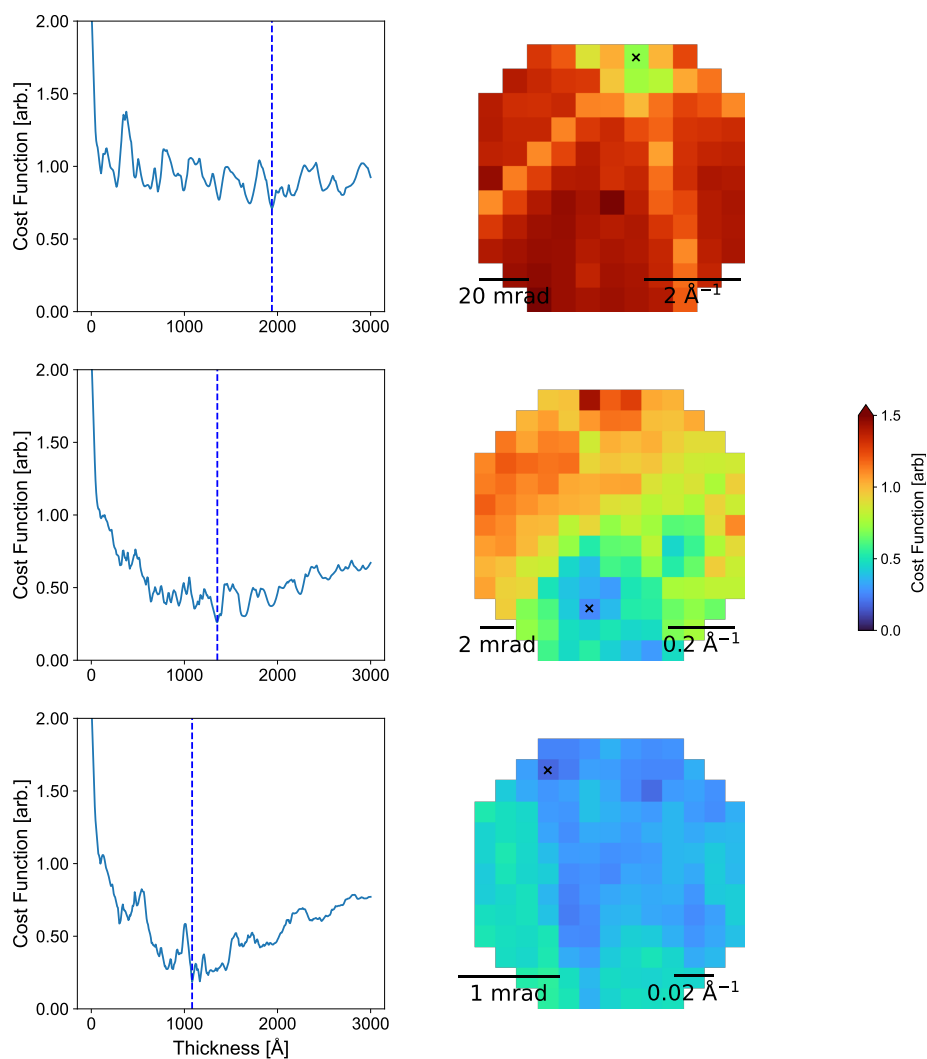


Figure 6.3: Projections of the 3D cost function onto thickness and orientation dimensions during progression of the iterative refinement procedure for a single diffraction pattern. Each row corresponds to one iteration, and the orientation search range is progressively narrowed around the best match from the previous search (indicated with a blue circle on the tilt projection). The first iteration searches a neighborhood centered on the orientation obtained from the kinematic library matching routine.

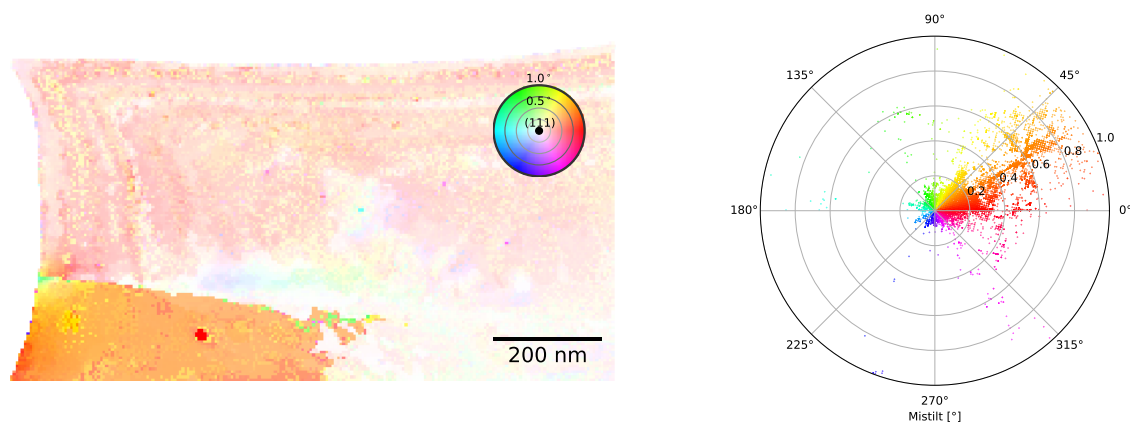


Figure 6.4: (left) Crystal orientations obtained from the dynamical refinement, shown as mistilt from the (111) crystal orientation (the maximum mistilt shown on the left panel is  $1^\circ$ ). (right) Histogram of misorientation of each pixel in the scan. The majority of the crystal is oriented within about  $0.2^\circ$  of the perfect zone axis orientation, and the defect area on the bottom left of the scan region is misoriented by approximately  $0.5^\circ$ .

or another iterative procedure are not appropriate as they are bound to get trapped in these local minima). The combination of thickness and tilt that minimize the cost function are marked on each iteration. As we progress, the range of tilts is made smaller and the sampling progressively finer in order to converge on the optimal match. We also note that the best matching thickness is highly dependent on the accuracy of the tilt value, as it changes from  $\approx 1950 \text{ \AA}$  at the first iteration to  $\approx 1100 \text{ \AA}$  at the final step (and finds a deeper minimum as the refinement progresses).

## Dynamical Refined Crystal Orientation

The refined orientation of the crystal is shown in Figure 6.4, where the saturation corresponds to the angle between the local orientation and an exact (111) zone axis orientation, and the color corresponds to the azimuthal angle of this misorientation. The histogram shows the distribution of misorientations from (111) on the same color scale as the orientation map. The majority of the crystal is found to be oriented within  $0.2^\circ$  of the zone axis orientation, and the defected region is misoriented by an average of about  $0.5^\circ$  from the rest of the scan region, with a total misorientation of about  $0.7^\circ$  from the zone axis. The majority of the scan region is tilted in the northeast direction, but some positions near the dislocation pileup at the center of the scan region tilt to the south or southwest, which may correspond to local lattice distortions produced by the defects. However, we cannot be certain of the source of these direction changes without high resolution imaging of the area or another independent measure of the local distortions. Instability in the indexing heuristic could cause some pixels

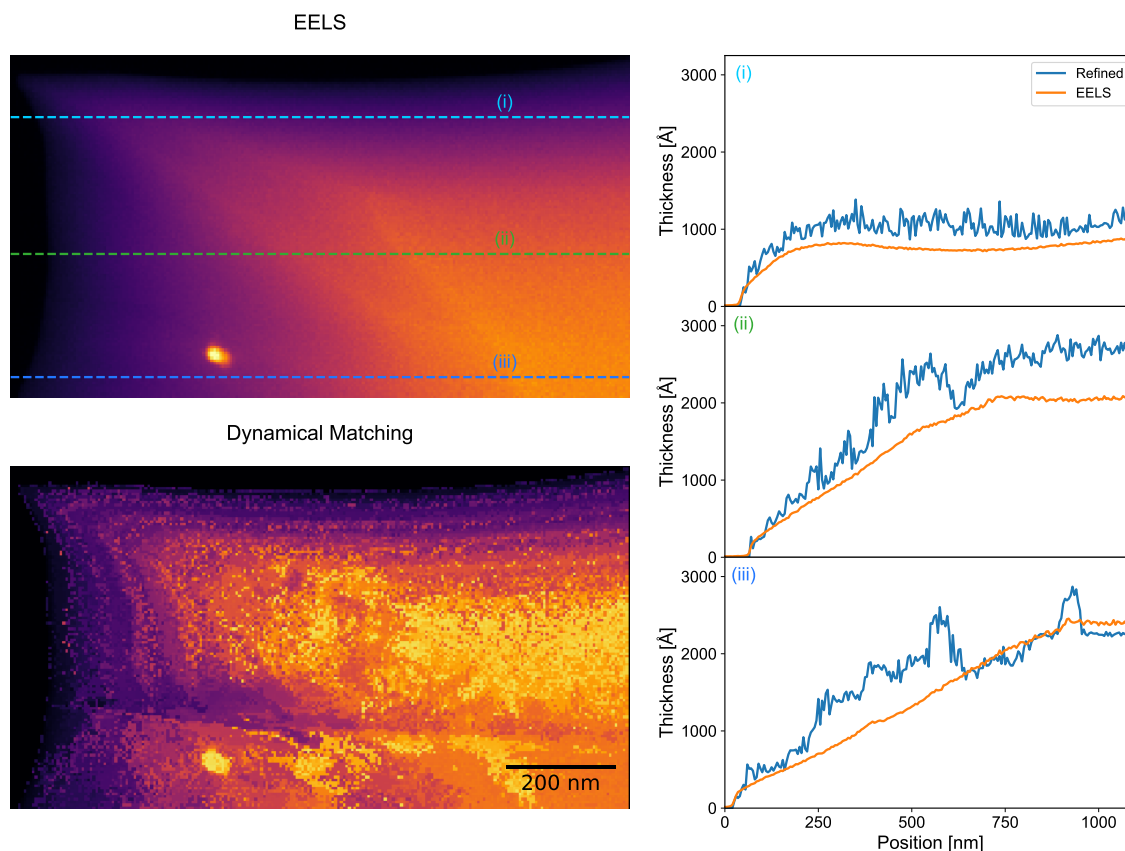


Figure 6.5: Maps of local sample thickness obtained from the dynamical refinement procedure and compared with low-loss EELS log ratio estimates of the thickness.

to be indexed differently and lead to a change in the apparent azimuthal direction which is a multiple of  $60^\circ$ .

## Crystal Thickness

The thickness matched at each scan position is shown in Figure 6.5, alongside with the sample thickness estimated from comparing the intensity of the first plasmon resonance in low-loss EELS. The thickness returned by the dynamical refinement procedure is in excellent qualitative agreement with the EELS result, showing the same wedge structure of the sample. However, line traces through three different locations on the crystal show that the dynamical

ical refinement procedure typically overestimates the sample thickness. There are several potential sources for this discrepancy: first, the model for the absorptive potential may be mismatched with the experiment. In particular, the Debye-Waller factor (for which we used an assumed room temperature value) may be incorrect, which would change how the modelled intensities decrease with thickness. The presence of an amorphous damage layer or a surface oxide on the sample may also be responsible for disagreement between the EELS and dynamical refinement. The presence of these surface layers would alter the effective  $Z$  of the sample, which would have a small effect on the EELS measurement, and miscalibration of the characteristic scattering length of the plasmon can contribute to the uncertainty in measured thickness [108]. Meanwhile, defect layers would contribute an amorphous background intensity to the diffraction patterns which reduce the total intensity in the Bragg disks, and this could lead to the refinement procedure overestimating the thickness (while still producing a highly accurate measurement of the orientation). Finally, the EELS data were acquired at the same zone axis orientation as the diffraction data in order to ensure good alignment between the two scans for pixelwise comparison, and channeling effects may alter the strength of the plasmon signal and thus the measured thickness [1].

## 6.5 Conclusion

In this work, we have developed a method for using dynamical diffraction simulations to improve the precision of orientation matching from nanobeam 4D-STEM experiments and to simultaneously obtain the thickness of the sample. This method relies on the extreme sensitivity of the diffracted intensities to small tilts and to the thickness of the sample, and uses an adaptive grid search to minimize the number of expensive diffraction calculations required to match the experimental data. We found that the dynamically refined orientations have precision of approximately  $0.2^\circ$ , and that the simultaneously measured thickness deviates from the thickness obtained from low-loss EELS by approximately 25%. This orientation precision is substantially better than what has been observed with kinematic library matching and allows for the identification of small-angle defects in crystals.

## Chapter 7

# Uncovering Polar Vortex Structures by Inversion of Multiple Scattering with a Stacked Bloch Wave Model

Nanobeam electron diffraction can probe local structural properties of complex crystalline materials including phase, orientation, tilt, strain, and polarization. Ideally, each diffraction pattern from a projected area of a few unit cells would produce clear a Bragg diffraction pattern, where the reciprocal lattice vectors can be measured from the spacing of the diffracted spots, and the spot intensities are equal to the square of the structure factor amplitudes. However, many samples are too thick for this simple interpretation of their diffraction patterns, as multiple scattering of the electron beam can produce a highly nonlinear relationship between the spot intensities and the underlying structure. Here, we develop a stacked Bloch wave method to model the diffracted intensities from thick samples with structure that varies along the electron beam. Our method reduces the large parameter space of electron scattering to just a few structural variables per probe position, making it fast enough to apply to very large fields of view. We apply our method to SrTiO<sub>3</sub>/PbTiO<sub>3</sub>/SrTiO<sub>3</sub> multilayer samples, and successfully disentangle specimen tilt from the mean polarization of the PbTiO<sub>3</sub> layers. We elucidate the structure of complex vortex topologies in the PbTiO<sub>3</sub> layers, demonstrating the promise of our method to extract material properties from thick samples.

---

The results presented in this chapter have been published as a preprint titled “Uncovering polar vortex structures by inversion of multiple scattering with a stacked Bloch wave model” as arXiv:2211.05842, by Steven E Zeltmann, Shang-Lin Hsu, Hamish G Brown, Sandhya Susarla, Ramamoorthy Ramesh, Andrew M Minor, and Colin Ophus [200]. The material is presented here with permission of the co-authors.



## 7.1 Introduction

Multiple scattering is often viewed as an unwanted and cumbersome artifact in electron microscopy as it is responsible for confounding effects such as contrast reversals in phase contrast images and electric field maps [191, 161], complex features inside nanobeam diffraction disks that hinder precise strain mapping [106, 34, 132, 199], and failure of many super-resolution imaging techniques [30].

In a naïve sense, an electric field built in to a thin sample causes a tilt of the electron wave as it propagates through the material. Depending on the optical setup and the length scale of the changes in the field this manifests as an intensity redistribution in a diffraction pattern when using a larger convergence angle [115], or a shift of the diffraction pattern when using a small convergence angle [31, 25]. In materials where the polarization is associated with a structural distortion, the diffracted intensities also change as a result of the change in the lattice electrostatic potential [38]. This is shown schematically in Fig 7.1a, which shows a sequence of diffraction patterns simulated at differing polarization for a thin sample of  $\text{PbTiO}_3$  (PTO).

When the sample thickness increases and multiple scattering occurs, however, the changes in the diffraction patterns become far more complex than this description. Deb et al. showed that when diffraction disks overlap, under the weak phase approximation there should be no contrast between opposing pairs of diffraction disks due to polarization; however, anomalous contrast between Friedel pairs arises when multiple scattering pathways are considered [42]. Mahr et al [107] showed that for an interface with an electric field due to a difference in mean inner potentials, dynamical scattering causes the measured electric field to oscillate wildly for most experimental setups, with beam precession providing the only partial remedy. Nguyen et al [118] also observed that when measuring chiral polarization domains via the diffraction intensity changes associated with a structural distortion, the chiral directions flip as a function of thickness. This case is shown schematically in Fig 7.1b for sample of PTO that is thick enough to cause multiple scattering of the electron probe.

Various approaches for reconstructing sample properties or structure under conditions of multiple scattering have been developed that utilize the  $\mathcal{S}$ -matrix description of the scattering process [22, 44, 133, 47, 21]. In these approaches, the electron scattering process is encapsulated in the “scattering matrix,” an object which contains the information about the material structure and thickness-propagation effects, such that the  $\mathcal{S}$ -matrix multiplies with a vector representing the incident electron wave to yield a vector representing the scattered wave [169]. As we will discuss in the theory section below, there is useful correspondence between the entries of the  $\mathcal{S}$ -matrix and the properties of the sample, such as its thickness, polarization, or tilt. This correspondence can be used in both the forward direction to simulate diffraction intensities given known structural parameters, as well as in the reverse direction to recover structural parameters from diffraction intensities. Much of the literature on this method is concerned with atomic-resolution reconstruction of the sample potential. Donatelli and Spence demonstrated a method for recovering the sample potential at high resolution from a tilt series of diffraction patterns by iterative inversion of the  $\mathcal{S}$ -matrix and

without prior knowledge of the sample thickness [44]. Brown and coworkers have developed methods for recovery of the  $\mathcal{S}$ -matrix, originally from focal series 4D-STEM data [22, 21] and later extended to a single defocus [47]. Their former method provided some 3-dimensional information about the sample, which they recovered from the  $\mathcal{S}$ -matrix by an optical sectioning approach. Alternatively, when the approximate structure of the material is known, and variation along the beam direction can be ascribed to a small number of parameters, the PRIMES (parameter retrieval and inversion from multiple electron scattering) family of methods can be used to obtain property variation along the beam direction [137, 136, 135, 139]. These methods all utilize a stacked  $\mathcal{S}$ -matrix model to represent parameter variation along the beam direction, which are refined against test CBED patterns using various numerical optimization schemes.

In this work, we analyze a complex sample consisting of multiple distinct layers through the thickness, and use a model of the electron multiple scattering to extract information about the material. The sample, consisting of 16 unit cells of SrTiO<sub>3</sub> (STO), 16 unit cells of PTO, and 16 units cells of STO, is shown in Fig 7.1c. The vortex structures in this material have been previously studied by plan-view and cross section transmission electron microscopy [190, 38] and by x-ray coherent diffractive imaging [163]. These vortex structures offer the promise of creating new electronic states of matter, with structured domains with nanometer-scale domain sizes. Similar to PRIMES, we do not aim to recover the full atomic structure of the material. Instead, we model the scattering matrix using parameters that represent small perturbations from an *a priori* known average structure of the material, and refine the model to match measured diffraction intensities. Parameterizing the scattering model allows us to choose a small and physically meaningful set of variables to refine against, and by computing gradients of the model semi-analytically we are able to dramatically accelerate the discovery of the model parameters. Our method is sufficiently fast to allow us to perform the parameter matching for each probe position in a four-dimensional scanning transmission electron microscopy (4D-STEM) scan, where a shallow converged electron probe is rastered across the sample surface.

## 7.2 Theory

To compute the dynamical diffraction intensities, we utilize the Bloch wave method, which is fully described in DeGraef [40]. In this method, the electron wave is written as a combination of Bloch states, and thus the Schrödinger equation for the fast electron wave is cast as an eigenvalue/eigenvector decomposition

$$\bar{\mathcal{A}}\mathcal{C} = 2k_n\gamma\mathcal{C}, \quad (7.1)$$

where the “structure” matrix  $\bar{\mathcal{A}}$  is determined by the crystal structure and orientation of the sample,  $\mathcal{C}$  is a matrix whose column vectors contain the Bloch wave coefficients,  $k_n$  is the normal component of the incident wavevector, and  $\gamma$  relative normal component of each




Figure 7.1: Nanobeam electron diffraction signals from lattices with varying in-plane polarization, indicated by the arrows. Diffraction pattern simulations of (a) thin  $\text{PbTiO}_3$ , (b) thick  $\text{PbTiO}_3$ , and (c) multilayer with 16:16:16 unit cells of  $\text{SrTiO}_3$ : $\text{PbTiO}_3$ : $\text{SrTiO}_3$ . Left to right, the in-plane  $\text{PbTiO}_3$  polarization varies smoothly from zero, full left-facing, zero full right-facing, and zero polarization.

of the Bloch waves. The entries of the structure matrix are given by

$$\bar{\mathcal{A}} = \begin{bmatrix} 0 & U_{-\mathbf{g}} & \cdots & U_{-\mathbf{h}} \\ U_{\mathbf{g}} & 2k_0 s_{\mathbf{g}} & \cdots & U_{\mathbf{g}-\mathbf{h}} \\ \vdots & \vdots & \ddots & \vdots \\ U_{\mathbf{h}} & U_{\mathbf{h}-\mathbf{g}} & \cdots & 2k_0 s_{\mathbf{h}} \end{bmatrix} \quad (7.2)$$

where  $U_{\mathbf{g}-\mathbf{h}}$  is the Fourier component of the sample electrostatic potential corresponding to the scattering vector  $\mathbf{g} - \mathbf{h}$  and  $s_{\mathbf{g}}$  is the excitation error for the beam  $\mathbf{g}$ . The excitation error is given by

$$s_{\mathbf{g}} = \frac{-\mathbf{g} \cdot (2\mathbf{k} + \mathbf{g})}{2|\mathbf{k} + \mathbf{g}| \cos \alpha} \quad (7.3)$$

where  $\mathbf{k}$  is the wavevector of the incident electron beam, and  $\alpha$  is the angle between the sample normal and the incident beam direction. To include absorption due to thermal diffuse or inelastic scattering, an imaginary component of the sample potential  $U'_{\mathbf{g}}$  is included in the matrix  $\bar{\mathcal{A}}$  by adding  $iU'_{\mathbf{g}}$  to the diagonal and  $iU'_{\mathbf{g}-\mathbf{h}}$  to the off-diagonal. By computing the eigenvalue/eigenvector decomposition of this matrix, we obtain the Bloch wave coefficients  $C_{\mathbf{g}}^{(j)}$  and the normal components  $\gamma^{(j)}$ , which are used to obtain the scattered wave amplitudes for a given crystal thickness. The electron wave  $\boldsymbol{\psi}$  at a depth  $z$  in the crystal is

$$\boldsymbol{\psi}(z) = \mathcal{C}\mathcal{E}(z)\mathcal{C}^{-1}\boldsymbol{\psi}_0 = \mathcal{S}(z)\boldsymbol{\psi}_0 \quad (7.4)$$

where  $\mathcal{C}$  is the matrix containing the eigenvectors and  $\mathcal{E}(z) = e^{2\pi i \gamma^{(j)} z} \delta_{ij}$  is a diagonal matrix which depends on the thickness and the Bloch wave normal components.  $\boldsymbol{\psi}_0$  is a vector containing the Fourier coefficients of the incident electron wave—in the case of plane wave illumination, it is a vector with the value of 1 at the index corresponding to the incident beam direction and zero elsewhere.

This transformation can be compactly represented by the scattering matrix  $\mathcal{S}$  which maps the vector representing the incident electron wave to the scattered wave at depth  $z$ . Writing the equation in this form, we see that the  $\mathcal{S}$ -matrix is the exponential of the structure matrix  $\mathcal{A}$  multiplied by  $2\pi iz$ .

## Stacked $\mathcal{S}$ -matrix model

In the description of the Bloch wave model above, we obtained a single  $\mathcal{S}$ -matrix which transformed a plane wave  $\boldsymbol{\psi}_0$  incident onto a crystal of some thickness  $z$  into the scattered wave  $\boldsymbol{\psi}$ . If the electron wave were to immediately enter another crystal, we can model this further scattering by simply using the  $\mathcal{S}$ -matrix of that layer to transform the previously scattered wave into the final scattered wave [165, 137]. This operation is equivalent to multiplying the complex  $\mathcal{S}$ -matrices together and applying it once to the original wave. From a numerical standpoint, it is more convenient to successively apply the  $\mathcal{S}$ -matrices to  $\boldsymbol{\psi}_0$  rather than multiply the  $\mathcal{S}$ -matrices first, as matrix-vector multiplications are cheaper than the matrix-matrix products needed to construct the total  $\mathcal{S}$ -matrix.

In this work, we will model the trilayer STO:PTO:STO sample using the product of 3  $\mathcal{S}$ -matrices:

$$\mathcal{S} = \mathcal{S}_{\text{STO}} \mathcal{S}_{\text{PTO}} \mathcal{S}_{\text{STO}} \quad (7.5)$$

For the trilayer sample we consider here, the epitaxial relationship between the layers considerably simplifies the use of the stacked  $\mathcal{S}$ -matrix approach, as the layers share the same lattice—this ensures all of the  $\mathcal{S}$ -matrices share the same set of beams. In a situation where the layers do not share the same lattice, one must ensure that all of the component  $\mathcal{S}$ -matrices include all of the beams from all of the layers. Further, in such case it is possible for a beam scattered by the top layer to be evanescent in a lower layer, with exponentially decaying intensity [165].

## Derivatives of the stacked $\mathcal{S}$ -matrix model

Using the Bloch wave method, we have obtained a scattering matrix which is the exponential of the structure matrix, whose entries are readily obtained from the crystal properties. The full scattering matrix describing the multilayer is represented as a product of three such  $\mathcal{S}$ -matrices. In order to match the parameters of the model to our experimental data, we will use an optimization procedure to minimize the error between the model and the experiment. Unfortunately, each evaluation of the scattered wave using this model requires (for each layer) building a new  $\mathcal{A}$ -matrix and diagonalizing it, which is computationally expensive and makes numerical optimization inefficient. Therefore we aim to obtain the derivatives of the scattering matrix with respect to the entries of the structure matrix, so that we can compute the gradient of the error without the large number of function evaluations necessitated by finite differences. Najfeld and Havel [117] define the directional derivative  $D_{\mathbf{V}}(t, \mathbf{A})$  of a matrix exponential  $e^{t\mathbf{A}}$  in the direction  $\mathbf{V}$  as

$$D_{\mathbf{V}}(t, \mathbf{A}) \equiv \lim_{h \rightarrow 0} \frac{1}{h} (e^{t(\mathbf{A}+h\mathbf{V})} - e^{t\mathbf{A}}) \quad (7.6)$$

We will also use the notation  $\frac{d\mathcal{S}}{d\theta}$  to refer to the derivative of a scattering matrix in the direction  $\mathbf{V}_{\theta}$ , where  $\theta$  is one of the structural perturbation parameters. For a matrix which has been spectrally decomposed as  $\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^{-1}$ , (i.e. its eigendecomposition has been computed, where the columns of  $\mathbf{U}$  contain the eigenvectors and the diagonal of  $\mathbf{\Lambda}$  contains the eigenvalues  $\lambda_i$ ) the directional derivative of its exponential can be computed as [117]

$$D_{\mathbf{V}}(t, \mathbf{A}) = \mathbf{U} ((\mathbf{U}^{-1}\mathbf{V}\mathbf{U}) \odot \mathbf{\Phi}(t)) \mathbf{U}^{-1}. \quad (7.7)$$

where  $\odot$  is the Hadamard, or elementwise, product. The entries of  $\mathbf{\Phi}(t)$  depend on the eigenvalues of  $\mathbf{A}$ :

$$\Phi_{ij}(t) = \begin{cases} (e^{t\lambda_i} - e^{t\lambda_j})/(\lambda_i - \lambda_j) & \text{if } \lambda_i \neq \lambda_j \\ te^{t\lambda_i} & \text{if } \lambda_i = \lambda_j \end{cases} \quad (7.8)$$

To compute the derivative of the total  $\mathcal{S}$ -matrix comprised of an ordered collection of  $N$  separate scattering matrices indexed with the superscript  $(j)$  (i.e.  $\mathcal{S} = \prod_j \mathcal{S}^{(j)}$ ), with respect

to a parameter  $\theta$  we use the product rule

$$\frac{d\mathcal{S}}{d\theta} = \sum_{i=0}^N \left[ \prod_{j=0}^{j=i-1} \mathcal{S}^{(j)} \cdot \frac{d\mathcal{S}^{(i)}}{d\theta} \cdot \prod_{k=i+1}^{k=N} \mathcal{S}^{(k)} \right]. \quad (7.9)$$

Due to our choice of parameters for the model, many of the derivatives  $\frac{d\mathcal{S}^{(j)}}{d\theta}$  will be zero. For the terms where there are nonzero derivatives, the derivative of the total  $\mathcal{S}$ -matrix comprises the scattering up to the layer affected by the parameter  $\theta$ , the change in scattering within that layer, and the further scattering of the wave by the following layers in the heterostructure.

In our model of the trilayer sample, there are two relevant classes of structural perturbation that we will attempt to recover: crystal tilt and structural distortion due to polarization. In the following section, we will derive the direction matrices  $\mathbf{V}_\theta$  for these types of perturbation.

### Crystal tilt

Because tilt of the crystal is included solely in the diagonal elements of the structure matrix via the excitation errors  $s_g$ , the derivative direction is

$$(\mathbf{V}_{\text{tilt}})_{g,h} = 2k_0 \frac{ds_g}{dk_0} \delta_{g-h} \quad (7.10)$$

The derivative of  $s_g$  with respect to the transverse ( $x$  and  $y$ ) components of the incident wavevector are approximately

$$\frac{ds_g}{dk_{\{x,y\}}} = -\frac{g_{\{x,y\}}}{|\mathbf{g} + \mathbf{k}_0|} + \frac{(g_{\{x,y\}} + k_{\{x,y\}})(\mathbf{g} \cdot (2\mathbf{k} + \mathbf{g}))}{2|\mathbf{g} + \mathbf{k}|^3}. \quad (7.11)$$

We have neglected the  $\cos \alpha$  term in the denominator of  $s_g$  in this derivation to greatly simplify the expression at the cost of some small error in the magnitude of  $\frac{ds}{dk}$ .

### Polarization

While it is possible under certain conditions to recover the locations of each atom in the unit cell by recovering the full  $\mathcal{S}$ -matrix from diffraction data, [44, 22] here we parameterize the model in terms of the polarization directly, and displace the Ti and O sites by interpolating between their positions in the canonical non-polar and polar structures. For each atom in the unit cell, we define displacement vectors  $\delta \mathbf{r}_a^{(j)}$  and  $\delta \mathbf{r}_b^{(j)}$  which takes the atom from its site  $\mathbf{r}^{(j)}$  in the non-polar structure to its site in the canonical polar structure, for polarizations in the  $a$  and  $b$  directions respectively. For intermediate or mixed-direction polarizations, we linearly scale the displacement vectors by the relative polarization,  $\rho_a$  and  $\rho_b$  for  $a$  and  $b$  polarizations respectively. The Fourier coefficients of the crystal potential are thus written as

$$U_{\mathbf{g}} = \frac{1}{\Omega} \sum_j f_e^{(j)} e^{2\pi i((\mathbf{r}^{(j)} + \rho_a \delta \mathbf{r}_a^{(j)} + \rho_b \delta \mathbf{r}_b^{(j)}) \cdot \mathbf{g})} \quad (7.12)$$

where  $\Omega$  is the unit cell volume, and the atomic form factors  $f_e^{(j)}$  are computed using the absorptive Weickenmeier-Kohl parameterization for isolated neutral atoms [183]. The derivative of the structure factor with respect to the relative  $a$ -axis polarization parameter,  $\rho_a$ , is then

$$\frac{dU_{\mathbf{g}}}{d\rho_a} = \frac{1}{\Omega} \sum_j 2\pi i f_e^{(j)}(\mathbf{g} \cdot \delta\mathbf{r}_a) e^{2\pi i((\mathbf{r}^{(j)} + \rho_a \delta\mathbf{r}_a^{(j)} + \rho_b \delta\mathbf{r}_b^{(j)}) \cdot \mathbf{g})} \quad (7.13)$$

and a similar expression arises for  $\frac{dU_{\mathbf{g}}}{d\rho_b}$ . The derivative direction matrix for polarization is simply filled with the derivatives of the Fourier coefficients, e.g.

$$\left( \mathbf{V}_{\rho_{\{a,b\}}} \right)_{\mathbf{g}, \mathbf{h}} = \frac{dU_{\mathbf{g}-\mathbf{h}}}{d\rho_{\{a,b\}}} \quad (7.14)$$

Note that we do not take into account expansion of the unit cell, so the mean inner potential term  $U_{\mathbf{0}}$  does not change, and thus polarization only affects the off-diagonal elements of the structure matrix.

## Numerical Optimization

To fit the stacked  $\mathcal{S}$ -matrix model to the experimental measurements, we implemented a version of the alternating direction method of multipliers [18]. At each iteration of the optimization algorithm, we first update the model parameters at each probe position using our previously derived gradients and taking a step along the direction of steepest descent. We then perform regularization of the fitted model parameters to ensure convergence to a physically sensible solution and enforce smoothness.

### Gradient Descent

The loss function  $\mathcal{L}$  is the sum squared difference between the simulated diffraction intensities from the model and the experimental intensities for each Bragg beam recorded

$$\mathcal{L} = \sum_{\mathbf{g}} (I_{\text{exp}}(\mathbf{g}) - \mu |\mathcal{S}\psi_0(\mathbf{g})|^2 - \nu)^2 \quad (7.15)$$

The modeled intensities have both an additive intensity offset  $\nu$  and multiplicative scaling  $\mu$ , which we found necessary in order to compensate for background noise and intensity variation in the experimental data. Note that the  $\mathcal{S}$ -matrix calculations are performed using a different, and larger, set of Bragg beams, in order to include scattering into the higher order beams (not recorded on the detector) in the forward model. Only the Bragg beams present in the experiment contribute to the loss function. The derivatives of the loss function

with respect to the intensity scale parameters are given as

$$\frac{\partial \mathcal{L}}{\partial \mu} = \sum_{\mathbf{g}} -2|\mathcal{S}\psi_0|^2 (I_{\text{exp}} - \mu|\mathcal{S}\psi_0|^2 - \nu) \quad (7.16)$$

$$\frac{\partial \mathcal{L}}{\partial \nu} = \sum_{\mathbf{g}} -2 (I_{\text{exp}} - \mu|\mathcal{S}\psi_0|^2 - \nu) \quad (7.17)$$

(where we have dropped the dependence on  $\mathbf{g}$  from the notation for compactness). The derivatives with respect to the structure perturbations involve the derivatives of the  $\mathcal{S}$ -matrix, and so are much more complicated expressions. For a generic parameter  $\theta$  that enters into the  $\mathcal{S}$ -matrices, the derivative of the loss function is

$$\frac{\partial \mathcal{L}}{\partial \theta} = -4\mu \Re \left[ \psi^* \frac{d\mathcal{S}}{d\theta} \psi_0 \right] \cdot [I_{\text{exp}} - \mu|\mathcal{S}\psi_0|^2 - \nu] \quad (7.18)$$

The gradients with respect to the tilt and polarization variables are obtained using the derivative directions in Eqs. 7.11 and 7.13, the product rule in Eq. 7.9, and the  $\mathcal{S}$ -matrix derivative method in Eq. 7.7. At each step of the optimization procedure, we update the parameters by taking a step along the negative gradient direction of this loss function.

## Regularization

In order to obtain physically sensible solutions to the optimization problem we found it necessary to apply several regularizers. Before performing the optimization, we de-noise the integrated disk intensities using principal component analysis, retaining the first 16 components.

At each iteration step, we apply further regularization. First, the estimated parameters are smoothed across the real-space dimensions of the scan using a Gaussian kernel. Since the intensity scale and offset and the tilts are expected to vary slowly across the field of view, we used a kernel size of 50 nm for the intensity parameters and 25 nm for the tilts. The polarization is expected to vary more rapidly, so we used a kernel size of 2 nm. Note that the experiment used a probe step size of 1 nm, giving equivalent values for the size of each kernel in terms of the number of probe positions.

In addition, we also clip the fitted parameters to be within set bounds, so that outliers do not excessively propagate error to their neighbors via the smoothing kernel. We note that we do not apply any explicit high-pass filtering to the fitted polarization values (on the contrary, they are Gaussian filtered, albeit with a very small kernel size). However, the strong smoothing regularization applied to the tilt and intensity signals can have the side effect of forcing all of the high frequency variation into the polarization channel.

## Atomic Form Factors

Wu et al [188] showed that 4D-STEM may be sensitive to the charge transfer between sites in ionic materials, using strontium titanate as a model system, which would imply that the



independent atom model for the crystal potential may not be valid for our computations. To test this possibility we used the GPAW density functional theory package [113] to simulate the charge transfer between species, and then used abTEM [102] to perform diffraction simulations for the simulated charge densities that match our experimental conditions. From these simulations we observed that the maximum deviation in the diffracted intensities between the DFT and IAM potentials was approximately 0.1% of the probe intensity, validating the use of the IAM model for our computations.

## 7.3 Methods

### Heterostructure Growth

We synthesized a trilayer structure consisting of 16 unit cells of SrTiO<sub>3</sub>, 16 unit cells of PbTiO<sub>3</sub>, and 16 unit cells of SrTiO<sub>3</sub>, on top of a SrRuO<sub>3</sub> buffer layer on a single crystal DyScO<sub>3</sub> substrate. The layers were grown at 610°C in a 100 mTorr oxygen atmosphere, using reflection high-energy electron diffraction (RHEED)-assisted pulsed laser deposition (PLD) with a KrF laser. The trilayer structure was removed from the substrate by chemical processing.

### 4D-STEM Experiments

We performed 4D-STEM measurements on the TEAM I microscope, an aberration-corrected Thermo Fisher Scientific Titan operated at 300 kV with a probe current of 100 pA. We used a STEM probe semiangle of 2 mrad, and a STEM probe step size of 1 nm. We recorded diffraction patterns using a Gatan K3 direct electron detector located beyond a Gatan Continuum energy filter. We operated the K3 detector in electron counting mode using a binning of 4x4 pixels, a camera length of 1.05 m, and an exposure time of 47 ms. We analyzed the 4D-STEM experiments using custom Python and Matlab code. The diffraction pattern simulations and Bloch wave calculations and optimizations have been implemented as part of the py4DSTEM analysis toolkit [158, 124, 198].

## 7.4 Results & Discussion

### Tilt/Polarization Confounding

Figure 7.2 shows how local mistilts of the sample from the perfect zone axis orientation can confound the measurement of local polarization when using a conventional metric based on Friedel pair asymmetry. The top row of Fig. 7.2a shows a sequence of simulated diffraction patterns for the STO/PTO trilayer sample at varying polarization in the  $x$ -direction. The inset numbers indicate the polarization signal measured using the anomalous contrast of the (200) Friedel pair [42], computed as  $(I_{200} - I_{\bar{2}00})/(I_{200} + I_{\bar{2}00})$  and normalized to the

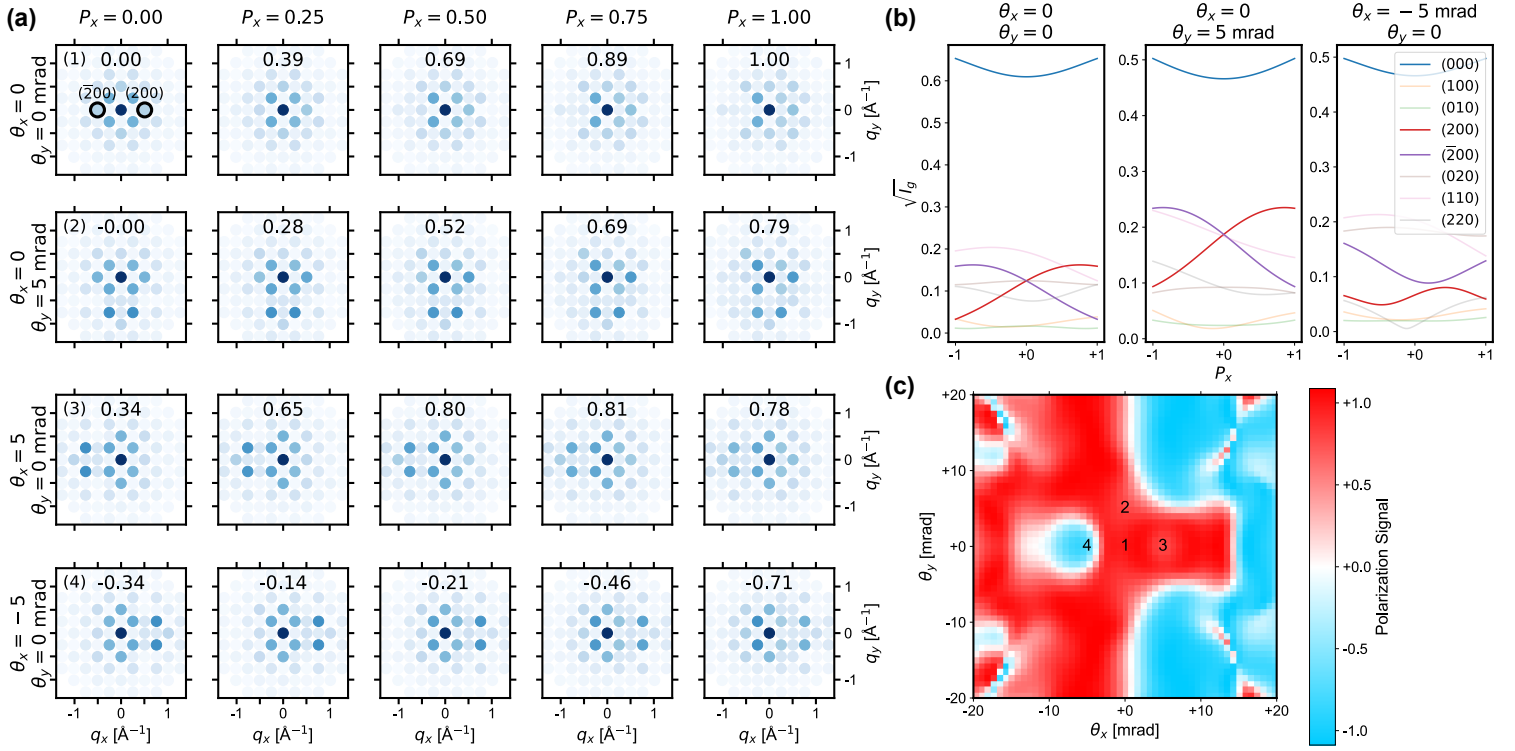


Figure 7.2: Confounding of tilt and polarization signals for a STO:PTO:STO trilayer. (a) Sequences of diffraction patterns simulated at relative  $x$ -direction polarization varying from 0–1 for various mistilts from a perfect [001] orientation. The two diffraction disks conventionally used to measure polarization, (200) and  $(\bar{2}00)$  are highlighted, and the inset text indicates the relative polarization measured as  $\frac{I_{200} - I_{\bar{2}00}}{I_{200} + I_{\bar{2}00}}$ . (b) Line traces of selected diffraction intensities for different mistilts, with the (200) and  $(\bar{2}00)$  reflections highlighted. (c) Map of apparent polarization signal for a trilayer with  $P_x = 1$  as measured from the asymmetry of the (200) and  $(\bar{2}00)$  disks for different mistilts. Contrast reversals in the polarization signal occur with at little as 5 mrad mistilt. The overlaid numbers indicate the tilt values corresponding to the rows of (a).

$\theta_x = \theta_y = 0$ ,  $P_x = 1$  value. The signal is monotonic with increasing polarization and approximately linear, indicating that in the ideal case the symmetry breaking of this pair of diffraction disks is a good measurement of the local polarization. As shown in the left panel of Fig. 7.2b, the intensities of these disks branch as a function of polarization. When tilting the incident beam towards the positive  $y$ -axis, as in the second row of Fig. 7.2a, the symmetry of (200) pair of diffraction disks is not broken, but the different excitation of these beams (shown in the center panel of Fig. 7.2b) causes the signal to be suppressed by approximately 20%. However, when the beam is tilted by 5 mrad ( $\approx 0.3^\circ$ ) towards the  $x$ -axis (as shown in the right panel of Fig. 7.2b), the effect of the tilting is to break the symmetry of the (200) disks. This slight mistilt in either direction along the  $x$ -axis completely destroys the polarization measurement. In the third row of Fig. 7.2a, where the beam is tilted towards the positive  $x$ -axis, there is an apparent nonzero polarization even when the material is not polarized. In the fourth row of Fig. 7.2a (and the right panel of Fig. 7.2b) a tilt towards the negative  $x$ -axis causes an inversion of the polarization signal when  $P_x > 0$ , non-monotonic behavior when  $P_x < 0$ , as well as an apparent nonzero polarization even when the material is not polarized. These effects are plotted as a function of  $x$  and  $y$  tilts in Fig. 7.2c.

## Gradients of the Diffracted Intensities

Computations of the derivatives of the diffraction disk intensities with respect to  $x$ -direction polarization and tilt are shown in Figure 7.3, evaluated over a range of tilts and polarizations. The inset numbers indicate the difference between the derivative of the (200) Friedel pair, which demonstrates the sensitivity of the signal derived from the anomalous contrast of those reflections to the chosen parameter. The difference in overall magnitude between the polarization and tilt derivatives is affected by the choice of units for the parameters; in the figure they have been scaled to be visually uniform, and in the optimization procedure the problem is rescaled to promote uniform convergence along all the parameter directions. In the on-zone, unpolarized case ( $P = 0$ ,  $\theta = 0$ ) in the top left, the (200) anomalous contrast signal will not distinguish between polarization of the crystal and tilt, as both cause the same anomalous contrast. However, other reflections respond in different ways to polarization and tilt. Thus, when considering all of the diffracted beams the gradient directions for tilt and polarization are approximately  $62^\circ$  separated. Since they are not orthogonal, an iterative optimization will be needed in order to solve the polarization and tilt. In the case of nonzero polarization and tilt, the gradients tend to become more parallel. In particular the gradients with respect to  $y$ -direction polarization and tilt, which are fully orthogonal to the  $x$ -direction parameter gradients at  $P = 0$ ,  $\theta = 0$ , will become partially coupled to the other direction when the crystal is tilted or polarized.

## Experimental Results

The fitted polarization and tilts of the STO:PTO:STO multilayer sample are plotted in Fig. 7.4 over the full field of view. Immediately, we can see several domains in the polarization

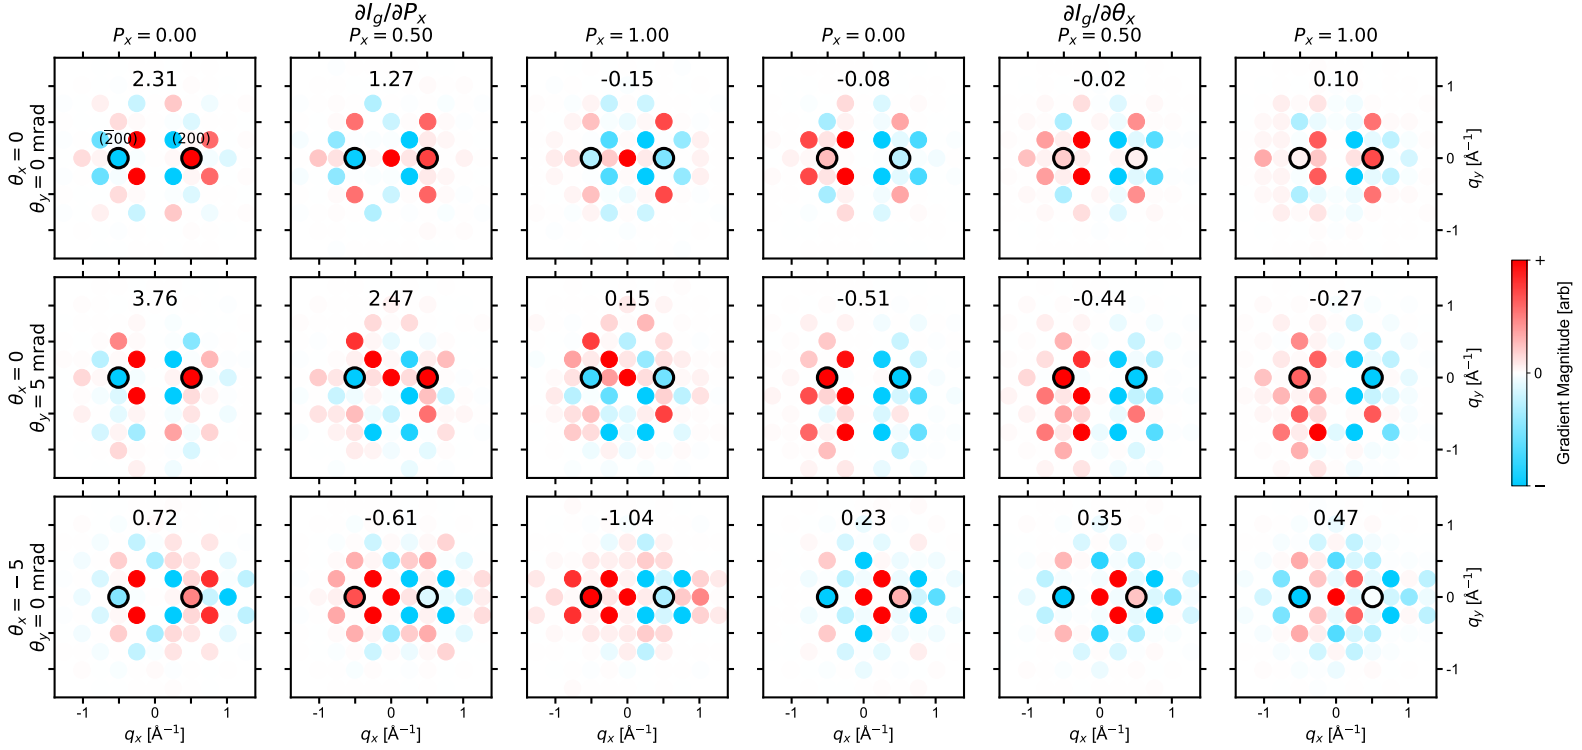


Figure 7.3: Derivatives of the diffracted intensities with respect to  $x$ -polarization  $P_x$  and tilt about the  $x$ -axis  $\theta_x$ . The inset numbers show the value of  $\frac{\partial I_{200}}{\partial \square} - \frac{\partial I_{500}}{\partial \square}$  in arbitrary units, where  $\square$  represents polarization for the left 3 columns, and tilt for the right 3 columns.

maps in Figs. 7.4a and b where the regular periodic structures show the tubular vortex structures. The “axial” polarization represents the PTO polarization along the directions parallel to the vortex cores, while the “lateral” polarization is perpendicular. Inside each vortex, the projected polarization is relatively constant in a given domain. Various domain boundaries are also visible, where the polarization abruptly changes sign, in either the axial or lateral directions, or both. These domain structures and domain walls are in good agreement with previous observations of STO/PTO multilayer samples [190, 38, 170]. The tilt maps shown in Figs. 7.4c and d show significant rotation from the ideal zone axis, especially in the  $y$ -axis direction. These maps demonstrate the need to include tilt in the modeling of the diffraction signals. The smoothness of the estimated tilt is due in part to the strong regularization applied during the reconstruction.

The polarization maps contain many complex domain and domain wall structures. We expect that the vortex cores will have alternating polarization signs in the axial direction. This alternating structure is visible in all domains in Fig. 7.4a, though interestingly we also observe a negative offset from zero mean axial polarization in the largest domain spanning the grain in the bottom half of the map. The grains at the top and bottom edges also

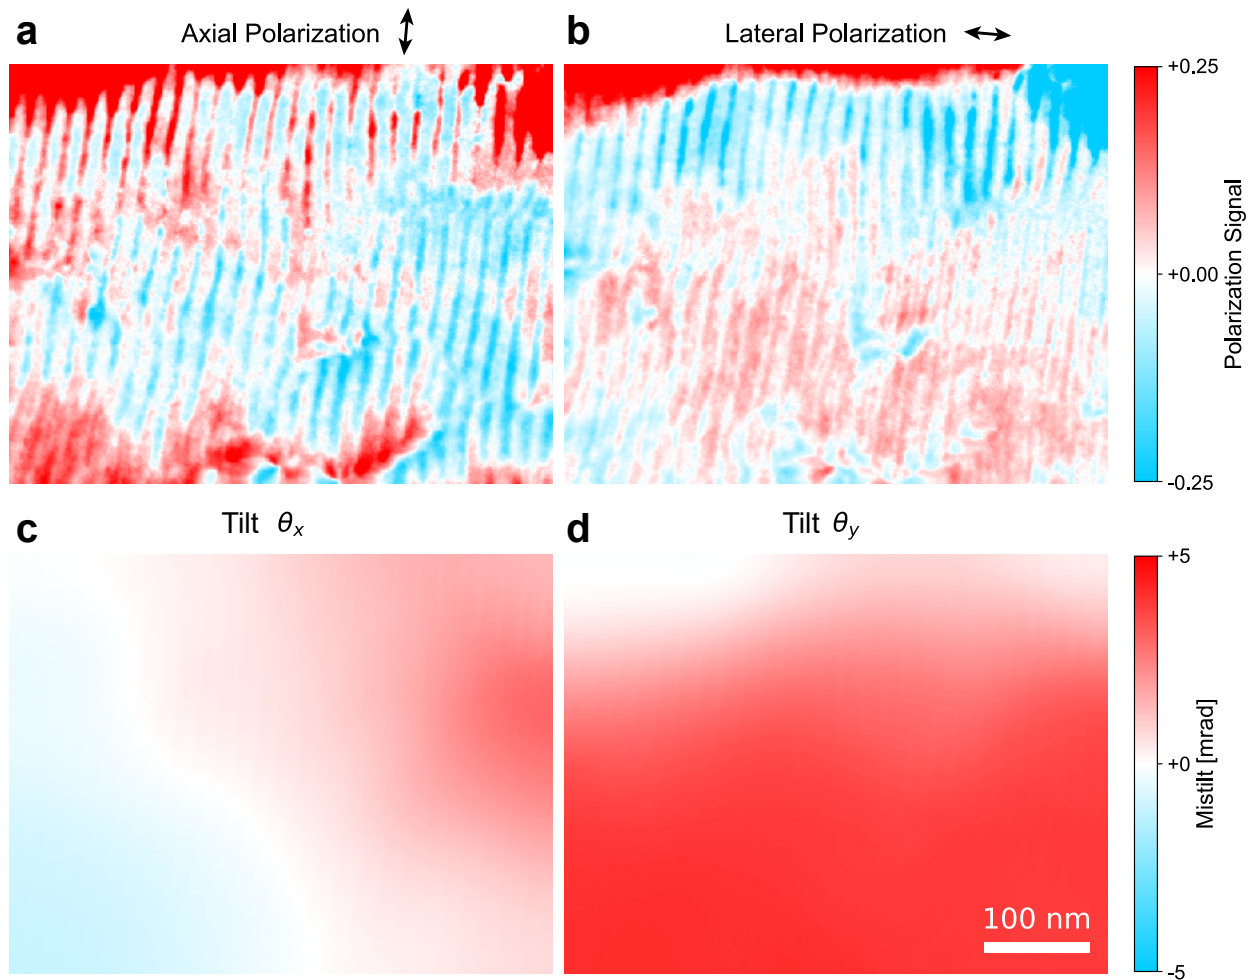


Figure 7.4: (a),(b) Polarization and (c),(d) tilt of the STO:PTO:STO sample recovered from the optimization procedure. Approximate polarization directions are labeled above.

show a significant positive offset from zero mean. These observations suggest that there may be a net axial polarization in many of the domains, which can't be directly observed from qualitative estimates of the polarization which have been high-pass filtered [170].

By contrast, in the lateral direction we expect oscillations in the polarization, but that each domain will have a larger net positive or negative polarization. This is because phase field predictions of the polarization structure of the PTO vortex phase predict that that every other vortex will be displaced towards one of the STO/PTO interfaces, while the remaining vortices will be displaced towards the other PTO/STO interface [170]. This in turn causes a net polarization flow to one of the lateral directions. Susarla et al. provide more phase field modeling and predicted vortex domain structures [170]. These net polarization features are indeed observed in Fig. 7.4b. Domains in the top third of the map and bottom left show significant polarization towards the negative direction, while the domains in the bottom two thirds show significant polarization in the positive direction. Various small domains are interspersed into the larger domains, but each shows a non-zero mean polarization. Overall, these observations provide a significant step forward in accurate modeling of the intensity of Bragg peaks when the beam undergoes significant multiple scattering and the sample has a large mistilt from the ideal zone axis.

## Future Directions

In this work we have shown a stacked  $\mathcal{S}$ -matrix model for scattering through a trilayer heterostructure parameterized with a single homogeneous polarization direction within the PTO layer. However, the samples we investigated are known to have polarization that varies in a complex manner through the thickness. A natural extension of Equation 7.5 to account for this is to model the PTO layer with a product of  $N_L$   $\mathcal{S}$ -matrices with distinct polarization:

$$\mathcal{S}_{\text{PTO}} = \prod_i^{N_L} \widehat{\mathcal{S}}_{\text{PTO}}(\rho_{a,i}, \rho_{b,i}; t = T/N_L) \quad (7.19)$$

where  $\rho_{a,i}$  and  $\rho_{b,i}$  are the relative polarization of the  $i$ -th layer in the  $a$  and  $b$  directions, and  $T$  is the total thickness. This modification to the model allows us to more accurately reproduce the physics of the scattering, at the cost of adding substantially more optimization variables. This added complexity can be mitigated somewhat by applying constraints to the variation in polarization with thickness. For example, if the polarization is constrained to vary linearly then the polarization variables at the  $i$ -th layer are expressed in terms of just two optimization variables,  $\rho_{\text{top}}$  and  $\rho_{\text{bottom}}$  regardless of the number of layers modeled

$$\rho_i = \rho_{\text{top}} + \frac{i}{N_L - 1}(\rho_{\text{bottom}} - \rho_{\text{top}}) \quad (7.20)$$

Pennington and Koch [135] have used a similar stacked model to solve for polarization changes along the beam direction, but their approach relies on having a “composite” CBED measurement that spans a large range of incident beam directions. This approach is both

experimentally more challenging and several orders of magnitude more computationally intensive, but the inclusion of many beam directions may be necessary in order to obtain 3D information. It may also be possible to use multibeam electron diffraction [66] to obtain several nanobeam diffraction patterns with large angular separation simultaneously.

## 7.5 Conclusions

In this work, we have constructed a model of the electron multiple scattering through a complex multi-layer sample, parameterized over the physically relevant variables, and utilized an optimization procedure to fit the model to experimental data. Simple models for measuring the polarization of materials from nanobeam electron diffraction patterns, using symmetry breaking of pairs of diffracted disks, break down in the presence of even small tilts of the crystal, causing contrast changes and reversals. Using a stacked  $\mathcal{S}$ -matrix approach, we are able to use all of the scattered beams to determine polarization and tilt simultaneously. In order to make the problem computationally feasible for a large area scan, we derived the analytic gradients of the diffraction intensities and used them to perform regularized gradient descent.

**Part III**  
**Conclusion**



# Chapter 8

## Summary and Future Work

### 8.1 Conclusions

In this dissertation, we began by overviewing the theory of diffraction in a TEM, how dynamical diffraction arises, and how diffraction patterns are computed using quantum mechanical methods. In Chapter 3, we demonstrated an algorithm for hyperspectral total variational denoising, with a high-performance open-source implementation, `cyTVDN`, that is specifically designed to handle very large 4D-STEM datasets. On simulated low-dose data we observed that the method could effectively reduce the noise in diffraction patterns and recovered Bragg disks that could be detected and utilized in structural analyses. In Chapter 4, we demonstrated how momentum-patterned electron probes can improve the precision of the Bragg disk detection that is used for calculating strain from 4D-STEM data. Our approach greatly improved the precision of strain measurements, in particular for thick samples with strong dynamical diffraction effects. In experimental strain maps we found that anomalous strain due to measurement errors was reduced from  $\epsilon = \pm 1.5\%$  to about  $\pm 0.1\%$  by the use of the patterned aperture, and in simulations we found the patterned probes to yield precise strain maps at lower dose than feasible with unpatterned probes. In Chapter 5 we introduced an efficient method to perform automated crystal orientation mapping using a sparse library matching procedure. This algorithm was implemented as part of the open-source analysis suite `py4DSTEM` and is freely available. We applied the method to a complex polycrystalline nanowire sample, and used it to determine the orientation relationships between neighboring grains. In Chapter 6 we extended the algorithm of the previous chapter to also include dynamical diffraction effects, using an adaptive grid refinement procedure and Bloch wave simulations. We found that the orientation precision was improved from the library matching precision of  $2^\circ$  to about  $0.2^\circ$ , while also simultaneously generating a measurement of the sample thickness. Finally, in Chapter 7 we extended the simulation matching approach even further to measure slight polar distortions in a complex multilayer sample. From a superlattice sample consisting of three distinct layers along the beam direction, we were able to separate the effects of tilting of the foil from the polarization of the center layer. These

effects are normally confounded when using a kinematic model of the diffraction process, but could be separated by our Bloch wave-based refinement procedure. To accelerate this procedure, we derived semi-analytic expressions for the gradients of the diffraction intensities with respect to the structural distortions, which allowed us to use the model to solve for the tilt and polarization of the material across a large-area scan.

## 8.2 Future Work

Several of the parts of this work have room for future development. In Chapter 5, our crystal orientation mapping method currently operates under the assumption that the sample consists of only one phase. However, many samples of interest consist of multiple distinct phases, such as in transition metal dichalcogenide heterostructures or precipitates in metallic samples. Adding the ability to distinguish between two phases is challenging for a number of reasons. For many samples, the phases present may have a symmetry relationship between them, and in many orientations they could be indistinguishable. Often, phases will have nearly equal lattice parameters, and then precise measurement of the disk positions becomes critical.

The work in Chapter 7 on fitting a dynamical diffraction model to measurements from a complex sample also has room for future development. In that work the complexity of the model had to be limited in order to keep the computational cost low so that large area scans could be fitted. This meant that we could not include and truly 3D information in the model despite the fact that the sample is expected to have property variation along the beam direction. Fitting methods like this also suffer greatly when there is a model mismatch, and in this case we believe that the thickness of the sample may have been varying over the field of view. An interesting future direction for this work would be to perform the fitting using the optimization procedure produced by the machine learning community. This could be performed by either using the minimization methods from the machine learning community to perform the fit to each of the diffraction patterns, or by training a model that can more rapidly predict the structural parameters from the diffraction intensities. With a more powerful fitting algorithm, it may become feasible to include more parameters in the model and obtain more robust measurements.

Several of the methods described here could also be adapted to work with experiments that include a variety of different beam tilts in the measurement. Multibeam diffraction [66], which uses a specially patterned condenser aperture to illuminate the sample with several converged probes separated by a few degrees from the optic axis, yields multiple diffraction patterns for each probe position. Orientation mapping, for instance, could be improved by considering the diffracted intensities arising from each of the tilted beams, and could overcome the inherent orientation matching limit imposed by the flatness of the Ewald sphere at high accelerating voltage. Similarly, precession diffraction [111] records diffraction patterns from a cone of incident angles and sums them incoherently. This produces more kinematic-like intensities which may match the library calculations more faithfully. In

addition, it tends to allow one to record higher-order reflections, which are more sensitive to tilt and so could improve the accuracy of the dynamical refinement. In order to include these composite diffraction techniques, the simulation pipelines would need to be modified, likely requiring longer computation time.

# Bibliography

- [1] AA Aliev, ZA Isakhanov, and MK Ruzibaeva. “A study of the orientation dependence of the electron energy loss spectra for single crystal films of copper and silver”. In: *Journal of Surface Investigation. X-ray, Synchrotron and Neutron Techniques* 6.1 (2012), pp. 149–153.
- [2] Leslie J Allen, SD Findlay, et al. “Modelling the inelastic scattering of fast electrons”. In: *Ultramicroscopy* 151 (2015), pp. 11–22.
- [3] *Anaconda Software Distribution*. Version Vers. 2-2.4.0. 2020. URL: <https://docs.anaconda.com/>.
- [4] Sylvain Arlot and Alain Celisse. “A survey of cross-validation procedures for model selection”. In: *Statistics Surveys* 4 (2010), pp. 40–79.
- [5] Hussein Banjak et al. “Evaluation of noise and blur effects with SIRT-FISTA-TV reconstruction algorithm: Application to fast environmental transmission electron tomography”. In: *Ultramicroscopy* 189 (2018), pp. 109–123.
- [6] Alvaro Barbero and Suvrit Sra. “Modular Proximal Optimization for Multidimensional Total-Variation Regularization”. In: *Journal of Machine Learning Research* 19.56 (2018), pp. 1–82. URL: <http://jmlr.org/papers/v19/13-538.html>.
- [7] Álvaro Barbero and Suvrit Sra. “Fast Newton-type Methods for Total Variation Regularization.” In: *ICML*. Ed. by Lise Getoor and Tobias Scheffer. Omnipress, 2011, pp. 313–320. URL: <http://dblp.uni-trier.de/db/conf/icml/icml2011.html#JimenezS11>.
- [8] Joshua Batson and Loic Royer. *Noise2Self: Blind Denoising by Self-Supervision*. 2019. arXiv: [1901.11365](https://arxiv.org/abs/1901.11365) [cs.CV].
- [9] A Béché et al. “Improved precision in strain measurement using nanobeam electron diffraction”. In: *Applied Physics Letters* 95.12 (2009), p. 123114.
- [10] A. Béché et al. “Strain measurement at the nanoscale: Comparison between convergent beam electron diffraction, nano-beam electron diffraction, high resolution imaging and dark field electron holography”. In: *Ultramicroscopy* 131 (2013), pp. 10–23. ISSN: 0304-3991. DOI: <https://doi.org/10.1016/j.ultramicro.2013.03.014>. URL: <http://www.sciencedirect.com/science/article/pii/S0304399113000843>.

- [11] Amir Beck and Marc Teboulle. “A fast iterative shrinkage-thresholding algorithm for linear inverse problems”. In: *SIAM journal on imaging sciences* 2.1 (2009), pp. 183–202.
- [12] SW Bedell, A Khakifirooz, and DK Sadana. “Strain scaling for CMOS”. In: *MRS Bulletin* 39.2 (2014), pp. 131–137.
- [13] Stefan Behnel et al. “Cython: The best of both worlds”. In: *Computing in Science & Engineering* 13.2 (2011), pp. 31–39.
- [14] Hans Bethe. “Theorie der beugung von elektronen an kristallen”. In: *Annalen der Physik* 392.17 (1928), pp. 55–129.
- [15] R Bierwolf et al. “Direct measurement of local lattice distortions in strained layer structures by HREM”. In: *Ultramicroscopy* 49.1-4 (1993), pp. 273–285.
- [16] Stephanie N Bogle et al. “Size analysis of nanoscale order in amorphous materials by variable-resolution fluctuation electron microscopy”. In: *Ultramicroscopy* 110.10 (2010), pp. 1273–1278.
- [17] Walter Borchardt-Ott. *Crystallography: an introduction*. Springer Science & Business Media, 2011.
- [18] Stephen Boyd et al. “Distributed optimization and statistical learning via the alternating direction method of multipliers”. In: *Foundations and Trends in Machine learning* 3.1 (2011), pp. 1–122.
- [19] EG Britton and WM Stobbs. “The analysis and application of dynamical effects in HOLZ patterns”. In: *Ultramicroscopy* 21.1 (1987), pp. 1–11.
- [20] Hamish Brown et al. “A python based open-source multislice simulation package for transmission electron microscopy”. In: *Microscopy and Microanalysis* 26.S2 (2020), pp. 2954–2956.
- [21] Hamish G Brown et al. “A Three-Dimensional Reconstruction Algorithm for Scanning Transmission Electron Microscopy Data from a Single Sample Orientation”. In: *Microscopy and Microanalysis* 28.5 (2022), pp. 1632–1640.
- [22] Hamish G Brown et al. “Structure retrieval at atomic resolution in the presence of multiple scattering of the electron probe”. In: *Physical Review Letters* 121.26 (2018), p. 266102.
- [23] G Brunetti et al. “Confirmation of the domino-cascade model by LiFePO<sub>4</sub>/FePO<sub>4</sub> precession electron diffraction”. In: *Chemistry of Materials* 23.20 (2011), pp. 4515–4524.
- [24] Karen C Bustillo et al. “4D-STEM of Beam-Sensitive Materials”. In: *Accounts of Chemical Research* (2021), pp. 860–866.
- [25] Michael C Cao et al. “Theory and practice of electron diffraction from single atoms and extended objects using an EMPAD”. In: *Microscopy* 67.suppl.1 (2018), pp. i150–i161.

- [26] Andrés-Felipe Castro-Méndez, Juanita Hidalgo, and Juan-Pablo Correa-Baena. “The role of grain boundaries in perovskite solar cells”. In: *Advanced Energy Materials* 9.38 (2019), p. 1901489.
- [27] Niels Cautaerts et al. *Free, flexible and fast: orientation mapping using the multi-core and GPU-accelerated template matching capabilities in the python-based open source 4D-STEM analysis toolbox Pyxem*. 2021. arXiv: [2111.07347](https://arxiv.org/abs/2111.07347) [[cond-mat.mtrl-sci](https://arxiv.org/archive/cond)].
- [28] Rohit Chandra et al. *Parallel programming in OpenMP*. Morgan kaufmann, 2001.
- [29] Yi Chang et al. “Anisotropic spectral-spatial total variation model for multispectral remote sensing image destriping”. In: *IEEE Transactions on Image Processing* 24.6 (2015), pp. 1852–1866.
- [30] Zhen Chen et al. “Electron ptychography achieves atomic-resolution limits set by lattice vibrations”. In: *Science* 372.6544 (2021), pp. 826–831.
- [31] L Clark et al. “Probing the limits of the rigid-intensity-shift model in differential-phase-contrast scanning transmission electron microscopy”. In: *Physical Review A* 97.4 (2018), p. 043843.
- [32] Colin B Clement, Matthew Bierbaum, and James P Sethna. “Image registration and super resolution from first principles”. In: *arXiv:1809.05583* (2018).
- [33] L Clément et al. “Strain measurements by convergent-beam electron diffraction: The importance of stress relaxation in lamella preparations”. In: *Applied Physics Letters* 85.4 (2004), pp. 651–653.
- [34] David Cooper et al. “Strain mapping of semiconductor specimens with nm-scale resolution in a transmission electron microscope”. In: *Micron* 80 (2016), pp. 145–165.
- [35] John M Cowley and A F Moodie. “The scattering of electrons by atoms and crystals. I. A new theoretical approach”. In: *Acta Crystallographica* 10.10 (1957), pp. 609–619.
- [36] Kostadin Dabov et al. “Image denoising with block-matching and 3D filtering”. In: *Image Processing: Algorithms and Systems, Neural Networks, and Machine Learning*. Vol. 6064. International Society for Optics and Photonics. 2006, p. 606414.
- [37] AD Darbal et al. “Nanoscale automated phase and orientation mapping in the TEM”. In: *Microscopy Today* 20.6 (2012), pp. 38–42.
- [38] S Das et al. “Observation of room-temperature polar skyrmions”. In: *Nature* 568.7752 (2019), pp. 368–372.
- [39] E De Castro and CJIT Morandi. “Registration of translated and rotated images using finite Fourier transforms”. In: *IEEE Transactions on pattern analysis and machine intelligence* 5 (1987), pp. 700–703.
- [40] Marc De Graef. *Introduction to conventional transmission electron microscopy*. Cambridge university press, 2003.

- [41] Marc De Graef et al. *EMsoft-org/EMsoft: EMsoft Release 5.0.0*. Version v5.0.0. Oct. 2019. DOI: [10.5281/zenodo.3489720](https://doi.org/10.5281/zenodo.3489720). URL: <https://doi.org/10.5281/zenodo.3489720>.
- [42] Pratiti Deb et al. “Imaging polarity in two dimensional materials by breaking Friedel’s law”. In: *Ultramicroscopy* 215 (2020), p. 113019.
- [43] PH Dederichs et al. “Lattice theory of point defects”. In: *Journal of Nuclear Materials* 69 (1978), pp. 176–199.
- [44] Jeffrey J. Donatelli and John C. H. Spence. “Inversion of Many-Beam Bragg Intensities for Phasing by Iterated Projections: Removal of Multiple Scattering Artifacts from Diffraction Data”. In: *Phys. Rev. Lett.* 125 (6 Aug. 2020), p. 065502. DOI: [10.1103/PhysRevLett.125.065502](https://link.aps.org/doi/10.1103/PhysRevLett.125.065502). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.125.065502>.
- [45] Alexander S Eggeman, Robert Krakow, and Paul A Midgley. “Scanning precession electron tomography for three-dimensional nanoscale orientation imaging and crystallographic analysis”. In: *Nature communications* 6.1 (2015), pp. 1–7.
- [46] P Ewald. “Die Berechnung optischer und elektrostatischer Gitterpotentiale”. In: *Ann. Phys* 64 (1921), p. 253.
- [47] Scott D Findlay et al. “Scattering Matrix Determination in Crystalline Materials from 4D Scanning Transmission Electron Microscopy at a Single Defocus Value”. In: *Microscopy and Microanalysis* 27.4 (2021), pp. 744–757.
- [48] Brent Fultz and James M Howe. *Transmission electron microscopy and diffractometry of materials*. Springer Science & Business Media, 2012.
- [49] J-J Fundenberger et al. “Polycrystal orientation maps from TEM”. In: *Ultramicroscopy* 96.2 (2003), pp. 127–137.
- [50] Pedro L. Galindo et al. “The Peak Pairs algorithm for strain mapping from HRTEM images”. In: *Ultramicroscopy* 107.12 (2007), pp. 1186–1193. ISSN: 0304-3991. DOI: <https://doi.org/10.1016/j.ultramicro.2007.01.019>. URL: <http://www.sciencedirect.com/science/article/pii/S0304399107000204>.
- [51] Christoph Gammer et al. “Local nanoscale strain mapping of a metallic glass during *in situ* testing”. In: *Applied Physics Letters* 112.17 (2018), p. 171905.
- [52] Josiah Willard Gibbs. *Elements of vector analysis: arranged for the use of students in physics*. Tuttle, Morehouse & Taylor, 1884.
- [53] Shaun S Gleason, Martin A Hunt, and William Bruce Jatko. “Subpixel measurement of image features based on paraboloid surface fit”. In: *Machine vision systems integration in industry*. Vol. 1386. International Society for Optics and Photonics. 1991, pp. 135–145.
- [54] Tom Goldstein and Stanley Osher. “The split Bregman method for L1-regularized problems”. In: *SIAM journal on imaging sciences* 2.2 (2009), pp. 323–343.

- [55] Berit H Goodge, David J Baek, and Lena F Kourkoutis. “Atomic-resolution elemental mapping at cryogenic temperatures enabled by direct electron detection”. In: *arXiv preprint arXiv:2007.09747* (2020).
- [56] Tim Grieb et al. “Optimization of NBED simulations for disc-detection measurements”. In: *Ultramicroscopy* 181 (2017), pp. 50–60.
- [57] Manuel Guizar-Sicairos, Samuel T. Thurman, and James R. Fienup. “Efficient sub-pixel image registration algorithms”. In: *Opt. Lett.* 33.2 (Jan. 2008), pp. 156–158. DOI: [10.1364/OL.33.000156](https://doi.org/10.1364/OL.33.000156). URL: <http://ol.osa.org/abstract.cfm?URI=ol-33-2-156>.
- [58] Hui Guo et al. “Electron-event representation data enable efficient cryoEM file storage with full preservation of spatial and temporal resolution”. In: *IUCrJ* 7.5 (2020), pp. 860–869.
- [59] Giulio Guzzinati et al. “Electron Bessel beam diffraction for precise and accurate nanoscale strain mapping”. In: *arXiv:1902.06979* (2019).
- [60] Angelika Hähnel et al. “Improving accuracy and precision of strain analysis by energy-filtered nanobeam electron diffraction”. In: *Microscopy and Microanalysis* 18.1 (2012), pp. 229–240.
- [61] CR Hall and Peter Bernhard Hirsch. “Effect of thermal diffuse scattering on propagation of high energy electrons through crystals”. In: *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 286.1405 (1965), pp. 158–177.
- [62] Y Han et al. “Strain Mapping of Two-Dimensional Heterostructures with Subpicometer Precision.” In: *Nano Letters* 18.6 (2018), p. 3746.
- [63] Charles R. Harris et al. “Array programming with NumPy”. In: *Nature* 585.7825 (Sept. 2020), pp. 357–362. DOI: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2). URL: <https://doi.org/10.1038/s41586-020-2649-2>.
- [64] H Hashimoto, Archibald Howie, and Michael John Whelan. “Anomalous electron absorption effects in metal foils: Theory and comparison with experiment”. In: *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* 269.1336 (1962), pp. 80–103.
- [65] Martin V. Holt et al. “Strain Imaging of Nanoscale Semiconductor Heterostructures with X-Ray Bragg Projection Ptychography”. In: *Phys. Rev. Lett.* 112 (16 Apr. 2014), p. 165502. DOI: [10.1103/PhysRevLett.112.165502](https://doi.org/10.1103/PhysRevLett.112.165502). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.112.165502>.
- [66] Xuhao Hong et al. “Multibeam electron diffraction”. In: *Microscopy and Microanalysis* 27.1 (2021), pp. 129–139.
- [67] FJ Humphreys. “Review grain and subgrain characterisation by electron backscatter diffraction”. In: *Journal of materials science* 36.16 (2001), pp. 3833–3854.



- [68] M.J. Hÿtch, E. Snoeck, and R. Kilaas. “Quantitative measurement of displacement and strain fields from HREM micrographs”. In: *Ultramicroscopy* 74.3 (1998), pp. 131–146. ISSN: 0304-3991. DOI: [https://doi.org/10.1016/S0304-3991\(98\)00035-7](https://doi.org/10.1016/S0304-3991(98)00035-7). URL: <http://www.sciencedirect.com/science/article/pii/S0304399198000357>.
- [69] Martin J Hÿtch and Andrew M Minor. “Observing and measuring strain in nanostructures and devices with transmission electron microscopy”. In: *MRS Bulletin* 39.2 (2014), pp. 138–146.
- [70] Anubhav Jain et al. “Commentary: The Materials Project: A materials genome approach to accelerating materials innovation”. In: *APL materials* 1.1 (2013), p. 011002.
- [71] Matthew T Janish et al. “Mapping cation disorder in irradiated Gd<sub>2</sub>Ti<sub>2</sub>O<sub>7</sub> pyrochlore by 4D-STEM”. In: *Microscopy and Microanalysis* 25.S2 (2019), pp. 1560–1561.
- [72] GCAM Janssen. “Stress and strain in polycrystalline thin films”. In: *Thin solid films* 515.17 (2007), pp. 6654–6664.
- [73] Jiwon Jeong et al. “Automated crystal orientation mapping by precession electron diffraction assisted four-dimensional scanning transmission electron microscopy (4D-STEM) using a scintillator based CMOS detector”. In: *arXiv preprint arXiv:2102.09711* (2021).
- [74] Rong-Qing Jia and Hanqing Zhao. “A fast algorithm for the total variation model of image denoising”. In: *Advances in Computational Mathematics* 33.2 (2010), pp. 231–241.
- [75] Yi Jiang et al. “Electron ptychography of 2D materials to deep sub-ångström resolution”. In: *Nature* 559.7714 (2018), p. 343.
- [76] Lewys Jones and Peter D Nellist. “Advances in 2D, 3D and 4D STEM image data analysis”. In: *Microscopy and Microanalysis* 19.S2 (2013), pp. 770–771.
- [77] PM Jones, GM Rackham, and John Wickham Steeds. “Higher order Laue zone effects in electron diffraction and their use in lattice parameter determination”. In: *Proceedings of the Royal Society of London A. Mathematical and Physical Sciences* 354.1677 (1977), pp. 197–222.
- [78] Earl J Kirkland. *Advanced computing in electron microscopy*. Springer Science & Business Media, 2010.
- [79] Earl J Kirkland. *Advanced computing in electron microscopy, 3rd edition*. Springer Science & Business Media, 2020.
- [80] David Kleiven and Jaakko Akola. “Precipitate formation in aluminium alloys: Multi-scale modelling approach”. In: *Acta Materialia* 195 (2020), pp. 123–131.
- [81] M Knoll and E Ruska. “The electron microscope. vol 78, p 318”. In: *Zeitschrift für Physik* 79.9-10 (1932), pp. 699–699.

- [82] A Kobler et al. “Combination of in situ straining and ACOM TEM: A novel method for analysis of plastic deformation of nanocrystalline metals”. In: *Ultramicroscopy* 128 (2013), pp. 68–81.
- [83] Christoph T Koch, V Burak Özdöl, and Peter A van Aken. “An efficient, simple, and precise way to map strain with nanometer resolution in semiconductor devices”. In: *Applied Physics Letters* 96.9 (2010), p. 091901.
- [84] L Fitting Kourkoutis et al. “Atomic-resolution spectroscopic imaging of oxide interfaces”. In: *Philosophical Magazine* 90.35-36 (2010), pp. 4731–4749.
- [85] Eric Lang et al. “Automated Crystal Orientation Mapping with a Liquid-Cell TEM”. In: *Microscopy and Microanalysis* 27.S1 (2021), pp. 2232–2233.
- [86] Ask Hjorth Larsen et al. “The atomic simulation environment—a Python library for working with atoms”. In: *Journal of Physics: Condensed Matter* 29.27 (2017), p. 273002. URL: <http://stacks.iop.org/0953-8984/29/i=27/a=273002>.
- [87] James M LeBeau et al. “Position averaged convergent beam electron diffraction: Theory and applications”. In: *Ultramicroscopy* 110.2 (2010), pp. 118–125.
- [88] JM LeBeau et al. “Quantitative STEM: Experimental Methods and Applications”. In: *Journal of Physics: Conference Series*. Vol. 371. 1. IOP Publishing. 2012, p. 012053.
- [89] Richard LeSar. “Simulations of dislocation structure and response”. In: *Annu. Rev. Condens. Matter Phys.* 5.1 (2014), pp. 375–407.
- [90] Ju Li, Zhiwei Shan, and Evan Ma. “Elastic strain engineering for unprecedented materials properties”. In: *MRS Bulletin* 39.2 (2014), pp. 108–114.
- [91] XY Li et al. “Constrained minimal-interface structures in polycrystalline copper with extremely fine grains”. In: *Science* 370.6518 (2020), pp. 831–836.
- [92] Martin Linck et al. “Chromatic aberration correction for atomic resolution TEM imaging from 20 to 80 kV”. In: *Physical review letters* 117.7 (2016), p. 076101.
- [93] ACY Liu et al. “Systematic mapping of icosahedral short-range order in a melt-spun  $Zr_36Cu_64$  metallic glass”. In: *Physical Review Letters* 110.20 (2013), p. 205505.
- [94] Amelia CY Liu et al. “Interpretation of angular symmetries in electron nanodiffraction patterns from thin amorphous specimens”. In: *Acta Crystallographica Section A: Foundations and Advances* 71.5 (2015), pp. 473–482.
- [95] Li Liu, Yongxin Han, and Mingwu Jin. “Fast alternating projection methods for constrained tomographic reconstruction”. In: *PLOS ONE* 12.3 (Mar. 2017), pp. 1–23. DOI: [10.1371/journal.pone.0172938](https://doi.org/10.1371/journal.pone.0172938). URL: <https://doi.org/10.1371/journal.pone.0172938>.
- [96] I Lobato and D Van Dyck. “An accurate parameterization for scattering factors, electron densities and electrostatic potentials for neutral atoms that obey all physical constraints”. In: *Acta Crystallographica Section A: Foundations and Advances* 70.6 (2014), pp. 636–649.

- [97] I Lobato and D Van Dyck. “MULTTEM: A new multislice program to perform accurate and fast electron diffraction and imaging simulations using Graphics Processing Units with CUDA”. In: *Ultramicroscopy* 156 (2015), pp. 9–17.
- [98] I. Lobato and D. Van Dyck. “A complete comparison of simulated electron diffraction patterns using different parameterizations of the electron scattering factors”. In: *Ultramicroscopy* 155 (2015), pp. 11–19. ISSN: 0304-3991. DOI: <https://doi.org/10.1016/j.ultramic.2015.03.018>. URL: <https://www.sciencedirect.com/science/article/pii/S0304399115000674>.
- [99] Alejandra Londoño-Calderon et al. “1D to 2D transition in tellurium observed by 4D electron microscopy”. In: *Small* 16.49 (2020), p. 2005447.
- [100] Alejandra Londoño-Calderon et al. “Intrinsic helical twist and chirality in ultrathin tellurium nanowires”. In: *Nanoscale* 13.21 (2021), pp. 9606–9614.
- [101] Ian MacLaren et al. “A comparison of a direct electron detector and a high-speed video camera for a scanning precession electron diffraction phase and orientation mapping”. In: *Microscopy and Microanalysis* 26.6 (2020), pp. 1110–1116.
- [102] J Madsen and T Susi. “The abTEM code: transmission electron microscopy from first principles [version 2; peer review: 2 approved]”. In: *Open Research Europe* 1.24 (2021). DOI: [10.12688/openreseurope.13015.2](https://doi.org/10.12688/openreseurope.13015.2).
- [103] Jacob Madsen, Timothy J Pennycook, and Toma Susi. “ab initio description of bonding for transmission electron microscopy”. In: *Ultramicroscopy* 231 (2021), p. 113253.
- [104] Arnulf Maeland and Ted B Flanagan. “Lattice spacings of gold–palladium alloys”. In: *Canadian journal of physics* 42.11 (1964), pp. 2364–2366.
- [105] Matteo Maggioni et al. “Nonlocal transform-domain filter for volumetric data denoising and reconstruction”. In: *IEEE transactions on image processing* 22.1 (2012), pp. 119–133.
- [106] Christoph Mahr et al. “Theoretical study of precision and accuracy of strain analysis by nano-beam electron diffraction”. In: *Ultramicroscopy* 158 (2015), pp. 38–48.
- [107] Christoph Mahr et al. “Towards the interpretation of a shift of the central beam in nano-beam electron diffraction as a change in mean inner potential”. In: *Ultramicroscopy* 236 (2022), p. 113503. ISSN: 0304-3991. DOI: <https://doi.org/10.1016/j.ultramic.2022.113503>. URL: <https://www.sciencedirect.com/science/article/pii/S0304399122000389>.
- [108] T Malis, SC Cheng, and RF Egerton. “EELS log-ratio technique for specimen-thickness measurement in the TEM”. In: *Journal of electron microscopy technique* 8.2 (1988), pp. 193–200.
- [109] Ankit Nalin Mehta et al. “Unravelling stacking order in epitaxial bilayer MX<sub>2</sub> using 4D-STEM with unsupervised learning”. In: *Nanotechnology* 31.44 (2020), p. 445702.

- [110] Yifei Meng and J-M Zuo. “Three-dimensional nanostructure determination from a large diffraction data set recorded using scanning electron nanodiffraction”. In: *IUCrJ* 3.5 (2016), pp. 300–308.
- [111] Paul A Midgley and Alexander S Eggeman. “Precession electron diffraction—a topical review”. In: *IUCrJ* 2.1 (2015), pp. 126–136.
- [112] P Moeck et al. “High spatial resolution semi-automatic crystallite orientation and phase mapping of nanocrystals in transmission electron microscopes”. In: *Crystal research and technology* 46.6 (2011), pp. 589–606.
- [113] J. J. Mortensen, L. B. Hansen, and K. W. Jacobsen. “Real-space grid implementation of the projector augmented wave method”. In: *Phys. Rev. B* 71 (3 Jan. 2005), p. 035109. DOI: [10.1103/PhysRevB.71.035109](https://doi.org/10.1103/PhysRevB.71.035109). URL: <https://link.aps.org/doi/10.1103/PhysRevB.71.035109>.
- [114] Knut Müller et al. “Scanning transmission electron microscopy strain measurement from millisecond frames of a direct electron charge coupled device”. In: *Applied Physics Letters* 101.21 (2012), p. 212110.
- [115] Knut Müller et al. “Strain measurement in semiconductor heterostructures by scanning transmission electron microscopy”. In: *Microscopy and Microanalysis* 18.5 (2012), pp. 995–1009.
- [116] Joydeep Munshi et al. “Improving Strain Mapping from Electron Diffraction Patterns Using Fourier Space Deep Learning Networks”. In: *manuscript in preparation* (2021).
- [117] I. Najfeld and T.F. Havel. “Derivatives of the Matrix Exponential and Their Computation”. In: *Advances in Applied Mathematics* 16.3 (1995), pp. 321–375. ISSN: 0196-8858. DOI: <https://doi.org/10.1006/aama.1995.1017>. URL: <https://www.sciencedirect.com/science/article/pii/S0196885885710172>.
- [118] Kayla X Nguyen et al. “Transferring orbital angular momentum to an electron beam reveals toroidal and chiral order”. In: *arXiv preprint arXiv:2012.04134* (2020).
- [119] Magnus Nord et al. “Fast pixelated detectors in scanning transmission electron microscopy. Part I: data acquisition, live processing, and storage”. In: *Microscopy and Microanalysis* 26.4 (2020), pp. 653–666.
- [120] CM O’Leary et al. “Phase reconstruction using fast binary 4D STEM data”. In: *Applied Physics Letters* 116.12 (2020), p. 124101.
- [121] Shyue Ping Ong et al. “Python Materials Genomics (pymatgen): A robust, open-source python library for materials analysis”. In: *Computational Materials Science* 68 (2013), pp. 314–319.
- [122] Colin Ophus. “A fast image simulation algorithm for scanning transmission electron microscopy”. In: *Advanced structural and chemical imaging* 3.1 (2017), pp. 1–11.

- [123] Colin Ophus. “Four-dimensional scanning transmission electron microscopy (4D-STEM): From scanning nanodiffraction to ptychography and beyond”. In: *Microscopy and Microanalysis* 25.3 (2019), pp. 563–582.
- [124] Colin Ophus et al. “Automated Crystal Orientation Mapping in py4DSTEM using Sparse Correlation Matching”. In: *Microscopy and Microanalysis* 28.2 (2022), pp. 390–403.
- [125] Colin Ophus et al. “Large-scale experimental and theoretical study of graphene grain boundary structures”. In: *Physical Review B* 92.20 (2015), p. 205402.
- [126] Stanley Osher et al. “An iterative regularization method for total variation-based image restoration”. In: *Multiscale Modeling & Simulation* 4.2 (2005), pp. 460–489.
- [127] VB Ozdol et al. “Strain mapping at nanometer resolution using advanced nano-beam electron diffraction”. In: *Applied Physics Letters* 106.25 (2015), p. 253107.
- [128] Ouliana Panova et al. “Orientation mapping of semicrystalline polymers using scanning electron nanobeam diffraction”. In: *Micron* 88 (2016), pp. 30–36.
- [129] Min-Ho Park et al. “Efficient Perovskite Light-Emitting Diodes Using Polycrystalline Core–Shell-Mimicked Nanograins”. In: *Advanced Functional Materials* 29.22 (2019), p. 1902017.
- [130] Gary W Paterson et al. “Fast Pixelated Detectors in Scanning Transmission Electron Microscopy. Part II: Post-Acquisition Data Processing, Visualization, and Structural Characterization”. In: *Microscopy and Microanalysis* 26.5 (2020), pp. 944–963.
- [131] Thomas C Pekin et al. “In situ nanobeam electron diffraction strain mapping of planar slip in stainless steel”. In: *Scripta Materialia* 146 (2018), pp. 87–90.
- [132] Thomas C Pekin et al. “Optimizing disk registration algorithms for nanobeam electron diffraction strain mapping”. In: *Ultramicroscopy* 176 (2017), pp. 170–176.
- [133] Philipp M Pelz et al. “Phase-contrast imaging of multiply-scattering extended objects at atomic resolution by reconstruction of the scattering matrix”. In: *Physical Review Research* 3.2 (2021), p. 023159.
- [134] Philipp M. Pelz et al. “Real-Time Interactive 4D-STEM Phase-Contrast Imaging From Electron Event Representation Data: Less computation with the right representation”. In: *IEEE Signal Processing Magazine* 39.1 (2022), pp. 25–31. DOI: [10.1109/MSP.2021.3120981](https://doi.org/10.1109/MSP.2021.3120981).
- [135] Robert S. Pennington and Christoph T. Koch. “A three-dimensional polarization domain retrieval method from electron diffraction data”. In: *Ultramicroscopy* 155 (2015), pp. 42–48. ISSN: 0304-3991. DOI: <https://doi.org/10.1016/j.ultramic.2015.04.002>. URL: <https://www.sciencedirect.com/science/article/pii/S0304399115000728>.

- [136] Robert S. Pennington and Christoph T. Koch. “Retrieving depth-direction information from TEM diffraction data under reciprocal-space sampling variation”. In: *Ultramicroscopy* 148 (2015), pp. 105–114. ISSN: 0304-3991. DOI: <https://doi.org/10.1016/j.ultramic.2014.10.006>. URL: <https://www.sciencedirect.com/science/article/pii/S0304399114001983>.
- [137] Robert S. Pennington, Wouter Van den Broek, and Christoph T. Koch. “Third-dimension information retrieval from a single convergent-beam transmission electron diffraction pattern using an artificial neural network”. In: *Phys. Rev. B* 89 (20 May 2014), p. 205409. DOI: [10.1103/PhysRevB.89.205409](https://doi.org/10.1103/PhysRevB.89.205409). URL: <https://link.aps.org/doi/10.1103/PhysRevB.89.205409>.
- [138] Robert S. Pennington, Feng Wang, and Christoph T. Koch. “Stacked-Bloch-wave electron diffraction simulations using GPU acceleration”. In: *Ultramicroscopy* 141 (2014), pp. 32–37. ISSN: 0304-3991. DOI: <https://doi.org/10.1016/j.ultramic.2014.03.003>. URL: <https://www.sciencedirect.com/science/article/pii/S0304399114000485>.
- [139] Robert S. Pennington et al. “Neural-network-based depth-resolved multiscale structural optimization using density functional theory and electron diffraction data”. In: *Phys. Rev. B* 97 (2 Jan. 2018), p. 024112. DOI: [10.1103/PhysRevB.97.024112](https://doi.org/10.1103/PhysRevB.97.024112). URL: <https://link.aps.org/doi/10.1103/PhysRevB.97.024112>.
- [140] Stephen J Pennycook. “A scan through the history of STEM”. In: *Scanning Transmission Electron Microscopy*. Springer, 2011, pp. 1–90.
- [141] Nicolas J Peter et al. “Segregation-induced nanofaceting transition at an asymmetric tilt grain boundary in copper”. In: *Physical review letters* 121.25 (2018), p. 255502.
- [142] Alan Pryor, Colin Ophus, and Jianwei Miao. “A streaming multi-GPU implementation of image simulation algorithms for scanning transmission electron microscopy”. In: *Advanced Structural and Chemical Imaging* 3.1 (2017), p. 15.
- [143] Haoyuan Qi et al. “Near-atomic-scale observation of grain boundaries in a layer-stacked two-dimensional polymer”. In: *Science advances* 6.33 (2020), eabb5976.
- [144] Alexander Rakowski et al. “A Complete Pipeline for Deep Learning Workflows in Transmission Electron Microscopy”. In: *manuscript in preparation* (2021).
- [145] Quentin M Ramasse. “Twenty years after: How “Aberration correction in the STEM” truly placed a “A synchrotron in a Microscope””. In: *Ultramicroscopy* 180 (2017), pp. 41–51.
- [146] Luis Rangel DaCosta et al. “Prismatic 2.0-Simulation software for scanning and high resolution transmission electron microscopy (STEM and HRTEM)”. In: *Micron* (2021), p. 103141.
- [147] Edgar F Rauch et al. “Automated nanocrystal orientation and phase mapping in the transmission electron microscope on the basis of precession electron diffraction”. In: *Zeitschrift für Kristallographie* 225.2-3 (2010), pp. 103–109.

- [148] EF Rauch and L Dupuy. “Rapid spot diffraction patterns identification through template matching”. In: *Archives of Metallurgy and Materials* 50 (2005), pp. 87–99.
- [149] EF Rauch and MJMC Véron. “Automated crystal orientation and phase mapping in TEM”. In: *Materials Characterization* 98 (2014), pp. 1–9.
- [150] Ian Robinson and Ross Harder. “Coherent X-ray diffraction imaging of strain at the nanoscale”. In: *Nature Materials* 8.4 (2009), p. 291.
- [151] H Rose. In: *Optik* 45.2 (1976).
- [152] C.J. Rossouw et al. “Quantitative absorption corrections for electron diffraction: Correlation between theory and experiment”. In: *Ultramicroscopy* 34.3 (1990), pp. 149–163. ISSN: 0304-3991. DOI: [https://doi.org/10.1016/0304-3991\(90\)90069-X](https://doi.org/10.1016/0304-3991(90)90069-X). URL: <https://www.sciencedirect.com/science/article/pii/030439919090069X>.
- [153] Jean-Luc Rouviere et al. “Improved strain precision with high spatial resolution using nanobeam precession electron diffraction”. In: *Applied Physics Letters* 103.24 (2013), p. 241913.
- [154] Jean-Luc Rouvière. *Method to facilitate positioning of diffraction spots*. US Patent App. 13/877,904. Aug. 2013.
- [155] David Rowenhorst et al. “Consistent representations of and conversions between 3D rotations”. In: *Modelling and Simulation in Materials Science and Engineering* 23.8 (2015), p. 083501.
- [156] Leonid I Rudin, Stanley Osher, and Emad Fatemi. “Nonlinear total variation based noise removal algorithms”. In: *Physica D: nonlinear phenomena* 60.1-4 (1992), pp. 259–268.
- [157] Benjamin Savitzky et al. *py4DSTEM: Open source processing and analysis of 4D-STEM data*. Version v0.5. July 2019. DOI: [10.5281/zenodo.3333960](https://doi.org/10.5281/zenodo.3333960). URL: <https://doi.org/10.5281/zenodo.3333960>.
- [158] Benjamin H Savitzky et al. “py4DSTEM: A software package for four-dimensional scanning transmission electron microscopy data analysis”. In: *Microscopy and Microanalysis* 27.4 (2021), pp. 712–743.
- [159] Jonathan Schwartz et al. “Dynamic compressed sensing for real-time tomographic reconstruction”. In: *Ultramicroscopy* 219 (2020), p. 113122. ISSN: 0304-3991. DOI: <https://doi.org/10.1016/j.ultramicro.2020.113122>. URL: <https://www.sciencedirect.com/science/article/pii/S0304399120302734>.
- [160] Robert A Schwarzer and J Sukkau. “Automated crystal orientation mapping (ACOM) with a computer-controlled TEM by interpreting transmission Kikuchi patterns”. In: *Materials Science Forum* 273 (1998), pp. 215–222.

- [161] Takehito Seki et al. “Linear imaging theory for differential phase contrast and other phase imaging modes in scanning transmission electron microscopy”. In: *Ultramicroscopy* 240 (2022), p. 113580.
- [162] R Serneels and R Gevers. “A Particular Many Beam Situation in Transmission Electron Diffraction”. In: *physica status solidi (b)* 45.2 (1971), pp. 493–504.
- [163] Ziming Shao et al. *Real-space imaging of polar and elastic nano-textures in thin films via inversion of diffraction data*. 2022. DOI: [10.48550/ARXIV.2211.01506](https://doi.org/10.48550/ARXIV.2211.01506). URL: <https://arxiv.org/abs/2211.01506>.
- [164] Ken Shoemake. “Animating rotation with quaternion curves”. In: *Proceedings of the 12th annual conference on Computer graphics and interactive techniques*. 1985, pp. 245–254.
- [165] S. Singh, M.J. Mills, and M. De Graef. “Dynamical scattering image simulations for two-phase  $\gamma$ - $\gamma'$  microstructures: A theoretical model”. In: *Ultramicroscopy* 185 (2018), pp. 32–41. ISSN: 0304-3991. DOI: <https://doi.org/10.1016/j.ultramic.2017.11.008>. URL: <https://www.sciencedirect.com/science/article/pii/S0304399117304138>.
- [166] R. Soummer et al. “Fast computation of Lyot-style coronagraph propagation”. In: *Opt. Express* 15.24 (Nov. 2007), pp. 15935–15951. DOI: [10.1364/OE.15.015935](https://doi.org/10.1364/OE.15.015935). URL: <http://www.opticsexpress.org/abstract.cfm?URI=oe-15-24-15935>.
- [167] JCH Spence. “On the accurate measurement of structure-factor amplitudes and phases by electron diffraction”. In: *Acta Crystallographica Section A: Foundations of Crystallography* 49.2 (1993), pp. 231–260.
- [168] Steven R Spurgeon et al. “Towards data-driven next-generation transmission electron microscopy”. In: *Nature materials* 20.3 (2021), pp. 274–279.
- [169] Lorenzo Sturkey. “The calculation of electron diffraction intensities”. In: *Proceedings of the Physical Society (1958-1967)* 80.2 (1962), p. 321.
- [170] Sandhya Susarla et al. “The emergence of three-dimensional chiral domain walls in polar vortices”. In: *Manuscript under review* (2022).
- [171] Toma Susi et al. “Efficient first principles simulation of electron scattering factors for transmission electron microscopy”. In: *Ultramicroscopy* 197 (2019), pp. 16–22.
- [172] Ming Tang, W Craig Carter, and Rowland M Cannon. “Diffuse interface model for structural transitions of grain boundaries”. In: *Physical Review B* 73.2 (2006), p. 024102.
- [173] Jing Tao et al. “Direct imaging of nanoscale phase separation in La 0.55 Ca 0.45 MnO 3: relationship to colossal magnetoresistance”. In: *Physical Review Letters* 103.9 (2009), p. 097202.



- [174] Carl V Thompson. “Structure evolution during processing of polycrystalline films”. In: *Annual review of materials science* 30.1 (2000), pp. 159–190.
- [175] Carl V Thompson and Roland Carel. “Texture development in polycrystalline thin films”. In: *Materials Science and Engineering: B* 32.3 (1995), pp. 211–219.
- [176] Koji Usuda et al. “Strain relaxation of strained-Si layers on SiGe-on-insulator (SGOI) structures after mesa isolation”. In: *Applied Surface Science* 224 (2004), pp. 113–116.
- [177] PM Voyles and DA Muller. “Fluctuation microscopy in the STEM”. In: *Ultramicroscopy* 93.2 (2002), pp. 147–159.
- [178] A Wang and M De Graef. “Modeling dynamical electron scattering with Bethe potentials and the scattering matrix”. In: *Ultramicroscopy* 160 (2016), pp. 35–43.
- [179] Feng Wang et al. “Noise2Atom: unsupervised denoising for scanning transmission electron microscopy images”. In: *Applied Microscopy* 50.1 (2020), pp. 1–9.
- [180] Yong Wang et al. “Chiral Transformation: From Single Nanowire to Double Helix”. In: *Journal of the American Chemical Society* 133.50 (2011), pp. 20060–20063.
- [181] Zhang-Jie Wang et al. “Sample size effects on the large strain bursts in submicron aluminum pillars”. In: *Applied Physics Letters* 100.7 (2012), p. 071906.
- [182] Geoff Wehmeyer et al. “Measuring temperature-dependent thermal diffuse scattering using scanning transmission electron microscopy”. In: *Applied Physics Letters* 113.25 (2018), p. 253101.
- [183] A Weickenmeier and H Kohl. “Computation of absorptive form factors for high-energy electron diffraction”. In: *Acta Crystallographica Section A: Foundations of Crystallography* 47.5 (1991), pp. 590–597.
- [184] Wikipedia contributors. *BITNET — Wikipedia, The Free Encyclopedia*. [Online; accessed 6-December-2022]. 2022. URL: <https://en.wikipedia.org/w/index.php?title=BITNET&oldid=1109177498>.
- [185] Stuart I Wright, Matthew M Nowell, and David P Field. “A review of strain analysis using electron backscatter diffraction”. In: *Microscopy and microanalysis* 17.3 (2011), pp. 316–329.
- [186] Stuart I Wright et al. “Introduction and comparison of new EBSD post-processing methodologies”. In: *Ultramicroscopy* 159 (2015), pp. 81–94.
- [187] Guilin Wu and Stefan Zaeferrer. “Advances in TEM orientation microscopy by combination of dark-field conical scanning and improved image matching”. In: *Ultramicroscopy* 109.11 (2009), pp. 1317–1325.
- [188] Lijun Wu, Qingping Meng, and Yimei Zhu. “Mapping valence electron distributions with multipole density formalism using 4D-STEM”. In: *Ultramicroscopy* 219 (2020), p. 113095. ISSN: 0304-3991. DOI: <https://doi.org/10.1016/j.ultramic.2020.113095>. URL: <https://www.sciencedirect.com/science/article/pii/S0304399120302461>.

- [189] Mingjian Wu et al. “Seeing Structural Evolution of Organic Molecular Nano-crystallites Using 4D Scanning Confocal Electron Diffraction”. In: *arXiv preprint arXiv:2110.02373* (2021).
- [190] AK Yadav et al. “Observation of polar vortices in oxide superlattices”. In: *Nature* 530.7589 (2016), pp. 198–201.
- [191] Hao Yang et al. “Electron ptychographic phase imaging of light elements in crystalline materials using Wigner distribution deconvolution”. In: *Ultramicroscopy* 180 (2017), pp. 173–179.
- [192] Renliang Yuan et al. “Training artificial neural networks for precision orientation and strain mapping using 4D electron diffraction datasets”. In: *Ultramicroscopy* (2021), p. 113256.
- [193] S Zaefferer. “A critical review of orientation microscopy in SEM and TEM”. In: *Crystal Research and Technology* 46.6 (2011), pp. 607–628.
- [194] Stefan Zaefferer and Robert A Schwarzer. “On-line interpretation of spot and Kikuchi patterns”. In: *Materials Science Forum*. Vol. 157. Trans Tech Publ. 1994, pp. 247–250.
- [195] Nestor J Zaluzec. “Computationally Mediated Experimental Science”. In: *Microscopy and Microanalysis* 9.S02 (2003), pp. 150–151.
- [196] Nestor J Zaluzec. “Quantitative measurements of magnetic vortices using position resolved diffraction in Lorentz STEM”. In: *Microscopy and Microanalysis* 8.S02 (2002), pp. 376–377.
- [197] SE Zeltmann. *cyTVDN/cyTVDN: alpha release*. Version v0,1-alpha. Mar. 2020. DOI: [10.5281/zenodo.3715582](https://doi.org/10.5281/zenodo.3715582). URL: <https://doi.org/10.5281/zenodo.3715582>.
- [198] Steven E Zeltmann, Andrew M Minor, and Colin Ophus. “4D-STEM Measurement of Thickness and Orientation by Bloch Wave Dynamical Diffraction Matching”. In: *Microscopy and Microanalysis* 28.S1 (2022), pp. 382–383.
- [199] Steven E Zeltmann et al. “Patterned probes for high precision 4D-STEM Bragg measurements”. In: *Ultramicroscopy* 209 (2020), p. 112890.
- [200] Steven E Zeltmann et al. “Uncovering polar vortex structures by inversion of multiple scattering with a stacked Bloch wave model”. In: *arXiv preprint arXiv:2211.05842* (2022).
- [201] Chenyu Zhang et al. “Denoising atomic resolution 4D scanning transmission electron microscopy data with tensor singular value decomposition”. In: *Ultramicroscopy* 219 (2020), p. 113123.
- [202] Peng Zhang et al. “Direct strain measurement in a 65 nm node strained silicon transistor by convergent-beam electron diffraction”. In: *Applied Physics Letters* 89.16 (2006), p. 161907.

- [203] Guo-Zhen Zhu, Guillaume Radtke, and Gianluigi A Botton. “Bonding and structure of a reconstructed (001) surface of SrTiO<sub>3</sub> from TEM”. In: *Nature* 490.7420 (2012), p. 384.
- [204] Menglin Zhu and Jinwoo Hwang. “Temperature Mapping with STEM Atomic Scale Debye-Waller Thermometry”. In: *Microscopy and Microanalysis* 28.S1 (2022), pp. 118–120.
- [205] Jian Min Zuo and John CH Spence. *Advanced transmission electron microscopy*. Springer, 2017.
- [206] Jian-Min Zuo and Xiurong Zhu. “Strategies for fast and reliable 4D-STEM orientation and phase mapping of nanomaterials and devices”. In: *Microscopy and Microanalysis* 27.S1 (2021), pp. 762–763.