

Imaging 3D Chemistry and Structure at Nanometer Resolution


by

Jonathan Schwartz

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Material Science and Engineering)
in the University of Michigan
2023

Doctoral Committee:

Associate Professor Robert Hovden, Chair
Professor Jeffrey A. Fessler
Associate Professor Ashwin Shahani
Professor Steven Yalisove



Jonathan Schwartz

jtschw@umich.edu

ORCID iD: 0000-0002-8063-6951

© Jonathan Schwartz 2023

DEDICATION

For my friends who motivated and supported me during my academic journey.

For my academic advisor who provided me guidance throughout the process.

For my parents and family that have always supported and loved me.

ACKNOWLEDGMENTS

When I started working on these specific subjects, it felt like the end would never come. It was apparently so far away. We knew completing this task was going to be difficult, but it came as unexpectedly as the Spanish Inquisition, to say the least. Yet it was possible, indeed, thanks to people ready to give advice and lend a hand. Recalling their efforts, I proceed here to tell them, "thank you." The process of the thesis to this point has been a long, windy road going up and down, lets focus on the ups.

I would like to thank the staff at MC2 who have also become my colleagues for more than five years now, for their help in the understanding and operation of the electron microscopes; especially Dr. Tao Ma, Dr. Kai Sun, Dr. Haiping and Bobby Kerns.

I also want to thank the staff scientists at Argonne and Lawrence Berkeley National Laboratory, specifically Dr. Yi Jiang, Dr. Huihuo Zheng Dr. Zichao Wendy Di, Dr. Colin Ophus, Dr. Peter Ercius and Prof. Mary Scott for their unbounded guidance in my early years in my Ph.D. Their guidance has provided me pivotal the assistance for developing software on supercomputers and patience when attempted challenging tomography experiments.

I would also like to express my deepest gratitude to Prof. Hovden and everyone in our research group. My gratitude is for all of them that helped teach me sample preparation, microscope operation and data analysis during the early days of my education.

I would also like to thank friends and family, several of which had helped make my long and windy road still be composed of fun adventures.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGMENTS	iii
LIST OF FIGURES	vi
LIST OF APPENDICES	xv
LIST OF ACRONYMS	xvi
ABSTRACT	xvii
CHAPTER	
1 Introduction and Background	1
1.1 Scanning Transmission Electron Microscopy	1
1.2 Electron Spectroscopy	4
1.2.1 Electron Energy Loss Spectroscopy	4
1.2.2 Energy Dispersive X-Ray Spectroscopy	5
1.3 Electron Tomography	7
1.3.1 Resolution and the Crowther Criterion	10
1.3.2 Resolution and the Missing Wedge	12
1.3.3 Electron Tomography as an Inverse Problem	13
1.3.4 Improving Electron Tomography with Sparse Optimization	15
1.3.5 Compressed Sensing Tomography with Total Variation Regularization	16
2 Removing Stripes, Scratches and Curtaining with Non-Recoverable Compressed Sensing	19
2.1 Introduction	19
2.2 Background	20
2.3 Results	22
2.4 Discussion and Conclusion	27
3 Imaging Atomic-Scale Chemistry from Fused Multi-Modal Electron Microscopy	28
3.1 Introduction	28
3.2 Principles of Multi-Modal Electron Microscopy	29
3.3 High-SNR Recovery of Nanomaterial Chemistry	30

3.4	Quantifying Chemical Concentration	35
3.5	Influence of Electron Dose	39
3.6	Discussion and Conclusion	42
4	Imaging 3D Chemistry at 1 nm Resolution with Fused Multi-Modal Electron Tomography	44
4.1	Introduction	44
4.2	Principles of Fused Multi-Modal Electron Microscopy	45
4.3	3D Chemistry at High-Resolution, Low-Dose	48
4.4	Sub-nanometer Chemical Resolution in 3D	54
4.5	Influence of Sampling	57
4.6	Discussion	61
4.7	Conclusion	62
5	Real-time 3D Analysis During Electron Tomography	64
5.1	Introduction	64
5.2	Real-Time Tomography Workflow	65
5.3	Real-Time 3D Visualization During Reconstruction	67
5.4	Dynamic Compressed Sensing for Real-Time Tomography	70
5.5	Convergence of Dynamic Compressed Sensing	72
5.6	Dynamic Parameter Tuning	74
5.7	The Live Tomography Software on tomviz	76
5.8	Parallelization and Performance	77
5.9	Discussion and Conclusion	79
6	Conclusion	80
	APPENDICES	82
	BIBLIOGRAPHY	97

LIST OF FIGURES

FIGURE

1.1	Aberration-Corrected STEM Microscope. Conceptual schematic of an aberration-corrected electron microscope with an HAADF image of a HgS nanoparticle shown on the right.	2
1.2	EELS spectra and inelastic scattering. a A sample EELS spectrum of MnO ₂ [adapted from eels.info]. b A classical (particle) view of electron scattering.	4
1.3	Spectroscopy in the STEM Microscope. a Conceptual schematic inelastic signal generation that occurs in the electron microscope. b-c Energy losses from transmitted electrons or characteristic energies from emitted X-ray correspond to chemical signatures that can be used to map the chemistry of materials (shown right).	6
1.4	Diagram of single-axis electron tomography. Projections are recorded at regular tilt increments of θ around a single fixed axis. A typical experiment contains 70 - 140 projection images from $\pm 70^\circ$ separated by consistent 1 or 2° increments.	8
1.5	Fourier Slice Theorem. A projection of a 2D object in real-space is equivalent to the central plane of the object's full Fourier transform at the same angle (θ).	9
1.6	Limited Angular Sampling. The Fourier slice theorem states that every angular projection ($p_\theta(t)$) contains full information about $F(k_x, k_y)$. A finite, regular sampling of t translate into a finite sampling along the projection lines (red lines). A limited sampling with only 20 views (top row) is compared a fully sampled tilt series with 180 views (bottom row). The empty space between adjacent planes negatively impacts the reconstruction quality.	11
1.7	The Missing Wedge and Limit Angular Sampling. The ground truth and sample tomography reconstruction of a Shepp-Logan Phantom using 40 projections with a 60° wide missing wedge (highlighted in blue), which suffers from elongation artifacts along the vertical direction.	12
1.8	Momentum Accelerated Optimization. Traditional gradient descent (GD) juxtaposed with Nesterov accelerated gradient descent (NAG). From comparable starting points, NAG converges ~ 3 times faster, requiring 9 steps to reach the minimum, while GD required 25 steps.	15
1.9	Illustration of Sparsity in the Gradient Magnitude Domain. A HAADF micrograph of ZnS-CuS nanoparticles and the magnitude of its spatial gradient shown right. While the HAADF image may not be sparse, its gradient image is.	17

2.1	Reconstruction of a scratched pearl surface. a) The original 10 keV BSE-SEM image of a scratched biomaterial with its FFT below. b) The output from deleting a wedge of information in Fourier Space. c) The TV minimization reconstruction and its Fourier Space image. d) The residual between images (b) and (a). e) The residuals between images (c) and (a).	22
2.2	Reconstructions of BF-TEM and FIB images. a) The original contrast reversal 300 keV BF-TEM image of an AlGa _N Quantum Well with horizontal intensity fluctuations and its FFT on the top right-hand corner. b) The reconstruction with the TV-minimization algorithm and its FFT. c) The residuals between images (a) and (b). d) The original 30 keV secondary electron (SE)-SEM image of an aluminum sample with curtaining collected during a FIB tomography experiment. e) The residuals between images (d) and (e). Contrast was reversed in (a) and (b) for clarity.	23
2.3	Reconstruction of a BSE-SEM micrograph under various horizontal (0°) ‘Missing’ Wedge widths. a) The original 10 keV SEM image of a pearl sample with vertical scratches highlighted by green arrows. b1-d1) The first row shows the images with a ‘missing’ wedge of information as the blurring artifacts are prominent in the yellow circles. b2-d2) The second row shows the TV minimization reconstructions with the removal of the blurring artifacts and loss of features for large wedges highlighted in the red circles. FFT insets shown lower left.	24
2.4	TV minimization reconstructions for decreasing SNR. a) The first row shows the images after Gaussian additive white noise is implemented and b2-d2) the second row shows the TV minimization reconstructions.	25
2.5	Quantitative Study of SNR and Wedge Size. a) A plot of the RMSE normalized by the error from various missing wedge sizes. Values below 1 indicate the reconstruction outperforms loss of information. Wedges below 8°, consistently achieves satisfactory performance at all SNR values above 10. b) DF-S/TEM micrograph (at 300 keV) of MBE grown InGa _N nanowires with platinum nanoparticles coated on the surface. c)	26
3.1	Consequence of Reconstructing Chemical Maps with Individual Components in Cost Function. a) The raw EDX maps and simultaneous HAADF image. b) Reconstructing the multi-modal dataset with $\Psi_1 + \Psi_2$ (model + data-consistency). The chemical maps are slightly improved but remain noisy. c) Reconstructing the multi-modal dataset with $\Psi_2 + TV$. This functional form is equivalent to a denoising problem; thus the resulting maps produce common staircase artifacts associated with TV. d) Reconstructing the raw EDX maps with fused Multi-Modal Electron Microscopy. Scale bar, 30 nm.	31
3.2	Nanoscale multi-modal chemical recovery of CoS catalysts using EDX + HAADF. a) Schematic highlighting the linked HAADF and EDX modalities collected in the microscope for every probe position. The algorithm links and correlates information between the two signals through an optimization process that produces chemical maps with higher SNRs. b) The raw EDX chemical maps for the Co, S, and O elemental distributions. c) The simultaneous HAADF micrograph of the CoS nanoparticle. d) The multi-modal reconstructions for the elemental distributions. e) EDX RGB overlay of the Co, S, and O maps. Scale bar, 30 nm.	32

3.3	Atomic-scale multi-modal chemical recovery of $\text{Co}_{3-x}\text{Mn}_x\text{O}_4$ supercapacitors using EELS + HAADF. a) Schematic highlighting the linked HAADF and EELS modalities collected in the microscope at every probe position. b) Raw EELS maps for the elemental distributions of Co, Mn - $L_{2,3}$ and O - K edges. c) The simultaneous HAADF micrograph of the $\text{Co}_{3-x}\text{Mn}_x\text{O}_4$ nanoparticle. d) The multi-modal reconstructions for the elemental distributions. e) EELS RGB overlay of the Co, S, and O maps. Scale bar, 2 nm.	33
3.4	Recovering chemistry in an atomically sharp $\text{ZnS-Cu}_{0.64}\text{S}_{0.36}$ heterointerface interface. a) The raw EELS maps for the Cu, S, and Zn $L_{2,3}$ edges. b) The multi-modal reconstructions for the elemental compositions. c) The simultaneous HAADF micrograph of the $\text{ZnS-Cu}_{0.64}\text{S}_{0.34}$ interface. d) Color overlay of the Zn, S, and Zn maps. The relative concentration for the constituent elements consist of $48\pm 5.9\%$ for Zn, $59.9\pm 3.2\%$ for Cu and $38\pm 2.6\%$ for S in the $\text{Cu}_{0.64}\text{S}_{0.36}$ layer and $48.9\pm 6\%$ in ZnS. e) Convergence plots for the three individual components in the cost function. Scale bar, 1 nm.	34
3.5	Inelastic FePt nanoparticle simulation. a) Ground truth EELS images generated from inelastic simulations. b) EELS maps degraded with Poisson shot noise. SNR shown on top right. c) Recovered atomic-resolution EELS maps for the Fe and Pt distributions. Estimated SNR shown on top right. d) Line profiles of the marked yellow bars (10 pixels in width) in (c) compares the Multi-Modal reconstruction and ground truth. Scale bar, 1 nm.	35
3.6	Histograms of intensities for ZrO_2 and NiO. Histogram of intensities for the Sulphur, Cobalt and Oxygen distributions. Due to the core-shell structure consisting of two phases, the histograms contain of two Gaussian distributions. The specimen core contains a relative concentration of $39\pm 1.6\%$, $42\pm 2.5\%$ and $13\pm 2.4\%$ and exterior shell composition of $26\pm 2.8\%$, $11\pm 2.0\%$, $54\pm 1.3\%$ for Co, S, O respectively. Scale bar, 30 nm.	36
3.7	Histograms of intensities for CoS chemical Maps. Histogram of intensities for the Sulphur, Cobalt and Oxygen distributions. Due to the core-shell structure consisting of two phases, the histograms contain of two Gaussian distributions. The specimen core contains a relative concentration of $39\pm 1.6\%$, $42\pm 2.5\%$ and $13\pm 2.4\%$ and exterior shell composition of $26\pm 2.8\%$, $11\pm 2.0\%$, $54\pm 1.3\%$ for Co, S, O respectively. Scale bar, 1 nm.	37
3.8	Measuring relative concentration for experimental and synthetic datasets. Pixel intensity histograms for an experimental Zr (green), Ni (blue) and synthetic Ga (red) concentration maps. The standard deviation (σ) for each element is reported. The raw and reconstructed EDX maps are illustrated inside of the plot. Ground truth concentrations are highlighted by the respective colored triangles above the top axis. Stable convergence for the three components in the cost function: model term (orange), data fidelity (magenta), and regularization (turquoise) are illustrated in the inset. Qualitatively the convergence is identical for all three example datasets. Zr and Ni scale bars: 5, 10 nm, respectively.	38

3.9	The influence of γ for measuring chemistry Histogram of intensities for the Sulphur, Cobalt and Oxygen distributions when γ varies between 1.0, 1.6, and 1.8. We see for values close to 1.7, the specimen's relative concentrations almost perfectly matches the distribution in Fig. S5. However if gamma is far off from what is experimentally feasible ($\gamma = 1.0$), the measured concentrations are incorrect. Scale bar, 30 nm.	39
3.10	Estimating dose requirements for accurate chemical recovery. a) A RMSE map representing the reconstruction error as a function of multiple spectroscopic and HAADF SNR. Brighter pixels denote results containing the incorrect concentrations from the ground truth. b) Visualization of three points on the phase diagram corresponding to increasing ADF / chemical electron dose. c) A plot of average standard error vs. RMSE demonstrating the two metrics are linearly correlated. d) The 3D model for generating synthetic chemical and ADF projections.	40
3.11	Uncertainty Maps for Synthetic CoS Recovery Visualization of the standard error (StdErr) maps for the synthetic CoS core shell simulation in Fig.. As we increase the electron dose, the average StdErr pixel values for all three elements decreases. $\langle \text{SNR} \rangle$ denotes the average SNR between the chemical and HAADF modality.	41
3.12	Recovering Missing Elements for $\text{Bi}_{0.35}\text{Sr}_{0.18}\text{Ca}_{0.47}\text{MnO}_3$ (BSCMO) a) Raw EELS spectrum for a BSCMO system shows Ca, O, and Mn core-loss edges, however, the Bi and Sr core edges were not measured in this energy range. b) The simultaneously collected HAADF image with the inset on the bottom right highlighting the A and B sites' composition. c) RGB overlay highlighting the chemical distribution. d) The raw EELS maps. e) The recovered chemical maps. The purple circles highlight Ca vacancies. Here, the Bi and Sr were treated as a single unknown signal since they are known a priori to occupy the same lattice sites. Although chemical maps of Bi and Sr are missing, multi-modal reconstruction shows considerable improvement to the SNR and quality of the Mn, Ca, and O distributions. Moreover, we see the recovered Bi + Sr map anti-correlates with Ca as expected, however the small intensity between lattice sites is a minor artifact that arises without spectroscopic measurement of Bi and Sr. Scale bar, 2 nm.	43
4.1	Nanoscale recovery of $\text{Au-Fe}_3\text{O}_4$ nanoparticle superlattice. a) Schematic highlighting the linked HAADF and EELS modalities for chemical tomography. HAADF projection images are collected at every tilt increment while core-loss EELS spectra are sparsely acquired every few tilts. b) The fused multi-modal reconstruction for the specimen's Fe $L_{2,3}$ (turquoise), O-K (turquoise), and gold $M_{4,5}$ edge (yellow). c) Chemical overlay of the superlattice nanoparticles over the entire 115 nm field of view. Scale cubes, 5 nm^3	46
4.2	Multi-Modal Reconstruction of $\text{Au-Fe}_3\text{O}_4$ Nanoparticles Inside a Carbon Support. The complete reconstruction of the $\text{Au-Fe}_3\text{O}_4$ nanoparticle superlattice inside the carbon matrix (highlighted in purple). Scale cube, 10 nm^3	47

4.3	<p>Nanoscale recovery of Co₃O₄-Mn₃O₄ core-shell nanoparticles. a-c Raw EELS reconstruction for the Co (blue-green) and Mn (orange) L_{2,3} core-loss edges. d-f The HAADF tomogram of Co₃O₄-Mn₃O₄ nanoparticle tracks the structure of the specimen but fails to describe materials chemistry in 3D. g-i The fused multi-modal reconstruction. Scale cubes, 25 nm³. a,d,g Representation in Fourier space of the projections used to reconstruct the tomograms. j Fused multi-modal tomogram of a single Co₃O₄-Mn₃O₄ nanoparticle. Scale cube, 10 nm³. e A line profile showing the average intensity across the diameter of the particle.</p>	49
4.4	<p>Multi-Modal EELS Tomography of ZnS - Cu_{0.64}S_{0.36} Heterostructured Nanoparticles. a The HAADF reconstruction and 2D slice are shown on the left. b The fused multi-modal reconstruction illustrating CuS or ZnS-rich nanoparticles and oxidized shells. c 2D slices of the chemical reconstructions with the noisy traditional reconstructions highlighted on the left of each image. Scale bar, 50 nm. d The HAADF and fused multi-modal chemical tomogram for a smaller field of view. e Representative EELS spectra for the S, C, Cu, and Zn core loss edges.</p>	51
4.5	<p>Multi-Modal EDX Tomography of Cu-SiC Nanoparticles. a The HAADF reconstruction with a few tilt micrographs is shown on right. b The fused multi-modal reconstruction highlighting Cu nanoparticles embedded inside the SiC support and raw EDX maps are shown on the right. Scale bar, 50 nm. c The PSF of an individual 3 nm nanoparticle inside the SiC. A few 2D slices of the reconstruction are shown on the right where we see the structure is sharp along the orthogonal axis perpendicular to the missing wedge and approximately a 20% reduction in resolution along the missing wedge direction. Scale bar, 2 nm. d EDX spectra for a single tilt.</p>	52
4.6	<p>SNR Dependency for Successful Fused Multi-Modal Recovery. a The initial corrupted chemical distributions for oxygen in the CuO-CoO synthetic dataset with increasing SNR. b A heat map expressing the relationship between average reconstruction error and SNR for either modality (HAADF or Chemistry) when 11 chemical maps ($\Delta\theta = +12^\circ$) and 141 HAADF projections ($\Delta\theta = +1^\circ$) are available. c SNR plot highlighting the average NRMSE as a function of chemical SNR for reconstructions without any regularization or fusion (traditional tomography), without fusion (compressed sensing) and within the multi-modal framework. 2D slices from oxygen is shown on right.</p>	53
4.7	<p>Resolution and Electron Dose for Multi-Modal Electron Tomography. a Resolution and dose relationship for electron tomography approximates the best achievable resolution for each material [116, 124, 92, 163, 30] – assuming an image contrast of 80%[105]. Fused multi-modal electron tomography results in a much higher dose-efficiency which enables higher resolution at any dose limit. b Multi-modal (green) and conventional chemical tomography (blue). This relationship between dose and resolution assumes sufficient tomographic sampling is achieved (i.e. many projections)—in practice the actual resolution will be much lower. Dose limited resolution assumes the material is adequately sampled (i.e. Crowther and Nyquist relations)</p>	54

4.8	Historical Demonstrations of Multi-Element Chemical Tomography. The reported dose and Nyquist limited resolutions for the fused multi-modal (MM) reconstructions reported in this manuscript are compared to previous multi-element chemical tomography (CT) experiments [95, 78, 134, 15, 155, 126, 79]. Note, the actual achieved 3D resolution of previously reported chemical tomography may be lower than the Nyquist resolution.	55
4.9	Resolution Analysis of Au-Fe₃O₄ superlattice nanoparticles. a Fused EELS tomograms of Au-Fe ₃ O ₄ nanoparticles. Scale cube, 2 nm ³ . b Power spectral density of the Fe reconstruction along the principal axial directions shown on the right. Scale bar, 0.5 nm ⁻¹ . c Power spectral density profiles for k_x - k_y and k_x - k_z directions. d Line profiles of a 2.5 nm Au nanoparticle gives a resolution of 0.8 nm, 0.8 nm, and 1.1 nm along the x, y, and z directions.	56
4.10	HAADF Resolution Analysis of Au-Fe₃O₄ superlattice nanoparticles. a Fused EELS tomograms of Au-Fe ₃ O ₄ nanoparticles. Power spectral density of the HAADF reconstruction along the principal axial directions shown on the right. Scale cube, 2 nm ³ . Scale bar, 0.5 nm ⁻¹ . b Power spectral density profiles for YZ and XY planes. c Line scan profiles of a 2.5 nm Au nanoparticle give a resolution of 1.00, 1.03, and 1.01 nm along the x, y, and z directions.	56
4.11	Estimating Sampling Requirements for Accurate Recovery with Synthetic CoO/CuO Nanocubes. a An NRMSE map representing the reconstruction error as a function of the number of HAADF and chemical tilts. Brighter pixels denote results containing incorrect reconstructions from the ground truth. b Visualization of three points corresponding to conventional chemical tomography (reconstruction without the HAADF modality), and low-dose fused multi-modal electron tomography. c The 3D models used for generating synthetic chemical and ADF projections. Scale bar, 75 nm.	58
4.12	Hyperparameter Estimation with Bayesian Optimization for the CuO-CoO Nanocubes. a Bayesian optimization optimizes the data fusion cost function (shown above) when provided a given number of chemical and HAADF tilts. b 3D visualization of the ground truth Au decorated CuO/CoO nanocubes. Scale bar, 75 nm. c Bayesian optimization parameter selection landscape where each black dot represents one of the many attempts to find the minimum NRMSE.	59
4.13	Measuring 3D Concentration for simulated CuO-CoO nanocubes. Histograms of chemical concentrations for each voxel in traditional and fused multi-modal tomography reconstructions. For traditional chemical tomography, the accuracy improves with SNR or increasing projections. Multi-modal tomography maintains low error, especially for experimentally realistic conditions.	60
4.14	Measuring 3D Stoichiometric Concentration of Au-Fe₃O₄ Superlattice Nanoparticles. 3D chemical reconstructions for each element are shown with their corresponding voxel intensity histograms. The mean values and standard deviations are 46.4±15.1%, 54.6±15.3%, 100±0% for Fe, O, and Au, respectively. The expected stoichiometry of this system is 42.9%, 57.1%, 100%.	61

4.15	Assessing Convergence and Selecting Hyperparameters with Pareto Front Curves. a Pareto fronts illustrates the relationship between reconstruction quality and regularization parameters for multi-modal electron tomography. Depicted are the tradeoffs from three reconstruction evaluation metrics: the multi-modal, self-consistency and average NRMSE across all elements. We see the highest quality reconstruction (lowest NRMSE) occurs around the inflection point of the pareto front. b The three individual components in the cost function plotted throughout the multi-modal electron tomography reconstruction process illustrates smooth asymptotic convergence. Convergence should be confirmed for accurate reconstruction.	62
5.1	Real-Time Electron Tomography Workflow of a helical nanoparticle visualized on tomviz. a , Specimen projections are sequentially collected in an electron microscope across an angular range ($< \pm 75^\circ$) and continually passed to tomviz for reconstruction and live 3D visualization. b , As projections accumulate during the experiment, the reconstruction updates in real-time and resolution improves. Scale bar, 100 nm c , A high-quality tomogram is available for data interpretation upon the end of an experiment.	66
5.2	Demonstration of iterative reconstruction algorithms. a-c , Visualization of the Co ₂ P nanoparticle early, mid, and at the end of the reconstruction process. At the beginning, the underlying structure can partially be seen behind the excess of background intensity. In the middle of the process, sharp features begin to form. The final iteration converges to a tomogram visually similar to the input tilt series. Scale bar, 50 nm. e-g , Visualization of an atomic resolution FePt nanoparticle. The atoms in the TV nanoparticle are resolved with increasing iteration and its periodicity demonstrated with the fast Fourier transform (FFT). Scale bar, 1 nm. d, h , A plot of the normalized residual to demonstrate convergence.	68
5.3	Real-Time Simulation for the Raw FePt Nanoparticle. a , The raw FePt tomographic projection images used for the real-time tomography simulation. Scale bar, 2 nm. b) As the simulation progresses, the reconstruction quality improves and atoms become visible as early as 50% of the experiment. c) An atomic-resolution tomogram is available after all the projections are reconstructed.	69
5.4	Demonstration of live-WBP. Live tomographic reconstruction in tomviz shown through freeze frames during progression of a weighted back projection algorithm (left to right). This unique capability allows users to interact and analyze the 3D structure throughout reconstruction. In the actual software the reconstruction updates in real time. a , Live volume rendering of Au/strontium titanate (STO) nanocubes. b , Live volume rendering of platinum (Pt) nanoparticles on a carbon support. Scale bar, 50nm.	70
5.5	External and internal architecture of tomviz GUI. The tomviz platform is composed of a multi-threaded pipeline that synchronously handles tomographic and 3D visualization on separate threads. a , Tomviz monitors for recently acquired tilt projections within a directory and b , automatically reads new data into the pipeline. c , As tomographic reconstructions proceed, visualizations dynamically update and remain interactive for analysis.	72

5.6	Dynamic CS convergence. a) 3D volume rendering with yellow indicating high intensity (Au) and purple representing low intensity (SrTiO ₃). These tomograms were constructed under the dynamic (left) and traditional (right) CS framework. b) Plots of data distance (DD), total variation (TV) and RMSE for both the dynamic algorithm (black) and conventional (blue). c) 2D cross-sections of the final output for each 3D reconstruction and the test object.	73
5.7	Convergence criteria. Stable (blue) and unstable (black) convergence of RMSE, DD, and TV for the synthetic Au/SrTiO ₃ nanoparticle. The blue curve shows proper convergence when the reconstruction completes many iterations (125) prior to appending more projections. The black curve shows incomplete convergence when insufficient intermediate iterations are not completed and TV may drift away from the true solution (TV ₀).	74
5.8	Dynamic manipulation of data tolerance parameter. a) Tightening (reducing) ϵ decreases the regularization weight which gradually produces sharper 3D tomograms. b) Loosening (increasing) ϵ allows for there to be more regularization and produce smoother 3D tomograms. c) Plots demonstrating the progression of DD (gray) and RMSE (magenta) vs iteration. As the calculation progresses, scientists can manipulate regularization without having to reset the algorithm. The visually desirable solution is obtained when $\epsilon = 0.0225$, highlighted in green. Reducing ϵ below this value produces noisy reconstructions. Overestimating ϵ blurs away fine features such as the Pt nanoparticles and internal voids.	75
5.9	External and internal architecture of tomviz GUI. The tomviz platform is composed of a multi-threaded pipeline that synchronously handles tomographic and 3D visualization on separate threads. a , Tomviz monitors for recently acquired tilt projections within a directory and b , automatically reads new data into the pipeline. c , As tomographic reconstructions proceed, visualizations dynamically update and remain interactive for analysis.	76
A.1	Determining optimal λ parameter for simulated compressed sensing electron tomography with Bayesian optimization. a The BO optimizer iteratively samples λ values and performs a tomographic reconstruction to estimate the RMSE landscape. BO strategically determines the next point to sample after each measurement. b The final estimated landscape of RSME vs λ . The blue star represents the Bayesian optimization's minimum RMSE solution, and the yellow and red represent the minima for the DD and TV metrics, respectively. c-e Corresponding reconstructions: c lies in the traditional regime, d is the RMSE minimizing reconstruction, and e is an over-smoothed reconstruction.	83
B.1	Load the sample dataset Tomviz comes packaged with sample datasets that include projection images taken across many specimen tilts (i.e. a tilt series).	86
B.2	View the tilt images. When projections are loaded into tomviz it is automatically displayed as orthogonal slices. The parameters of each visualization mode can be set in the 'Properties' panel in the lower left. To view images along different tilts, use the 'Slice' slider (red box) or manually enter the image you want to view.	87

B.3	Align the Tilt Images. The sample tilt series is slightly misaligned. We can automatically correct misalignments by centering the images with the Center of Mass method (#1), rotating the tilt axis (#2), and shifting its location (#3) so its center and parallel to x-axis.	88
B.4	Select the reconstruction algorithm for live visualization. The tomography menu contains many algorithms for electron tomography. The reconstruction sub-section contains all the iterative and direct algorithms for tomography. Here we will be selecting the Simultaneous Iterative Reconstruction Technique (SIRT).	88
B.5	Enter parameters for the SIRT algorithm	89
B.6	Visualizing Elements in the Data ‘Pipelines’. Once a dataset has been loaded and live reconstruction begins, the pipeline can be populated with visualization modules of choice. a , By default, tomviz will continue displaying the input dataset after the reconstruction is initialized. b-c , To visualize updates for the live volumetric process, visualization modules should be present below the ‘Reconstruction’ dataset icon. . . .	90
B.7	Visualizing a Live-SIRT reconstruction. The tomviz graphical user interface for 3D visualization contains a variety of visualization tools for analyzing the 3D structure of specimens. 2D and 3D visualization modules can be selected from the top left menubar. Data transformations and visualizations are recorded in the Data Pipelines column on the left side for reproducible workflows. A histogram of voxel intensities is displayed on the top center where the black line represents the opacity map. Users can exit the reconstruction early by pressing the ‘Cancel’ or ‘Ok’ (Done) button. . . .	91
B.8	Load an empty dataset to begin real-time tomography.	92
B.9	Initialize a Real-Time Tomography Reconstruction.	92
B.10	GUI for real-time tomographic reconstruction.	93
B.11	Copy individual projections into monitor directory. Tomviz responds to changes in monitored folders once files are saved into the directory.	94
B.12	Specimen drift from eucentric height. a , If the eucentric height for the goniometer is incorrect, tilting the specimen holder causes the specimen to sweep along an arc around the tilt axis. b , After correcting the stage position, the specimen drift is reduced significantly.	95

LIST OF APPENDICES

A Automated Hyperparameter Tuning With Bayesian Optimization 82

B User Manual for Real Time Tomography on tomviz 85

LIST OF ACRONYMS

ART	Algebraic Reconstruction Technique
BF	Bright Field
CTF	Contrast Transfer Function
CS	Compressed Sensing
EELS	Electron Energy Loss Spectroscopy
EDX	Energy Dispersive Spectroscopy
FFT	Fast Fourier Transform
GUI	Graphic User Interface
HAADF	High-Angle Annular Dark-Field Detector
MBIR	Model Based Iterative Reconstruction
NRMSE	Normalized Root Mean Square Error
RGB	Red Green Blue
RMSE	Root Mean Square Error
SIRT	Simultaneous Iterative Reconstruction Technique
SNR	Signal-to-Noise Ratio
STEM	Scanning Transmission Electron Microscope
TEM	Transmission Electron Microscope
TV	Total Variation
TV_{min}	Total Variation Minimization
WBP	Weighted Backprojection

ABSTRACT

Microscopists relentlessly strive to observe structure at the highest possible resolution. Over the past century, there has been a transition away from photonic wave sources, including light and X-rays, towards harnessing the potentials of electron optics. Whereas optical microscopes encountered their inherent far-field diffraction constraints long ago, thereby limiting characterization to the micron or nanoscale, electron microscopes can now routinely image materials by using accelerated electrons.

The pursuit of achieving high-resolution micrographs has significantly advanced through the development of aberration correctors. These devices utilize multi-pole magnetic lenses to rectify imperfections in the electron probe. As a result, electron probes can now reach sub-angstrom length scales, improving spatial resolution. This relentless pursuit of improving resolution played a vital role in expanding our understanding of materials and biological structures at the atomic and nanoscale levels.

Despite the remarkable progress in corrector hardware, there remains an increasingly need for advanced algorithms capable of efficiently extract valuable information from sparse low-dose acquisitions. The challenges are particularly pronounced in both atomic-scale measurements and 3D imaging, where substantial radiation dosages may exceed the dose tolerance of materials. This necessity assumes an even greater significance for chemically sensitive measurements, wherein the likelihood of capturing inelastic scattering events markedly diminishes.

To mitigate beam damage, reducing electron exposure is a common strategy but can lead to a decrease in signal-to-noise ratio (SNR) – hindering the analysis process. The extent radiative damage is particularly problematic for tomography experiments, where designing experimental configurations suitable for a material's tolerance is crucial for successfully obtaining accurate 3D measurements.

Over the past few decades, efforts to enhance image SNR have occurred simultaneously with developments in detector sensitivity and algorithms that faithfully reveals hidden structure lost in noisy signals. By using the concept of 'compressive sensing,' it is possible to reduce dose through a computational procedure that assumes signals should be sparse in some domain. More importantly, compressive sensing provides theoretical guarantees for recovering signals. This thesis extends this framework into the field of electron microscopy and uncovers several applications that benefit.

The dissertation introduces image processing algorithms that enables such sparse chemical and 3D structural imaging of materials. These algorithms are designed to maintain maximal quality with minimal radiative dosages. The development of several computational tools designed for the visualization of high-resolution of material's 3D structure and chemistry are presented in this thesis; these tools are applied to several physical systems including catalytic, semi-organic and oxide systems from projections micrographs collected by scanning transmission electron microscopes. Most of the methods presented in this thesis have already been published in several journals include *Nature Communications*, *npj Computational Materials* and *Ultramicroscopy*. Overall, this research aims to push the boundaries of electron microscopy and unlock new insights into the 3D structure and chemistry for a wider class of materials.

The first half of this thesis relates to image recovery and restoration. Chapter 2 provides an introductory overview of electron microscopy, image formation, convex optimization, and tomography algorithms. Chapter 3 presents the framework to eliminate structured stripe artifacts that typically are introduced by mechanical polishing sample techniques for the preparation transmission electron experiments. This elimination is achieved through constrained optimization using total variation regularizers. Chapter 4 explores imaging denoising recovery within the context of chemical maps. This is performed by the novel introduction of data fusion, which correlates structural and chemical features from simultaneously acquired signals collected in electron microscopes during spectroscopy experiments. We demonstrate image recovery at various dimensions spanning from micron to atomic length scales.

The latter half of the thesis centers on contributions to 3D imaging. Chapter 5 presents the fused multi-modal electron tomography framework for low dose 3D chemical imaging of materials. We demonstrate high resolution chemical tomography on several hard and semi-organic materials. Beyond application, we also explored the dose limitations and sampling requirements for this novel form of chemical imaging. Chapter 6 discusses the software and algorithmic developments for enabling real-time 3D imaging of materials. We demonstrate real-time tomography on helical chiral nanoparticles and the open-source platform that enables this experimentation for any researcher.

Conclusions and future work is presented in Chapter 7 with efforts for measuring the atomic scale distribution of matter for objects larger than the electron probes. This experimental technique enabled the measurement of hundreds of thousands of atoms, thus a new accomplishment for high resolution 3D imaging.

CHAPTER 1

Introduction and Background

Characterization lies at the heart of material science research. The discovery of crystal symmetries, quantification of defects, and inspection of new materials all demand reliable methods that measure structural properties across scales ranging from the microscopic to the atomic level. Modern scanning transmission electron microscopes (STEM) have revolutionized our ability to explore atomic-scale structure in a diverse range of materials (e.g., III-V Nitrides or perovskites) by using high-energy electrons that can resolve individual bond-lengths. This chapter provides an introduction for image formation in electron microscopes and its extension for 3D tomographic imaging. Additionally, we delve into the principles of convex optimization and its practical applications in image recovery.

1.1 Scanning Transmission Electron Microscopy

Electron microscopes deliver imaging resolutions that surpass optical microscopes. This advantage arises from fundamental differences between the utilization of light and electron waves. Optical microscopes rely on visible light waves whose smallest wavelengths exceed 400 nm. This inherently prevents light from resolving features smaller than half a micron. In contrast, the resolving power of scattered electrons fundamentally lies in its picometer wavelengths. In the STEM, a coherent electron source emitting from a sharp metal tip (electron gun) generates free electrons that are accelerated to voltages of 20-300 keV. To put this into perspective, at 80 keV the electron's velocity is roughly half the speed of light ($\lambda = 4.18$ pm) and it reaches about 80% of the speed of light at 300 keV ($\lambda = 1.97$ pm). However the electromagnetic lenses used to focus the electrons are often plagued by deviations that prevent the maximal achievable resolution.

Despite the advantages from accelerated electron waves, achieving maximal imaging resolutions is constrained by imperfections inherent to electron magnetic lenses. These lenses play a pivotal role in focusing accelerated electron waves, but frequently fall short of meeting ideal expectations. In 1936, Otto Scherzer demonstrated a fundamental resolution limit exists for rota-

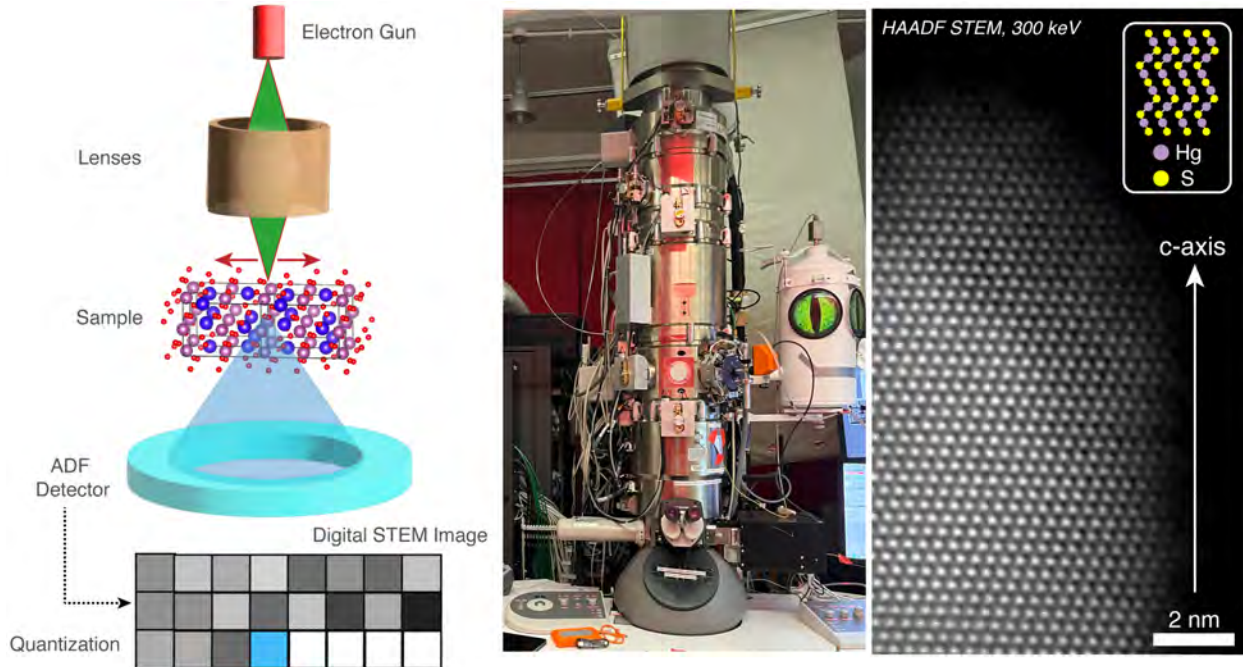


Figure 1.1: **Aberration-Corrected STEM Microscope.** Conceptual schematic of an aberration-corrected electron microscope with an HAADF image of a HgS nanoparticle shown on the right.

tionally symmetric magnetic lenses from the inherent spherical aberrations they introduce [140]. In general, geometric aberrations are connected to the aperture cutoff (described by the convergence semi-angle α) which constrains the extent to which a wave can be focused. Thus, geometric aberrations, characterized by phase errors across the aperture plane prevent the maximal usable convergence semi-angle (α). Simply put, the aperture cutoff defines the highest spatial frequencies that can be faithfully transferred to the resulting image while omitting dominant high convergent angle aberrations. In 1947 Otto outlined methods to circumvent this limitations by including non-rotationally symmetric lenses that can correct and minimize geometric aberrations, thus leading the pathway for modern microscopes that achieve sub-angstrom resolutions [141]. Nonetheless focusing and aberration correction occur in lenses positioned above the specimen, however the collection of scattering events provide image contrast in STEM micrographs.

In a STEM, the magnetic lenses focus and raster scan a small electron probe converged to a focal diameter of $0.5 - 50 \text{ \AA}$ across a sample. The electrons transmit through the sample and scattering events are subsequently collected by a detector centered on the optical axis located in the diffraction plane (Fig. 1.1a). At each probe position, the detector integrates all scattered electrons within a defined exposure time, converting the measured collision events into a digital pixel intensity. Two detectors commonly found in STEM microscopes are a disc geometry that integrates the forward scattered electrons at low angles to produce a bright-field signal and a hollow annulus ge-

ometry that integrates the high elastically scattered electrons yielding an annular dark-field (ADF) image. ADF-STEM is a popular imaging mode because it produces directly interpretable signals (Fig. 1.1b).

The contrast observed in ADF-STEM images can be approximated using an incoherent linear imaging model. With the incoherent model, the resulting ADF-STEM image ($I(\mathbf{x})$) is characterized as a convolution between the focused probe with an object ($f(\mathbf{x})$). Mathematically this is expressed as: $I(\mathbf{x}) = |\psi(\mathbf{x})|^2 \otimes f(\mathbf{x})$. The object function represents the specimen described through the probability of scattering to large angles in the ADF detector: $f(\mathbf{x}) \sim \int D(\mathbf{k}) \frac{\partial\sigma(\mathbf{x})}{\partial k_s} d^2k_s$ where $\frac{\partial\sigma(\mathbf{x})}{\partial k_s}$ is the partial cross section of the specimen for scattering to angle k_s and $D(\mathbf{k})$ is the function that expresses the annulus detector:

$$D(\mathbf{k}) = \begin{cases} 1, & \text{for } \alpha_{\min} \leq \lambda|\mathbf{k}| \leq \alpha_{\max} \\ 0 & \text{otherwise} \end{cases} \quad (1.1)$$

To approximately achieve incoherent imaging, the key requirement is to have a sufficiently large detector radius. This minimum collection angle is primarily dependent on the electron probe. Adjusting the collection range to a desired amount can be easily accomplished by modifying the compression of the projector lenses. As a rough guideline, if $\alpha_{\min} > 3\alpha$, then the image will generally be incoherent [68]. Practically, the inner detector angle can be larger, extending to $\alpha_{\min} > 5\alpha$ to fully suppress diffraction contrast. This mode is often referred to as high-angle ADF (HAADF) imaging.

Consequently, the image contrast for ADF-STEM imaging is influenced by both the electron probe shape and material's partial cross section. Specifically, the partial cross section describes the angular distribution of incident electrons that undergo scattering due to the Coulomb potential of a stationary atomic nucleus. The earliest model for elastic scattering was used by Rutherford to explain the scattering of charged particles through the unscreened electrostatic field of a nucleus. Notably, in the scenario of Rutherford scattering, image contrast can exhibit variations of up to Z^2 . However, Rutherford scattering neglects the impact of electron cloud screening within the nuclear electrostatic field, leading to an overestimation of scattering at smaller angles.

ADF micrographs produce directly interpretable images of atoms, where intensity is related to Z -number. However the intensity only provides insight into relative chemistry. When specific chemical identification is required, we rely on inelastic scattering events.

1.2 Electron Spectroscopy

Measuring material chemical signatures are accessible through inelastic scattering events that occur between the microscope's fast incident and core electrons that surround atoms within the specimen. Whereas an elastic collision preserves kinetic energy and momentum, an inelastic collision conserves the total energy and momentum, a part of the kinetic energy being converted to atom–electron excitation. Within the scanning transmission electron microscope, two chemically sensitive spectroscopic methods are available: electron energy loss spectroscopy (EELS) and energy-dispersive X-ray Spectroscopy (EDX).

1.2.1 Electron Energy Loss Spectroscopy

High-energy electrons that pass through a material can engage in several inelastic processes that induce change the transmitted electron's kinetic energy. Electron energy loss spectroscopy (EELS) measures the interactions between the energetic incident electrons and solid-state materials that are thin. A magnetic prism below the sample deflects scattered electrons, yielding a spectrum that quantifies the number of transmitted electrons with specific energy losses (Fig. 1.3a). The collective energy distribution of inelastically scattered electrons provides the local chemical characteristics of the specimen.

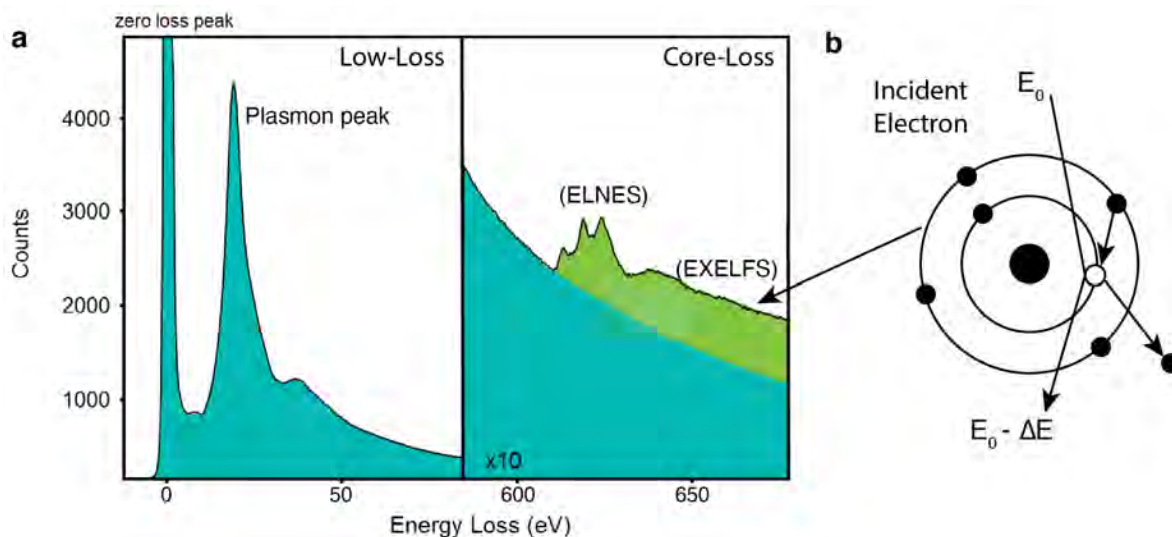


Figure 1.2: **EELS spectra and inelastic scattering.** **a** A sample EELS spectrum of MnO₂ [adapted from eels.info]. **b** A classical (particle) view of electron scattering.

The EELS spectrum contains three primary regions: the zero-loss peak, the low-loss region, and the core-loss region (Fig. 1.3a). The zero-loss peak encompasses the elastically scattered electrons that might have incurred energy losses below the detectable threshold. Analyzing the zero-loss

peak peak's width facilitates the determination of the probe's energy resolution (typically ranging from 0.1 - 2 eV) and can also be utilized to estimate the specimen's thickness. Coulomb interactions between the high-energy incident electrons and the specimen results in inelastic scattering that extend beyond the zero-loss peak. This region is characterized by scattering events involving outer shell electrons, thus leading to collective electronic excitations of plasmons and phonons. These surface plasmons describe the valence electron density and unfortunately can not measure chemistry in specimen.

Information about the elements present are accessible by measuring the material's core electrons transitioning to unoccupied states above the Fermi level. During this process, the shell electrons of atoms in the specimen move outside the attractive nucleus field leading to ionization of the atom (Fig. 1.2b). The core loss edges are represented by an increase in intensity – representing the ionization threshold, the energy that is approximately equivalent to its inner-shell binding energy. Core-loss edges can be divided into two regions: the energy loss near-edge structure, extending 30-50 eV above the edge onset and extended energy loss fine structure which extends 200-300 eV above the near-edge region. Both regions capture dipole transitions from a selected core-orbital to unoccupied states above the Fermi level. While generally extended energy loss is rarely applied in studies due to weak signal modulations, the fine structure from the onset peaks provides insights into the unoccupied density of states [59, 85].

1.2.2 Energy Dispersive X-Ray Spectroscopy

EDX is the measurement of X-rays emitted from specimens. When the focused electron beam strikes the specimen, it prompts the excitation of an inner-shell electron. This excited electron is ejected from its orbital, resulting in the formation of an electron vacancy or “positively charged hole.” An electron from a higher-energy outer shell can release energy in the form of X-rays to drop into this vacancy. These emitted X-rays carry the distinct energy signature, closely tied to the atom's elemental identity and chemical concentration.

A solid-state detector captures and records the emitted X-ray energies. Positioned at an angle relative to the specimen, this detector enables the separation and identification of X-rays based on their distinct energies (Fig.1.3b). Current EDX detectors commonly offers total solid angles of 0.9 steradian, allowing 7% of the emitted x-rays to be analyzed. Modern microscopes can use several detectors in unison, thus scaling the total solid angle accordingly. Notably, EDX serves a dual purpose of qualitative and quantitative analysis, offering insights into the chemical composition of a specimen. It empowers analysts to not only determine the presence of specific elements but also ascertain the relative concentration of each element within the sample. The EDX image intensity (I_a) is proportional to the sample thickness (t) and concentration (C_a) of element a . This

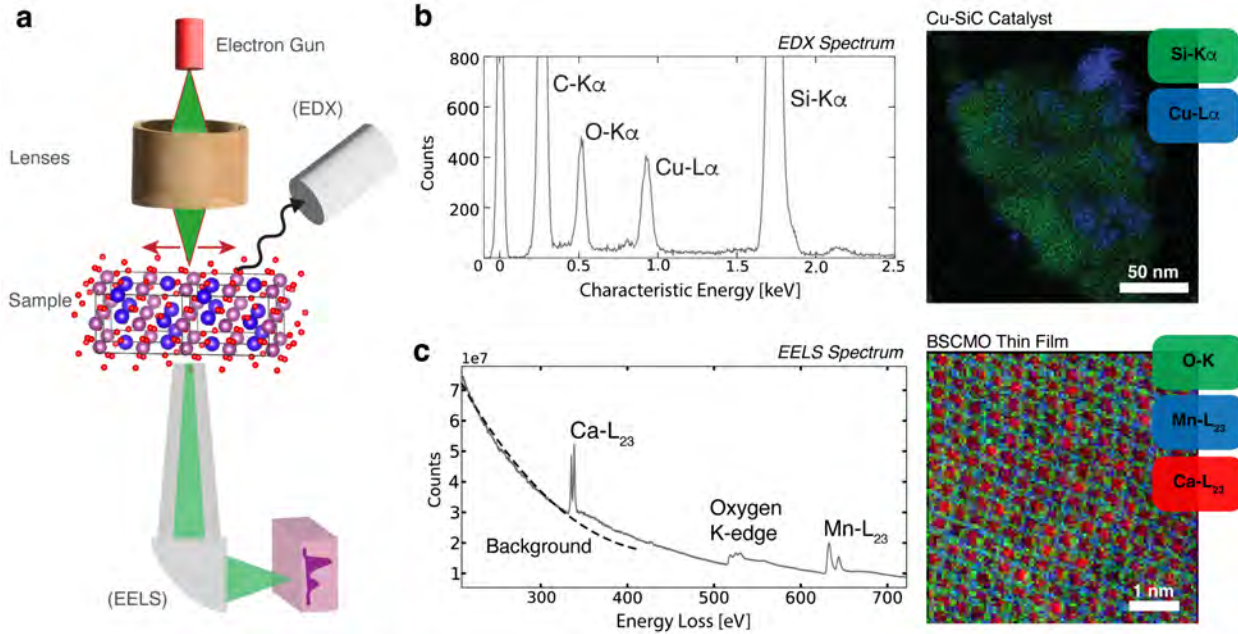


Figure 1.3: **Spectroscopy in the STEM Microscope.** **a** Conceptual schematic inelastic signal generation that occurs in the electron microscope. **b-c** Energy losses from transmitted electrons or characteristic energies from emitted X-ray correspond to chemical signatures that can be used to map the chemistry of materials (shown right).

relationship is estimated as: $I_a = DC_{apt}/\zeta_a$ where D_a is the total electron dose, ρ is the specimen density, and ζ_a is a sensitivity factor that is intrinsically linked to the ionization cross section and characteristic X-ray generation probability [166].

While EDX exhibits relatively poor energy resolution (approximately 120 eV compared to 0.1 eV for EELS), it is frequently favored due to its straightforward quantification methodology. However, this preference comes with the drawback of reduced sensitivity to lighter elements, as the X-ray fluorescence yield falls below 1% for $Z < 10$ [49]. Complexities arise when overlapping peaks can obscure meaningful EDX analysis. Especially in the case of light element quantification (e.g. K -peaks for carbon or oxygen), where heavier elements containing L -peaks within 100 eV can interfere. Nonetheless these edges could still overlap for EELS, but they can still be distinguishable from the better energy resolution.

A fundamental challenge in either method (EDX or EELS) lies in dealing with the low signal-to-noise ratio (SNR) that occurs from relatively small inelastic cross sections. The spectral measurements are limited by a Poisson process as detections are triggered from individual electron or X-ray counts arriving at the detector. Strategies aimed at improving SNR often involve increasing the electron dose or prolonging the acquisition time per pixel to increase the counts. However, these approaches carry the potential of accelerating irradiation damages or spatial drift leading to

image distortions.

In combination both inelastic and elastic signals (EELS, EDX and HAADF) can provide a comprehensive understanding of the atomic-scale structure of materials. This signals can be simultaneously acquired across various length scales, spanning from micron to sub-angstrom resolutions, to obtain projections at unique viewing angles.

1.3 Electron Tomography

Three dimensional characterization and visualization of nano-sized objects can be achieved through electron tomography. The genesis of tomography traces back to the mathematical underpinnings of reconstructing an object from its lower-dimensional projections, a concept first formulated by Radon in 1917 [129]. The technique enables the recovery of specimen structures by rotating the specimen around a fixed eucentric axis while acquiring projection images at several viewing angles (commonly referred to as a tilt-series). A 3D reconstruction of the specimen's volumetric structure is reconstructed from a series of 2D projections. Although the foundational mathematical principles were established early on, the first practical application emerged in 1956 when Bracewell employed them to reconstruct two-dimensional maps of microwave emissions from the sun [12]. Over time, this approach found its place within the electron microscopy community, where De Rosier and Klug reconstructed the helical structure of the T4 bacteriophage tail from a single TEM projection by leveraging prior knowledge about its helical symmetry [133]. Expanding on this progress, spherical symmetry was harnessed in 1970 to successfully reconstruct the negatively stained human wart virus and tomato bushy virus from multiple projections [32]. The milestone achievement in the evolution of electron tomography came with the 3D reconstruction of the low-symmetry fatty acid synthetase molecule in 1974 [72, 74].

To obtain a directly interpretable 3D volume (i.e. tomogram), each recorded image must satisfy the projection requirement. This principle requires that the image intensities must exhibit a monotonic correlation with specimen thickness [69]. In the case of bright field STEM, this criterion fails as signal intensity can dramatically change with tilt from diffraction contrast. On the other hand, ADF-STEM and chemically sensitive spectral maps effectively mitigate phase and diffraction contrast, resulting in image intensities that reflect variations in projected mass-thickness. Each of these modalities, traditionally acquired independently, can be employed to measure a specimen's 3D structure or chemistry. An illustration of a HAADF tomography experiment is presented in Figure 1.4.

Projection images can be mathematically described by the Radon transform. For simplicity, let's begin by reconstructing a 2D image from a 3D object represented by the function $f(x, y)$. The projection data ($p_\theta(t)$) for the sample tilted by angle θ and ray at scanning position $t =$




Figure 1.4: **Diagram of single-axis electron tomography.** Projections are recorded at regular tilt increments of θ around a single fixed axis. A typical experiment contains 70 - 140 projection images from $\pm 70^\circ$ separated by consistent 1 or 2° increments.

$x \cos \theta + y \sin \theta$ is given by:

$$p_\theta(t) = \int_L f(x, y) ds = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) \delta(x \cos \theta + y \sin \theta - t) dx dy \quad (1.2)$$

To fully appreciate the inherent challenges to accurately reconstruct our sample with sufficient sampling, we can concisely visualize our tomography experiment in Fourier space. The Fourier slice theorem is the fundamental principle for tomography, bridging the connection between the Radon and Fourier transform. It states that projections are equivalent to central planes in Fourier space (illustrated in Fig. 1.2). The Fourier transform of the projection is given by:

$$\begin{aligned} \mathcal{F}[p_\theta(t)] &= \int_{-\infty}^{\infty} p_\theta(t) \exp(-2\pi i \omega t) dt \\ &= P_\theta(\omega) \end{aligned} \quad (1.3)$$

Where ω is the frequency variable corresponding to t . Now let us consider the simplest form of the

Fourier slice theorem, in the case where $\theta = 0$:

$$\begin{aligned}
 F(k_x, 0) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) \exp(-2\pi i x k_x) dx dy \\
 &= \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} f(x, y) dy \right] \exp(-2\pi i x k_x) dx \\
 &= \int_{-\infty}^{\infty} p_{\theta=0}(x) \exp(-2\pi i x k_x) dx \\
 &= P_{\theta=0}(k_x)
 \end{aligned}
 \tag{1.4}$$

In the general case, we can rotate the coordinates and obtain:

$$\begin{aligned}
 P_{\theta}(\omega) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) \exp(-2\pi i \omega t) dx dy \\
 &= F(\omega \cos \theta, \omega \sin \theta) \\
 &= F(\omega, \theta)
 \end{aligned}
 \tag{1.5}$$

Essentially, the Fourier slice theorem tells us that $P_{\theta}(\omega)$ is $F(k_x, k_y)$ in polar coordinates ($k_x = \omega \cos \theta$ and $k_y = \omega \sin \theta$). By acquiring projections at various angles, we can effectively fill in the object's information in Fourier space, enabling us to reconstruct the true structure with a single inversion operation ($\mathcal{F}^{-1}[F(k_x, k_y)]$). Nonetheless a central challenge in this method arises from to the disparity that measurements exist on a radial slices, while the inverse Fourier transform requires data on Cartesian coordinates.




Figure 1.5: **Fourier Slice Theorem.** A projection of a 2D object in real-space is equivalent to the central plane of the object's full Fourier transform at the same angle (θ).

To circumvent the complications of direct Fourier inversion, tomography can leverage the Radon transform to produce tomograms. A fundamental linkage between Radon and Fourier space exists, as both domains provide theoretically equivalent approaches to tomographic reconstruction [162]. The inverse Radon transform (commonly referred to as back-projection) is expressed as follows:

$$\hat{f}(x, y) = \int_0^\pi p_\theta(t) d\theta \quad (1.6)$$

In this process, the intensity projected onto individual pixel is linearly integrated into the recovered object (\hat{f}) along a ray corresponding to the relevant angle. This integration contributes to the formation of object density. However, using the back-projection or inverse Fourier transform alone often produces inaccurate reconstructions, because the density of samples diminishes toward higher frequencies thus impeding resolution. We can combat the uneven sampling distribution with the weighted back-projection (WBP) method [18]. As its name implies, WBP encompasses two sequential steps: filtering and subsequently back-projecting the projections. Unfortunately while filtering can preserve fine features, the consequence from incomplete sampling still has a treacherous impact on resolution.

1.3.1 Resolution and the Crowther Criterion

Practical experimental limitations for electron tomography imposes significant limitations on achieving perfect sampling in Fourier space. Firstly, due to radiation dose limitations, we only have the liberty to acquire a finite number of measurements before the sample begins to degrade. Notably, this finite measurements manifests as an radial imbalance between the low- and high-frequency components. This limitation leads to reconstructions characterized by blurred features, as sample density diminishes at higher frequencies. Ultimately, imbalanced sampling in Fourier space adversely impacts resolution as fine spatial features are irrecoverable.

The spatial resolution of tomograms obtained from a single-axis tilt series exhibits anisotropic characteristics. Along the tilt axis, the resolution can be equivalent to the pixel size as defined by the Nyquist limit (assuming perfect alignment). However, in directions perpendicular to the tilt axis (y) the resolution can be reduced due to the discrete number of acquired projections (N). The best achievable resolution estimate was geometrically derived by Crowther as follows [33]:

$$d_y = \frac{\pi D}{N} \quad (1.7)$$

Eq 1.7 is known as the ‘‘Crowther criterion’’ which illustrates that achieving higher resolution for a given object size necessitates an increased number of uniformly sampled projections. Figure 1.6

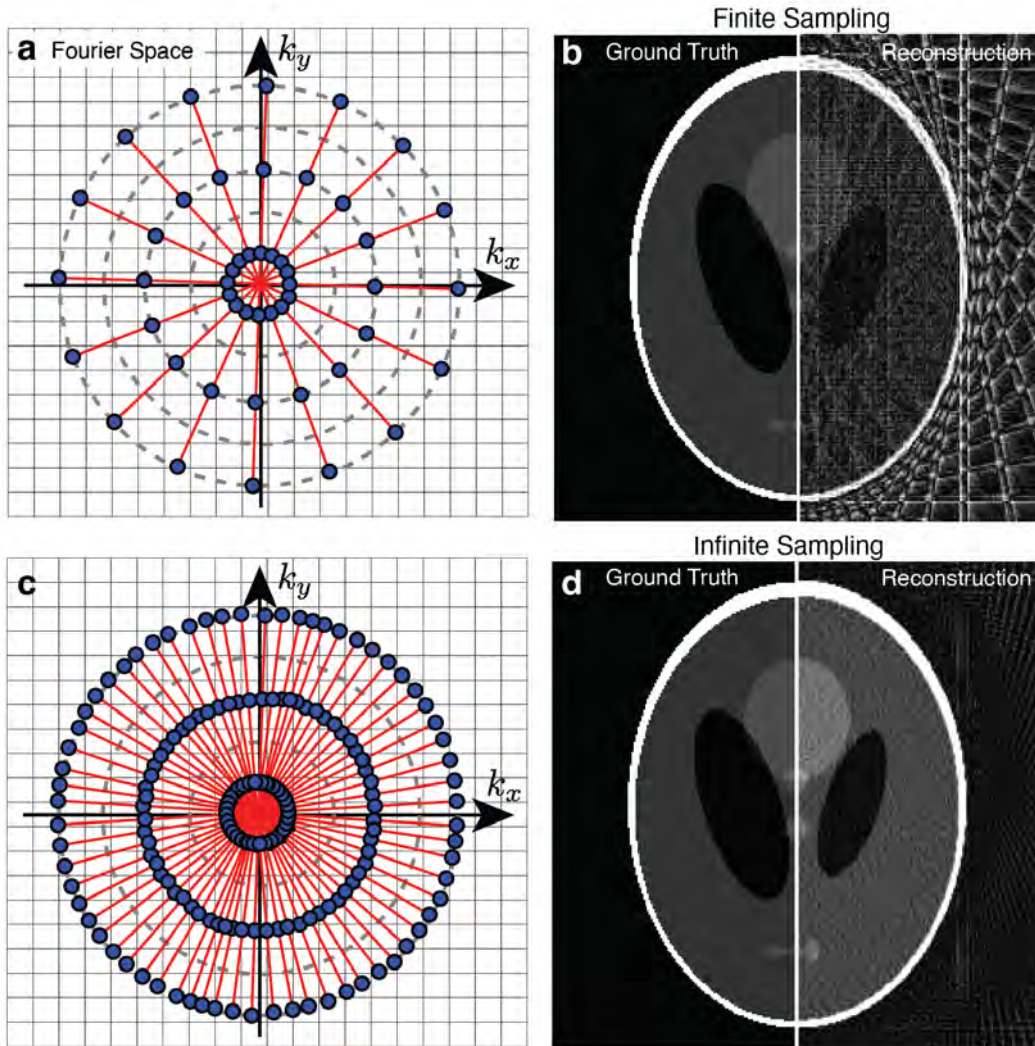


Figure 1.6: **Limited Angular Sampling.** The Fourier slice theorem states that every angular projection ($p_\theta(t)$) contains full information about $F(k_x, k_y)$. A finite, regular sampling of t translate into a finite sampling along the projection lines (red lines). A limited sampling with only 20 views (top row) is compared a fully sampled tilt series with 180 views (bottom row). The empty space between adjacent planes negatively impacts the reconstruction quality.

demonstrates the effect of limited angular sampling on a Shepp-Logan phantom. When compared to the reconstruction with uniform sampling (Fig. 1.6c,d), the lack of information in Fourier space from the limited views cripples the quality – leading to the presence of streak artifacts and boundary distortions that obscures the visibility of small pores (Fig. 1.6a,b).

1.3.2 Resolution and the Missing Wedge

Achieving a full tilt range within an electron microscope is often hindered by hardware limitations. These limitations result in an information gap along the k_z direction when measuring the 3D structures of materials. The confined space between lenses restricts extensive sample rotations, as electrons cannot transmit through the increased thickness at high tilt angles. In most scenarios, researchers can attain a maximum tilt angle (α) of 75° . Ultimately this incomplete sampling prevents the use of direct inversion methods as object recovery in electron tomography becomes an under-measured problem. As a result, features along the z -axis are stretched by (e_{yz}) that reduces the z -axis resolution [128]. In the case of 2D tomography, a missing wedge along the k_x -axis (Fig. 1.7a) manifests as elongation along the y -axis (Fig. 1.7b).

$$d_z = d_y e_{yz} = d_y \sqrt{\frac{\alpha - \sin \alpha \cos \alpha}{\alpha + \sin \alpha \cos \alpha}} \quad (1.8)$$

Beyond limitations posed by the missing wedge, resolution can also be influenced by various experimental factors including sample drift, alignment discrepancies, microscope stability, and beam-induced damage. Fortunately practical resolution estimation has proven to be more an intricate process due to the advancement of iterative algorithms designed to enhance overall quality by formulating the process as an inverse problem.

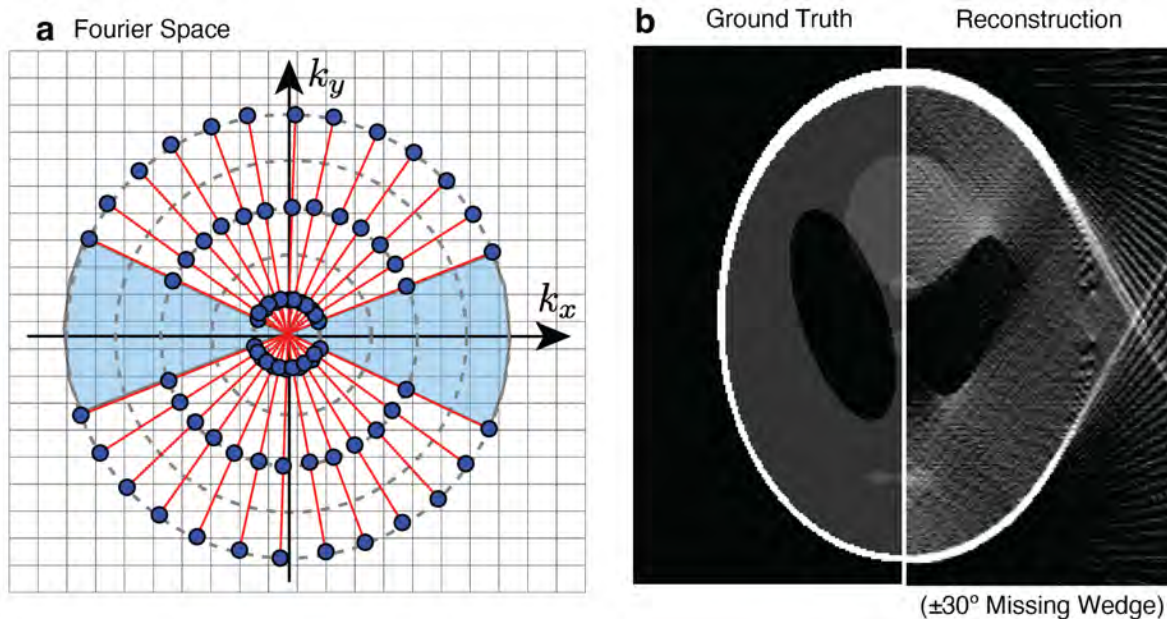


Figure 1.7: **The Missing Wedge and Limit Angular Sampling.** The ground truth and sample tomography reconstruction of a Shepp-Logan Phantom using 40 projections with a 60° wide missing wedge (highlighted in blue), which suffers from elongation artifacts along the vertical direction.

1.3.3 Electron Tomography as an Inverse Problem

Inverse problems is the process of recovering a signal from a collection of observations by modeling the physics for generating measurements. Tomography and image restorations (e.g. denoising or denoising) are two examples of inverse problems, where objectives are to either recover the 2D or 3D specimen morphology or remove noise in an image. We can model the inverse problem as a system of equations that expresses the outcome of measurements, given a complete description of the physical system of interest. In case of tomography, this would involve generating projections from the recovered object.

Iterative algorithms prove to be highly effective for solving large systems of linear equations, particularly those requiring regularized to address noisy data. In the context of linear systems of equations, our discrete inverse problem can be expressed as: $\mathbf{Ax} = \mathbf{b}$ where \mathbf{x} signifies the recovered object, \mathbf{b} represents the measured projections and $\mathbf{A} \in \mathbb{R}^{m \times n}$ stands as the matrix modeling the measurement process, often referred to as the sensing or measurement matrix. In the case of tomography, \mathbf{x} would be the recovered volume or image of our specimen and \mathbf{A} is the Radon transform.

Unfortunately insufficient measurements for electron tomography experiments complicates the reconstruction process. In cases where the problem is fully described, (i.e. sufficient projections over the entire angular range), we can execute simple matrix inversion ($\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$). The direct matrix inversion approach accurately recovers the object when the inverse problem is well-posed. Hadamard's conditions establish the prerequisites for a properly described problem, entailing the existence of a unique and stable solution resilient to alterations in data. However, problems failing to meet these conditions (e.g., \mathbf{A} having fewer rows than columns) lead to nontrivial nullspaces, thus invalidating the uniqueness condition. Thus, when insufficient projections are available, the inverse problem is ill-posed.

Handling an ill-posed inverse problem necessitates a modification that yields a unique and stable solution. One approach involves transforming the recovery process from a linear system ($\mathbf{Ax} = \mathbf{b}$) into a least-squares problem, which results in solutions approximating the unknown solution ($\hat{\mathbf{x}}$):

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{Ax} - \mathbf{b}\|_2^2 \quad (1.9)$$

To solve the optimization problem stated in Eq. 1.9, we need to use numerical algorithms that determine the given solution. Within the literature, a variety of optimization algorithms are available. Among the simplest are the first-order gradient descent algorithms, which follow the steepest gradient direction. Extending this approach, the Simultaneous Iterative Reconstruction Technique (SIRT) incorporates a projection operator to enforce constraints that align with realistic conditions. The most common constraint is non-negativity which assumes the recovered object possesses pos-

itive and real-valued attributes. The pseudo-code for SIRT is given as:

```

Initialize:  $\mathbf{x}_0 = \mathbf{0}, L_A = \|\mathbf{A}\|_2^2$ 
for  $k = 1, \dots, N_{\text{iter}}$  do ▷ Main Loop
     $\mathbf{x}_k = P_{\mathcal{C}}\left(\mathbf{x}_{k-1} - \frac{1}{L_A}\mathbf{A}^\top(\mathbf{A}\mathbf{x}_{k-1} - \mathbf{b})\right)$ 
end for
Return:  $\mathbf{x}_{N_{\text{iter}}}$  ▷ Return the Final Reconstruction

```

where $P_{\mathcal{C}}$ is a projection onto a convex set \mathcal{C} like non-negativity and L_A is the Lipschitz constant. Computation of the Lipschitz constant is useful for estimating appropriate step sizes that ensure monotonic decay. Variants of SIRT consist of normalizing the gradient by the rows or columns of \mathbf{A} (e.g., Cimmino’s method or component averaging). An alternative approach, the algebraic reconstruction technique (ART), also known as the Kaczmarz algorithm, involves only one row from the measurement matrix (a_i) per iteration. While each SIRT iteration entails a single matrix vector multiplication operation, ART follows sweeping process that either passes through all the rows sequentially or randomly. The pseudo-code for ART is given as:

```

Initialize:  $\mathbf{x}_0 = \mathbf{0}$ 
for  $k = 1, \dots, N_{\text{iter}}$  do ▷ Main Loop
    for  $i = 1, m$  do ▷ Cycle Through all the Rows
         $x^{k,i} = P_{\mathcal{C}}\left(x^{k,i-1} + \omega \frac{b_i - a_i^T y^{k,i-1}}{\|a_i\|_2^2} a_i\right)$ 
    end for
end for
Return:  $\mathbf{x}_{N_{\text{iter}}}$  ▷ Return the Final Reconstruction

```

Unfortunately, these first-order methods are known to exhibit slow convergence rates ($\mathcal{O}(1/k)$) thus requiring several hundreds or thousands of iterations to obtain the final solution. Here, the big-O notation for worst-case convergence means a decay rate proportional to $1/k$. Modifications to the first-order gradient methods can accelerate the convergence rate through the incorporation of momentum. At a high level, momentum leverages previous steps during reconstruction as prior information for the current step to provide a multiplicative boost in convergence. We can include momentum into standard gradient descent optimization through the addition of an extrapolation step after the gradient update.

Conceptually the momentum term nudges each iterate further down the parabolic landscape by descending further along the previous update’s direction and dampens potential oscillations. While the gradient descent algorithm is simple and guarantees monotonic improvement for every iteration, the fixed learning rate can slowly progress toward the local solution. Increasing the learning rate poses the risk of oscillating in regions of high curvature. Momentum-based descent adaptively scales the step size in response to the local curvature per iteration. The advantages of momentum

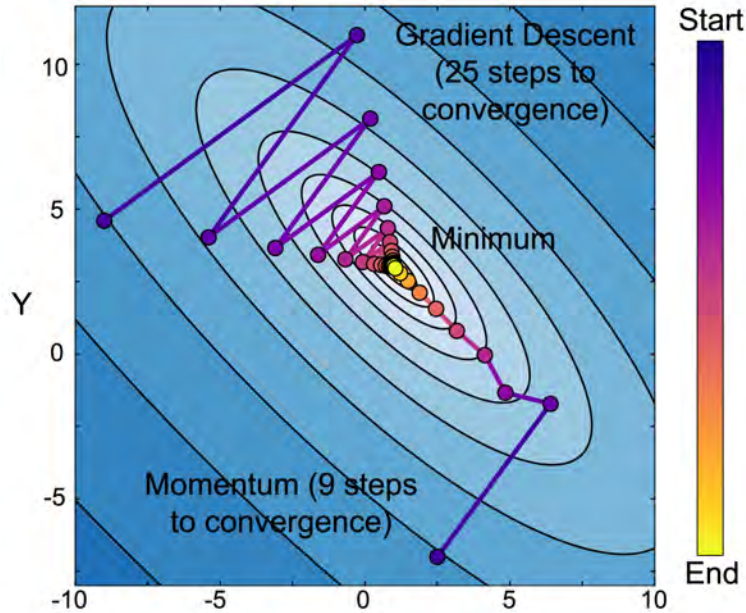


Figure 1.8: **Momentum Accelerated Optimization.** Traditional gradient descent (GD) juxtaposed with Nesterov accelerated gradient descent (NAG). From comparable starting points, NAG converges ~ 3 times faster, requiring 9 steps to reach the minimum, while GD required 25 steps.

can be seen by seeking the global minimum of a synthetic landscape known as Booth’s Function (Fig. 1.8)—a convex function often used to test optimization performance [80].

For the Booth landscape demonstration, standard gradient descent requires roughly 25 iterations to reach the minimum due to a slowdown near the flat minimum. With momentum, the descent retains speed and only requires $\sim 40\%$ of steps to reach the solution (9 in total). This simple modification remarkably achieves an optimal quadratic convergence rate [6]. Specifically, Nesterov has shown analytically that the convergence rate of these classic first-order methods can be sped up from $\mathcal{O}(1/k)$ to $\mathcal{O}(1/k^2)$ after k iterations with the incorporation of momentum [118].

1.3.4 Improving Electron Tomography with Sparse Optimization

Although iterative methods offer a versatile approach for tackling large-scale linear systems, the challenge lies in the ill-posed nature of the measurement matrix – especially in experiments where minimal projections are available. In such cases, the resulting solution (\mathbf{x}) becomes highly susceptible to errors present in the raw measurements (\mathbf{b}). Worse, the complexity is exacerbated by the fact that measurements are often not perfect, but rather are contaminated by noise and error ($\mathbf{b} = \hat{\mathbf{b}} + \epsilon$).

To address the intricacies of ill-posed problems, introducing regularity helps promote the production of unique and stable solutions. Regularization helps facilitate accurate recovery from

reduced measurements by incorporating any available information about the object prior to acquiring any data. The regularizer should be chosen such that desirable recovery is encouraged. For instance, the regularizer can be a p -norm $\left(\|\mathbf{x}\|_p = \left(\sum_{i=1}^n |x_i|^p\right)^{1/p}\right)$ of the signal itself or involve applying a transformative operation (e.g., wavelet transform or finite difference operator) on the signal. The choice for the p -norm influences the type of signals that are favored, whether they emphasize differences in appearance or prioritize smoothness. Overall, the unconstrained regularized cost function comes out as the following form:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{Ax} - \mathbf{b}\|_2^2 + \lambda \mathcal{R}(\mathbf{x}) \quad (1.10)$$

Where λ is a regularization parameter that balances the trade-off between the data fidelity and regularization ($\mathcal{R}(\mathbf{x})$). Alternatively, we can express Eq. 1.10 as a constrained optimization where a restriction is placed on the data fidelity term. If for example the noise level is known, the ε can be selected directly on this knowledge.

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \mathcal{R}(\mathbf{x}) \text{ s.t. } \frac{1}{2} \|\mathbf{Ax} - \mathbf{b}\|_2^2 < \varepsilon \quad (1.11)$$

1.3.5 Compressed Sensing Tomography with Total Variation Regularization

Over the past two decades there has been a growing interest in regularizers designed to enforce sparse recovery. This interest has been driven by the mathematical principal of compressive sensing (CS) – which guarantees perfect (or sometimes near-perfect) signal recovery [40]. In this context, signal sparsity refers when the number of non-zero values is significantly lower than the total number of elements – often quantified using the ℓ_0 -norm ($\|\mathbf{x}\|_0$). To solve the CS problem, we modify Eq. 1.11 with the ℓ_0 -norm as the regularizer ($\|\mathbf{x}\|_0$) and the hard constraint:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \|\mathbf{x}\|_0 \text{ s.t. } \mathbf{Ax} = \mathbf{b} \quad (1.12)$$

Unfortunately testing all possible combinations that satisfy this constraint is computationally intractable and can be shown to be NP -hard. Nonetheless, this intractability can be circumvented by relaxing the ℓ_0 -norm using a more manageable approximation, commonly the ℓ_1 -norm, and allowing some degree of data inconsistency within the unregularized formulation (Eq. 1.10). The ℓ_1 -norm is renowned for promoting sparsity and convexity, rendering the optimization problem computationally tractable.

Tomograms can be expressed as sparse in various ways. One prevalent approach involves representing tomograms in the gradient magnitude domain, where relatively uniform regions are separated by distinct boundaries. An example of sparsity enforced by the gradient magnitude operator

is illustrated in Figure 1.9 on a HAADF micrograph of ZnS-CuS nanocrystals. While the HAADF image is composed of relatively large homogeneous regions, its gradient image is sparse as its composed of 14.9% nonzero pixels.




Figure 1.9: **Illustration of Sparsity in the Gradient Magnitude Domain.** A HAADF micrograph of ZnS-CuS nanoparticles and the magnitude of its spatial gradient shown right. While the HAADF image may not be sparse, its gradient image is.

Total variation (TV) has emerged as a widely used regularizer in low-dose tomography imaging challenges. This is due to its ability to sparsely represent the reconstructed object in the gradient magnitude domain. Originally introduced by Rudin in 1992 for image denoising within a continuous-domain framework [136], total variation has found utility in edge-preserving regularization and sparse recovery within the field of tomography. Noteworthy contributions in this domain include the works of Sidky and Pan in medical CT, as well as Leary and Midgely in electron tomography [154, 160].

TV norm penalizes abrupt transitions, thereby promoting smoother transitions within the signal. Two variants of total variation norms (isotropic and anisotropic) are commonly employed in image recovery tasks. The key distinction lies in how the isotropic norm treats all directions equally, thus yielding a globally smooth solution. While the anisotropic total variation preserves strong edges by adapting its smoothing based on the local orientation of edges. Consequently, edges and features are retained more effectively, as the directional aspects of gradients are taken into account.

The problem addressed throughout a significant portion of this thesis pertains to formulating the unconstrained optimization (Eq. 1.10) as the Total Variation Minimization (TVMin) problem:

$$\arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \lambda \|\mathbf{x}\|_{\text{TV}} \quad (1.13)$$

where $\|\mathbf{x}\|_{\text{TV}} = \|\nabla \mathbf{x}\|_{2,1} = \|\sqrt{|D_x \mathbf{x}|^2 + |D_y \mathbf{x}|^2}\|_1$ and D_i is the finite-difference approximation

along the axis i . Note this formulation expresses the isotropic total variation for a 2D image; in 3D an additional finite difference operation would be applied along the z -axis.

Alternatively we can use the anisotropic total variation which can be expressed as $\|\mathbf{x}\|_{\text{ATV}} = \|\nabla \mathbf{x}\|_1 = |D_x \mathbf{x}| + |D_y \mathbf{x}|$. In either case, we can express the total variation norm as $\|\mathbf{x}\|_{\text{TV}}$. We can use FISTA to solve Eq. 1.13 (shown below).

```

Initialize:  $\mathbf{x}_0 = \mathbf{y}_0 = \mathbf{0}, t_0 = 1, L_A = \|\mathbf{A}\|_2^2, \bar{\lambda} := \lambda/L_A$ 
for  $k = 1, N_{iter}$  do
     $\mathbf{y}_k = \mathbf{y}_{k-1} - \frac{1}{L_A} \mathbf{A}^\top (\mathbf{A} \mathbf{y}_{k-1} - \mathbf{b})$ 
     $\mathbf{x}_k = \arg \min_{\mathbf{y}} \|\mathbf{y}\|_{\text{TV}} + 1/(2\bar{\lambda}) \|\mathbf{y}_k - \mathbf{y}\|_2^2$ 
     $t_k = \frac{1}{2} (1 + \sqrt{1 + 4t_{k-1}})$ 
     $\mathbf{y}_{k+1} = \mathbf{x}_k + \frac{t_{k-1}-1}{t_k} (\mathbf{x}_k - \mathbf{x}_{k-1})$ 
end for
Return:  $\mathbf{x}_{Niter}$ 

```

▷ Main Loop

▷ Gradient Descent

▷ Proximal Operation for TV

▷ Momentum weighting

▷ Momentum update

▷ Return the Final Reconstruction

CHAPTER 2

Removing Stripes, Scratches and Curtaining with Non-Recoverable Compressed Sensing

1

Highly directional image artifacts such as ion mill curtaining, mechanical scratches, or image striping from beam instability degrade the interpretability of micrographs. These unwanted, aperiodic features extend the image along a primary direction and occupy a small wedge of information in Fourier space. Deleting this wedge of data replaces stripes, scratches, or curtaining, with more complex streaking and blurring artifacts—known within the tomography community as ‘missing wedge’ artifacts. Here, we overcome this problem by recovering the missing region using total variation minimization, which leverages image sparsity-based reconstruction techniques—colloquially referred to as compressed sensing—to reliably restore images corrupted by stripe-like features. Our approach removes beam instability, ion mill curtaining, mechanical scratches, or any stripe features and remains robust at low signal-to-noise. The success of this approach is achieved by exploiting compressed sensing’s inability to recover directional structures that are highly localized and missing in Fourier Space.

2.1 Introduction

Streaks, stripes, scratches and curtaining artifacts commonly degrade image quality in microscopy datasets. This broad class of highly directional artifacts arise from varying conditions during image scanning—or may be artifacts inherent to the specimen, but artificially introduced during sample preparation such as curtaining during ion beam milling or mechanical scratches from polishing techniques. As a result, these artifacts can plague micrographs across any length scale. At the mesoscale, 3D focused ion beam (FIB) tomography is limited by the streaks induced by mill curtaining [174]. At high-resolution, stripes appear in scanning transmission electron microscopy

¹This chapter is based on results from [147].

(STEM) from beam instability and is most noticeable when signal is low relative to the background—common to bright-field (BF) detectors, imaging thick specimens or beam current fluctuation. Even more broadly, stripe artifacts are seen in atomic force microscopes (AFM) [26], light sheet fluorescence microscopy (LSFM) [98], and even globally at km length scales in planetary satellite imaging [130]. When possible, these linear artifacts are best mitigated experimentally, however experimental solutions are often difficult or unavoidable.

Over the years, a few methods have been demonstrated for destriping images outside of electron microscopy. Statistical-based methods developed for multiple-sensor imaging systems in planetary satellites assume the distribution of digital numbers in each sensor should be consistent (i.e. histogram or moment matching) [130, 55]. However, these matching-based methods are highly limited by the similarity assumption and fail on single-sensor imaging systems. Alternatively, filtering-based methods suppress the presence of stripe noise by constructing a filter on a transformed domain with a Fourier transform [26, 25] or wavelet analysis [161, 115]. Unfortunately, filtering methods risk removing or suppressing useful structural information falling within the filter. The third approach treats the destriping issue as an ill-posed inverse problem. Prior knowledge is used to regularize an optimization problem [11] and separate the unidirectional stripes from the image [24, 1, 139]. A similar class of research, known as compressed sensing (CS), has become highly successful toward solving inverse problems with incomplete data by finding maximally sparse solutions — but has yet been applied to remove scratch and stripe artifacts.

In this chapter, we describe a compressed sensing inspired approach that can remove highly-directional artifacts and demonstrate applications for ion mill curtaining, mechanical scratches, and beam instability. Here, a wedge of information containing the stripe artifacts is removed in Fourier space and the specimen’s information is recovered using total variation (TV) minimization, which maximizes sparsity of the image’s gradient magnitude and preserves sharp edges. A data constraint is imposed to produce a stripe free image with near identical appearance in signal-to-noise. The algorithm effectively removes striping when the missing wedge encompasses all stripe artifacts (typically $5^\circ \sim 10^\circ$) and is relatively insensitive to noise.

2.2 Background

The simplest way to remove streaks and stripes is to delete corresponding planes (or wedges) of information in Fourier space, however in turn, this also degrades the image. To illustrate, Figure 2.1 shows a backscatter electron (BSE) image of a biomineral surface scratched during mechanical polishing. The aperiodic scratches that extend a vertical direction (Fig. 2.1a) are confined in an angular range (i.e., $\sim 5^\circ$) in Fourier space. Removing an information wedge (Figure 2.1b) can better estimate the true object. However, deleting information in Fourier Space introduces smearing,

elongation, and blurring (highlighted in the yellow circles). These missing wedge artifacts are well known to the electron tomography community where larger wedge sizes exacerbate smearing and elongation artifacts [110].

Recently, CS inspired approaches have been applied to tomography as a tool for recovering information in the missing wedge. It is possible to reconstruct models with high data-fidelity from sparse projections under the CS framework. TV minimization is widely used in image restoration because of its ability to preserve edges [127]. The optimization problem for this algorithm can be written as:

$$\arg \min_{\mathbf{x}} \|\nabla \mathbf{x}\|_1 \text{ s.t. } \Phi \mathbf{x} = \mathbf{b}$$

where \mathbf{x} and \mathbf{b} represent the reconstructed image and measured data, Φ is the measurement matrix and ∇ transforms the image to the gradient-magnitude (sparse) domain.

Compressed sensing has demonstrated high-quality signal recovery is possible from a minimal number of measurements (\mathbf{b}) by assuming maximal sparsity [40, 21]. CS solves inverse problems ($\Phi \mathbf{x} = \mathbf{b}$) by seeking the sparsest representation of the original object (\mathbf{x}) via ℓ_1 - norm optimization [117, 160]. The theory of CS requires two assumptions to be true: (1) the object must have a sparse representation in a known transform domain (i.e. to be compressible), (2) there must be incoherence (i.e., contain a high level of dissimilarity) between the sensing and sparse basis [103, 21]. Conveniently, experimental data is often sparse in certain domains, as reflected by the compressibility of real-world information [42]. Incoherence expresses the level of dissimilarity between the sensing basis and the sparse basis [93] and depends on both how information is sampled (or missing) and the structure of the information (i.e. specimen).

One popular sensing basis used in electron microscopy is the Fourier basis. Each plane in Fourier space represents a projection of lines or stripes in real-space [12]. Thus, extended stripe-like structures in an image are confined to a plane in Fourier space. This is illustrated in the vertical scratches of Figure 1a, which exists in a horizontal plane of the FFT. If there is a small angular range to the stripes, the planes will broaden out to a wedge. Unlike non-directional image features that are typically spread-out in Fourier space, all knowledge of stripes becomes lost by removing a well-chosen plane (or wedge) of information. Thus, without sufficient sampling of particular Fourier planes, linear types of features cannot be retrieved by the standard CS sampling strategy.

In this work we demonstrate that TV minimization advantageously fails to restore stripes and scratches confined within a missing wedge (low incoherence) – but recovers structural information that is relatively spread in Fourier space (high incoherence). This provides the basis to our destriping approach.

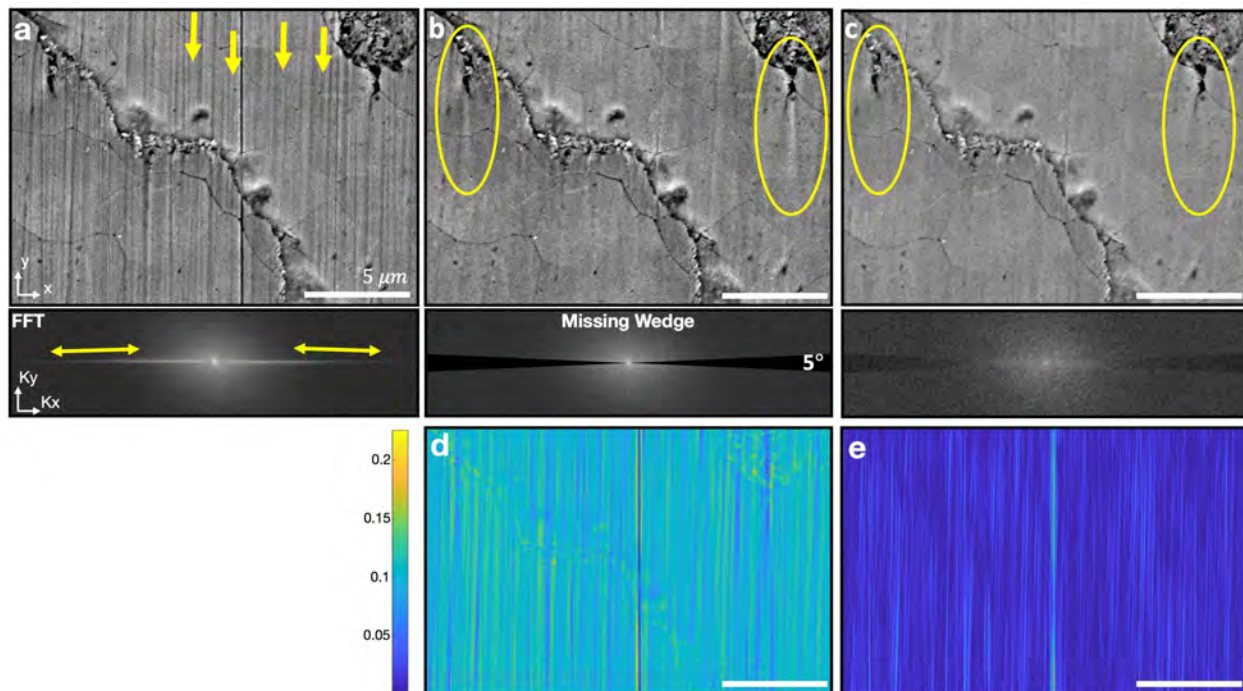


Figure 2.1: **Reconstruction of a scratched pearl surface.** **a)** The original 10 keV BSE-SEM image of a scratched biomaterial with its FFT below. **b)** The output from deleting a wedge of information in Fourier Space. **c)** The TV minimization reconstruction and its Fourier Space image. **d)** The residual between images (b) and (a). **e,** The residuals between images (c) and (a).

2.3 Results

Here we destripe images, by first removing the information in Fourier space containing unwanted artifacts, thus creating a ‘missing’ wedge. This ‘missing’ wedge of information is then recovered by minimizing the image’s TV using a gradient descent approach. Simply deleting information within the ‘missing’ wedge in Fourier Space (Figure 2.1b) provides a poor estimate of the clean image and creates elongation and blurring artifacts perpendicular to the ‘missing’ wedge (Fig. 2.1b). The residuals reveal the quality of the decomposition by calculating the absolute difference between the output and the original image. The residuals from a wedge Fourier filter, for Figure 1b, shows the removal of both unwanted stripes and useful structural information (Fig. 2.1d), a typical problem associated with filtering. However, recovering information in the ‘missing’ wedge with TV minimization produces an image without these artifacts (Fig. 2.1c). Figure 1e shows TV minimization only removed features that strictly pertain to the scratches. The gradient descent incorporates a convergence parameter (a) and requires many iterations to converge. Large scale simulations to investigate the optimal conditions for convergence found that for both experimental and simulated images the algorithm normally converges after ~ 150 iterations when $a = 0.1$.




Figure 2.2: **Reconstructions of BF-TEM and FIB images.** **a)** The original contrast reversal 300 keV BF-TEM image of an AlGaIn Quantum Well with horizontal intensity fluctuations and its FFT on the top right-hand corner. **b)** The reconstruction with the TV- minimization algorithm and its FFT. **c)** The residuals between images **(a)** and **(b)**. **d)** The original 30 keV secondary electron (SE)-SEM image of an aluminum sample with curtaining collected during a FIB tomography experiment. **e)**, The residuals between images **(d)** and **(e)**. Contrast was reversed in **(a)** and **(b)** for clarity.

Other stripe artifacts, such as beam instability common to BF-STEM images (Fig. 2.2a) or curtaining in FIB micrographs (Fig. 2.2d) can also be removed with our approach. Figure 2.2a. shows an atomic resolution image of an $\text{Al}_x\text{Ga}_{1-x}\text{N}$ quantum well suffering from unidirectional horizontal stripes caused by current fluctuations in the beam. The stripes prevent assessment of interface sharpness of the bright Al-rich layer. Similarly, waterfall/curtaining effects (Fig. 2.2d) — typically arise in FIB-tomography due to irregular milling rates caused by specimen inhomogeneity [71]. The presence of curtaining artifacts reduces interpretability for segmentation or object recognition [8]. Figures 2.2b-c and 2.2e-f demonstrate that our approach can separate the corrupted image into the clean and stripped components. For beam instability, the stripes are unidirectional (i.e. perfectly horizontal) and in this case, the missing wedge is along the x-axis only (no angular spread). In this limiting case, our approach converges to a one-dimensional problem with some similarity to that proposed by [11, 24] for Earth satellite and FIB data.

An optimal wedge size must be just large enough to remove all the stripe artifacts. Figure 2.5 shows the performance of removing scratches for 4° , 8° , and 15° missing wedges. If too small




Figure 2.3: **Reconstruction of a BSE-SEM micrograph under various horizontal (0°) ‘Missing’ Wedge widths.** **a)** The original 10 keV SEM image of a pearl sample with vertical scratches highlighted by green arrows. **b1-d1)** The first row shows the images with a ‘missing’ wedge of information as the blurring artifacts are prominent in the yellow circles. **b2-d2)** The second row shows the TV minimization reconstructions with the removal of the blurring artifacts and loss of features for large wedges highlighted in the red circles. FFT insets shown lower left.

of an angle is chosen, then scratch features remain (Fig. 2.5b1) before and after TV minimization recovery (Fig. 2.5b2). Thus, a sufficiently large missing wedge should be used to ensure that all information relating to the stripe, scratch, or curtaining has been removed in Fourier space (Fig 2.5c) and thus cannot be recovered with TV minimization. Increasing the missing wedge of information exacerbates image degradation before TV minimization recovery (Fig. 2.5b1-d1) but less noticeably after reconstruction. However, one must be cautious not to make the wedge unnecessarily large, as specimen features elongated along the direction of the scratches are susceptible to alteration. The vertical domain boundary highlighted by red circles in Figure 2.5 is preserved using an optimal angle of 8° but becomes blurred for much larger angles 15° .

Our approach is robust even down to low signal to noise ratios (SNR) and has the capability to preserve fine features from the original image. The algorithm consistently produced similar results to the ideal image in Fig. 2.5c2, after random Gaussian white noise was added to Fig. 2.4a. The standard deviation (σ) of Gaussian variance of the noise is related to SNR by $\sigma = \frac{\mu}{\text{SNR}}$ where μ is the image’s mean. Even though low signal-to-noise (Fig. 2.4d1) makes the scratches less




Figure 2.4: **TV minimization reconstructions for decreasing SNR.** **a)** The first row shows the images after Gaussian additive white noise is implemented and **b2-d2)** the second row shows the TV minimization reconstructions.

visible, the algorithm continues to recover the object. Noise preservation is expected from the strict data constraint ($\Phi \mathbf{x} = \mathbf{b}$) that preserves all information in Fourier Space outside of the missing wedge. The goal of this algorithm is not to filter or reduce noise, but to reproduce the original image free from scratches and stripes. Softening the data constraint has been used in electron tomography to smooth tomograms and reduce noise [83]. However, the reconstructed image can be over smoothed by TV minimization if the data constraint is too relaxed. Typically, in electron tomography experiments, missing wedges $\geq 25^\circ$ are recovered with CS. However, here we are showing that highly linear features may not be properly recovered. Because a hard data constraint only allows noise reduction within the small region of the missing wedge (no more than 6% in Fourier Space), the noise structure looks nearly identical to the test images as shown by the reconstructions in Fig. 2.4.

To further understand the relationship between SNR and the missing wedge, a quantitative study was performed on a DF-S/TEM micrograph of InGaN nanowires (Fig. 5b). Figure 5a shows a plot of the reconstruction's root mean square error (RMSE) normalized by the RMSE of the images before reconstruction (i.e. with missing wedge). Pixels with values below 1 indicate improvement from the reconstruction. The typical blurring and elongation artifacts are highlighted by the red circles in Fig. 5c. Similarly to the previous figure, the test image was exposed to random Gaus-




Figure 2.5: **Quantitative Study of SNR and Wedge Size.** **a)** A plot of the RMSE normalized by the error from various missing wedge sizes. Values below 1 indicate the reconstruction outperforms loss of information. Wedges below 8° , consistently achieves satisfactory performance at all SNR values above 10. **b)** DF-S/TEM micrograph (at 300 keV) of MBE grown InGaN nanowires with platinum nanoparticles coated on the surface. **c)**

sian white noise to verify the algorithm's performance across multiple SNRs. Our approach best reconstructs the test object for small missing wedges ($\leq 12^\circ$) and SNR values above 10. The reconstructions are visualized for two wedge sizes and SNR values (Fig. 5d-e) to verify the test image (Fig. 5b) is accurately reconstructed while preserving its noise.

2.4 Discussion and Conclusion

In this study, we demonstrate a compressed sensing-based approach to remove highly directional artifacts that commonly occur from ion mill curtaining, mechanical scratches, and beam instability. These highly directional aperiodic features can be removed without introducing blurring or elongation artifacts by removing wedges of information in Fourier space and recovering with TV minimization. Furthermore, our approach remains robust at low SNR. Overall, the qualitative results demonstrate that our technique achieves successful recovery, especially when small wedges are implemented ($\leq 10^\circ$). Within electron microscopy, these artifacts may become more common with the rising popularity of FIB sectional tomography that can contain curtaining or monochromated STEM where lens instability causes current fluctuations and image banding. The destriping technique investigated in this manuscript may also have application to a broader field of imaging techniques—such as atomic force microscopy or Raman spectroscopy.

Moreover, this destriping study provides insight to the recoverability of missing wedges in Fourier Space using compressed sensing. Specifically, we show highly directional features that exist within the missing wedge are not recoverable using TV minimization. This has implications for electron tomography, where incomplete experimental measurement commonly results in a missing wedge. Thus, we expect the recovery of highly extended, unidirectional features—such as interfaces—can become difficult for tomography. Acquiring projections perpendicular to interfaces or linear features should overcome this limitation.

CHAPTER 3

Imaging Atomic-Scale Chemistry from Fused Multi-Modal Electron Microscopy

1

Efforts to map atomic-scale chemistry at low doses with minimal noise using electron microscopes are fundamentally limited by inelastic interactions. Here, fused multi-modal electron microscopy offers high signal-to-noise ratio (SNR) recovery of material chemistry at nano- and atomic- resolution by coupling correlated information encoded within both elastic scattering (high-angle annular dark field (HAADF)) and inelastic spectroscopic signals (electron energy loss (EELS) or energy-dispersive x-ray (EDX)). By linking these simultaneously acquired signals, or modalities, the chemical distribution within nanomaterials can be imaged at significantly lower doses with existing detector hardware. In many cases, the dose requirements can be reduced by over one order of magnitude. This high SNR recovery of chemistry is tested against simulated and experimental atomic resolution data of heterogeneous nanomaterials.

3.1 Introduction

Modern scanning transmission electron microscopes (STEM) can focus sub-angstrom electron beams on and between atoms to quantify structure and chemistry in real space from elastic and inelastic scattering processes. The chemical composition of specimens is revealed by spectroscopic techniques produced from inelastic interactions in the form of energy dispersive X-rays (EDX) [37, 88] or electron energy loss (EELS) [156, 114]. Unfortunately, high-resolution chemical imaging requires high doses (e.g., $> 10^6$ e/Å²) that often exceed the specimen limits—resulting in chemical maps that are noisy or missing entirely [67, 35]. Substantial effort and cost to improve detector hardware has brought the field closer to the measurement limits set by inelastic processes [106, 89]. Direct interpretation of atomic structure at higher-SNR is provided by elastically scattered

¹The results presented in this chapter lead to a publication in *npj Computational Materials* [144].

electrons collected in a high-angle annular dark field detector (HAADF); however, this signal under-describes the chemistry [94]. Reaching the lowest doses at the highest SNR ultimately requires fusing both elastic and inelastic scattering modalities.

Currently, detector signals—such as HAADF and EDX/EELS—are analyzed separately for insight into structural, chemical, or electronic properties [157]. Correlative imaging disregards shared information between structure and chemistry and misses opportunities to recover useful information. Data fusion, popularized in satellite imaging, goes further than correlation by linking the separate signals to reconstruct new information and improve measurement accuracy [64, 91, 38]. Successful data fusion designs an analytical model that faithfully represents the relationship between modalities, and yields a meaningful combination without imposing any artificial connections [19].

Here we introduce fused multi-modal electron microscopy, a technique offering high SNR recovery of nanomaterial chemistry by linking correlated information encoded within both HAADF and EDX / EELS. We recover chemical maps by reformulating the inverse problem as a non-linear optimization which seeks solutions that accurately match the actual chemical distribution in a material. Our approach substantially improves SNRs for chemical maps, often around 300-500%, and can reduce doses over one order of magnitude while remaining consistent with original measurements. We demonstrate on EDX/EELS datasets at sub-nanometer and atomic resolution. Moreover, fused multi-modal electron microscopy recovers a specimen’s relative concentration, allowing researchers to measure local stoichiometry with less-than 15% error without any knowledge of the inelastic cross sections. Convergence and uncertainty estimates are identified along with simulations that provide ground-truth assessment of when and how this approach can fail.

3.2 Principles of Multi-Modal Electron Microscopy

Fused multi-modal electron microscopy recovers chemical maps by solving an optimization problem seeking a solution that strongly correlates with (1) the HAADF modality containing high SNR, (2) the chemically sensitive spectroscopic modality (EELS and / or EDX), and (3) encourages sparsity in the gradient domain producing solutions with reduced spatial variation. The overall optimization function results as following:

$$\begin{aligned} \arg \min_{\mathbf{x}_i \geq 0} \quad & \frac{1}{2} \left\| \sum_i (Z_i \mathbf{x}_i)^\gamma - \mathbf{b}_H \right\|_2^2 + \\ & \lambda_1 \sum_i \left(\mathbf{1}^T \mathbf{x}_i - \mathbf{b}_i^T \log(\mathbf{x}_i + \varepsilon) \right) + \lambda_2 \sum_i \|\mathbf{x}_i\|_{\text{TV}}, \end{aligned} \quad (3.1)$$

where λ are regularization parameters, \mathbf{b}_H is the measured HAADF, \mathbf{b}_i and \mathbf{x}_i are the measured and reconstructed chemical maps for element i , ε herein prevents $\log(0)$ issues but can also account for background, the \log is applied element-wise to its arguments, superscript T denotes vector transpose, and $\mathbf{1}$ denotes the vector of $n_x n_y$ ones, where $n_x \times n_y$ is the image size.

The three terms in (3.1) define our multi-modal approach to surpass traditional dose limits for chemical imaging. First, we assume a forward model where the simultaneous HAADF is a linear combination of elemental distributions (\mathbf{x}_i^γ where $\gamma \in [1.4, 2]$). The incoherent linear imaging approximation for elastic scattering scales with atomic number as Z_i^γ where γ is typically around 1.7 [68, 90, 76]. This γ is bounded between 2 for Rutherford scattering from bare nuclear potentials to 4/3 as described by Lenz-Wentzel expressions for electrons experiencing a screened coulombic potential [31, 164]. Second, we ensure the recovered signals maintain a high-degree of data fidelity with the initial measurements by using maximum negative log-likelihood for spectroscopic measurements dominated by low-count Poisson statistics [39, 120]. In a higher count regime, this term can be substituted with a simple least-squares error. Lastly, we utilize channel-wise total variation (TV) regularization to enforce a sparse gradient magnitude, which reduces noise by promoting image smoothness while preserving sharp features [136]. This sparsity constraint, popularized by the field of compressed sensing (CS), is a powerful yet minimal prior toward recovering structured data [40, 20]. When implementing, each of these three terms can and should be weighted by an appropriately selected coefficients that balances their contributions. All three terms are necessary for accurate recovery (Fig. 3.1).

3.3 High-SNR Recovery of Nanomaterial Chemistry

Figure 3.2 demonstrates high-SNR recovery for EDX signals of commercial cobalt sulfide (CoS) nano-catalysts for oxygen-reduction applications—a unique class with the highest activity among non-precious metals [135]. Figure 3.2a illustrates the model that links the two modalities (EDX and HAADF) simultaneously collected in the electron microscope. The low detection rate for characteristic X-rays is due to minimal emission (e.g., over 50% for $Z > 32$ and below 2% for $Z < 11$) and collection yield ($< 9\%$) [142]. For high-resolution EDX, the low count rate yields a sparse chemical image dominated by shot noise (Fig. 3.2b). However, noise in the fused multi-modal chemical map is virtually eliminated (Fig. 3.2d) and recovers chemical structure without a loss of resolution—including the nanoparticle core and oxide shell interface. The chemical maps produced by fused multi-modal EM quantitatively agree with the expected stoichiometry—the specimen core contains a relative concentration of $39 \pm 1.6\%$, $42 \pm 2.5\%$ and $13 \pm 2.4\%$ and exterior shell composition of $26 \pm 2.8\%$, $11 \pm 2.0\%$, $54 \pm 1.3\%$ for Co, S, O respectively. The dose for this dataset was approximately $\sim 10^5 \text{ e}/\text{\AA}^{-2}$ and a 0.7 sr EDX detector was used; however, these quantitative

$$\arg \min_{\mathbf{x}_i \geq 0} \underbrace{\frac{1}{2} \left\| \mathbf{b}_H - \sum_i (Z_i \mathbf{x}_i)^\gamma \right\|_2^2}_{\Psi_1(\mathbf{x})} + \underbrace{\lambda_1 \sum_i \left(\mathbf{1}^T \mathbf{x}_i - \mathbf{b}_i^T \log(\mathbf{x}_i + \varepsilon) \right)}_{\Psi_2(\mathbf{x})} + \lambda_2 \sum_i \|\mathbf{x}_i\|_{\text{TV}}$$




Figure 3.1: **Consequence of Reconstructing Chemical Maps with Individual Components in Cost Function.** **a)** The raw EDX maps and simultaneous HAADF image. **b)** Reconstructing the multi-modal dataset with $\Psi_1 + \Psi_2$ (model + data-consistency). The chemical maps are slightly improved but remain noisy. **c)** Reconstructing the multi-modal dataset with $\Psi_2 + \text{TV}$. This functional form is equivalent to a denoising problem; thus the resulting maps produce common staircase artifacts associated with TV. **d)** Reconstructing the raw EDX maps with fused Multi-Modal Electron Microscopy. Scale bar, 30 nm.

estimates remained consistent when the dose was reduced to $\sim 10^4 \text{ e}/\text{\AA}^{-2}$.

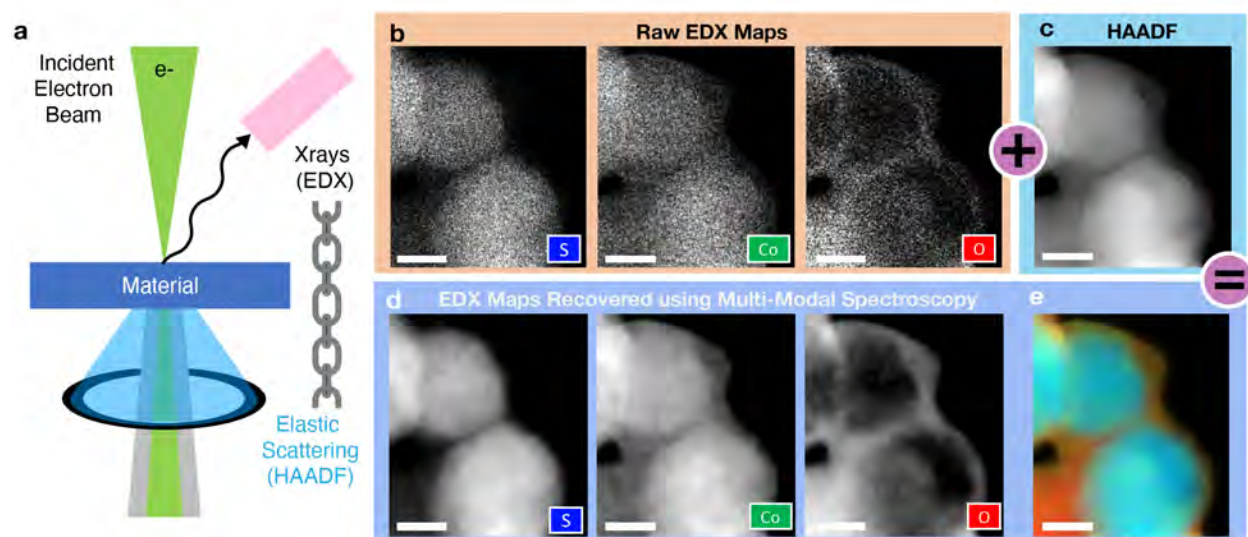


Figure 3.2: Nanoscale multi-modal chemical recovery of CoS catalysts using EDX + HAADF. **a)** Schematic highlighting the linked HAADF and EDX modalities collected in the microscope for every probe position. The algorithm links and correlates information between the two signals through an optimization process that produces chemical maps with higher SNRs. **b)** The raw EDX chemical maps for the Co, S, and O elemental distributions. **c)** The simultaneous HAADF micrograph of the CoS nanoparticle. **d)** The multi-modal reconstructions for the elemental distributions. **e)** EDX RGB overlay of the Co, S, and O maps. Scale bar, 30 nm.

Fused multi-modal electron microscopy accurately recovers chemical structure down to atomic length scales—demonstrated here for EELS spectroscopic signals. EELS derived chemical maps for $\text{Co}_{3-x}\text{Mn}_x\text{O}_4$ ($x = 1.49$) high-performing super-capacitor nanoparticles [125] are substantially improved by fused multi-modal electron microscopy in Figure 3.3. This composite Co-Mn oxide was designed to achieve a synergy between cobalt oxide’s high specific capacitance and manganese oxide’s long life cycle [125, 9]. While the $\text{Co}_{3-x}\text{Mn}_x\text{O}_4$ nanoparticle appears chemically homogeneous in the HAADF projection image along the [100] direction (Fig. 3.3c), core-shell distinctions are hinted at in the raw EELS maps (Fig. 3.3b). Specifically, these nanoparticles contain a Mn-rich center with a Co shell and homogeneous distribution of O. However the raw EELS maps are excessively degraded by noise, preventing analysis beyond rough assessment of specimen morphology. The multi-modal reconstructions (Fig. 3.3d) confirm the crystalline Co-rich shell and map the Co/Mn interface in greater detail (Fig. 3.3e). In the presence of cobalt and manganese, the HAADF image lacks noticeable contrast from oxygen; the resulting oxygen map lacks detail and benefits mostly from regularization.

Fused multi-modal electron microscopy accurately recovers chemical structure down to atomic length scales—demonstrated here for EELS spectroscopic signals. EELS derived chemical maps

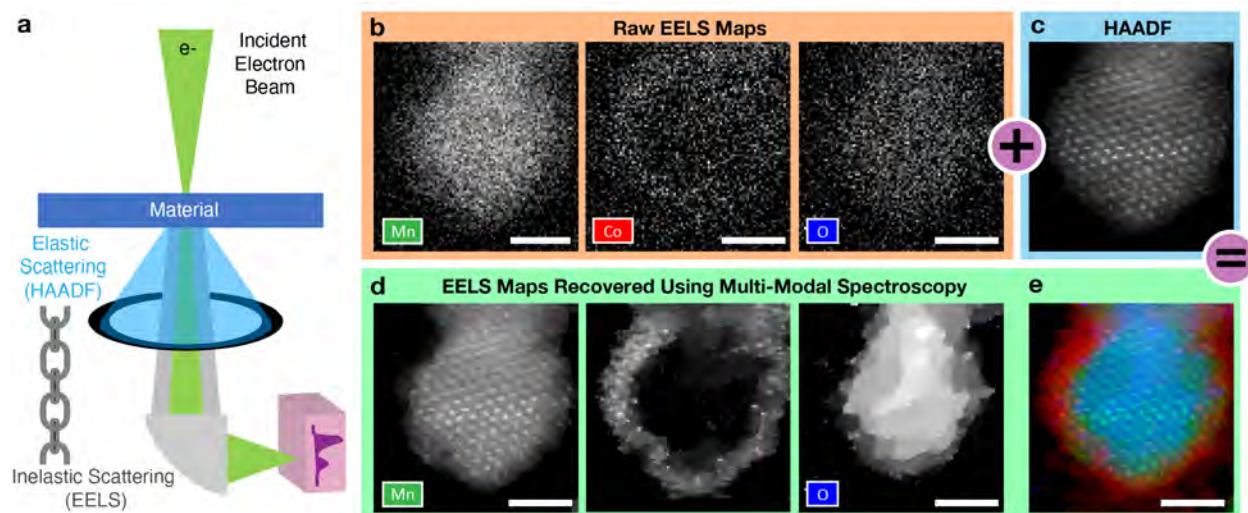


Figure 3.3: Atomic-scale multi-modal chemical recovery of $\text{Co}_{3-x}\text{Mn}_x\text{O}_4$ supercapacitors using EELS + HAADF. **a)** Schematic highlighting the linked HAADF and EELS modalities collected in the microscope at every probe position. **b)** Raw EELS maps for the elemental distributions of Co, Mn - $L_{2,3}$ and O - K edges. **c)** The simultaneous HAADF micrograph of the $\text{Co}_{3-x}\text{Mn}_x\text{O}_4$ nanoparticle. **d)** The multi-modal reconstructions for the elemental distributions. **e)** EELS RGB overlay of the Co, S, and O maps. Scale bar, 2 nm.

for $\text{Co}_{3-x}\text{Mn}_x\text{O}_4$ ($x = 1.49$) high-performing super-capacitor nanoparticles [125] are substantially improved by fused multi-modal electron microscopy in Figure 3.3. This composite Co-Mn oxide was designed to achieve a synergy between cobalt oxide's high specific capacitance and manganese oxide's long life cycle [125, 9]. While the $\text{Co}_{3-x}\text{Mn}_x\text{O}_4$ nanoparticle appears chemically homogeneous in the HAADF projection image along the [100] direction (Fig. 3.3c), core-shell distinctions are hinted at in the raw EELS maps (Fig. 3.3b). Specifically, these nanoparticles contain a Mn-rich center with a Co shell and homogeneous distribution of O. However the raw EELS maps are excessively degraded by noise, preventing analysis beyond rough assessment of specimen morphology. The multi-modal reconstructions (Fig. 3.3d) confirm the crystalline Co-rich shell and map the Co/Mn interface in greater detail (Fig. 3.3e). In the presence of cobalt and manganese, the HAADF image lacks noticeable contrast from oxygen; the resulting oxygen map lacks detail and benefits mostly from regularization.

Figure 3.4 exhibits fused multi-modal electron microscopy at atomic resolution on copper sulphur heterostructured nanocrystals with zinc sulfide caps with potential applications in photovoltaic devices or battery electrodes [62]. The copper sulfide properties are sensitive to the Cu-S stoichiometry and crystal structure at the interface between ZnS and $\text{Cu}_{0.64}\text{S}_{0.36}$. Figure 3.4 shows high-resolution HAADF and EELS characterization of a heterostructure $\text{Cu}_{0.64}\text{S}_{0.36}$ -ZnS interface. Fused multi-modal electron microscopy maps out the atomically sharp $\text{Cu}_{0.64}\text{S}_{0.36}$ -ZnS interface

and reveals step edges between the two layers. The labeled points on the RGB chemical overlay (Fig. 3.4d) shows the chemical ratios produced by multi-modal EM for the $\text{Cu}_{0.64}\text{S}_{0.36}$ and ZnS regions—values which are consistent with the reported growth conditions. Figure 3.4e shows the algorithm convergence for each of the three terms in the optimization function (Eq. 3.1)—smooth and asymptotic decay is an indicator of reliable reconstruction.

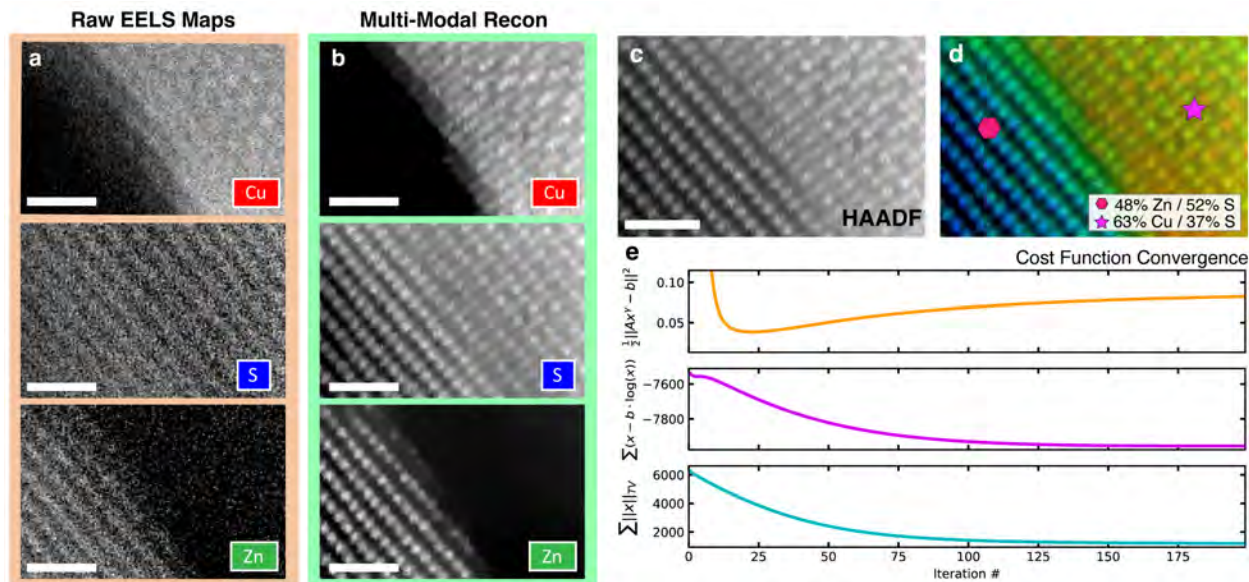


Figure 3.4: **Recovering chemistry in an atomically sharp ZnS- $\text{Cu}_{0.64}\text{S}_{0.36}$ heterointerface interface.** **a)** The raw EELS maps for the Cu, S, and Zn $L_{2,3}$ edges. **b)** The multi-modal reconstructions for the elemental compositions. **c)** The simultaneous HAADF micrograph of the ZnS- $\text{Cu}_{0.64}\text{S}_{0.34}$ interface. **d)** Color overlay of the Zn, S, and Zn maps. The relative concentration for the constituent elements consist of $48 \pm 5.9\%$ for Zn, $59.9 \pm 3.2\%$ for Cu and $38 \pm 2.6\%$ for S in the $\text{Cu}_{0.64}\text{S}_{0.36}$ layer and $48.9 \pm 6\%$ in ZnS. **e)** Convergence plots for the three individual components in the cost function. Scale bar, 1 nm.

Fused multi-modal imaging of Fe and Pt distributions from inelastic multislice simulations (Fig. 3.5) provide ground truth solutions to validate recovery at atomic resolution under multiple scattering conditions of an on-axis ~ 8 nm nanoparticle. Here, we applied Poisson noise (Fig. 3.5b) containing electron doses of $\sim 10^9$ $e \text{ \AA}^{-2}$, to produce chemical maps with noise levels resembling experimental atomic-resolution EELS datasets ($\text{SNR} \simeq 5$). We estimated SNR improvements by measuring peak-SNR for the noisy and recovered chemical maps [75]. Qualitatively, the recovered chemical distributions (Fig. 3.5c) match the original images. Fig. 3.5d illustrates agreement of the line profiles as the atom column positions and relative peak intensities between the ground truth and multi-modal reconstruction are almost identical.

Simulating EELS chemical maps is computationally demanding as every inelastic scattering event requires propagation of an additional wavefunction [43, 3]—scaling faster than the cube of

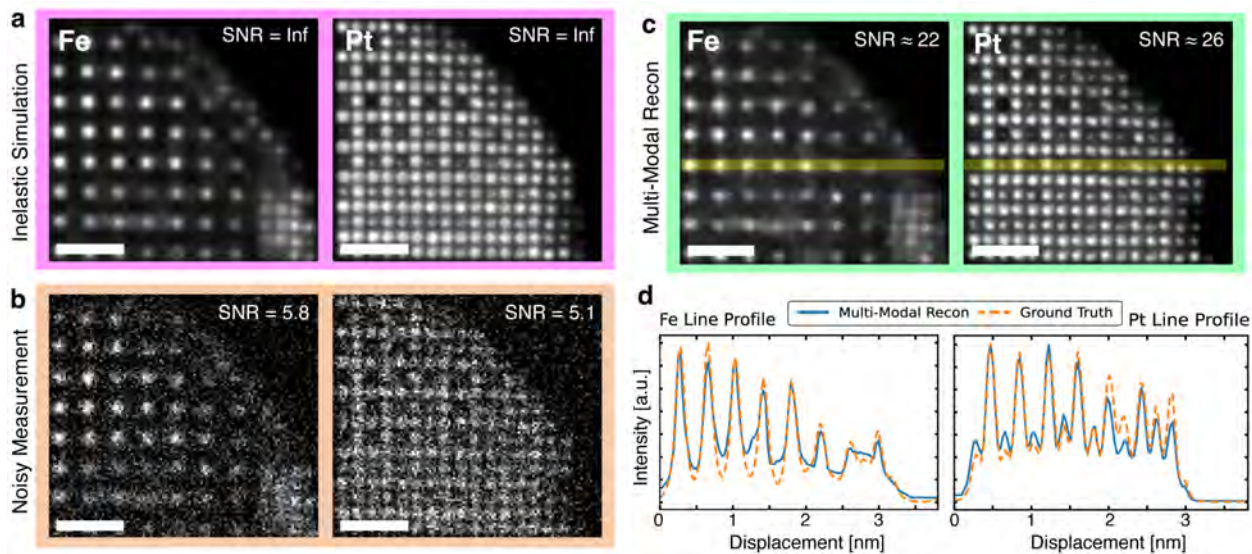


Figure 3.5: **Inelastic FePt nanoparticle simulation.** **a)** Ground truth EELS images generated from inelastic simulations. **b)** EELS maps degraded with Poisson shot noise. SNR shown on top right. **c)** Recovered atomic-resolution EELS maps for the Fe and Pt distributions. Estimated SNR shown on top right. **d)** Line profiles of the marked yellow bars (10 pixels in width) in (c) compares the Multi-Modal reconstruction and ground truth. Scale bar, 1 nm.

the number of beams, $O(N^3 \log N)$. Inelastic transition potentials of interest (in this case the $L_{2,3}$ Fe and $M_{4,5}$ Pt edges) were calculated from density function theory. Long computation times (nearly 4,000 core-hours) result from a large number of outgoing scattering channels corresponding to the many possible excitations in a sample. For this reason, there is little precedence for inelastic image simulations. We relaxed the runtime by utilizing the PRISM STEM-EELS approximation, achieving over a ten-fold speedup [14]. Future work may explore the effects of smaller ADF collection angles with increased coherence lengths and crystallographic contrast [68, 179], or thicker specimens where electron channeling becomes more concerning [5, 77].

3.4 Quantifying Chemical Concentration

Fused multi-modal electron microscopy can produce stoichiometrically meaningful chemical maps without specific knowledge of inelastic cross sections. Here, the ratio of pixel values in the reconstructed maps quantify elemental concentration. We demonstrate quantifiable chemistry on experimental metal oxide thin films with known stoichiometry: NiO [45] and ZrO_2 . A histogram of intensities from the recovered chemical maps are fit with Gaussian distributions to determine the average concentration (Fig. 3.6). The recovered pixel values highlighted in Figure 3.6b,d followed a single Gaussian distribution where the Zr and Ni concentrations are centered about $35 \pm 5.8\%$ and

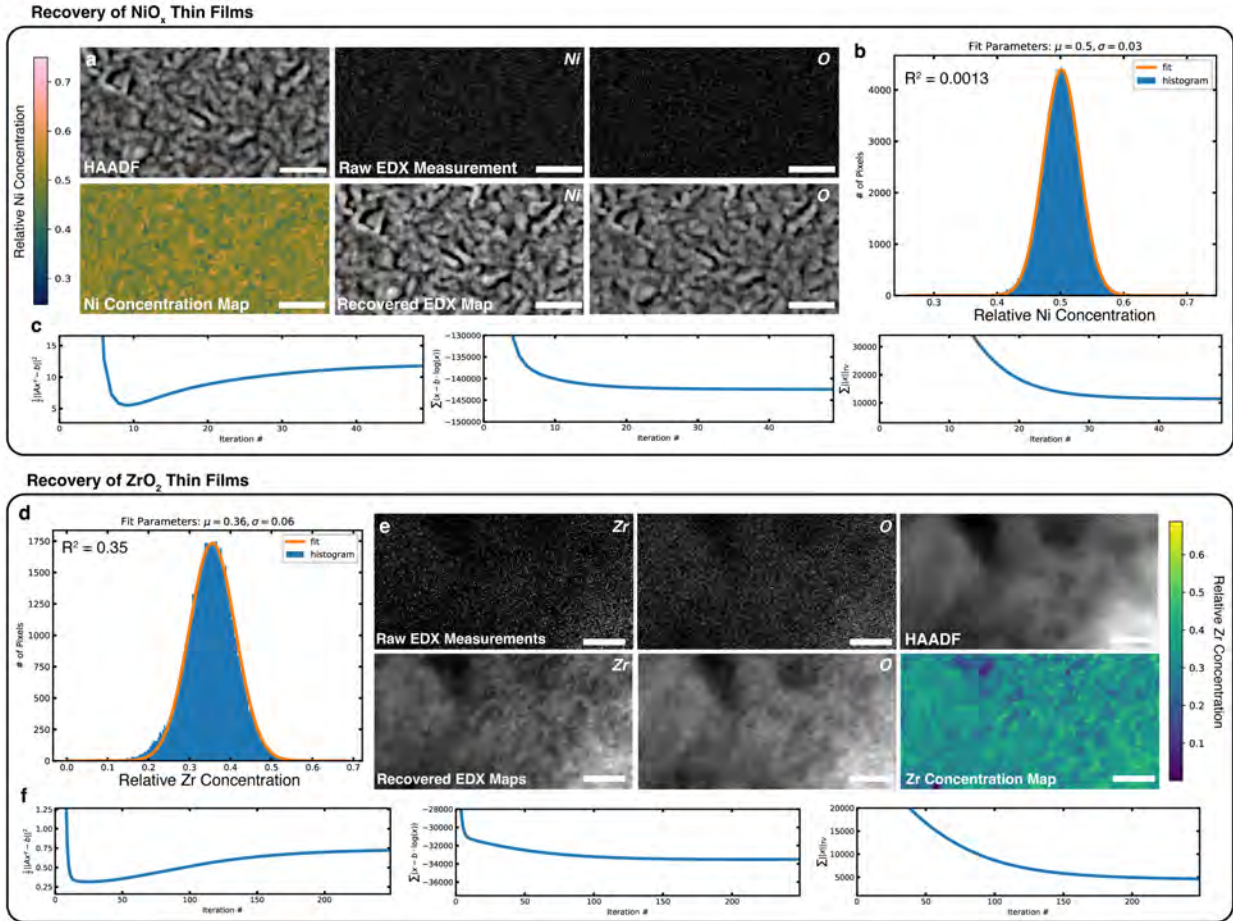


Figure 3.6: **Histograms of intensities for ZrO₂ and NiO.** Histogram of intensities for the Sulphur, Cobalt and Oxygen distributions. Due to the core-shell structure consisting of two phases, the histograms contain of two Gaussian distributions. The specimen core contains a relative concentration of $39 \pm 1.6\%$, $42 \pm 2.5\%$ and $13 \pm 2.4\%$ and exterior shell composition of $26 \pm 2.8\%$, $11 \pm 2.0\%$, $54 \pm 1.3\%$ for Co, S, O respectively. Scale bar, 30 nm.

$50 \pm 2.9\%$. More importantly, in the context of EDX, minimal preprocessing is required as stoichiometry is decomposed from the HAADF micrographs (Fig. 3.6a,d). In both cases, the average Ni and Zr relative concentration is approximately equivalent to the expected ratio from the crystal stoichiometry: 33% and 50%. Convergence still needs to be reliably assessed (Fig. 3.6c,f) to ensure the reconstructions traverse to the optimal minima defined by Eq. 3.1.

Alternatively, the heterogeneous CoS nanoparticles in Fig. 3.2 follows a bi-modal distribution for the core and shell phases (Fig. 3.7). More importantly, measuring relative chemical concentration remains accurate as we reduce our fields of view down to atomic length scales. As the ZnS-CuS interface, the recovered relative chemistries matches expected crystal stoichiometry.

We further validate stoichiometric recovery on a synthetic gallium oxide crystal (Fig 3.8) where

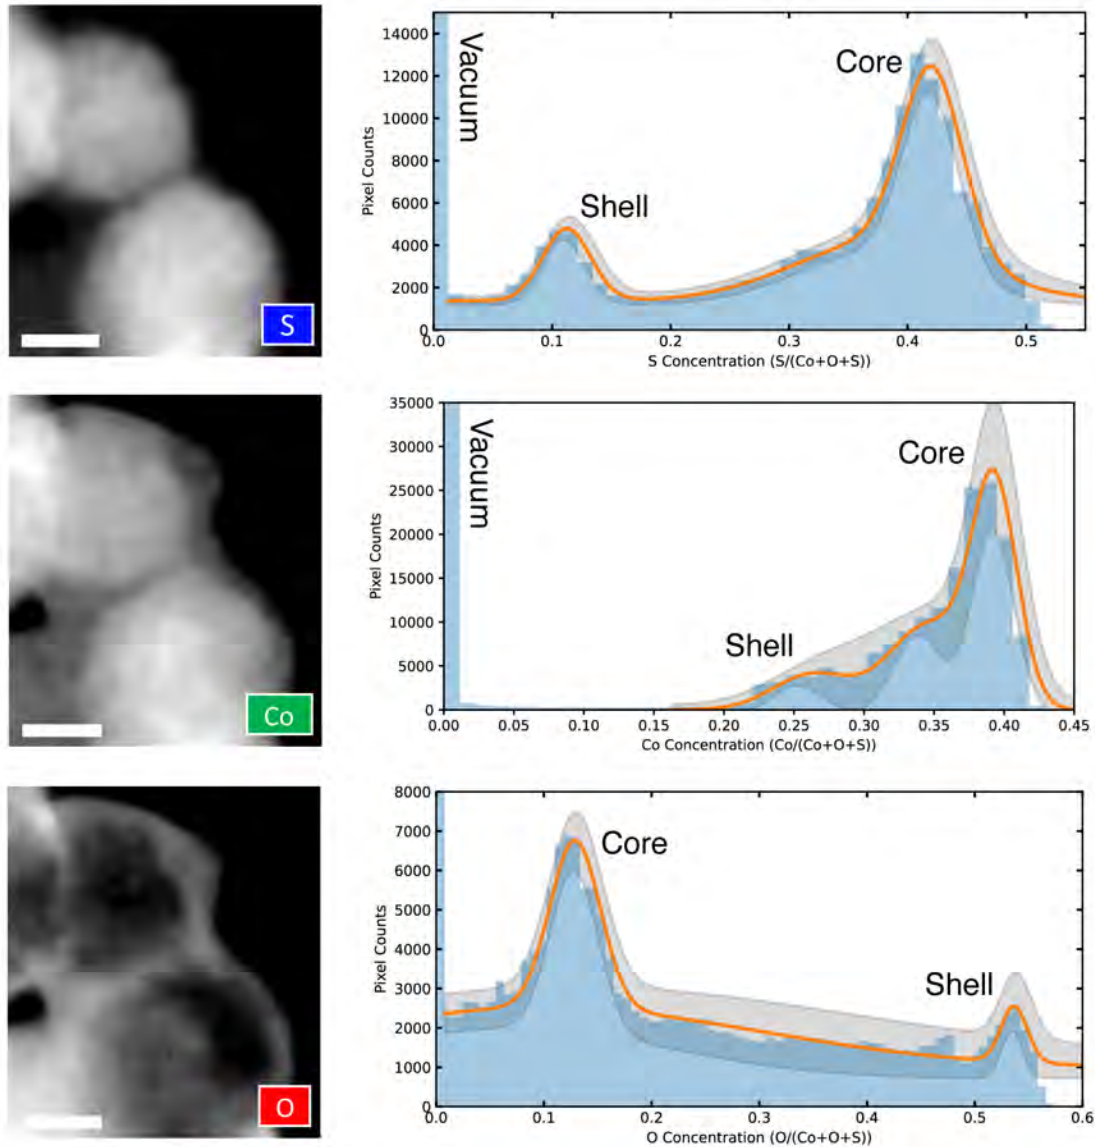


Figure 3.7: **Histograms of intensities for CoS chemical Maps.** Histogram of intensities for the Sulphur, Cobalt and Oxygen distributions. Due to the core-shell structure consisting of two phases, the histograms contain of two Gaussian distributions. The specimen core contains a relative concentration of $39\pm 1.6\%$, $42\pm 2.5\%$ and $13\pm 2.4\%$ and exterior shell composition of $26\pm 2.8\%$, $11\pm 2.0\%$, $54\pm 1.3\%$ for Co, S, O respectively. Scale bar, 1 nm.

two overlapping Ga and O thin films of equal thickness have a stoichiometry of Ga_2O_3 . The simulated HAADF signal is proportional to $\sum_i (x_i Z_i)^\gamma$ where x_i is the concentration for element i and Z_i is the atomic number. As shown by the histogram, the simulated results agree strongly with the prior knowledge and successfully recovers the relative Ga concentration. The Gaussian distribution is centered about $40\pm 0.4\%$ when the ground truth is 40% . The inset shows convergence

plots.

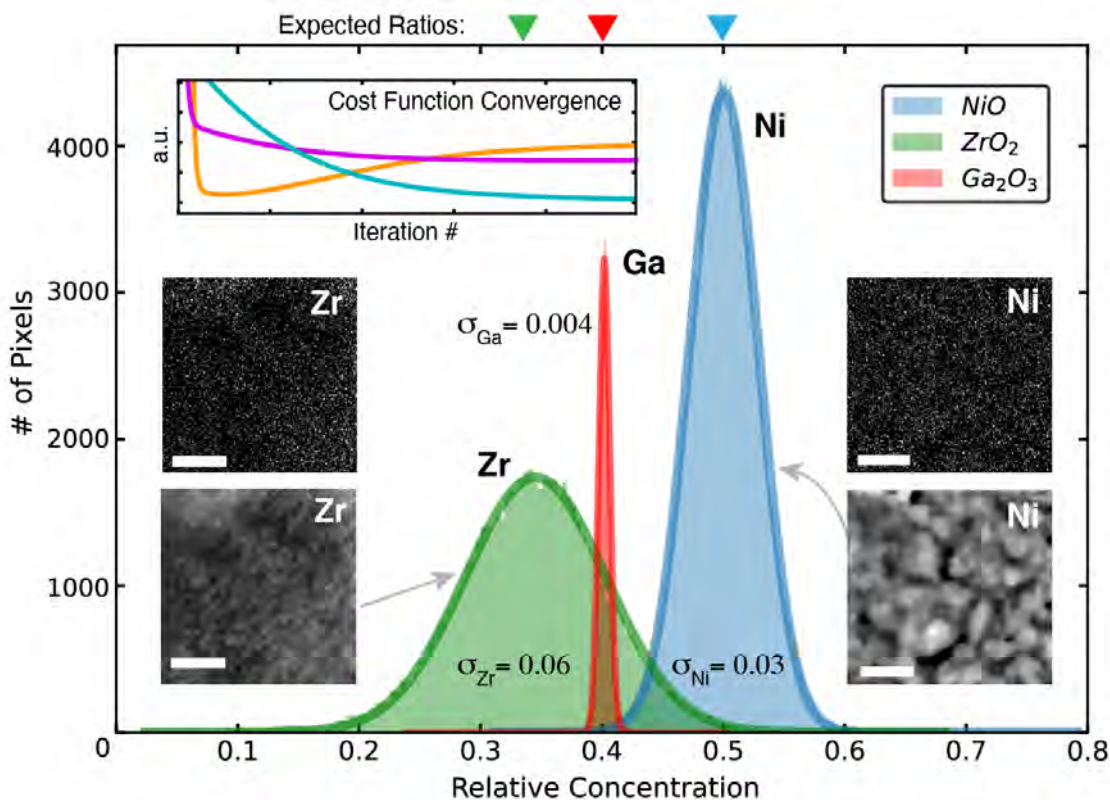


Figure 3.8: **Measuring relative concentration for experimental and synthetic datasets.** Pixel intensity histograms for an experimental Zr (green), Ni (blue) and synthetic Ga (red) concentration maps. The standard deviation (σ) for each element is reported. The raw and reconstructed EDX maps are illustrated inside of the plot. Ground truth concentrations are highlighted by the respective colored triangles above the top axis. Stable convergence for the three components in the cost function: model term (orange), data fidelity (magenta), and regularization (turquoise) are illustrated in the inset. Qualitatively the convergence is identical for all three example datasets. Zr and Ni scale bars: 5, 10 nm, respectively.

We estimate a stoichiometric error of less-than 15% for most materials based on the relative concentration's standard deviation ($\pm 7\%$) added in quadrature with the variation of solutions ($\pm 6\%$). Although the algorithm shows stable convergence, the overall quantitative conclusions are slightly sensitive to the selection of hyperparameters. We estimate incorrect selection of hyperparameters could result in variation of roughly $\pm 6\%$ from the correct prediction in stoichiometry even when the algorithm converges. We found measuring stoichiometry is robust across a range of γ values close to 1.7. In cases where γ is far off (e.g., $\gamma = 1.0$), the quantification is systematically incorrect (Fig. 3.9).

This error is comparable to estimating chemical concentrations directly from EELS / EDX

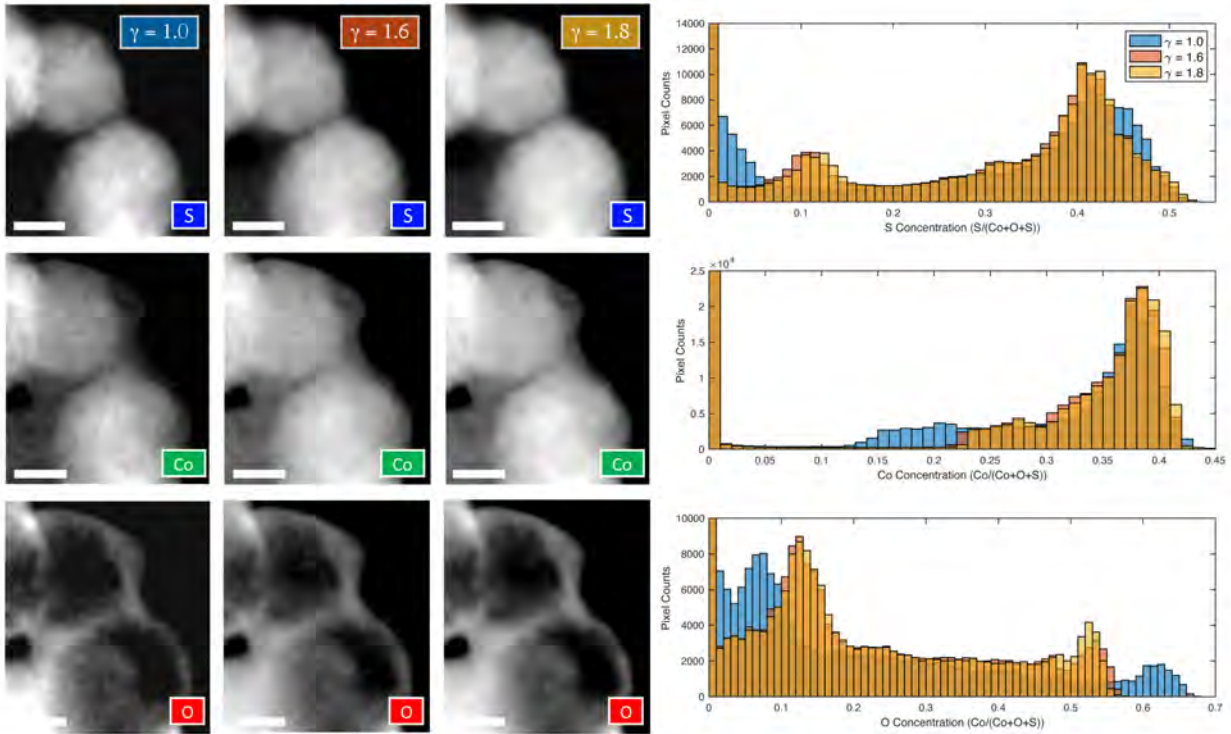


Figure 3.9: **The influence of γ for measuring chemistry** Histogram of intensities for the Sulphur, Cobalt and Oxygen distributions when γ varies between 1.0, 1.6, and 1.8. We see for values close to 1.7, the specimen's relative concentrations almost perfectly matches the distribution in Fig. S5. However if gamma is far off from what is experimentally feasible ($\gamma = 1.0$), the measured concentrations are incorrect. Scale bar, 30 nm.

spectral maps from the ratio of scattering cross section against core-loss intensity [131]. However, traditional approaches require accurate knowledge of all experimental parameters (e.g., beam energy, specimen-thickness, collection angles) and accurate calculation of the inelastic cross-section typically to provide errors roughly between 5-10% [44].

3.5 Influence of Electron Dose

To better understand the accuracy of fused multi-modal electron microscopy at low doses, we performed a quantitative study of normalized root-mean-square error (RMSE) concentrations for a simulated 3D core-shell nanoparticle (CoS core, CoO shell). Figure 3.10 shows the fused multi-modal reconstruction accuracy across a wide range of HAADF and chemical SNR. The simulated projection images were generated by simple linear incoherent imaging model of the 3D chemical compositions highlighted in Fig. 3.10d—here the probe's depth of focus is much larger than the object. Random Poisson noise corresponding to different electron dose levels was applied to vary

the SNR across each pixel.

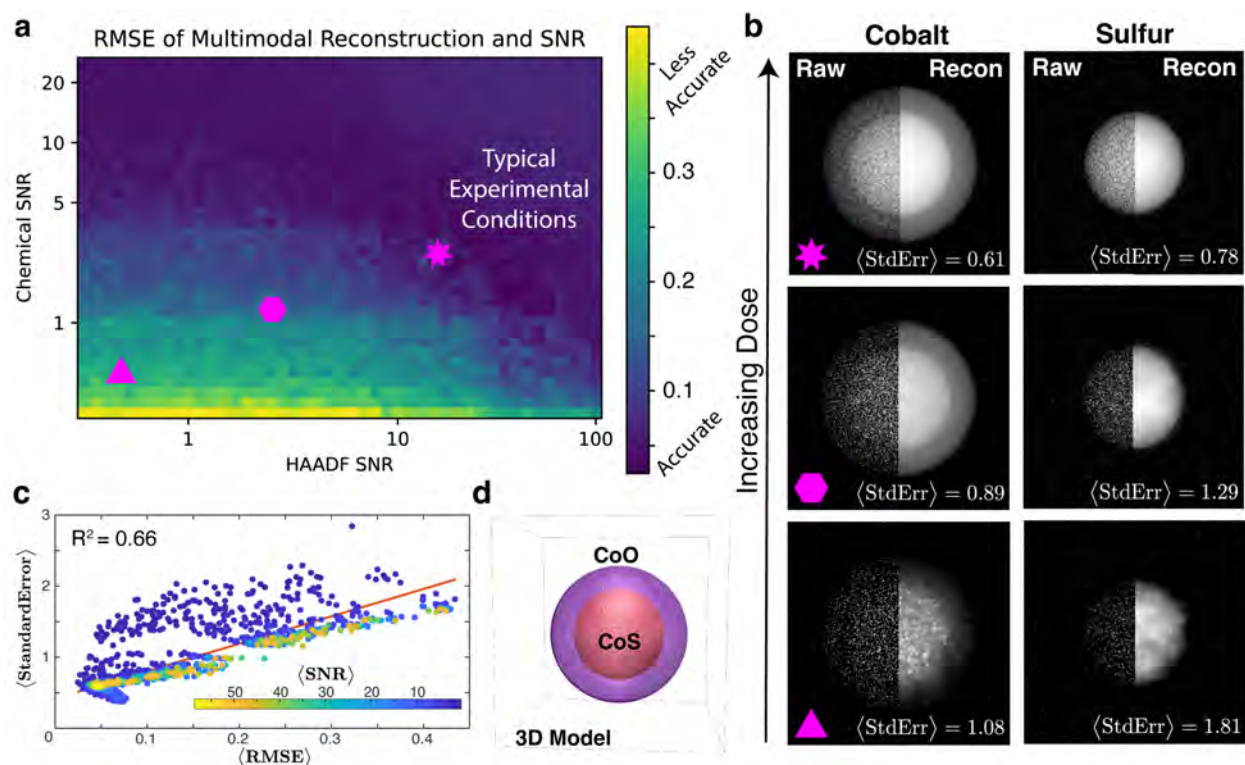


Figure 3.10: **Estimating dose requirements for accurate chemical recovery.** **a)** A RMSE map representing the reconstruction error as a function of multiple spectroscopic and HAADF SNR. Brighter pixels denote results containing the incorrect concentrations from the ground truth. **b)** Visualization of three points on the phase diagram corresponding to increasing ADF / chemical electron dose. **c)** A plot of average standard error vs. RMSE demonstrating the two metrics are linearly correlated. **d)** The 3D model for generating synthetic chemical and ADF projections.

Overall, the RMSE simulation map (Fig. 3.10a) shows the core-shell nanoparticle chemical maps are accurately recovered at low-doses (HAADF SNR $\gtrsim 4$ and chemical SNR $\gtrsim 2$); however, they become less accurate at extremely low doses. The RMSE map for multimodal reconstruction shows a predictably continuous degradation in recovery as signals diminish. The degraded and reconstructed chemical maps for various noise levels are highlighted in Figure 3.10b. The Co map closely mirrors the Z-contrast observed in HAADF (not shown) simply because it is the heaviest element present. Usually researchers will perform spectroscopic experiments in the top right corner of Fig. 3.10a (e.g., HAADF SNR > 20, chemical SNR > 3), which for this simulation, provides accurate recovery.

In actual experiments, the ground truth is unknown and RMSE cannot be calculated to assess fused multi-modal electron microscopy. However we can estimate accuracy by calculating an average standard error of our recovered image from the Hessian of our model. The standard error

reflects uncertainty at each pixel in a recovered chemical map by quantifying the neighborhood size for similar solutions (Figure 3.11). The average standard error across all pixels in a fused multi-modal image provides a single value metric of the reconstruction accuracy (see Methods). Figure 3.10c shows that RMSE and average standard error correlate, especially at higher doses (SNR > 10).

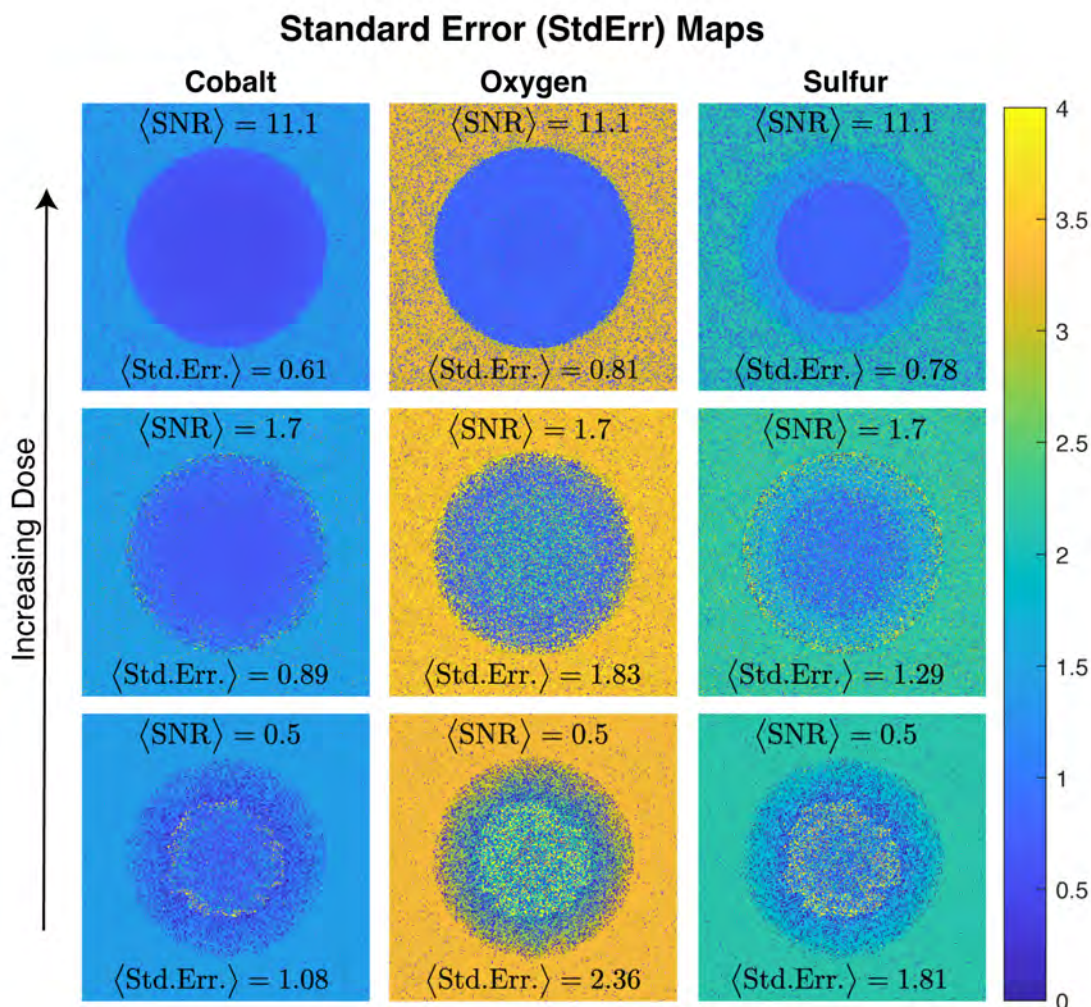


Figure 3.11: **Uncertainty Maps for Synthetic CoS Recovery** Visualization of the standard error (StdErr) maps for the synthetic CoS core shell simulation in Fig.. As we increase the electron dose, the average StdErr pixel values for all three elements decreases. $\langle \text{SNR} \rangle$ denotes the average SNR between the chemical and HAADF modality.

3.6 Discussion and Conclusion

While this chapter highlights the advantages of multi-modal electron microscopy, the technique is not a black-box solution. Step sizes for convergence and weights on the terms in the cost function (Eq. 3.1) must be reasonably selected. This manuscript illustrates approaches to assess the validity of concentration measurements using confidence estimation demonstrated across several simulated and experimental material classes. Standard spectroscopic pre-processing methods become ever more critical in combination with multi-modal fusion. Improper background subtraction of EELS spectra or overlapping characteristic X-ray peaks that normally causes inaccurate stoichiometric quantification also reduces the accuracy of fused multi-modal imaging.

Fused multi-modal electron microscopy offers little advantage in recovering chemical maps for elements with insignificant contrast in the HAADF modality. This property is limiting for analyzing specimens with low-Z elements in the presence of heavy elements (e.g., oxygen and lutetium). Future efforts could resolve this challenge by incorporating an additional complementary elastic imaging mode where light elements are visible, such as annular bright field (ABF) [54]. However in some instances, fused multi-modal electron microscopy may recover useful information for under-determined chemical signals. For example, in a $\text{Bi}_{0.35}\text{Sr}_{0.18}\text{Ca}_{0.47}\text{MnO}_3$ (BSCMO) system [138], only the Ca, Mn, and O EELS maps were obtained, yet multimodality remarkably improves the SNR of measured maps despite missing two elements (Figure. 3.12).

Although fused multi-modal chemical mapping appears quite robust at nanometer or sub-nanometer resolution, we found atomic-resolution reconstructions can be challenged by spurious atom artifacts which require attention. However, this is easily remedied by down-sampling to frequencies below the first Bragg peaks and analysing a lower resolution chemical map. Alternatively, recovery with minimal spurious atom artifacts is achieved when lower resolution reconstructions are used as an initial guess.

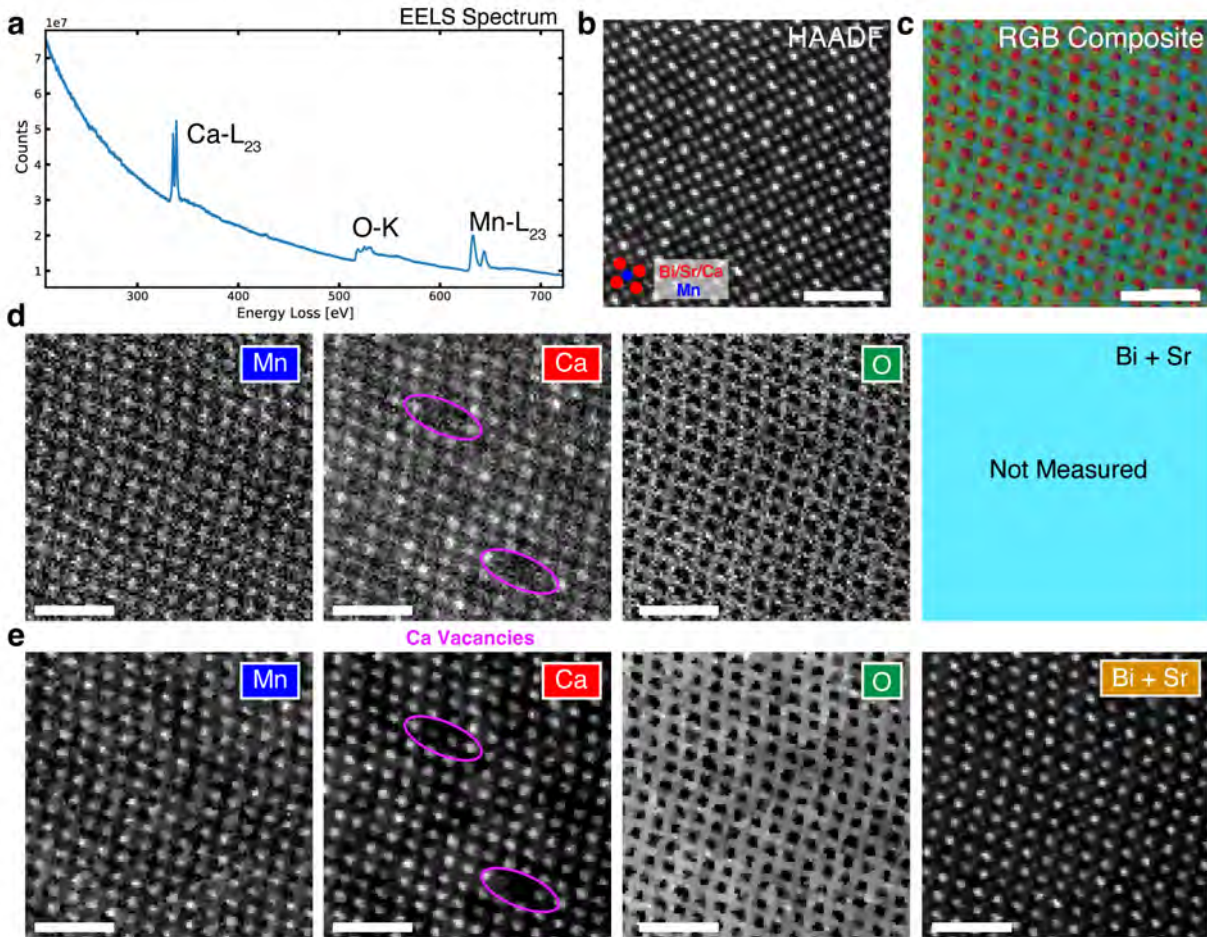


Figure 3.12: **Recovering Missing Elements for $\text{Bi}_{0.35}\text{Sr}_{0.18}\text{Ca}_{0.47}\text{MnO}_3$ (BSCMO)** a) Raw EELS spectrum for a BSCMO system shows Ca, O, and Mn core-loss edges, however, the Bi and Sr core edges were not measured in this energy range. b) The simultaneously collected HAADF image with the inset on the bottom right highlighting the A and B sites' composition. c) RGB overlay highlighting the chemical distribution. d) The raw EELS maps. e) The recovered chemical maps. The purple circles highlight Ca vacancies. Here, the Bi and Sr were treated as a single unknown signal since they are known a priori to occupy the same lattice sites. Although chemical maps of Bi and Sr are missing, multi-modal reconstruction shows considerable improvement to the SNR and quality of the Mn, Ca, and O distributions. Moreover, we see the recovered Bi + Sr map anti-correlates with Ca as expected, however the small intensity between lattice sites is a minor artifact that arises without spectroscopic measurement of Bi and Sr. Scale bar, 2 nm.

CHAPTER 4

Imaging 3D Chemistry at 1 nm Resolution with Fused Multi-Modal Electron Tomography

1

Measuring the three-dimensional (3D) distribution of chemistry in nanoscale matter is a long-standing challenge for metrological science. The inelastic scattering events required for 3D chemical imaging are too rare, requiring high beam exposure that destroys the specimen before an experiment completes. Even larger doses are required to achieve high resolution. Thus, chemical mapping in 3D has been unachievable except at lower resolution with the most radiation-hard materials. Here, high-resolution 3D chemical imaging is achieved near or below one nanometer resolution in a Au-Fe₃O₄ metamaterial, Co₃O₄ - Mn₃O₄ core-shell nanocrystals, and ZnS-Cu_{0.64}S_{0.36} nanomaterial using fused multi-modal electron tomography. Multi-modal data fusion enables high-resolution chemical tomography often with 99% less dose by linking information encoded within both elastic (HAADF) and inelastic (EDX / EELS) signals. Now sub-nanometer 3D resolution of chemistry is measurable for a broad class of geometrically and compositionally complex materials.

4.1 Introduction

Knowing the complete chemical arrangement of matter in all dimensions is fundamental to engineering novel nanomaterials [109]. Although electron tomography provides comprehensive 3D structure at resolutions below 1 nm using elastic scattering signals [150, 172, 96], chemical tomography obtained from inelastic scattering remains largely out of reach. Several demonstrations of chemical tomography using electron energy loss or x-ray energy spectroscopy (EELS / EDX) accompanied the introduction of scanning transmission electron microscope (STEM) tomography and provide a milestone for 3D imaging [113, 111, 119, 95]. However, chemical tomography from core-excitation spectroscopy demands high electron doses that almost always exceed the specimen

¹This chapter is currently in the review process.

limits (e.g., $> 10^7$ e/Å²) [35, 67, 29]. If attempting chemical tomography, researchers must sacrifice resolution by collecting few specimen projections (e.g., 5-10) and constrain the total dose (e.g., $< 10^6$ e/Å²). Consequently, 3D resolution is penalized from undersampling and noisy chemical maps [33]. Therefore, a paradigm shift is necessary for high-resolution chemical tomography.

We show achieving high-resolution 3D chemistry at lower dose requires fusing both elastic and inelastic scattering signals. Typically these detector signals are analyzed separately and correlated [60, 102, 99]. However, correlative imaging disregards shared but also complementary information between structure and chemistry and misses opportunities to recover useful information [157]. Data fusion, popularized in satellite imaging, goes further than correlation by linking separate signal modalities to reconstruct new information and improve measurement accuracy [64, 91, 39]. Recent developments in multi-modal data fusion paved new opportunities for high-resolution chemical imaging by substantially reducing the dose requirements to successfully acquire an atomic-resolution map [144]. In alignment with the principles of fused multi-modal electron microscopy, we extend its algorithmic framework into the third dimension.

Here we introduce fused multi-modal electron tomography, which offers high signal-to-noise (SNR) and high-resolution recovery of material chemistry in 3D by linking information encoded within both elastic high-angle annular dark field (HAADF) and inelastic (EDX / EELS) scattering signals. Multi-modal electron tomography reconstructs the volumetric chemical structure of specimens by solving a 3-term inverse problem that fuses signals from multiple detectors. This framework enables new sampling strategies that minimize dose by measuring a high number of HAADF projections alongside far fewer chemical projections—dose reductions of one-hundred fold are readily achieved. Although the chemical structure is severely underdetermined, fusing the two modalities fills in missing information, notably improving resolution and reconstruction quality. Our approach demonstrates that researchers can measure 3D chemistry at 1 nm resolution using electron doses as low as 10^4 e/Å² and as few as 9 spectroscopic maps while remaining consistent with original measurements. Multi-modal tomography is validated across multiple material systems, including Au-Fe₃O₄ superlattice clusters, core-shell Co₃O₄-Mn₃O₄ [121], ZnS-Cu_{0.64}S_{0.36} heterostructures [62], Cu-SiC nanoparticles and a range of simulated specimens. By fusing modalities, chemical tomography is now possible at sub-nanometer resolution for a wider class of material systems.

4.2 Principles of Fused Multi-Modal Electron Microscopy

High-resolution 3D chemical imaging is achieved using the multi-modal electron tomography framework illustrated in Fig. 4.1a for a binary Au-Fe₃O₄ nanoparticle superlattice within a carbon-based matrix. In multi-modal electron tomography, projections of the specimen structure are mea-

sured from a HAADF detector and the specimen chemistry is extracted from spectroscopy (EELS or EDX). These two detector modalities are fused during the reconstruction process to provide the complete 3D chemical distribution of a specimen at high resolution and SNR. Figure 4.1b shows the 3D reconstruction of each individual chemistry: larger 10.2 ± 1.1 nm Fe nanoparticles (blue) and smaller 3.9 ± 0.4 nm Au nanoparticles (orange). Both chemistries are visualized simultaneously in Fig 4.1c to show the self-organization of the chemical superlattice. The light-element, carbon matrix is shown in Fig. 4.2.

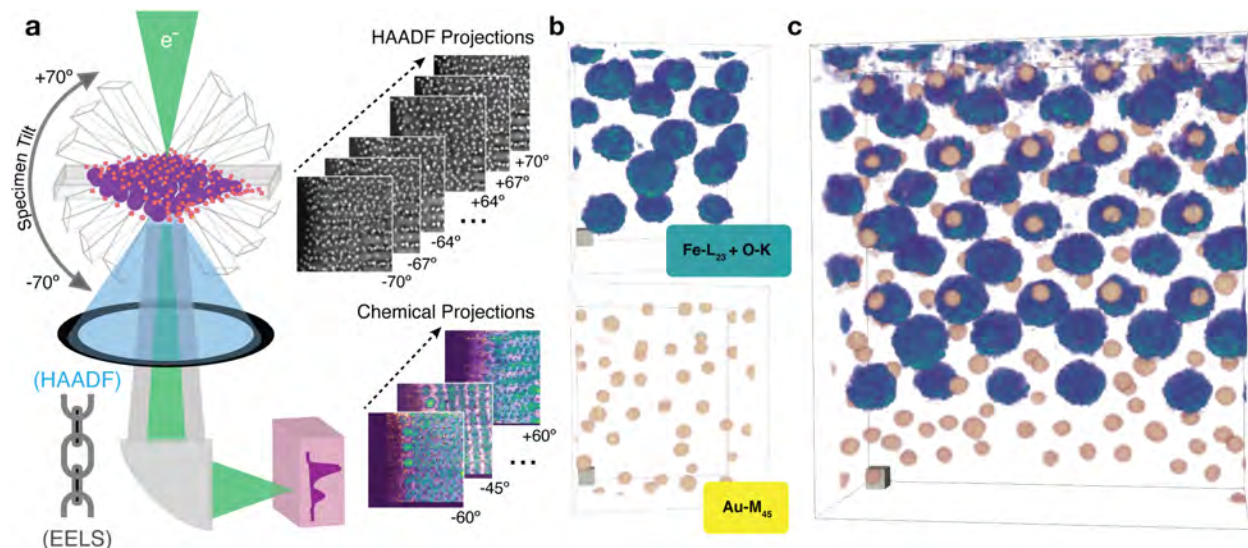


Figure 4.1: **Nanoscale recovery of Au-Fe₃O₄ nanoparticle superlattice.** **a** Schematic highlighting the linked HAADF and EELS modalities for chemical tomography. HAADF projection images are collected at every tilt increment while core-loss EELS spectra are sparsely acquired every few tilts. **b** The fused multi-modal reconstruction for the specimen's Fe L_{2,3} (turquoise), O-K (turquoise), and gold M_{4,5} edge (yellow). **c** Chemical overlay of the superlattice nanoparticles over the entire 115 nm field of view. Scale cubes, 5 nm³.

In multi-modal tomography, the number of structural HAADF projections usually exceeds the chemical projections. In this first demonstration, only 9 chemical maps ($\Delta\theta = 15^\circ$) are measured from the Fe-L_{2,3} and Au-M_{4,5} core-excitation edges in an EELS spectrum whereas 47 HAADF images ($\Delta\theta = 3^\circ$) are collected over a $\pm 70^\circ$ specimen tilt range. Linking both modalities into the reconstruction enables a clear distinction between Fe₃O₄ and Au nanoparticles at high resolution from just a few EELS maps and a total electron dose of 5×10^5 e/Å²—roughly two orders of magnitude lower total electron dose than an equivalent conventional approach.

Fused multi-modal electron tomography reconstructs three-dimensional chemical models by solving an optimization problem seeking a solution that strongly agrees with (1) the HAADF modality containing high SNR, (2) the chemically sensitive spectroscopic modality (EELS and / or EDX), and (3) encourages sparsity in the gradient domain producing solutions with reduced




Figure 4.2: **Multi-Modal Reconstruction of Au-Fe₃O₄ Nanoparticles Inside a Carbon Support.** The complete reconstruction of the Au-Fe₃O₄ nanoparticle superlattice inside the carbon matrix (highlighted in purple). Scale cube, 10 nm³.

spatial variation. The overall optimization function is as follows:

$$\arg \min_{\mathbf{x}_i \geq 0} \frac{\lambda_1}{2} \left\| \mathbf{A}_h \sum_i (Z_i \mathbf{x}_i)^\gamma - \mathbf{b}_H \right\|_2^2 + \lambda_2 \sum_i \left(\mathbf{1}^T \mathbf{A}_c \mathbf{x}_i - \mathbf{b}_i^T \log(\mathbf{A}_c \mathbf{x}_i + \varepsilon) \right) + \lambda_3 \sum_i \|\mathbf{x}_i\|_{\text{TV}}, \quad (4.1)$$

\mathbf{x}_i is the reconstructed 3D chemical distributions for element i , \mathbf{b}_i is the measured 2D chemical maps for element i , \mathbf{b}_H is the measured HAADF micrographs, \mathbf{A}_h and \mathbf{A}_c are forward projection operators for HAADF and chemical modalities, λ are regularization parameters, ε herein prevents $\log(0)$ issues but can also account for background, the \log is applied element-wise to its arguments, superscript T denotes vector transpose, and $\mathbf{1}$ denotes the vector of $N_{\text{chem}}^{\text{proj}} n_y n_i$ ones, where n_y is the number of pixels, n_i is the number of elements present, and $N_{\text{chem}}^{\text{proj}}$ is the number of projections for the chemical modality.

The three terms in Equation 4.1 define our fused multi-modal framework designed to surpass traditional limits for chemical tomography. First, we assume a forward model where the simultaneous HAADF is a linear combination of the reconstructed 3D elemental distributions (\mathbf{x}_i^γ where $\gamma \in [1.4, 2]$). The incoherent linear imaging approximation for elastic scattering scales with atomic number as Z_i^γ , where experimentally γ is typically around 1.7 [68, 90, 76]. This γ is bounded between $4/3$ as described by Lenz-Wentzel expressions for electrons passing through a screened coulombic potential and 2 for Rutherford scattering from bare nuclear potentials [31, 164]. Second, we ensure the recovered 3D distributions maintain a high degree of data fidelity with the initial measurements by using the log-likelihood for spectroscopic measurements dominated by low-count Poisson statistics [39, 120]. In a higher count regime, this term can be substituted with a least-squares discrepancy ($\|\mathbf{Ax} - \mathbf{b}\|_2^2$) [34]. Lastly, we include channel-wise isotropic total variation (TV) regularization to enforce a sparse gradient magnitude, which reduces noise by promoting image smoothness while preserving sharp features [136]. This sparsity constraint, popularized by the field of compressed sensing (CS), is a powerful yet modest prior for recovering structured data [40, 20]. When solving Equation 4.1, each of these three terms should be weighted appropriately by determining coefficients (λ) that balance their contributions. Ultimately, optimization of all three terms is necessary for accurate recovery.

The improvement in reconstruction quality with fused multi-modal chemical tomography (Fig. 4.3i) is dramatic when compared to traditional chemical tomography (Fig. 4.3c).

4.3 3D Chemistry at High-Resolution, Low-Dose

In tomography, 3D resolution is described by the Crowther criterion, which states resolution is limited by the object size and the number of specimen projections measured [87] – higher resolution requires more projections [169]. For traditional chemical tomography, few chemical projections are collected and the Crowther relation devastates resolution in 3D. This limitation occurs from the high-dose requirements of chemical mapping (i.e., EDX, EELS) where only a few projections can be collected before radiation damage alters the specimen structure.

Figure 4.3 shows how specimen projections from each modality are superimposed as planes of information in Fourier space. Chemical tomography is sparsely sampled in Fourier space (Fig. 4.3a), which results in a tomographic reconstruction containing artifacts and low SNR (Fig. 4.3b,c). Despite the poor quality, traditional chemical tomography tracks the chemical distribution, and the Mn shell (orange) can be seen surrounding the Co core (blue-green). In contrast, elastically scattered electrons collected by the HAADF detector provide high signals at lower doses and allow many projections to be collected—in practice, HAADF sampling is five to ten times more finely spaced than chemical (Fig. 4.3d) [76]. The dose required for a single HAADF projection

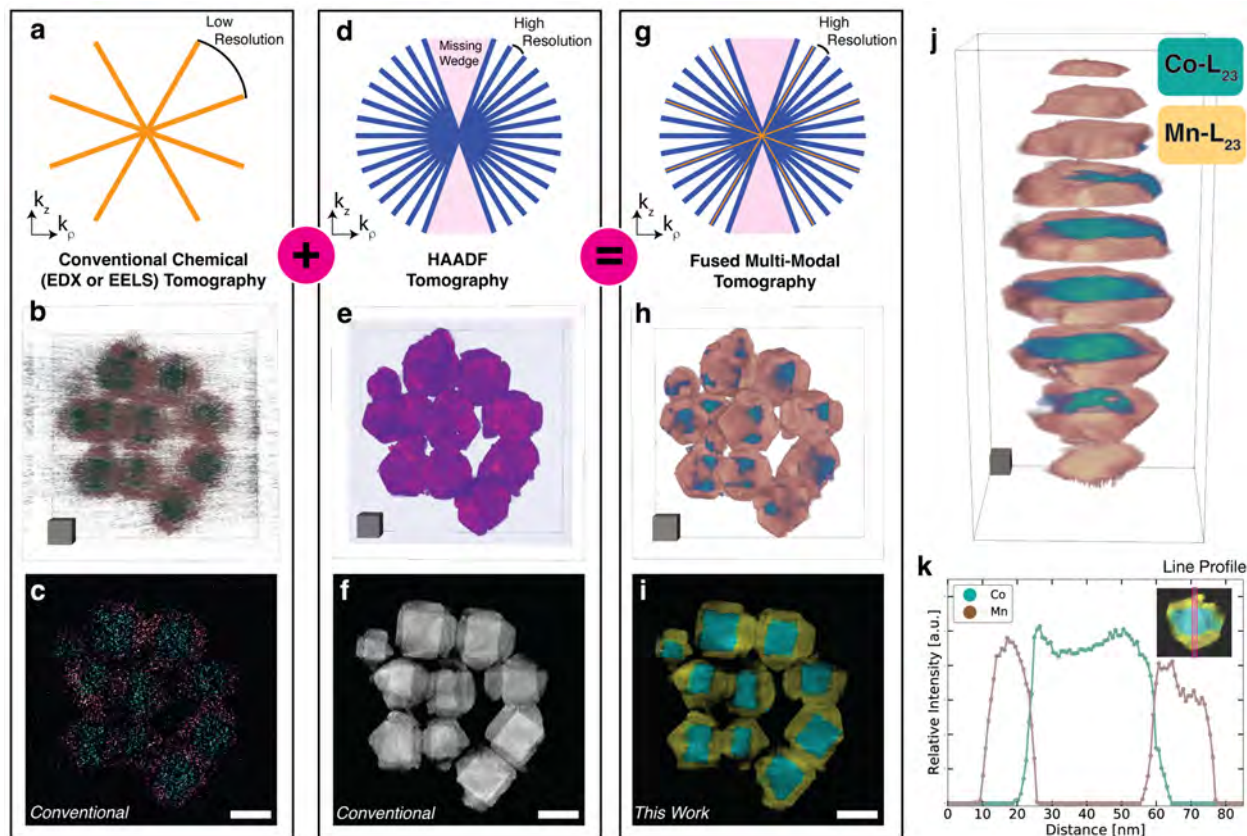


Figure 4.3: **Nanoscale recovery of $\text{Co}_3\text{O}_4\text{-Mn}_3\text{O}_4$ core-shell nanoparticles.** **a-c** Raw EELS reconstruction for the Co (blue-green) and Mn (orange) $L_{2,3}$ core-loss edges. **d-f** The HAADF tomogram of $\text{Co}_3\text{O}_4\text{-Mn}_3\text{O}_4$ nanoparticle tracks the structure of the specimen but fails to describe materials chemistry in 3D. **g-i** The fused multi-modal reconstruction. Scale cubes, 25 nm^3 . **a,d,g** Representation in Fourier space of the projections used to reconstruct the tomograms. **j** Fused multi-modal tomogram of a single $\text{Co}_3\text{O}_4\text{-Mn}_3\text{O}_4$ nanoparticle. Scale cube, 10 nm^3 . **e** A line profile showing the average intensity across the diameter of the particle.

is $10^2\text{-}10^3$ times lower than a chemical projection acquired using core-energy loss spectroscopy. Thus, it is favorable to acquire more HAADF images and achieve higher resolution. Although HAADF tomography permits high-resolution and high-SNR reconstructions of structure, it lacks chemical specificity. This is seen in Figure 4.3e,f where the structure is well defined with low noise but the Co and Mn regions are not identifiable.

Exploiting shared information in both modalities, multimodal tomography achieves a chemical resolution in 3D comparable to high-resolution HAADF reconstructions. Although few chemical measurements pose a severely underdetermined problem, fusing with the HAADF modality fills in missing chemical information. This is reflected in Figure 4.3g where many HAADF projections (e.g., 50-180) are measured while far fewer chemical projections (e.g., 5-15) are intermittently measured. In this reconstruction, 9 EELS maps and 45 HAADF projections (50-200 mrad detector

inner and outer semi-angles) were collected over a $\pm 60^\circ$ tilt range using a 2.4 Å probe with a 24.3 nm depth of focus (300 keV acceleration voltage, 10 mrad convergence angle). High-resolution of 3D chemistry is visible in the the core shell $\text{Co}_3\text{O}_4\text{-Mn}_3\text{O}_4$ using multi-modal tomography in Figure 4.3h,i.

Fused multi-modal electron tomography provides unique insight for studying heterostructured nanocrystals with unprecedented geometries. In the case of Co_3O_4 - Mn_3O_4 nanocrystals, the manganese oxide shell is divided into several ordered grains that grow on each surface plane for the cobalt oxide nanocube core [121]. However the core and shell interface can vary per plane driven by the growth interplay between strain and surface energy, resulting in the formation of grain boundaries [151]. The complete 3D distribution of Co and Mn at the surface and interface is difficult to discern with 2D projected EELS maps or HAADF reconstructions. Fortunately, the fused chemical distributions reveals surface coverage of the shell grains and cross-sections quantify the shell thickness and interface chemistry.

In cases where the HAADF distribution fails to provide any insight into the unique chemical phases of transition metal heterostructures, fused multi-modal still succeeds. Figure 4.4 highlights fused chemical tomography on heterostructured nanocrystals with applications in photovoltaic devices and battery electrodes [62]. The copper sulfide properties are sensitive to stoichiometry and crystal structure at the interface. The HAADF tomograms (Fig. 4.4a,d) fails to discern the Cu and Zn rich phases due to minimal difference in atomic numbers between the two elemental species. The voxel and pixel intensity across the interface is homogenous. The raw EELS spectra (Fig. 4.4e) highlights in difference in SNR between the metal and light small atomic number species (carbon and sulphur) where the number of counts is orders of magnitude smaller for the Cu and Zn core-loss edges. This occurs because the excitation events for metal core edges is far fainter than light elements. We can visualize the improvements from the raw and fused chemical reconstructions where now the rich unique 3D phases have been recovered (Fig. 4.4b,c,d).

Fused multi-modal chemical tomography is flexible with any form of inelastic scattering spectral signals where the measured chemistry is available. Figure 4.5 highlights the 3D chemical distribution of Cu-supported silicon carbon (SiC) catalysts designed for the production of fuels and chemicals from biomass. These silica supported catalysts efficiently convert ethanol into acetaldehyde because of their high selectivity and stability [97]. The EDX spectra confirms the composition of these nanoparticles, with additional Au background signal excited by the microscope holder. The HAADF reconstruction reveals the specimen morphology, while the contrast between the Cu and SiC composition is small (roughly 30% difference). These nanoparticles are impregnated with Cu particles that vary in size from 6 to 40 nm in diameter. Through the fused EDX tomogram, the rich Cu size distribution is entirely recovered at a far higher SNR than the raw EDX maps (Fig. 4.5b). In addition, the small nanoparticles are resolved at near the Nyquist limit

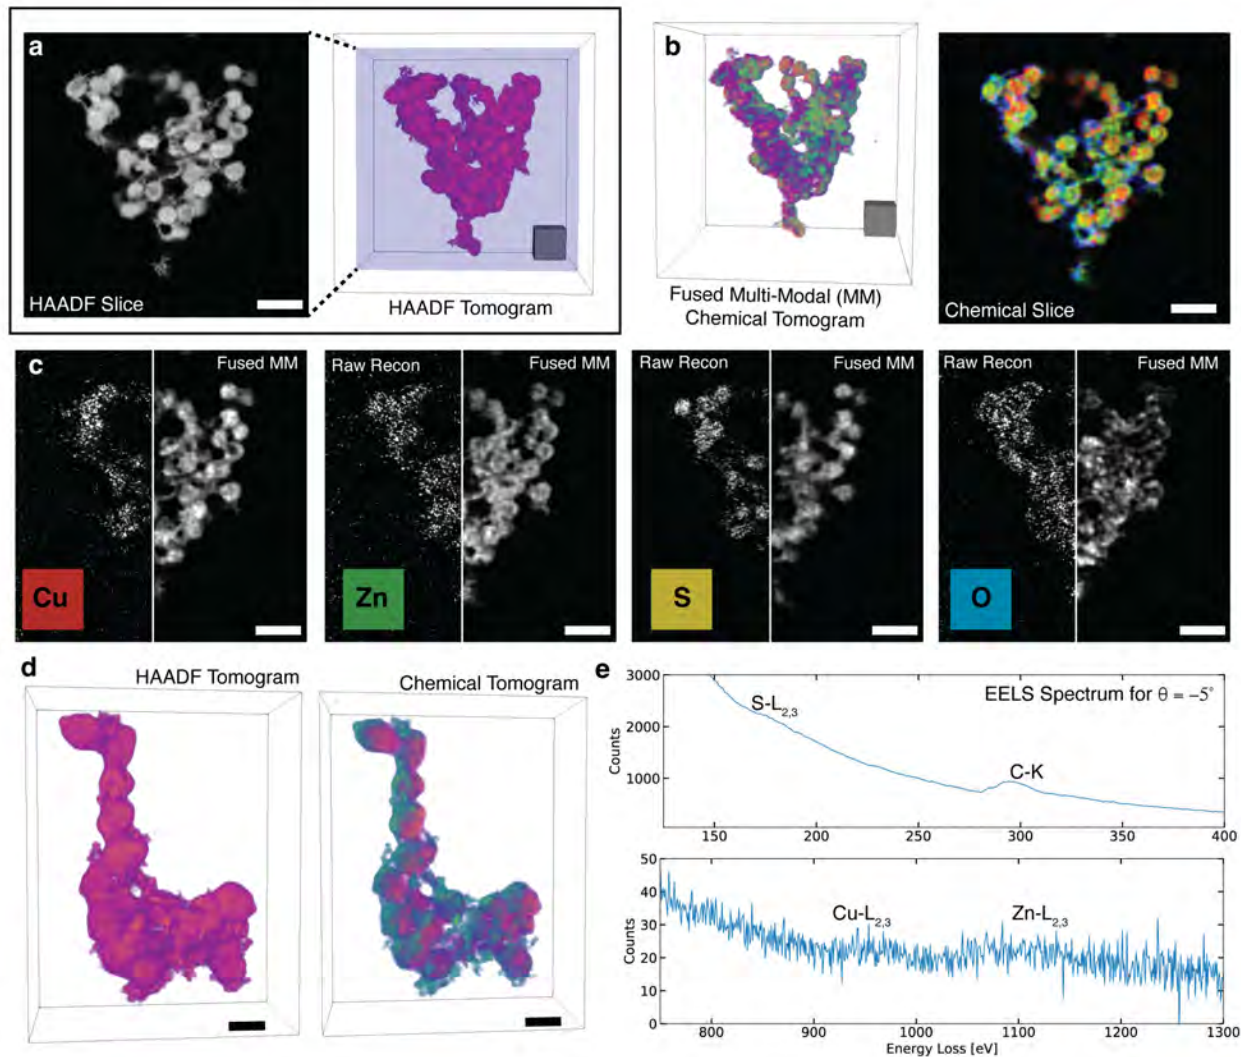


Figure 4.4: **Multi-Modal EELS Tomography of ZnS - Cu_{0.64}S_{0.36} Heterostructured Nanoparticles.** **a** The HAADF reconstruction and 2D slice are shown on the left. **b** The fused multi-modal reconstruction illustrating CuS or ZnS-rich nanoparticles and oxidized shells. **c** 2D slices of the chemical reconstructions with the noisy traditional reconstructions highlighted on the left of each image. Scale bar, 50 nm. **d** The HAADF and fused multi-modal chemical tomogram for a smaller field of view. **e** Representative EELS spectra for the S, C, Cu, and Zn core loss edges.

(Fig. 4.5c) from the sampling of this experiment (1.47 nm pixel size).

Data fusion eliminates noticeable noise in the final 3D chemical reconstruction without a loss of resolution. This noise reduction accompanies a dose reduction of roughly one-hundred fold. Linking the chemical projections to the high SNR HAADF signals dose-efficiently boosts the chemical specificity. Figure 4.6 highlights the relationship between fused reconstruction accuracy with SNR on either modality. Even at modest HAADF signals (e.g., SNR \simeq 10), multi-modal tomography notably outperforms traditional chemical tomography (Fig. 4.6b). To illustrate, in

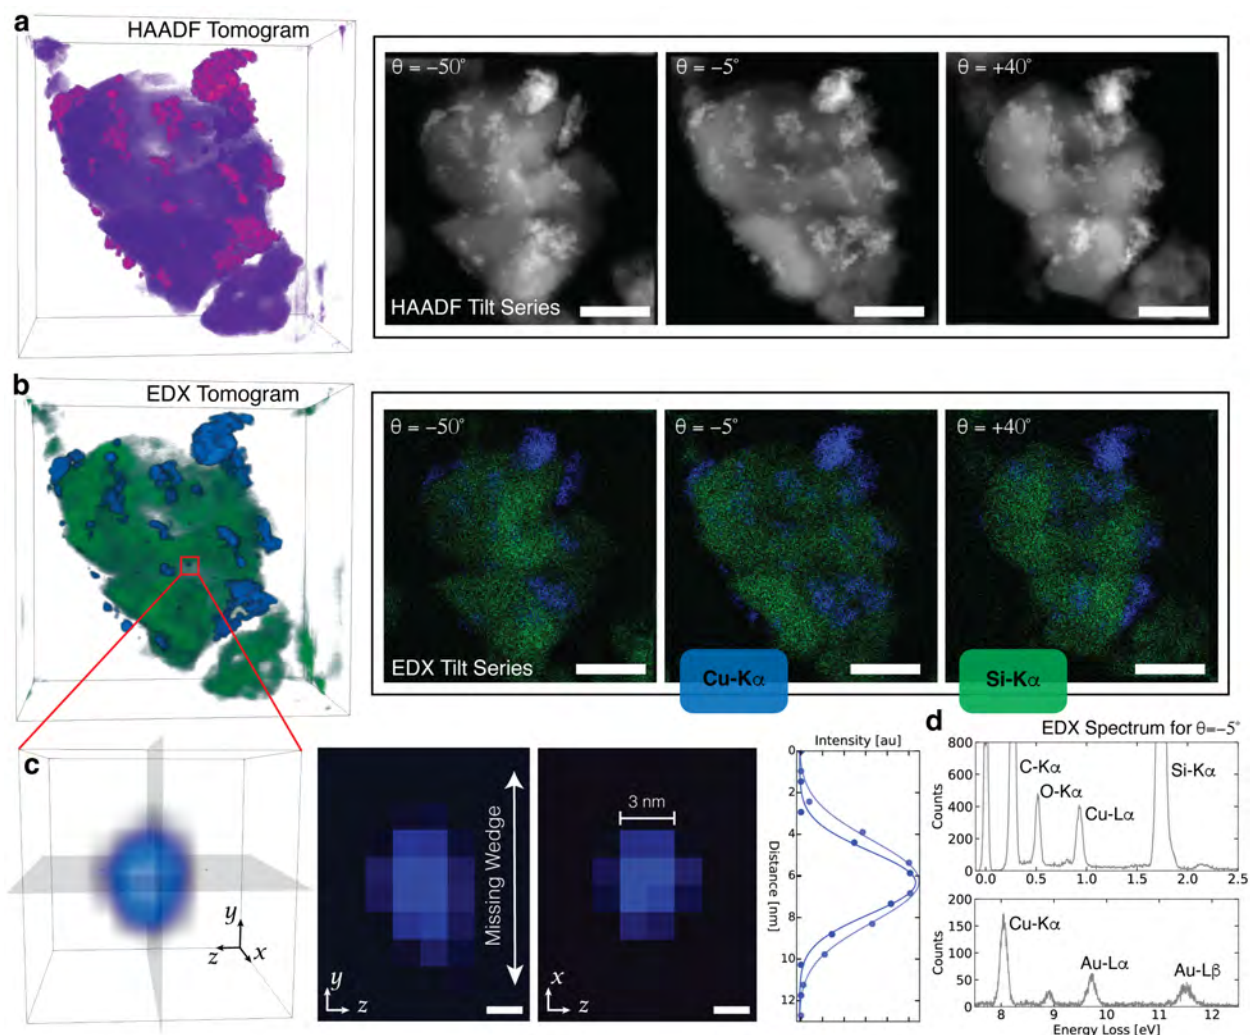


Figure 4.5: Multi-Modal EDX Tomography of Cu-SiC Nanoparticles. **a** The HAADF reconstruction with a few tilt micrographs is shown on right. **b** The fused multi-modal reconstruction highlighting Cu nanoparticles embedded inside the SiC support and raw EDX maps are shown on the right. Scale bar, 50 nm. **c** The PSF of an individual 3 nm nanoparticle inside the SiC. A few 2D slices of the reconstruction are shown on the right where we see the structure is sharp along the orthogonal axis perpendicular to the missing wedge and approximately a 20% reduction in resolution along the missing wedge direction. Scale bar, 2 nm. **d** EDX spectra for a single tilt.

Figure 4.3, matching the resolution of fused multi-modal chemical tomography using traditional methods would require 45 EELS maps—a five-fold dose increase. However, the SNR of each chemical projection would still fall short and requires roughly twenty-times additional dose. In total, multi-modal chemical tomography performs well at one-hundredth the dose requirement of traditional methods. Moreover, we can compare the performance of fused tomography against conventional algorithms to see how chemical SNR can influence the accuracy. Figure 4.6c shows fused multi-modal tomography consistently outperforms traditional and regularized chemical to-

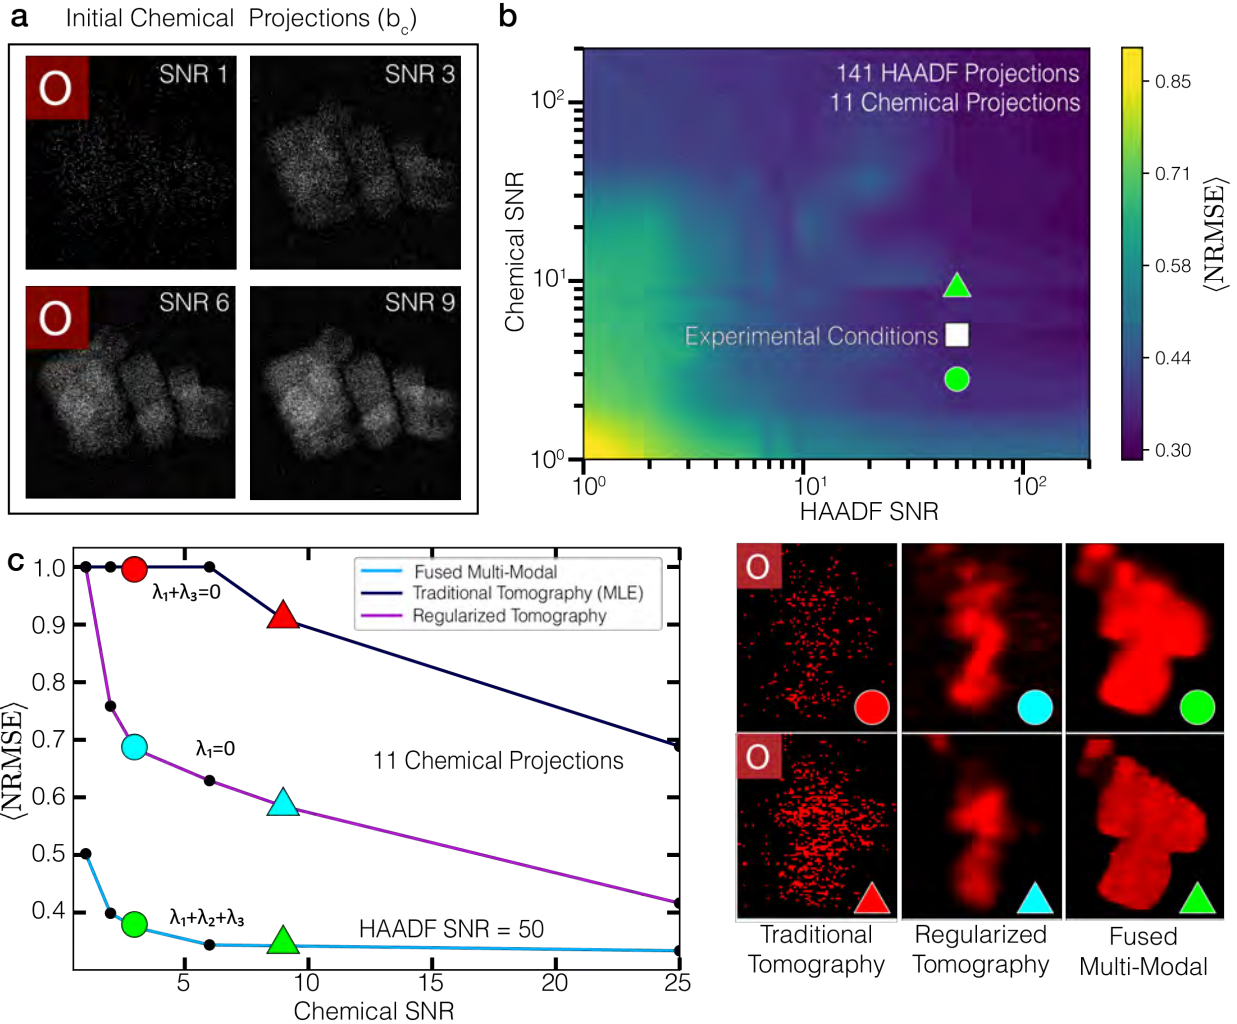


Figure 4.6: SNR Dependency for Successful Fused Multi-Modal Recovery. **a** The initial corrupted chemical distributions for oxygen in the CuO-CoO synthetic dataset with increasing SNR. **b** A heat map expressing the relationship between average reconstruction error and SNR for either modality (HAADF or Chemistry) when 11 chemical maps ($\Delta\theta = +12^\circ$) and 141 HAADF projections ($\Delta\theta = +1^\circ$) are available. **c** SNR plot highlighting the average NRMSE as a function of chemical SNR for reconstructions without any regularization or fusion (traditional tomography), without fusion (compressed sensing) and within the multi-modal framework. 2D slices from oxygen is shown on right.

mography (single modal compressed sensing algorithms) by 30-60% across all SNR. Eventually regularized tomography reaches comparable average NRMSE at an SNR of 50.

Reduction of electron beam dose produce irreplaceable advantages for electron tomography—both in terms of accessible resolution and the range of materials classes that can be imaged in 3D. Dose requirements for tomography scale quickly with higher resolution (resolution \propto dose $^{-4}$) [70, 105]. In specific, For 3D chemical imaging, multi-modal electron tomography notably im-

proves the sampling and dose constraints that limit resolution across a range of radiation sensitive materials (Fig. 4.7).




Figure 4.7: **Resolution and Electron Dose for Multi-Modal Electron Tomography.** **a** Resolution and dose relationship for electron tomography approximates the best achievable resolution for each material [116, 124, 92, 163, 30] – assuming an image contrast of 80% [105]. Fused multi-modal electron tomography results in a much higher dose-efficiency which enables higher resolution at any dose limit. **b** Multi-modal (green) and conventional chemical tomography (blue). This relationship between dose and resolution assumes sufficient tomographic sampling is achieved (i.e. many projections)—in practice the actual resolution will be much lower. Dose limited resolution assumes the material is adequately sampled (i.e. Crowther and Nyquist relations)

Ultimately, through the multi-modal data-fusion framework chemical tomography is now achievable with orders of magnitude of less dose compared to conventional approaches. Figure. 4.8 highlights the amount of dose used to perform chemical tomography with conventional approaches (CT) and with our multi-modal framework (MM). We see in all cases, our experiments shown herein use at least 2 orders of magnitude less dose and achieves lower Nyquist limited resolutions than ever perform.

4.4 Sub-nanometer Chemical Resolution in 3D

3D resolution of the chemical distribution in Au-Fe₃O₄ nanoparticle superlattice (Fig. 4.9a) is demonstrated at or below 1 nm using multi-modal tomography. The achieved resolution is quanti-




Figure 4.8: **Historical Demonstrations of Multi-Element Chemical Tomography.** The reported dose and Nyquist limited resolutions for the fused multi-modal (MM) reconstructions reported in this manuscript are compared to previous multi-element chemical tomography (CT) experiments [95, 78, 134, 15, 155, 126, 79]. Note, the actual achieved 3D resolution of previously reported chemical tomography may be lower than the Nyquist resolution.

fied in real and reciprocal space. In real space, the resolution limit is verified by visually inspecting a single 3 nm Au nanoparticle (Fig. 4.9d). The edge sharpness between the reconstructed nanoparticle and vacuum is visibly less than 1 nm. From line profiles, the half pitch resolution is 0.8 nm \times 0.8 nm \times 1.1 nm along the x , y , and z directions respectively. Along optimal directions (x , y) the resolution is comparable to the Nyquist frequency (8.05 Å). The real-space resolution is consistent with reciprocal space estimates of the cutoff frequency at which the signal drops to the noise floor [109]. Figure 4.9b highlights power spectral density variations projected on three orthogonal planes. Measured power spectral density along the k_x - k_y and k_x - k_z directions shows information transfer roughly occurring at 0.99 nm and 1.02 nm respectively (Fig. 4.9c). These directions conservatively represent the 3D resolution from an average of the high-resolution and low-resolution (z -axis) directions. This 3D chemical resolution nearly matches the 3D HAADF resolution 1.00 nm, 1.01 nm in Figure 4.9 (Supplementary Fig. 4.10).

For fused multi-modal chemical tomography, the HAADF 3D resolution provides a new upper bound to the highest obtainable 3D chemical resolution. A reduction of resolution along the z -axis is expected from the incomplete tilt range that creates a missing wedge of information in Fourier space [93]. Avoiding this anisotropic resolution loss has been demonstrated by acquiring a full tilt

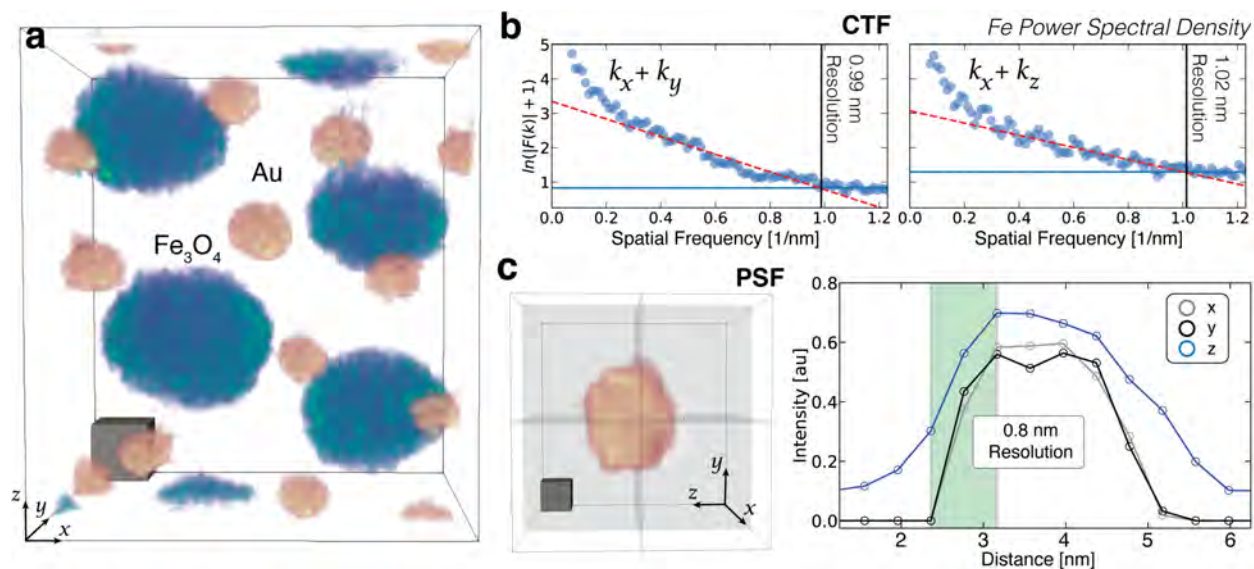


Figure 4.9: **Resolution Analysis of Au-Fe₃O₄ superlattice nanoparticles.** **a** Fused EELS tomograms of Au-Fe₃O₄ nanoparticles. Scale cube, 2 nm³. **b** Power spectral density of the Fe reconstruction along the principal axial directions shown on the right. Scale bar, 0.5 nm⁻¹. **c** Power spectral density profiles for k_x-k_y and k_x-k_z directions. **d** Line profiles of a 2.5 nm Au nanoparticle gives a resolution of 0.8 nm, 0.8 nm, and 1.1 nm along the x, y, and z directions.

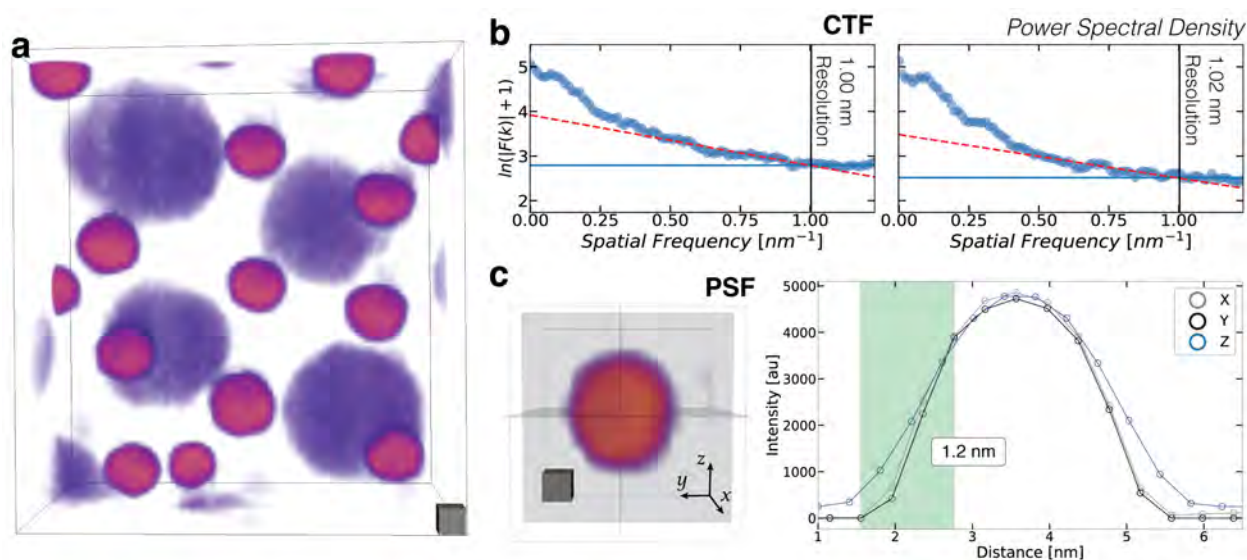


Figure 4.10: **HAADF Resolution Analysis of Au-Fe₃O₄ superlattice nanoparticles.** **a** Fused EELS tomograms of Au-Fe₃O₄ nanoparticles. Power spectral density of the HAADF reconstruction along the principal axial directions shown on the right. Scale cube, 2 nm³. Scale bar, 0.5 nm⁻¹. **b** Power spectral density profiles for YZ and XY planes. **c** Line scan profiles of a 2.5 nm Au nanoparticle give a resolution of 1.00, 1.03, and 1.01 nm along the x, y, and z directions.

range ($\pm 90^\circ$) through the preparation of needle wire samples or preparing nanoparticles on carbon nanofibers [167, 123]. Here, we observe approximately a 25% reduction in resolution along the missing wedge direction of the multi-modal chemical reconstruction.

4.5 Influence of Sampling

Electron tomography simulations show a 3-5 fold improvement in the normalized root mean square error ($\langle \text{NRMSE} \rangle$) averaged across all elements when multi-modal tomography is used over conventional chemical tomography. In Figure 4.11 synthetic gold decorated CoO / CuO nanocubes inspired by real experimental data [123] provide a ground truth comparison to assess the accuracy of fused multi-modal tomography. Simulated projection images are generated from a simple linear incoherent imaging model of the 3D chemical composition with added Poisson noise (See Methods). The specimen tilt range is limited to $\pm 70^\circ$ to better match typical experimental conditions. The advantages of multi-modal tomography are clearly visible in the 2D slices (Fig. 4.11b) taken from 3D reconstructions obtained by conventional chemical tomography ($\langle \text{NRMSE} \rangle = 1.301$) and fused multi-modal tomography ($\langle \text{NRMSE} \rangle = 0.33$). For all chemistries (Au, O, Cu, Co,) fused multi-modal tomography is more consistent with the ground truth with higher resolution and reduced noise.

For any number of chemical projections acquired, we see a notable reduction in NRMSE when HAADF projections are integrated into the chemical reconstruction. Figure 4.11 shows the improved fused multi-modal reconstruction accuracy across a wide range of HAADF and chemical projections for the gold-decorated CoO / CuO nanocubes. The reconstruction error (average NRMSE) across most of the multi-modal parameter space is less than 0.40 compared to values around 1.2 for conventional tomography. Pixel values on the diagram (Fig. 4.11a) represent the average NRMSE across all of the elements. This NRMSE map shows data fusion strongly benefits by increasing the HAADF information available. It requires substantially less dose to increase the HAADF projections (i.e. moving vertically on the map) compared to increasing the chemical projections (i.e. moving horizontally on the map).

Conventional chemical tomography does not use HAADF projections (bottom row, Fig. 4.11a) resulting in an average reconstruction error larger than the entire multi-modal regime. In practice fused multi-modal tomography is performed in the regime with equal or more HAADF projections than chemical (i.e. top-left triangle). Multi-modal tomography also performs well when the chemical projections exceed the number of HAADF projections, however, this is not practical since HAADF signals can be acquired simultaneously with EDX and EELS.

Determining optimal regularization parameters for the phase diagram (Fig. 4.11a) is computationally expensive to explore due to its variability across sampling conditions. While grid search

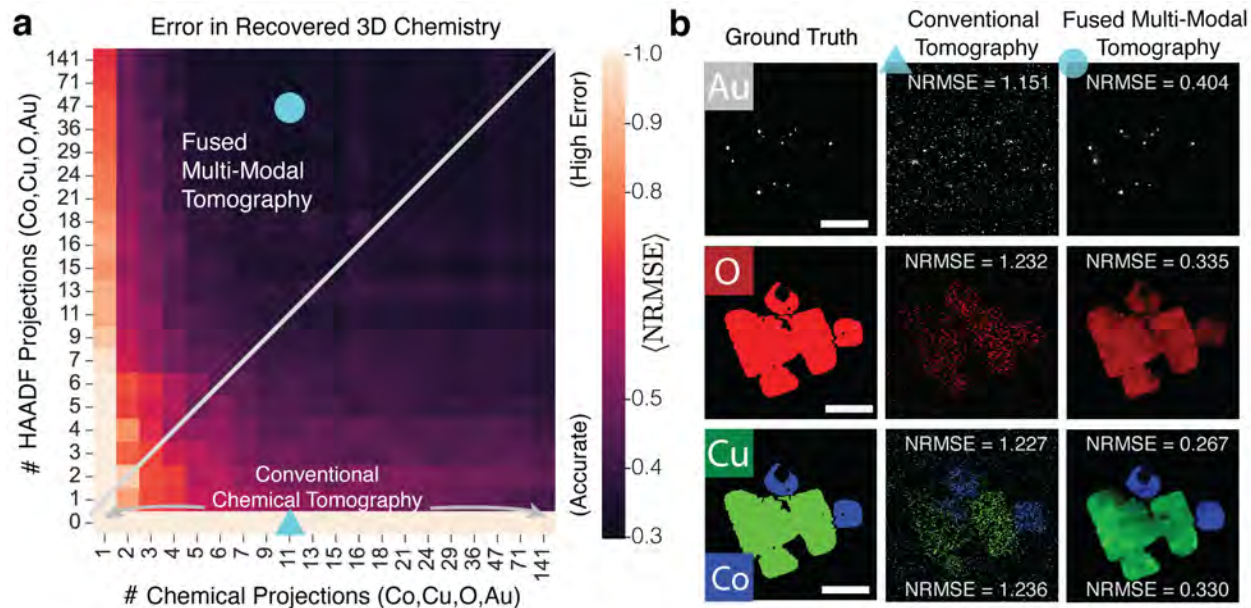


Figure 4.11: **Estimating Sampling Requirements for Accurate Recovery with Synthetic CoO/CuO Nanocubes.** **a** An NRMSE map representing the reconstruction error as a function of the number of HAADF and chemical tilts. Brighter pixels denote results containing incorrect reconstructions from the ground truth. **b** Visualization of three points corresponding to conventional chemical tomography (reconstruction without the HAADF modality), and low-dose fused multi-modal electron tomography. **c** The 3D models used for generating synthetic chemical and ADF projections. Scale bar, 75 nm.

could find the best parameters by exhaustively exploring all possible candidate values, the computation time would be expensive as each map would take approximately 125 days to complete on a single GPU.

We efficiently explored the parameter space with Bayesian optimization (BO) — a machine learning framework known for optimizing expensive unknown objective functions with minimal evaluations [178, 23]. It works by building a probabilistic model of the objective function with Gaussian processes (GP) regression. GP not only estimates our function of interest but also provides the uncertainty measurements to guide future predictions. BO takes into account past evaluations when determining future hyperparameter selections via an acquisition function [84]. For our simulations, we carried out BO with GP with the Matern kernel and GP Hedge acquisition strategy [13]. By exploiting BO with GP, we are able to provide an atlas of balanced hyperparameters for Eq. 4.1 with the CoCuO synthetic datasets. Overall, these maps can guide future scientists to produce multi-modal reconstructions with reasonable hyperparameter selections.

Asynchronous parallel BO on supercomputing resources allowed us to efficiently run several reconstructions simultaneously on a single node. This form of parallel computing resulted in sev-

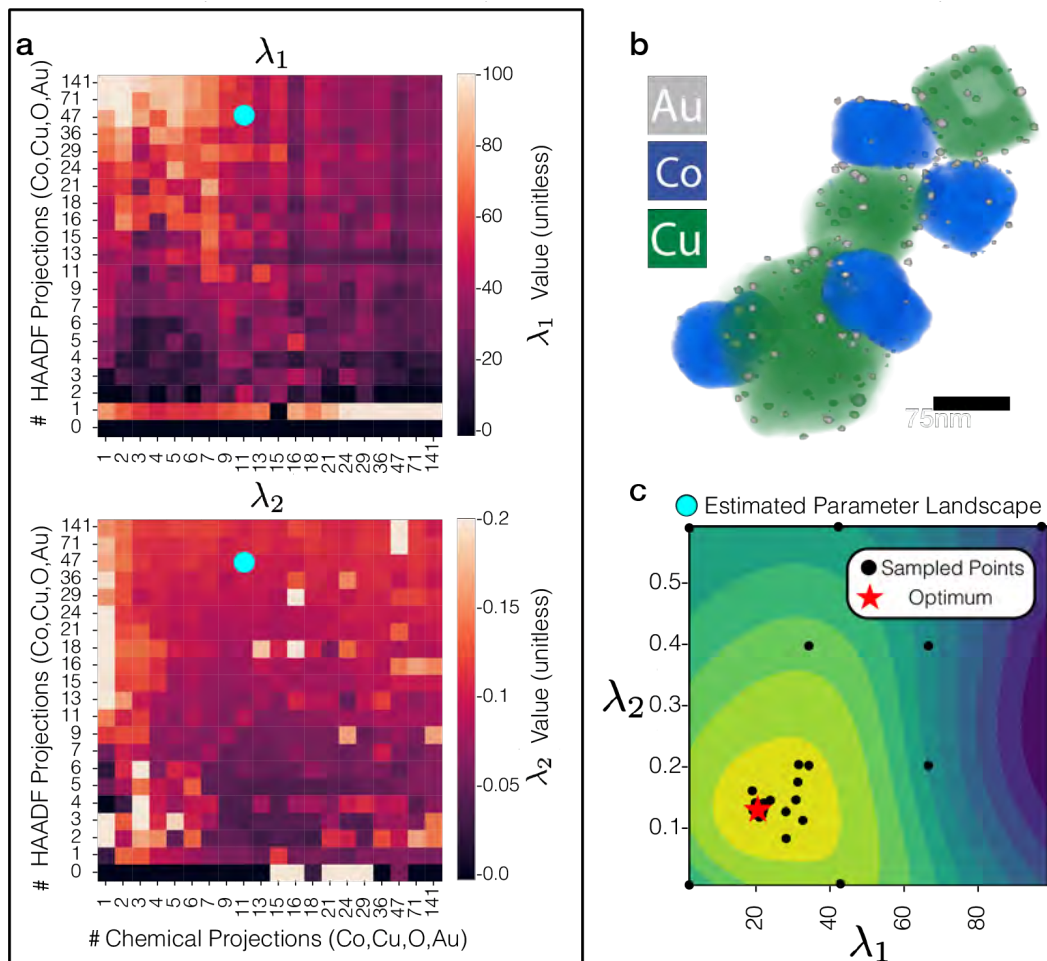


Figure 4.12: **Hyperparameter Estimation with Bayesian Optimization for the CuO-CoO Nanocubes.** **a** Bayesian optimization optimizes the data fusion cost function (shown above) when provided a given number of chemical and HAADF tilts. **b** 3D visualization of the ground truth Au decorated CuO/CoO nanocubes. Scale bar, 75 nm. **c** Bayesian optimization parameter selection landscape where each black dot represents one of the many attempts to find the minimum NRMSE.

eral factors of computational speed up as multiple GPUs received unique experimental parameters (e.g., SNR or sampling) to reconstruct concurrently amongst each other. Specifically, the computation time to generate an NRMSE map was reduced by 99.8% – taking less than a day to complete (18 hours) using GPU cluster resources available through Argonne and Oak Ridge National Laboratories (Summit-OLCF and ThetaGPU-ALCF).

Fused multi-modal electron tomography provides direct measurement of materials 3D chemistry. This means the ratio of voxels intensity reveals the specimens 3D stoichiometry without knowledge of inelastic cross-sections. For the simulated CoO-CuO nanocubes, values agree with the ground truth (Fig. 4.13)—concentrations of Cu, Co, and O are centered at the expected value of 0.50. Here, stoichiometric precision of multi-modal tomography ($\sigma = 0.04$) is four-times better

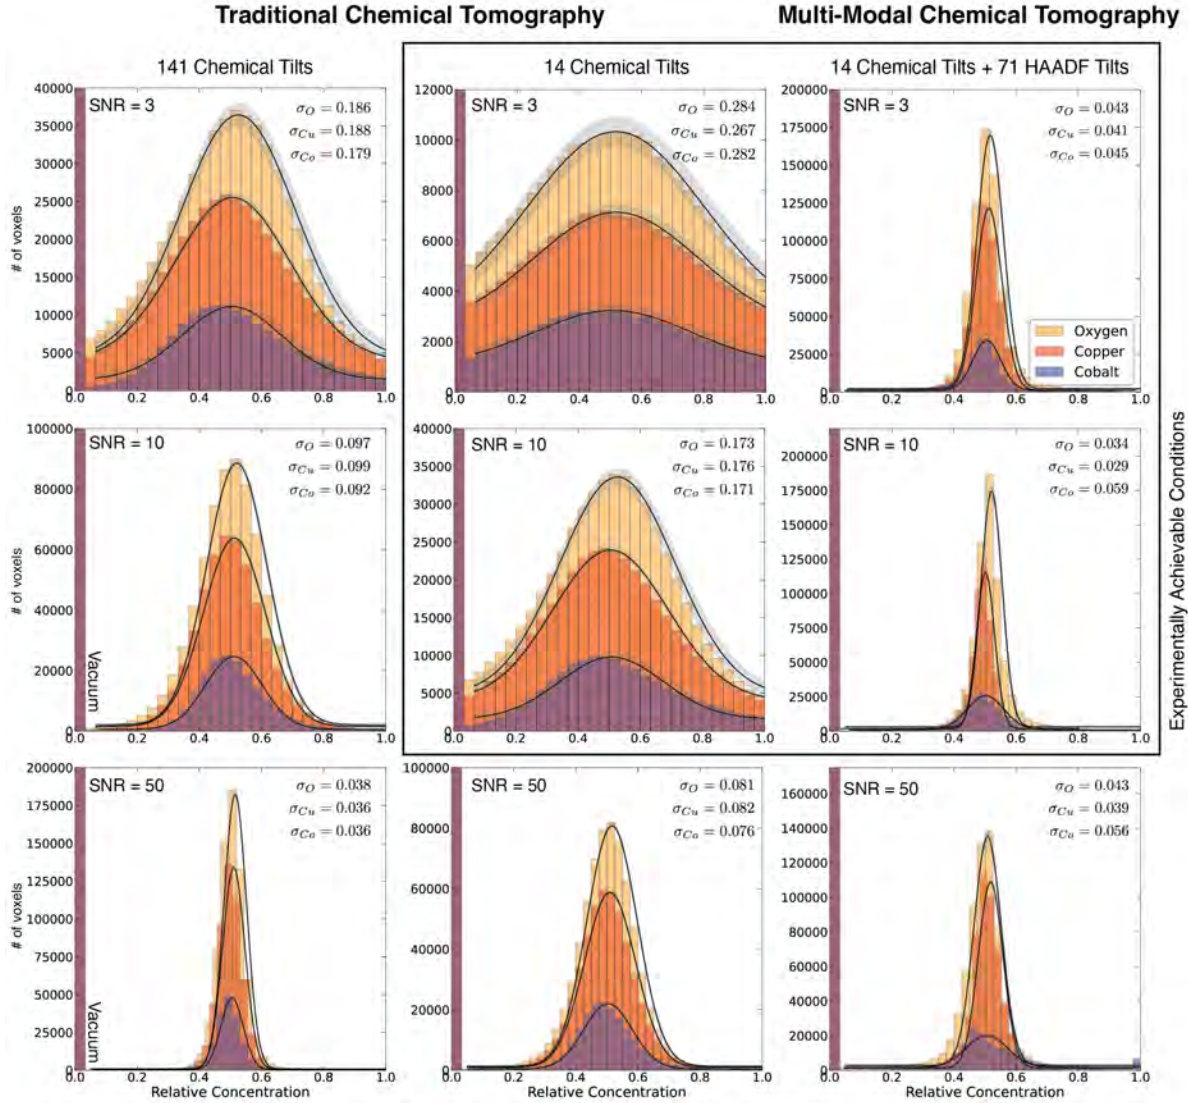


Figure 4.13: **Measuring 3D Concentration for simulated CuO-CoO nanocubes.** Histograms of chemical concentrations for each voxel in traditional and fused multi-modal tomography reconstructions. For traditional chemical tomography, the accuracy improves with SNR or increasing projections. Multi-modal tomography maintains low error, especially for experimentally realistic conditions.

than traditional chemical tomography ($\sigma = 0.17$).

For experimental Fe_3O_4 nanoparticles (Fig. 4.14), multi-modal tomography produces an average Fe concentration of 0.46 (0.43 expected) with a standard deviation of 0.15. Note, determining stoichiometry using traditional chemical tomography also requires accurate calculation of the inelastic cross-sections for each experiment [47].

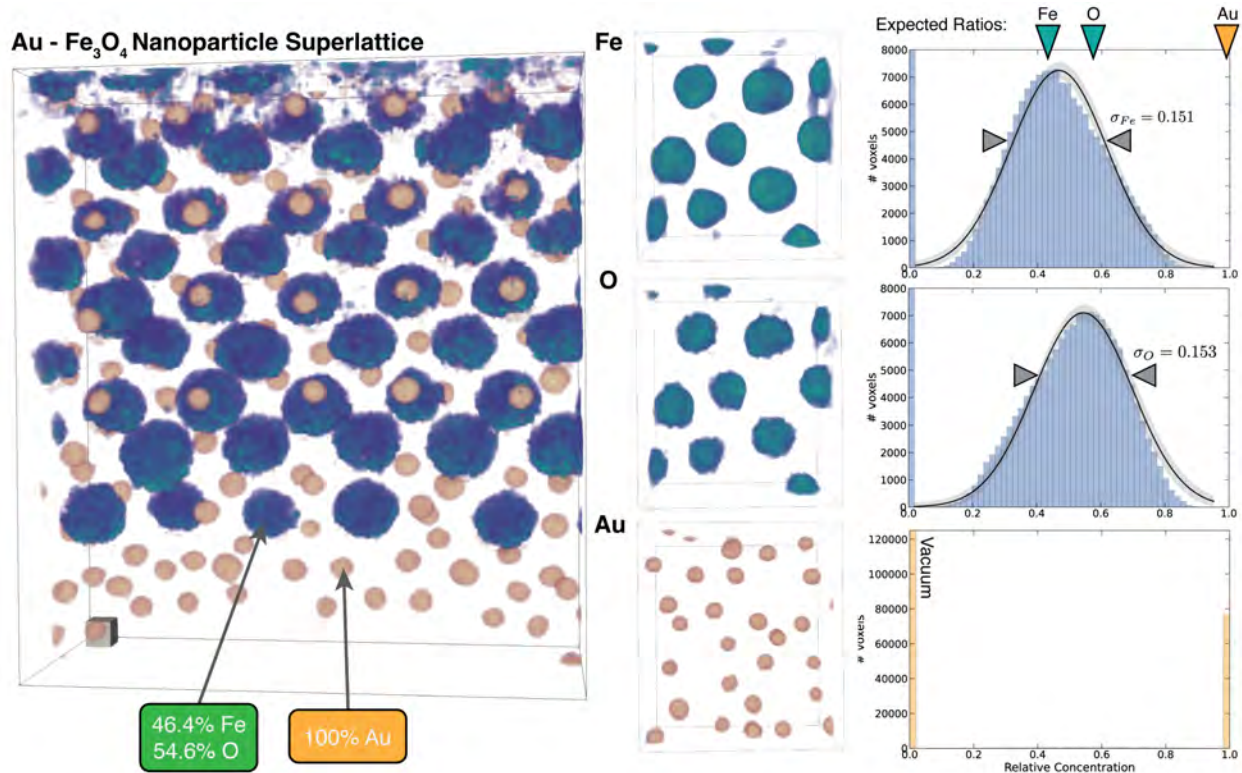


Figure 4.14: **Measuring 3D Stoichiometric Concentration of Au-Fe₃O₄ Superlattice Nanoparticles.** 3D chemical reconstructions for each element are shown with their corresponding voxel intensity histograms. The mean values and standard deviations are $46.4 \pm 15.1\%$, $54.6 \pm 15.3\%$, $100 \pm 0\%$ for Fe, O, and Au, respectively. The expected stoichiometry of this system is 42.9%, 57.1%, 100%.

4.6 Discussion

While this chapter highlights the advantages of fused multi-modal electron tomography, the technique is not a simple black-box solution. Step sizes for convergence and weights on the terms in the cost function (Eq. 4.1) must be reasonably selected (Fig. 4.15).

Standard spectroscopic pre-processing methods become ever more critical in combination with multi-modal fusion. Improper background subtraction of EELS spectra [114] or overlapping characteristic X-ray peaks that normally cause inaccurate stoichiometric quantification also reduces the accuracy of fused multi-modal tomography. Although lighter elements have smaller elastic cross-sections, they tend to have larger inelastic cross-sections which benefits chemical tomography. For example, the K-shell cross-section (chemical spectroscopic signal) of Carbon ($Z=6$) is over 20-fold larger than Germanium ($Z=32$) [46, 48]. EELS is advantageous for discerning lighter elements whereas overlapping peaks may occur in EDX. In general, electron tomography is favorable for measuring volumes in the range of $(10 \text{ nm})^3$ to $(1000 \text{ nm})^3$ at resolutions around 3

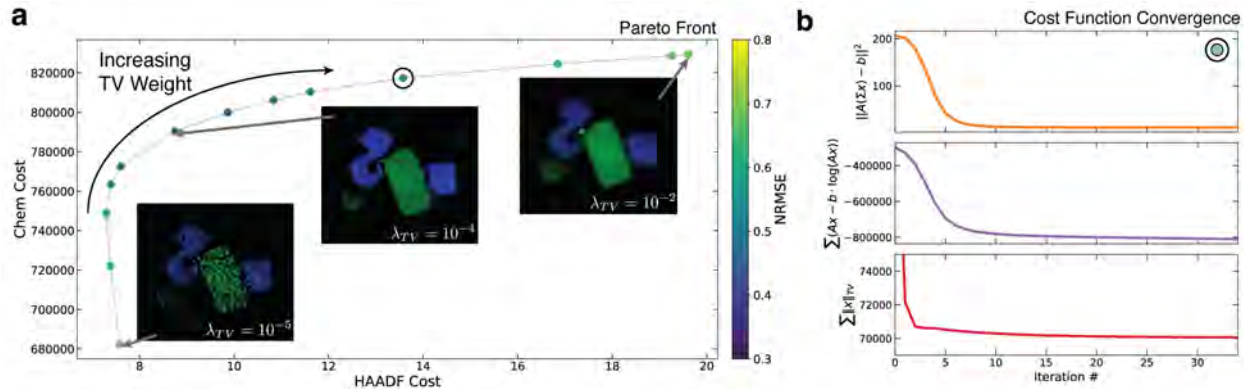


Figure 4.15: **Assessing Convergence and Selecting Hyperparameters with Pareto Front Curves.** **a** Pareto fronts illustrates the relationship between reconstruction quality and regularization parameters for multi-modal electron tomography. Depicted are the tradeoffs from three reconstruction evaluation metrics: the multi-modal, self-consistency and average NRMSE across all elements. We see the highest quality reconstruction (lowest NRMSE) occurs around the inflection point of the pareto front. **b** The three individual components in the cost function plotted throughout the multi-modal electron tomography reconstruction process illustrates smooth asymptotic convergence. Convergence should be confirmed for accurate reconstruction.

to 30 Å [96]. Thick specimens with dimensions that far exceed the mean free path of the electron can produce inversion contrast that will cause electron tomography to fail [53]—also causing failure for multi-modal electron tomography. Electron tomography performs best for thicknesses less than three times the incident electron’s mean free path (e.g., < 550 nm for Silicon at 300 keV) [108]. In all electron tomography experiments, beam convergence angles should be chosen to match the desired resolution and depth of focus [169]. As shown for 2D fused multi-modal electron microscopy [144], fused multi-modal tomography works best when elements have discernible contributions to the HAADF contrast and all chemical elements have been imaged. Multi-modal tomography leverages compressed sensing (e.g., TV min.) which assumes incoherence (i.e., a high level of dissimilarity) between the sensing and sparsifying transform [22, 103, 147]—although this assumption typically holds as demonstrated for the datasets presented herein.

4.7 Conclusion

In summary, we present fused multi-modal electron tomography that enables chemically-sensitive 3D reconstruction of matter with nanometer resolution at high SNR. Researchers no longer must choose between measuring 3D structure without chemical detail or characterizing chemistry along a single viewing direction. By linking signals from elastic (HAADF) and inelastic (EDX / EELS) scattering processes, the traditional dose limits of chemical tomography are substantially sur-

passed. In some cases, a one-hundred fold reduction in dose is estimated. To demonstrate, the complete volumetric density of each chemistry was mapped in several systems including Au-Fe₃O₄, Co₃O₄-Mn₃O₄, ZnS-Cu_{0.64}S_{0.36}, and Cu-SiC nanomaterials. In both synthetic and experimental datasets, fused multi-modal electron tomography shows substantial advantages in the accuracy of 3D chemical imaging. This approach enables chemical tomography of a wide range of previously inaccessible materials with moderate radiation sensitivity. At chemical resolutions of 1 nm, fused multi-modal electron tomography opens up new understanding of geometrically complex materials—from 3D semiconductor gate stacks [100], clean energy materials [10, 175], or photoluminescence quantum dot nanoparticles [165].

Here, fused multi-modal tomography used commonly available STEM detectors (HAADF, EDX, and EELS), however, this approach can be extended to other modalities in development—including pixel-array detectors [158], annular bright field [54], ptychography [82], low-loss EELS [63], etc. Furthermore, the tremendous potential of multi-modal data fusion as a paradigm readily enhances deep learning to capitalize on the unique advantages from both domains [65]. One can imagine a future wherein all scattered and emitted signals in an electron microscope are collected and fused for maximally efficient characterization of matter in all dimensions.

CHAPTER 5

Real-time 3D Analysis During Electron Tomography

1

The demand for high-throughput electron tomography is rapidly increasing in biological and material sciences. However, this 3D imaging technique is computationally bottlenecked by alignment and reconstruction which runs from hours to days. We demonstrate real-time tomography with dynamic 3D tomographic visualization to enable rapid interpretation of specimen structure immediately as data is collected on an electron microscope. Using geometrically complex chiral nanoparticles, we show volumetric interpretation can begin in less than 10 minutes and a high-quality tomogram is available within 30 minutes. Real-time tomography is integrated into tomviz, an open-source and cross-platform 3D data analysis tool that contains intuitive graphical user interfaces (GUI), to enable any scientist to characterize biological and material structure in 3D.

5.1 Introduction

Three-dimensional (3D) characterization across the nanoscale is now possible using scanning / transmission electron microscopes (S/TEM) [32, 107, 110, 149, 171]. In an electron tomography experiment, volumetric structure of biological or materials specimens are reconstructed from high-resolution projection images acquired across many viewing angles [133, 73]. Unfortunately, tomographic reconstructions can take one to several days to complete depending upon the dataset size or algorithm(s) employed. Even worse, the reconstruction occurs offline, long after all the data has been collected, preventing immediate interpretation during an ongoing experiment. While advancements in detector hardware have boosted throughput with digital data collection [143], substantial human effort and computational resources are still required to process the raw data before visualization. It has been a longstanding goal to begin 3D analysis of specimens in real-time to allow immediate assessment of nanoscale structure and data quality [17].

¹The results presented in this chapter lead publications in *Ultramicroscopy* and *Nature Communications* [148, 146].

Here we present facile 3D visualization of specimens during an electron or cryo-electron tomography experiment using the tomviz platform (tomviz.org). Our platform now provides interactive 3D material or biological structure in real-time to enhance high-throughput specimen interpretation. Tomviz offers multiple real-time reconstruction algorithms integrated into a fully graphical interface that presents the user with immediate visualization during data collection. Achieving high-throughput electron tomography requires an integrated pipeline that links the microscope hardware to optimized reconstruction algorithms and efficient 3D volumetric visualization. A multi-threaded data analysis pipeline runs dynamic visualizations that update as new data is collected or reconstruction algorithms proceed. Iterative reconstruction algorithms efficiently accommodate new data and keep pace with typical experimental acquisition rates. Scientists can interactively analyze 3D specimen structure concurrent with a tomographic reconstruction after or during an experiment. The robust graphical interface allows for 3D specimens to be rendered as shaded contours or translucent volumes that can be rotated, cropped, or sliced as the reconstruction occurs. In favorable cases, structural interpretation can begin as early as 10 minutes and a high-resolution volume is available after only 60% of data is acquired (~ 30 minutes). The latest tomviz release (v 2.0), is now packaged with real-time 3D analysis for electron tomography, is available as an open-source cross-platform tool with compiled binaries certified for Linux, Mac, and Windows.

5.2 Real-Time Tomography Workflow

The real-time tomography workflow is illustrated in Figure 5.1: electron micrographs are collected, passed to tomviz for reconstruction, and visualized as an interactive 3D rendering. This process runs simultaneously and continuously while the electron microscope is being operated. During experimental acquisition, tomviz monitors when new projections are collected (Fig. 5.1a) and appends new data into the reconstruction process. Importantly, tomograms are reconstructed in parallel with data acquisition. Real-time algorithms accommodate the arrival of new data without restarting the reconstruction process. Iterative reconstruction methods are made efficient for real-time processes by utilizing dynamic descent parameters.

Modifications to the common implementation for SIRT and TVmin were made to account for the dynamic addition of input projections throughout an experiment. SIRT seeks the minimal error between the reconstruction and experimental data: $\arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$ where \mathbf{A} is the measurement matrix, \mathbf{b} are the experimental projections and \mathbf{x} is the tomogram. We can further regularize the process through the assumption that our volumes should be piece-wise smooth and minimize its total variation $\|\mathbf{x}\|_{TV}$. Iterative algorithms require rescaling of the descent parameter based on the number of projections sampled. SIRT can easily estimate the descent parameter through calculation of the Lipschitz constant ($L = \|\mathbf{A}^T \mathbf{A}\|_2$). The Lipschitz constant can be

estimated by using the power method [86]. The descent parameter for TVmin is scaled by a dampening envelope that ensures its magnitude decays linearly [148]. Non-iterative algorithms such as WBP do not require rescaling of descent parameters and simply needs to reinitialize the computation with the new projection images collected.

Dynamic reconstructions maintain pace with typical experimental acquisitions (e.g. $512^3 - 1024^3$ voxels) using a personal computer. The intermediate reconstructions are rendered in 3D and immediately presented to the scientist (Fig. 5.1b). Thus, the tomogram dynamically improves with time as both the reconstruction algorithm converges and additional specimen information arrives. High-quality 3D reconstructions are available before the end of the experiment (Fig. 5.1c).

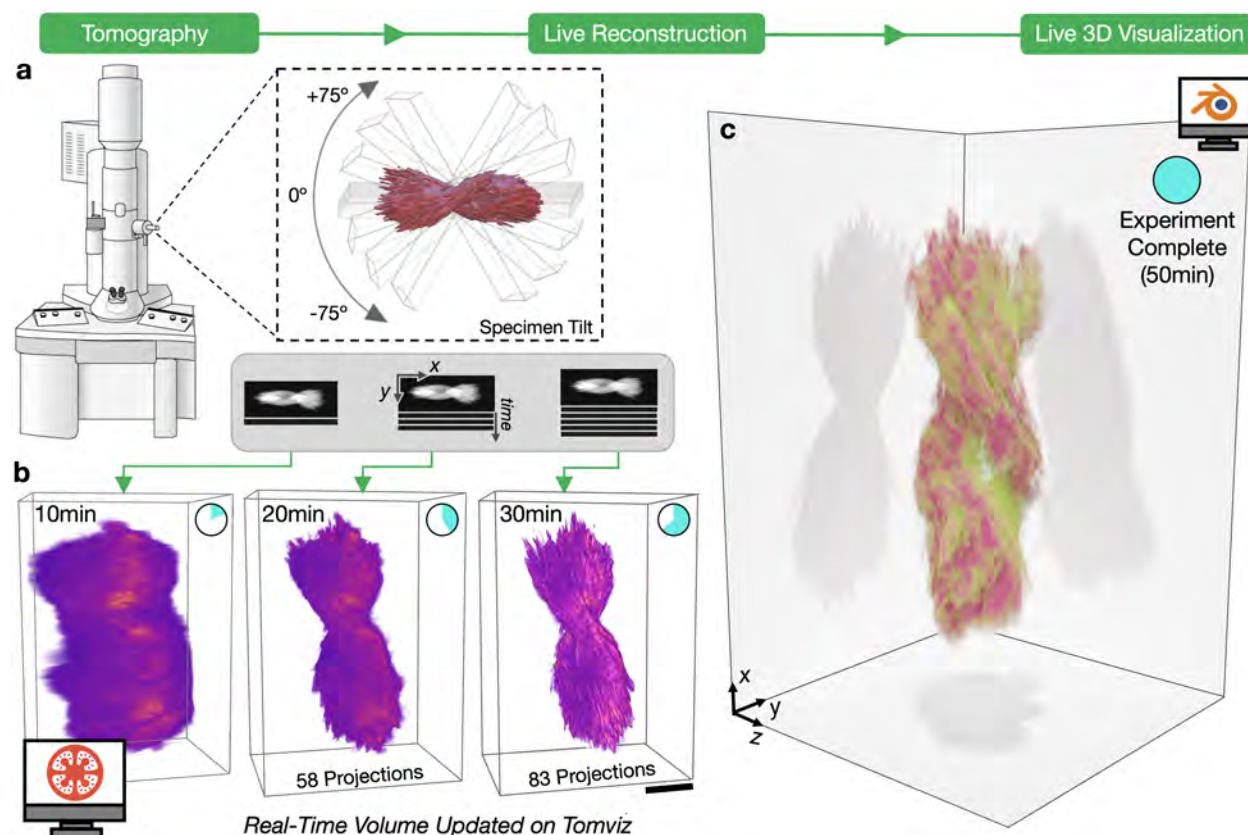


Figure 5.1: **Real-Time Electron Tomography Workflow of a helical nanoparticle visualized on tomviz.** **a**, Specimen projections are sequentially collected in an electron microscope across an angular range ($< \pm 75^\circ$) and continually passed to tomviz for reconstruction and live 3D visualization. **b**, As projections accumulate during the experiment, the reconstruction updates in real-time and resolution improves. Scale bar, 100 nm **c**, A high-quality tomogram is available for data interpretation upon the end of an experiment.

Direct visualization of a specimen's 3D structure enables immediate identification of morphological and internal information shortly after a tomography experiment begins. We demonstrate real-time tomography on a helical nanoparticle comprised of a chiral dipeptide Cystine amino-acid

coordinated with Cadmium (Cyst/Cd). The bowtie-shaped particles were synthesized using a size limited self-assembly process described by Yan *et. al.* [170]. These semiconducting nanoparticles contain strong tunable chiroptical properties due to a twisted geometry [170]. As shown in Figure 5.1b, the overall morphology for the Cd/Cyst nanoparticle can be observed in as early as 10 minutes and fine details are visible after 20-30 minutes of the experiment (roughly half-completion). The specimen's right handed chirality cannot be determined from a single projection image and requires 3D imaging (Fig. 5.1c). With real-time tomography the material's chirality and symmetry were identified within the first third of data acquisition (~ 15 minutes). This immediate feedback can save researchers days of effort as reconstructions are no longer processed offline. Moreover, real-time visualization allows quick adjustment and optimization of reconstruction parameters that can greatly influence the reconstruction quality. Ultimately, scientists can efficiently investigate 3D nanostructure during imaging to guide experiments and redefine scientific objectives while simultaneously operating the microscope.

5.3 Real-Time 3D Visualization During Reconstruction

Currently, the best tomographic reconstructions are obtained from algorithms that are slow and iterative. In practice, electron tomography experiments are limited by a finite and restricted angular range (e.g. $< \pm 70^\circ$) resulting in incomplete information that degrades resolution in 3D [122]. Iterative algorithms can recover tomograms with high spatial resolution and minimal reconstruction error [4]. While these algorithms better estimate 3D structure from under-determined measurements, they come at the expense of computational time [56]. Fortunately, using the tomviz tool, iterative reconstructions can be visualized in real-time throughout the arduous computation.

Real-time tomography greatly alleviates the wait-time by visualizing the intermediate 3D structure between algorithm iterations—beneficial during an experiment or analysis. Figure 5.2 demonstrates interactive visualizations of the Simultaneous Iterative Reconstruction Technique (SIRT) [57] for a cobalt phosphide (Co_2P) hyperbranched nanoparticle [176] (512^3 pixels volume reconstructed across the 363.52 nm full field of view). SIRT tomograms begin with a loose estimate [168] (Fig. 5.2a) and develop sharper, high frequency information with each increasing iteration (Fig. 5.2b-c). Compressed sensing algorithms such as total variation minimization (TVmin) seek maximally sparse solutions to recover high-resolution, low-noise structure using fewer projections than conventional methods [154, 137]. Figure 5.2e-f demonstrates an interactive 3D visualization using TVmin reconstruction of a iron platinum (FePt) nanoparticle at atomic resolution (256^3 pixels volume reconstructed across the 9.536 nm full field of view)—data provided and pre-processed by Yang *et. al.* [173]. This work replicates the atomic resolution tomogram using independent pre-processing and reconstruction methods. Recent developments in dynamic

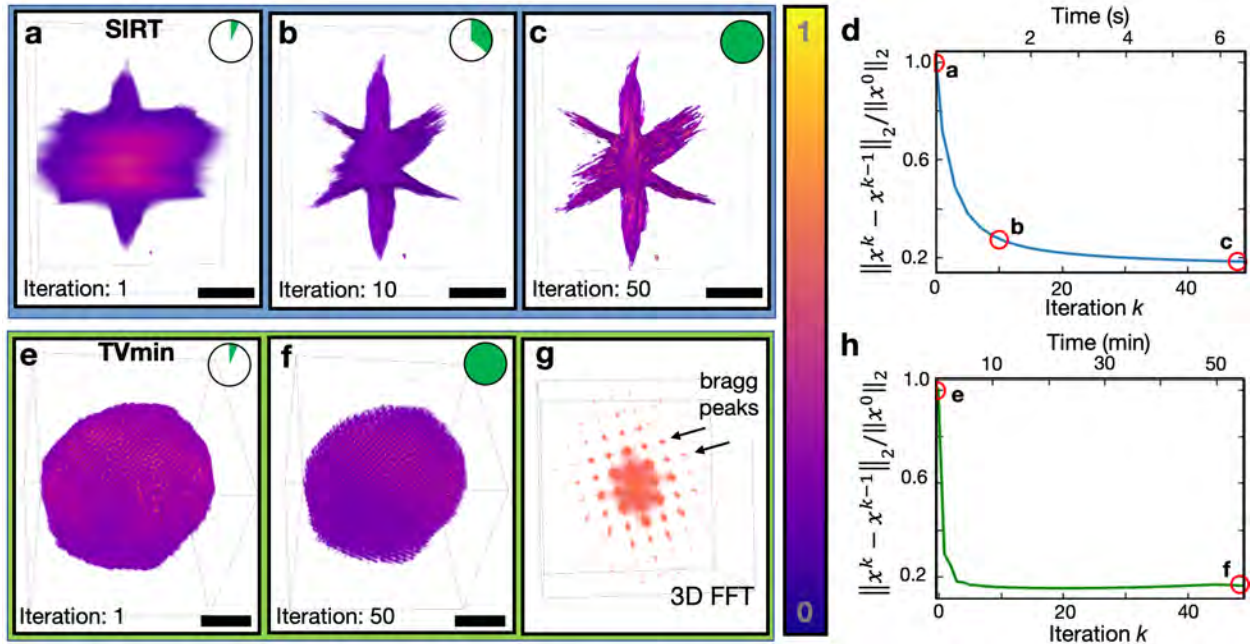


Figure 5.2: **Demonstration of iterative reconstruction algorithms.** **a-c**, Visualization of the Co₂P nanoparticle early, mid, and at the end of the reconstruction process. At the beginning, the underlying structure can partially be seen behind the excess of background intensity. In the middle of the process, sharp features begin to form. The final iteration converges to a tomogram visually similar to the input tilt series. Scale bar, 50 nm. **e-g**, Visualization of an atomic resolution FePt nanoparticle. The atoms in the TV nanoparticle are resolved with increasing iteration and its periodicity demonstrated with the fast Fourier transform (FFT). Scale bar, 1 nm. **d, h**, A plot of the normalized residual to demonstrate convergence.

compressive sensing [148] have also been incorporated into tomviz to accommodate the arrival of new projections during an experiment.

In addition to early estimates of specimen structure, real-time tomography allows assessment of the reconstruction convergence. This is observed qualitatively in the 3D visualization (Fig 5.2e-f) and quantitatively plotted in the residuals (Figure 5.2d,h). Watching the convergence provides visual inspection and intuition to how hyperparameters influence the final 3D structure and ensures proper convergence. For example, compressed-sensing inspired reconstruction methods are sensitive to regularization weights and require visual inspection to assess accuracy [83]. Furthermore, these advanced reconstruction algorithms do not exhibit predictable or monotonic convergence a priori and require monitoring to optimize convergence and determine when to terminate[51, 52]. Even for traditional algorithms where convergence is more predictable, they are often slow and changes become marginal—the scientist need not wait to begin interpreting the 3D structure. Lastly, practical issues such as misalignment, spurious values in data (e.g. hot pixels), in-plane rotations, and other pre-processing artifacts alter or degrade a reconstruction, however, these prob-

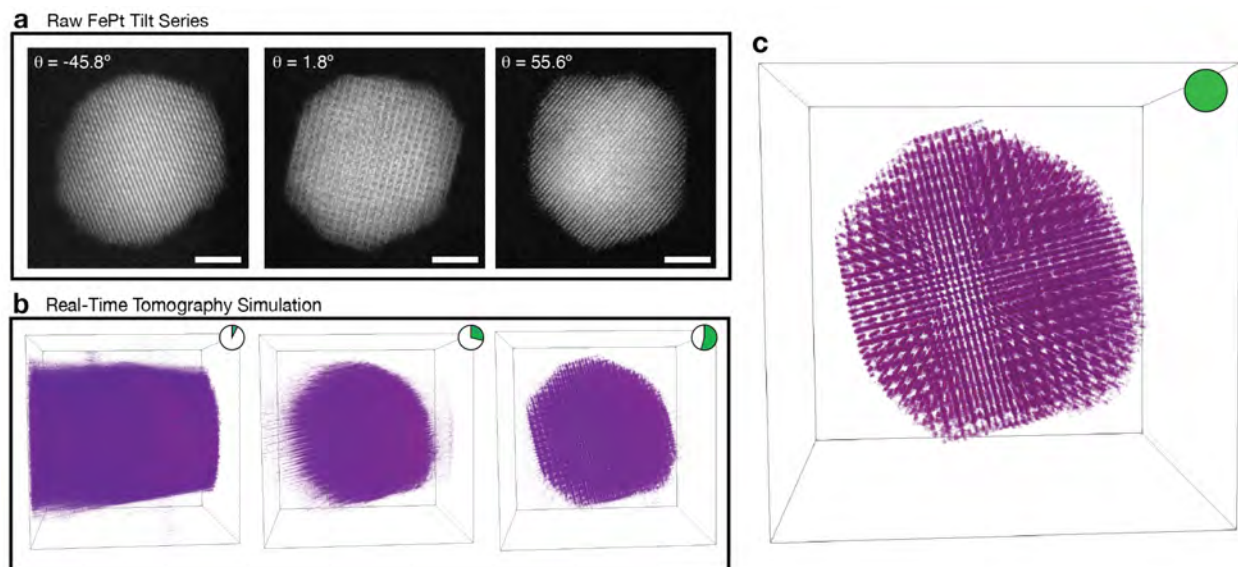


Figure 5.3: Real-Time Simulation for the Raw FePt Nanoparticle. **a**, The raw FePt tomographic projection images used for the real-time tomography simulation. Scale bar, 2 nm. **b**) As the simulation progresses, the reconstruction quality improves and atoms become visible as early as 50% of the experiment. **c**) An atomic-resolution tomogram is available after all the projections are reconstructed.

lems are diagnosable without completing a full reconstruction. Real-time assessment saves researchers time by providing early feedback and optimizing reconstruction parameters to serve the longstanding goal of high-throughput tomography.

Real-time tomography is agnostic to spatial resolutions and can faithfully reconstruct specimens at atomic length scales. Figure 5.3 illustrates the progressive convergence of the FePt nanoparticle while using the raw unaligned tilt series on tomviz. New unaligned HAADF projections (Fig. 5.3a) are incrementally fed into the ART algorithm every minute – acquisition speeds representative of experimental conditions. We see atomic scale structure becomes visible within 50% of the experiment (Fig. 5.3b), and individual atomic columns are resolved by completion (Fig. 5.3c). Further improvement can be achieved by integrating compressed sensing or regularized algorithms to remove low SNR and sampling artifacts.

Alternatively, weighted back projection (WBP) reconstructions are ideal for quick assessment of specimen morphology due to their fast, non-iterative computation [12, 16]. Figure 5.4 shows screenshots taken from a live WBP reconstruction visualized using tomviz—time proceeds from left to right. Figure 5.4a is a tomogram of gold (Au) nanoparticles on a strontium titanate (STO) nanocubes. Figure 5.4b shows platinum (Pt) nanoparticles on a carbon (C) support with the rotation axis along the x-direction. For WBP of single-axis tomography, partial volumetric updates are provided slice by slice along the direction parallel to the rotation axis. In the software, the 3D

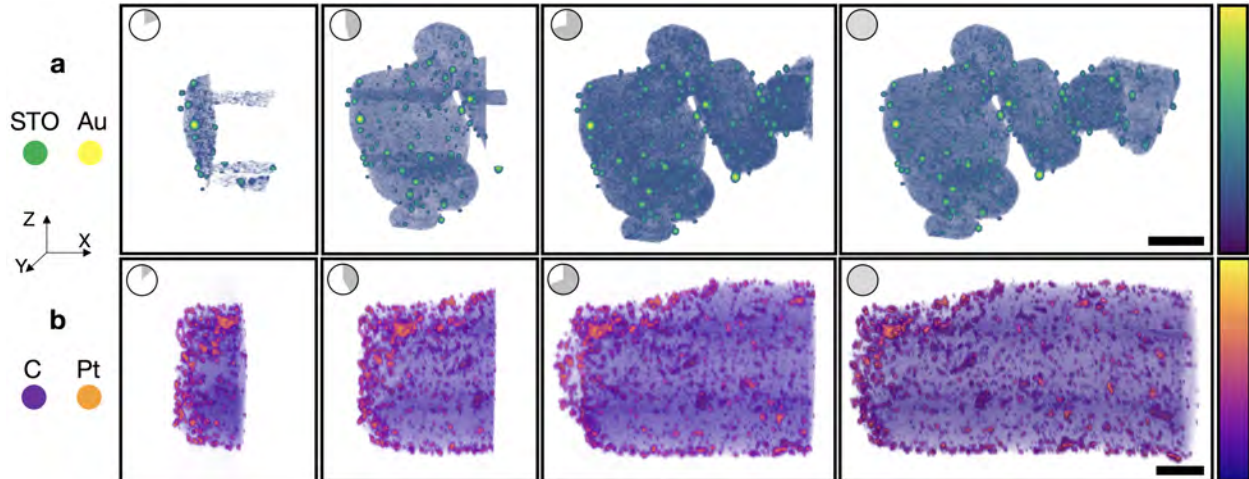


Figure 5.4: **Demonstration of live-WBP.** Live tomographic reconstruction in tomviz shown through freeze frames during progression of a weighted back projection algorithm (left to right). This unique capability allows users to interact and analyze the 3D structure throughout reconstruction. In the actual software the reconstruction updates in real time. **a**, Live volume rendering of Au/stromtium titanate (STO) nanocubes. **b**, Live volume rendering of platinum (Pt) nanoparticles on a carbon support. Scale bar, 50nm.

visualizations dynamically grow along one direction throughout the computation.

5.4 Dynamic Compressed Sensing for Real-Time Tomography

In addition to software advancements, we developed a dynamic CS framework that offers 3D specimen reconstruction in real-time as projection data is collected. It enables direct feedback and on the fly optimization of experimental parameters. The reconstruction algorithm begins immediately upon acquiring the first projection and dynamically updates the 3D structure as new projections arrive—unlike traditional schemes which start after the experiment is complete. This means researchers can start analysis and characterization with high-fidelity tomograms before an experiment is complete. Using scanning transmission electron microscope (STEM) tomography [110], we demonstrate our method accelerates the final convergence by a factor of 2-3 over conventional CS and provides insight into 3D nanostructure within 62% of the total experimental acquisition time. We demonstrated dynamic reconstruction reduces reconstruction error for synthetic Au-SrTiO₃ nanoparticles by 27% and converges 100% faster than a traditional approaches. Moreover, we show our iterative framework enables dynamic manipulation of the data-tolerance throughout the reconstruction to efficiently explore tunable parameters without having completely reset the algorithm. Implementing dynamic CS required complete parallelization that includes the 3D total variation regularization for the isotropic norm. Tomograms ($\sim 512^3$ voxels) can reconstructed

dynamically on modest multi-core laptops during an electron tomography experiment and larger reconstructions ($\sim 2048^3$) are achievable with high performance computing.

The notion of sparsity has become widely used in signal processing and image reconstruction as a prior knowledge to regularize solutions in underdetermined problems. It was greatly popularized by the theory of compressed sensing [153] that demonstrates the possibility to accurately recover the 3D structure of specimens ($\hat{\mathbf{x}}$) from an insufficient number of projections (b) with ℓ_1 -norm optimizations. One of the most representative sparsity-exploiting algorithms is the total variation minimization (TV-min), which was originally proposed for image denoising [136] and widely used to reduce tomographic artifacts for reconstructions from a limited number of projections [152, 93, 61]. The technique can effectively remove noisy features while preserving the edges of the object by minimizing its gradient magnitude. In this work, we consider a constrained optimization problem defined as:

$$\hat{\mathbf{x}}^* = \arg \min_{\hat{\mathbf{x}} \geq 0} \|\hat{\mathbf{x}}\|_{\text{TV}} \text{ s.t. } \|\mathbf{A}\hat{\mathbf{x}} - \mathbf{b}\|_2 \leq \epsilon \quad (5.1)$$

where $\|\cdot\|_{\text{TV}}$ and $\|\cdot\|_2$ denote the TV and ℓ_2 norms, and ϵ is a data-tolerance parameter that controls the trade-off between regularization (smoothness) and data fidelity. Here, a tomography experiment is formulated as an inverse problem, $\mathbf{A}\hat{\mathbf{x}} = \mathbf{b}$, where A is the projection matrix that models the measurement physics. The constrained optimization problem can be solved with a combination of adaptive steepest-descent (ASD) to minimize TV and projection onto convex sets (POCS) to enforce data constraints—commonly referred to as (ASD-POCS) [154]. This constrained optimization provides physical meaning to the tunable regularization parameter, ϵ , which can be initially estimated from the data quality [101]. Across all ‘flavors’ of compressed sensing tomography, the optimization process begins after all data has been collected. The iterative process can take thousands of iterations and runs from hours to a full day before converging to the designed solution. Moreover, the regularization parameter (here, ϵ) is often task-dependent and needs to be adjusted to produce the best image quality, further increasing the computation time for the reconstruction process.

Dynamic reconstruction during data collection allows researchers early insight into 3D structure throughout a tomography experiment. Figure 5.5 highlights the overall framework for the dynamic CS algorithm. Instead of starting at the end of an experiment, the reconstruction task begins immediately when the first projection is available. As more projections are experimentally acquired, the new information is accommodated as additional constraints in the optimization process (Fig. 5.5a) and improves the reconstructed tomogram quality (Fig. 5.5b). Because the reconstruction process is continuously running throughout the entire data acquisition, which typically take several hours, dynamic compressed sensing is able to produce a high-quality reconstruction before or upon

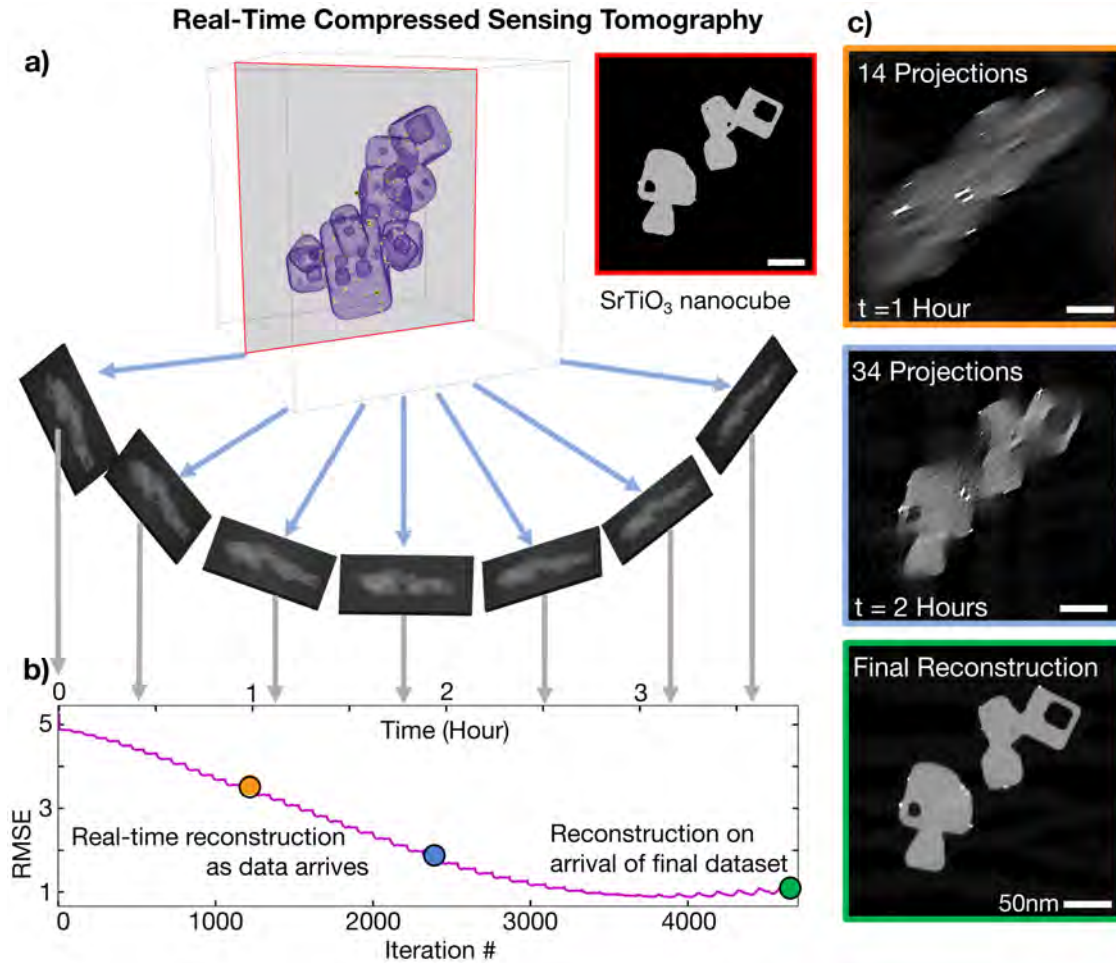


Figure 5.5: **External and internal architecture of tomviz GUI.** The tomviz platform is composed of a multi-threaded pipeline that synchronously handles tomographic and 3D visualization on separate threads. **a**, Tomviz monitors for recently acquired tilt projections within a directory and **b**, automatically reads new data into the pipeline. **c**, As tomographic reconstructions proceed, visualizations dynamically update and remain interactive for analysis.

arrival of the final projection (Fig. 5.5c).

5.5 Convergence of Dynamic Compressed Sensing

As dynamic CS progresses, in both iterations and the number of projections, it reduces RMSE but eventually diverges from the optimal solution with minimal error and approaches a solution defined by the optimization problem (Eq. 5.1). Iterative algorithms (e.g. Kaczmarz, Landweber, or Cimmino Method) typically deviate from solutions with minimum RMSE when applied to noisy data [50]. While RMSE is a useful quantitative measure to assess reconstruction performance, it

does not match the visually desired solution [83]. Moreover, computing RMSE requires knowledge of the true 3D specimen structure.

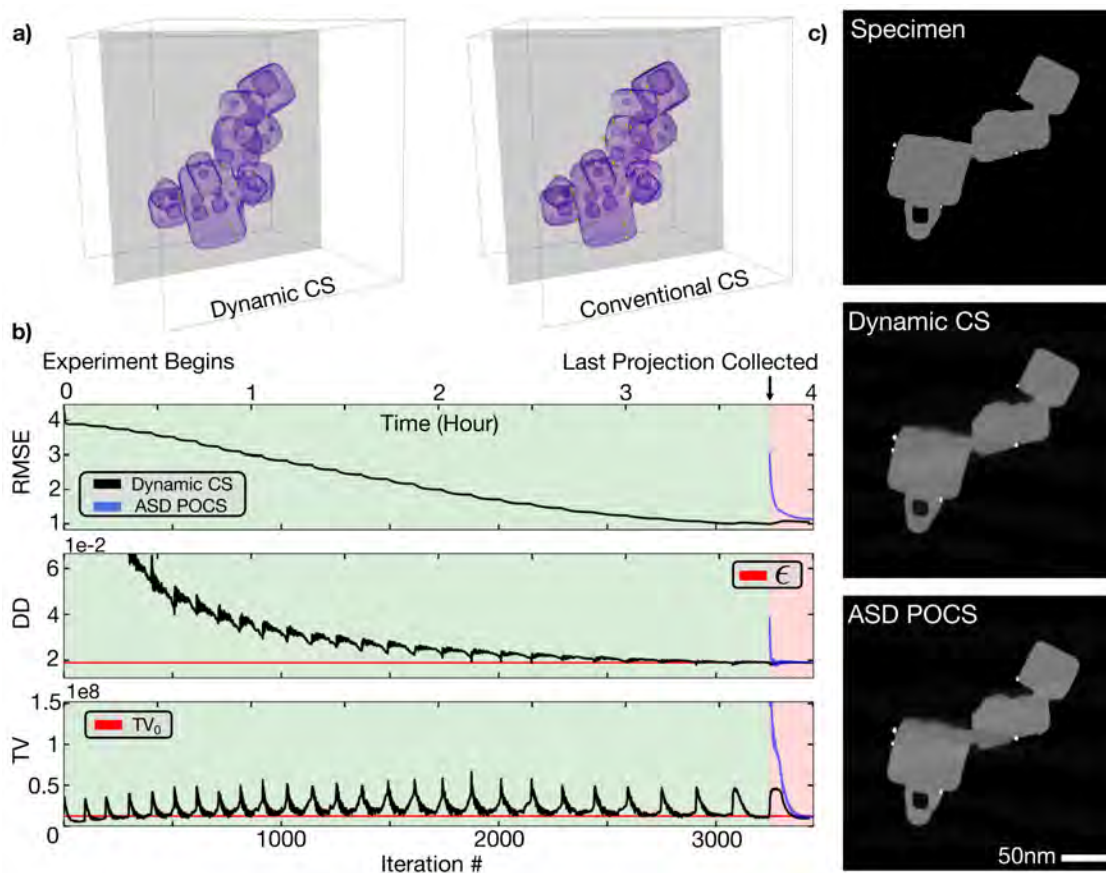


Figure 5.6: **Dynamic CS convergence.** a) 3D volume rendering with yellow indicating high intensity (Au) and purple representing low intensity (SrTiO_3). These tomograms were constructed under the dynamic (left) and traditional (right) CS framework. b) Plots of data distance (DD), total variation (TV) and RMSE for both the dynamic algorithm (black) and conventional (blue). c) 2D cross-sections of the final output for each 3D reconstruction and the test object.

In real experiments the true solution is unknown and a tomogram's RMSE cannot be measured. Instead, the progression of a reconstruction's TV and data distance ($\text{DD} = \|A\hat{x} - b\|_2$) can be utilized to assess the convergence towards an optimal solution. Figure 5.6 plots RMSE, DD and TV vs. time for the Au- SrTiO_3 phantom nanoparticle during a dynamic compressed sensing reconstruction. Throughout an experiment (shaded green), the data distance trends downward to the specified data tolerance, ϵ (red line) indicating stable convergence. The incorporation of new data creates sharp discontinuities in DD and TV. Unlike RMSE which drops with the addition of new projections, DD and TV momentarily rise sharply because ASD-POCS is attempting to minimize the distance between DD and ϵ by iteratively adjusting the weights between data fidelity and regularization. After the arrival of new data, the algorithm will sufficiently converge to a solution

within ~ 125 iterations. Dynamic CS performs best when there are enough iterations to satisfy its data tolerance constraint ($DD \simeq \epsilon$) before new projections are introduced. If additional projections are added too quickly, the overall convergence may drift (See Fig. 5.7) and the algorithm will be unable to reach its optimal solution by experimental completion.




Figure 5.7: **Convergence criteria.** Stable (blue) and unstable (black) convergence of RMSE, DD, and TV for the synthetic Au/SrTiO₃ nanoparticle. The blue curve shows proper convergence when the reconstruction completes many iterations (125) prior to appending more projections. The black curve shows incomplete convergence when insufficient intermediate iterations are not completed and TV may drift away from the true solution (TV_0).

5.6 Dynamic Parameter Tuning

Selecting ϵ often requires computing several reconstructions and ultimately relies on the scientist’s judgement. Here we show dynamic compressed sensing allows ϵ to be tightened (decreased) or loosened (increased) mid-reconstruction. Furthermore, the data-constraint can be reversibly adjusted—reflecting stable and convex convergence. Generally speaking, selection of the data consistency constraint, ϵ , depends on the SNR as it accommodates all sources of data inconsistency (e.g. noise) and ensures re-projections are within a given ℓ_2 distance from the actual (experimental) data [177]. Dynamic parameter tuning allows researchers to more efficiently dial in the optimal pa-

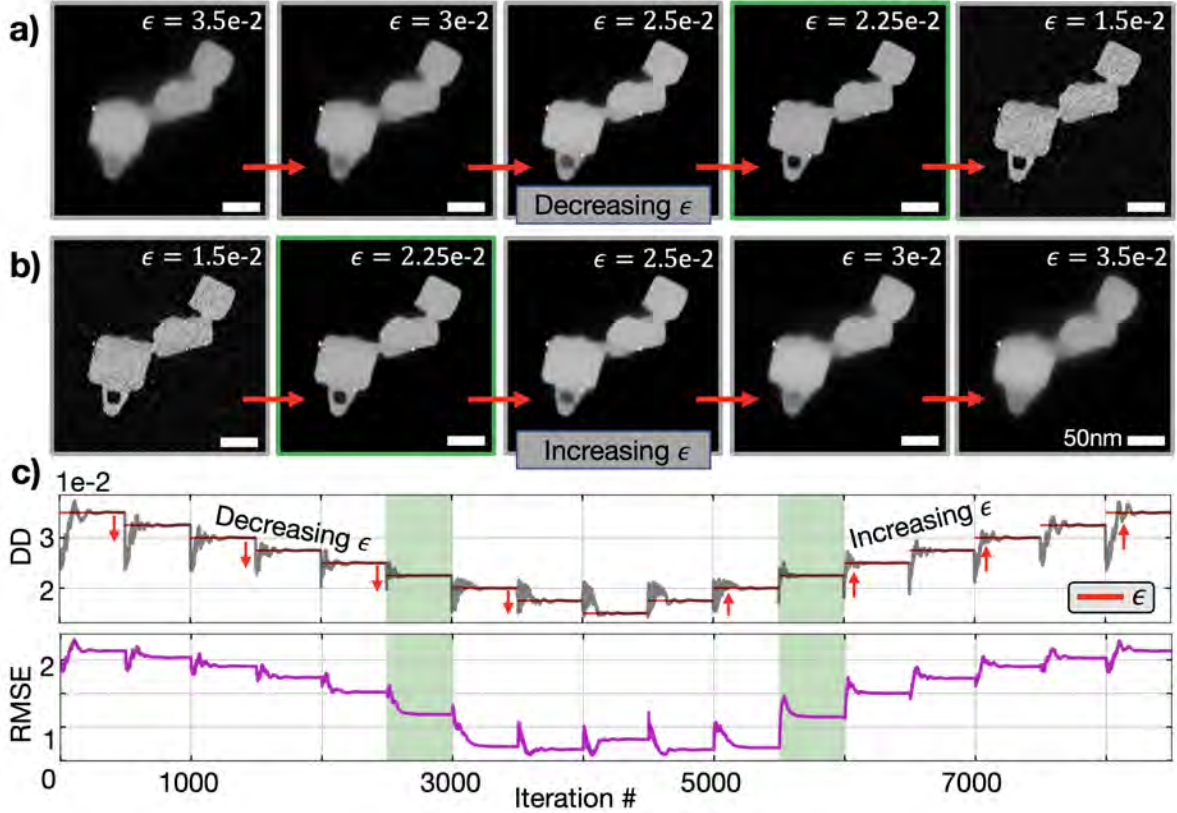


Figure 5.8: Dynamic manipulation of data tolerance parameter. a) Tightening (reducing) ϵ decreases the regularization weight which gradually produces sharper 3D tomograms. b) Loosening (increasing) ϵ allows for there to be more regularization and produce smoother 3D tomograms. c) Plots demonstrating the progression of DD (gray) and RMSE (magenta) vs iteration. As the calculation progresses, scientists can manipulate regularization without having to reset the algorithm. The visually desirable solution is obtained when $\epsilon = 0.0225$, highlighted in green. Reducing ϵ below this value produces noisy reconstructions. Overestimating ϵ blurs away fine features such as the Pt nanoparticles and internal voids.

parameter value. If ϵ is too low the reconstruction appears noisy; if ϵ is too high, detail and resolution is degraded.

Figure 5.8 demonstrates that manipulating ϵ throughout the reconstruction consistently converges in stable results. Here, we reconstructed a fully sampled Au/SrTiO₃ tilt series and waited until all of the optimization parameters (RMSE, DD, and TV) were fully converged prior to perturbing the data constraint. During iterations: 0–4,500 we reduced ϵ eight times by -0.025 before increasing by $+0.025$ eight times during iterations 4,500–8,500. For this dataset, we found small perturbations ($|\Delta\epsilon| \leq 0.05$) guarantee convergence within ~ 100 iterations. Dynamic CS reliably converges to solutions defined by the final ϵ chosen. Whether incrementally increased or decreased, the final value of ϵ determines a nearly unique solution, which can be seen both visually

(Fig 5.8a,b) and quantitatively in the RMSE 5.8c. Note that the minimal RMSE ($\epsilon \simeq 0.0175$) retains grainy artifacts and does not produce a desirable reconstruction due to Au particle's high intensities. As discussed by Jiang, *et. al.*, the visually appealing result is generally obtained from a slightly larger ϵ [83]. We observed similar data-tolerance properties, the visually desirable reconstruction (highlighted in green Fig. 5.8) occurs at $\epsilon = 0.025$.

5.7 The Live Tomography Software on tomviz

The latest tomviz release (v. 2.0) includes real-time tomography capabilities, is entirely open-source (BSD License), runs on all operating systems (OSX, Windows, Linux) with certified installers, and can be implemented on rudimentary TEMs available at most institutions.

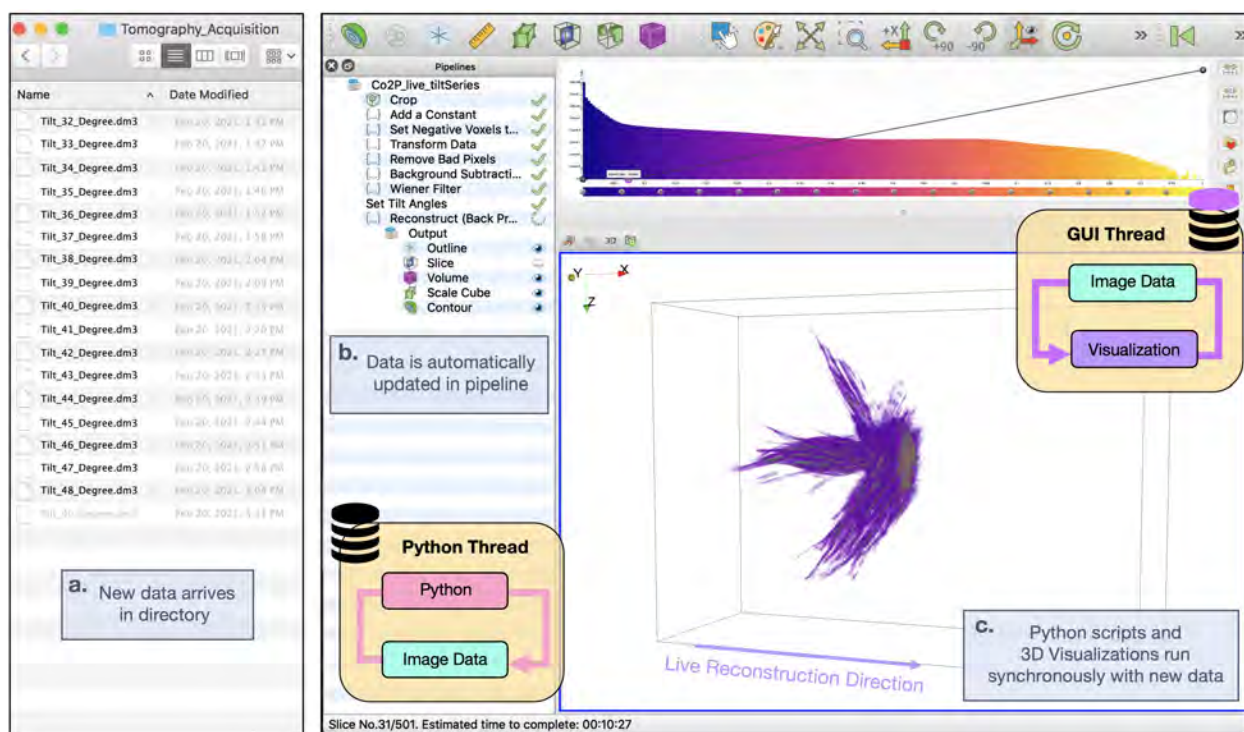


Figure 5.9: **External and internal architecture of tomviz GUI.** The tomviz platform is composed of a multi-threaded pipeline that synchronously handles tomographic and 3D visualization on separate threads. **a**, Tomviz monitors for recently acquired tilt projections within a directory and **b**, automatically reads new data into the pipeline. **c**, As tomographic reconstructions proceed, visualizations dynamically update and remain interactive for analysis.

The tomviz graphical user interface (GUI) (Fig. 5.9) provides an intuitive tomography tool that allows scientists to focus on 3D specimen interpretation [145]. Tomviz monitors data directories for the arrival of new projections during an experiment (Fig. 5.9 a) and visualizes the 3D recon-

struction as it dynamically updates (Fig. 5.9 c). During a real-time tomographic reconstruction users can zoom, rotate, slice, and segment the object to highlight regions of interest as the algorithm runs independently. Each voxel in the 3D render (i.e. volumes or isometric contours) are assigned a color and opacity controlled by the color-opacity transfer function overlaid on the histogram visible at the top of the GUI (Fig. 5.9, top-right). Users can intuitively define voxel transparency by selecting points on the curve and dragging it between transparent or opaque. The data ‘Pipeline’ retains all transformations and modules performed to produce visualizations, all of which can be saved in a state file for sharing and reproducibility.

A seamless user experience is enabled by an underlying multi-threaded framework of Python / C++ interactions. As the reconstruction occurs, algorithms written in Python trigger signals to notify the application that a new volume is available. Tomographic reconstructions can either run on basic computer infrastructure found on any laptop or scaled across multiple GPUs to process large volumes ($> 1024^3$ voxels). Live reconstructions without performance degradation requires a tripling of memory requirements. One data copy resides on the visualization (GUI) thread, another on the reconstruction (Python) thread, plus a temporary copy for efficient staging and handoff. The temporary copy allows the reconstruction to run unhindered during the handoff process. The total memory usage for real-time reconstruction is usually well within a consumer grade computer (c.a. 0.4 – 16 GB). After the reconstruction is complete, all copies are released from memory and only the final reconstruction remains. Analytical reconstruction methods such as WBP can run slice-by-slice with new reconstructed slices appended along a single reconstruction direction. For iterative methods, we recommend updating the entire volume either every iteration or every few (depending on the speed of computation)—especially for complex sampling schemes such as dual or multiple-axis tomography which lacks a single rotation axes. Enhancements to the underlying 3D rendering (VTK) within tomviz were made to improve interactive visualization and analysis throughout the reconstruction process [66].

5.8 Parallelization and Performance

Computational efficiency is key to the success of real-time tomography dynamic compressed sensing. Several software advancements were required to develop an asynchronous platform that handles visualization and tomographic reconstruction simultaneously to provide a seamless user experience. The multithreaded pipeline within the tomviz application executes long-running jobs while simultaneously offering real-time visualization of the progress. As the reconstruction occurs, algorithms written in Python can trigger signals to notify the application that a new volume is available. A slot on the C++ side listens for this signal, using a mutual exclusive lock (mutex) on the image data to secure access to the updated volume. The new data is copied into the foreground thread

(main GUI), and once it is available the mutex is released. Once the application receives a signal indicating that the output has been updated, downstream data operations can then be re-executed and any connected visualization modules will also be notified. As an effect, the histogram is recalculated in another background thread while all the current visualization modules display the rendered representation. In the case of the contour module, this will necessitate the recalculation of the surface mesh or the update will be uploaded to the GPU for volume rendering.

Modifications to the common implementation for SIRT and TVmin were made to account for the dynamic addition of input projections throughout an experiment. SIRT seeks the minimal error between the reconstruction and experimental data: $\arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$ where \mathbf{A} is the measurement matrix, \mathbf{b} are the experimental projections and \mathbf{x} is the tomogram. We can further regularize the process through the assumption that our volumes should be piece-wise smooth and minimize its total variation $\|\mathbf{x}\|_{\text{TV}}$. Iterative algorithms require rescaling of the descent parameter based on the number of projections sampled. SIRT can easily estimate the descent parameter through calculation of the Lipschitz constant ($L = \|\mathbf{A}^T \mathbf{A}\|_2$). The Lipschitz constant can be estimated by using the power method [86]. The descent parameter for TVmin is scaled by a dampening envelope that ensures its magnitude decays linearly [148]. Non-iterative algorithms such as WBP do not require rescaling of descent parameters and simply needs to reinitialize the computation with the new projection images collected.

As discussed previously (Section 5.5), the reconstruction process should reach stable solutions before incorporating new projections. We emphasize that the overall computation should always be faster than the data acquisition, so that by performing the image reconstruction on the fly, we can obtain the reconstructed image almost right after the experiment finishes. This is seen in all the previous cases we have presented, in which RMSE converges to a plateau value before the arrival of a new projection. However, as the size of the object increases, the computational complexity grows as $O(N^3)$, where N is the size of the object in each dimension; the experimental time, however, grows only as $O(N^2)$. Therefore, for large physical systems ($>2048^3$), single laptop/desktop or workstation is not powerful enough for dynamic compressed sensing. To overcome this problem, we deployed high performance computing (HPC) resources at Theta, a Cray XC40 11.69 petaflops supercomputer at Argonne Leadership Computing Facility.

We use Message Passing Interface (MPI) [28] to parallelize our code across different nodes. In particular, we distribute the 2D slices to different MPI processes. There are two dominant computations involved in our algorithm: (1) ART – for minimizing $\|\mathbf{A}\hat{\mathbf{x}} - \mathbf{b}\|_2$; (2) TV – for calculating the TV gradient of the object, $\|\hat{\mathbf{x}}\|_{\text{TV}}$. For ART, the computation is independent for different slices and no data exchange is needed among the processes. For TV, each process only needs to send the first and last slices owned by that process to the two nearby processes respectively. The communication overhead thus is minimal. We expect our algorithm to scale efficiently in HPC

supercomputers. Besides using MPI for the inter-node parallelism, for the intranode parallelism, we use OpenMP [36] to parallelize the computation intensive loops.

5.9 Discussion and Conclusion

We demonstrate real-time visualization of electron tomography reconstructions as they proceed during or after an experiment using tomviz, an open-source cross-platform tool compatible with all electron microscopes. We achieved real-time electron tomography by integrating dynamic volumetric data analysis tools, data input / output, processing, reconstruction and visualization into a single software tool. In the actual software, the 3D visualizations are dynamically updated in parallel with computation. This means that scientists need not wait for a reconstruction to complete, or all data to be collected before beginning the interpretation of results. Continuous feedback provides high-throughput and early diagnoses of 3D specimens, opportunities to optimize experimental parameters, or investigate multiple regions of interest. Although dose is fundamentally set by the experimental acquisition parameters (e.g. dwell time, beam current, sampling rate or tilt increment), in practice real-time tomography may reduce dose by streamlining acquisition and allowing the possibility of early termination if the reconstruction requirements are met. Optimized, threaded pipelines and the iterative nature of tomographic methods allows tomviz to show intermediate results with minimal impact on performance. This enables interactive 3D analysis of the current reconstruction state while the reconstruction proceeds on a separate thread. A robust graphical interface allows objects to be rendered as shaded contours or volumetric projections and these objects can be rotated, cropped, or sliced. This capability opens radically new possibilities for developing high-throughput, real-time tomographic reconstruction algorithms for geometrically complex inorganic [81] or biological materials. Ultimately, interactive real-time visualization goes beyond high-throughput and allows researchers to make early judgments to answer or identify new scientific questions during experimentation.

CHAPTER 6

Conclusion

As the field of material science is increasingly reliant on electron microscopy, underscoring the growing demand for innovative techniques in advanced low-dose imaging, with expanding applications in organic matter. Notably, the integration of compressed sensing algorithms has been a hallmark of inverse problem accomplishments in the past two decades, particularly through the incorporation of sparse modeling. This dissertation focused on solving the challenges associated with low-dose spectroscopy and real-time 3D imaging.

The relevance of sparse signal recovery was initially demonstrated in Chapter 2, where the restoration of electron micrographs by removing structured stripe artifacts were presented. This was achieved through a constrained cost function that minimizes the image's total variation while preserving information in Fourier Space. By capitalizing on the shortcomings of compressed sensing, which struggles to recover highly localized directional structures missing in Fourier Space, we achieved significant progress. In application, destriping has the most substantial impact on soft matter (e.g., nacreous pearls) which are particularly susceptible to damage from mechanical polishing methods [58]. Consequently researchers can allocate less emphasis on sample preparation, thus accelerating workflows and directing more emphasis on data analysis. Furthermore, this destriping study provides insight into the recoverability of extended, unidirectional features (e.g., interfaces), which can be problematic for inverse problems beyond electron microscopy.

In Chapters 3 and 4, a novel framework that facilitates data fusion between complementary elastic and inelastically scattered signals within microscopes was introduced. This approach transforms the exploitation of tangential features into a multi-objective inverse problem, promoting synergy between available modalities. Remarkably, thoughtful design considerations yielded dose reductions of over 99%, equating to two orders of magnitude, while enabling 2D and 3D chemistry measurements at sub-angstrom and sub-nanometer resolutions, respectively. This fidelity enhancement of chemical distributions, alongside material integrity retention, was consistently evident. Multi-modal data fusion not only improves SNR but opens a pathway for the low-dose chemical imaging of radiation-sensitive materials.

Ultimately, this framework offers a new form of experimentation that efficiently utilizes all the scattered and emitted signals in electron microscope to characterize the atomic structure of matter. We can extend our work to additional modalities including phase-contrast signals like integrated Differential Phase Contrast (iDPC) or electron ptychography [104, 27]. With phase contrast signals, we can expanding the scope of multi-modal imaging to fully encompass the light elements in oxides and soft materials as intensity would be linearly proportional to Z . Thereby improving simultaneous high-resolution chemical imaging of heavy and light elements.

Chapter 5 introduces software and algorithmic strides that accelerate electron tomography, rendering near real-time 3D imaging feasible. Through optimizations, the time required to display volumetric visualizations was dramatically reduced from over 24 hours to under 20 minutes on a software platform concurrently managing 3D rendering and reconstruction through separate threads. Algorithmic developments further bolster this advancement by adaptively handling the inflow of new projection views during microscope data acquisition. Intermediate visualization of specimen structure boosts high-throughput, as researchers can assess the specimen's integrity during the experiment and diagnose practical issues that could hinder success. For instance, this work enhances the likelihood of reconstructing highly beam-sensitive organic matter, which is prone to degradation under the electron beam. Currently, the temporal resolution of real-time tomography is partially constrained by the slow speed of the goniometer. The synergy of future improvements in automated rapid acquisition with faster reconstruction speeds on GPU clusters holds promise for rapidly visualizing 3D material structures within seconds. In the future, such rapid experiments combined with the real-time tomography presented herein, can empower the investigation of nanoscale dynamic processes in materials subjected to external stimuli, such as heat or liquid environments in 3D [2].

APPENDIX A

Automated Hyperparameter Tuning With Bayesian Optimization

Bayesian optimization (BO) with Gaussian processes (GP) has emerged at the forefront of expensive black-box optimization due to its efficient exploration capabilities. BO is designed to identify the global optima of unknown functions and is well suited for reliably identifying preferred parameters in an unknown objection function f :

$$\hat{\boldsymbol{x}} = \arg \min_{\boldsymbol{x} \in \mathcal{X}} f(\boldsymbol{x}) \quad (\text{A.1})$$

Where \mathcal{X} denotes the search space and $\hat{\boldsymbol{x}}$ is the optimum. Bayesian optimization framework can be applied to more unusual search spaces that involve categorical or conditional inputs. Because the Bayesian optimization framework is very data efficient, it is particularly helpful in situations where evaluations of f are costly and is non-differentiable or convex. In these situations, BO accounts previous observations to efficiently guide the search process for future measurements.

Bayesian optimization involves two core components: (1) develop a posterior probability distribution of the parameter objective with GP regression and (2) specified control the exploration of future measurements [84]. BO uses GP to model the objective as a joint distribution of functions whose landscape (e.g. smoothness) is defined by a kernel or covariance function. In addition to modeling the objective, the posterior distribution quantifies the uncertainty (standard deviation) over all unsampled points. The acquisition function proposes sampling points while balancing the importance between exploitation of extrema and exploration in regions with large uncertainty. Both phases act as a safeguard against convergence to local optimum by exploring the entire parameter space [23].

Figure A.1 shows the Bayesian optimization process to systematically determine the optimal λ values from estimates of the RMSE landscape. Here we are trying to reconstruct synthetic AuSTO nanocubes using the unregularized TV cost function (Eq. 1.13) with 47 low SNR projections solved with the FISTA algorithm. In the case of this cost function, we only are trying to optimize

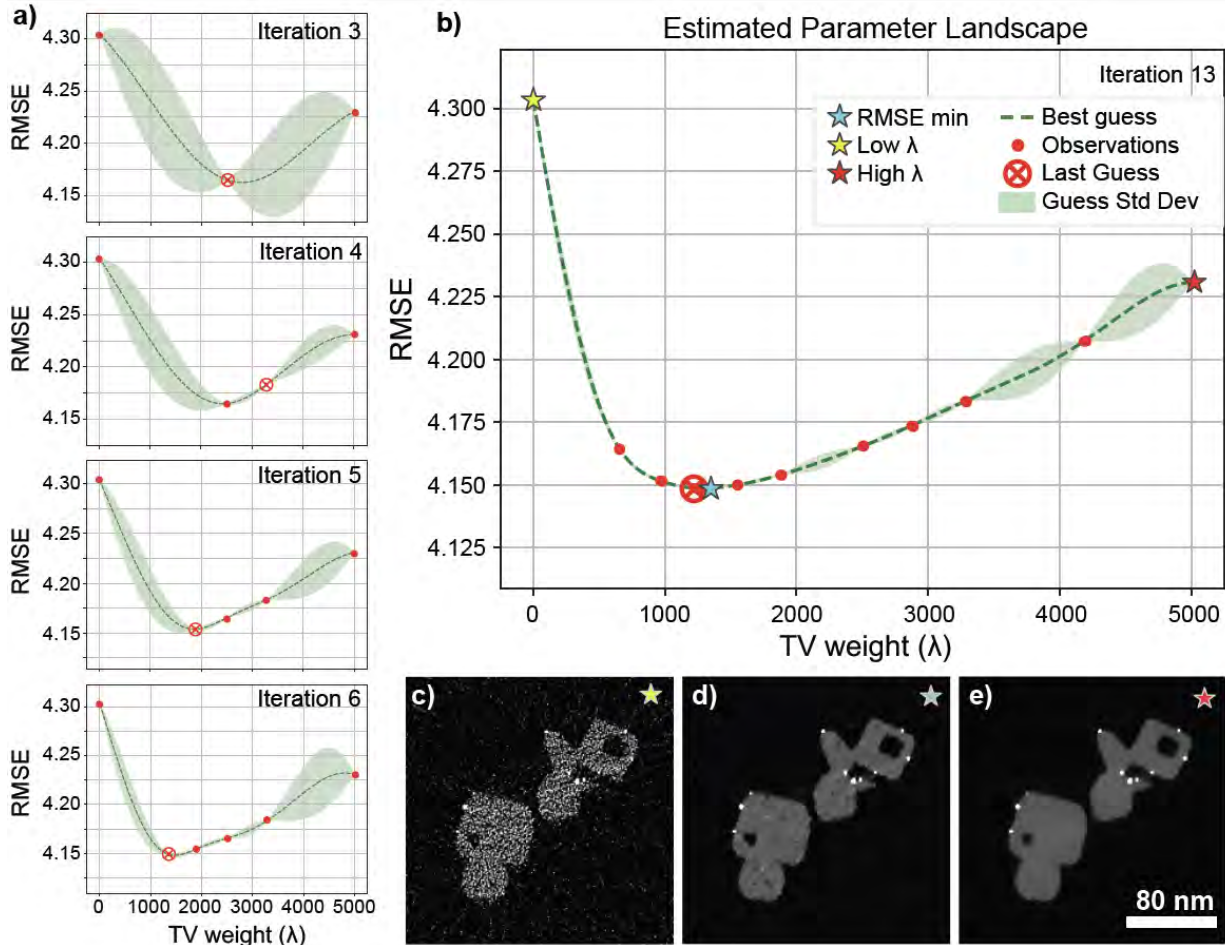


Figure A.1: Determining optimal λ parameter for simulated compressed sensing electron tomography with Bayesian optimization. **a** The BO optimizer iteratively samples λ values and performs a tomographic reconstruction to estimate the RMSE landscape. BO strategically determines the next point to sample after each measurement. **b** The final estimated landscape of RSME vs λ . The blue star represents the Bayesian optimization’s minimum RMSE solution, and the yellow and red represent the minima for the DD and TV metrics, respectively. **c-e** Corresponding reconstructions: **c** lies in the traditional regime, **d** is the RMSE minimizing reconstruction, and **e** is an over-smoothed reconstruction.

our cost function in relation to our regularization parameter (λ). As more points are gathered, certainty of the RMSE landscape improves. The variance in its estimate of the landscape is used to predict the global minimum. The estimated landscape is updated after each prediction until a certain number of iterations is performed or the solution converges. It is critical that the global minimizer is within the bounds provided to the optimizer.

For simulated data, autonomous parameter tuning for iterative algorithms electron tomography is achievable using Bayesian optimization with Gaussian processes. BO with GP autonomously

reduces the computational time required to discover optimal tomography parameters without any user intervention. This work is essential for large-scale tomographic simulations that require automatic parameter selection and efficient use of computational resources. Specifically, BO was particularly useful for carrying out the big scale fused multi-modal tomography simulations (Fig. 4.11) which required several observations to determine the optimal hyperparameters in Eq. 4.1. In the case of 2D optimization, the estimated parameter landscape transitions from a simple 1D curve to a surface as shown in Fig. 4.12c.

APPENDIX B

User Manual for Real Time Tomography on tomviz

Electron tomography and cryogenic electron tomography (cryo-ET) generate three-dimensional (3D) reconstructions of native biological and material specimens. However computational bottlenecks in tomographic reconstruction stymie the 3D investigation of specimens. This set of tutorials provides scientists with expedited analysis by enabling real-time tomography with the ability to visualize intermediate volumetric results while reconstruction algorithms or data collection is ongoing. The first two tutorials are quick demonstrations that illustrate real-time 3D visualization capabilities without the need for an electron microscope.

The real-time electron tomography toolset is built using the publicly available tomviz platform. Tomviz is a full featured tomography toolset with real-time analysis and reconstruction capabilities. The software contains a multithreaded pipeline that enables interactive 3D visualization of the current reconstruction state with minimal impact on performance. Thus, scientists can go beyond superficial inspection to quantify specimen features or internal structure while simultaneously operating the microscope. This immediate feedback can save researchers days of effort as reconstructions are no longer processed offline. Real-time tomography also improves offline analysis by dynamically visualizing iterative tomographic reconstructions as they progress. Whether the computation runs online or offline, tomviz users can evaluate 3D specimen structure and optimize the reconstruction accuracy in real-time. These features are highlighted herein as four tutorials: the first two are quick demonstrations of real-time tomography using data pre-bundled with tomviz, two more that demonstrate real-time tomography on an electron microscope, and the last which shows how custom scripts with real-time 3D visualization can be written and used within tomviz.

B.1 Quick Demo: Live 3D Visualization During Reconstruction

In this demonstration, we will visualize a live tomographic reconstruction after all the data has been collected (i.e. offline analysis). Live visualization provides insight into the specimen structure as the computation evolves. This tutorial demonstrates live tomography using projection images collected on a Co₂P hyper-branched star nanoparticle—the dataset is prepackaged with tomviz and does not require additional downloads.

(1) **Load the tilt series data from the drop-down menu.** In this section we're going to start by loading the projection images (i.e. a 'tilt series') from the Sample Data drop-down menu (Fig. B.1). The dataset is a collection of projection images acquired across $\pm 75^\circ$ at a $+2^\circ$ increment of a complex nanoparticle. When the data is loaded, it will appear in the 'Pipeline' panel along with visualization modules automatically displayed on the 'RenderView' panel (Fig. B.2). To view the projection images at each tilt, select the 'Slice' module from the 'Pipelines' column, use the 'Slice' slider located below in the 'Properties' panel (highlighted in red) and select the XY plane. These projection images are intentionally slightly misaligned, fortunately our automatic algorithms will implement the necessary corrections.

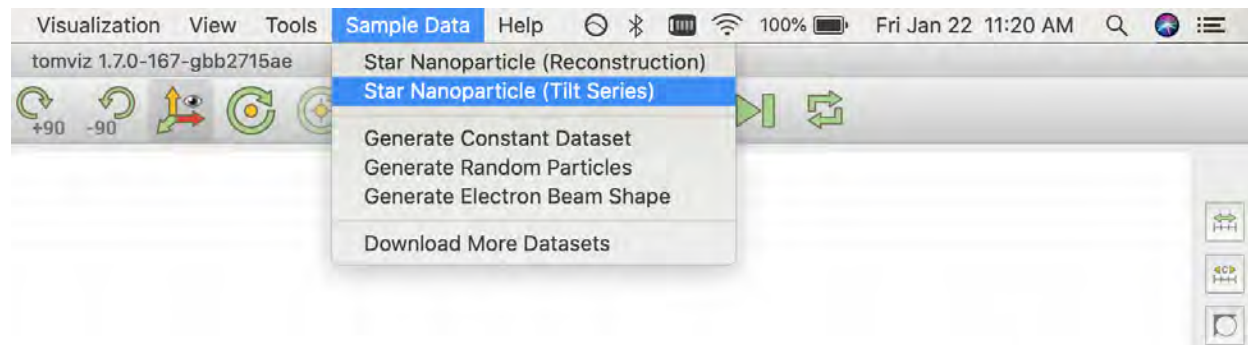


Figure B.1: **Load the sample dataset** Tomviz comes packaged with sample datasets that include projection images taken across many specimen tilts (i.e. a tilt series).

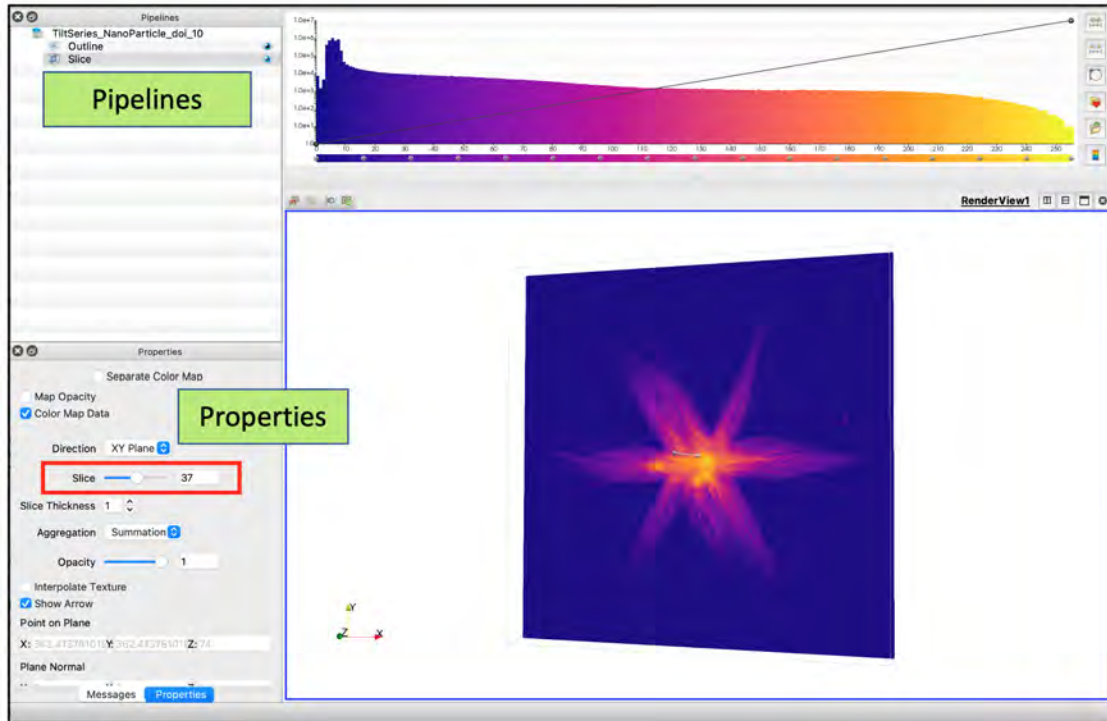


Figure B.2: **View the tilt images.** When projections are loaded into tomviz it is automatically displayed as orthogonal slices. The parameters of each visualization mode can be set in the ‘Properties’ panel in the lower left. To view images along different tilts, use the ‘Slice’ slider (red box) or manually enter the image you want to view.

(2) **Preprocess and Align the Tilt Series.** For BF-TEM, contrast inversion and CTF correction is often applied. The data can be inverted by ‘Invert Data’ in Data Transforms and CTF correction can be accessed in the Tomography dropdown menu. After the CTF of the instrument is specified [132, 7, 159] the image data will be reweighted in Fourier space [112, 41].

The performance of tomography reconstruction algorithms depends on alignment quality. We will begin the alignment process by shifting the projection images so the specimen’s center of mass is located at the origin. Select ‘Image Alignment (Auto: Center of Mass)’ from the Tomography drop-down menu (#1 in Fig. B.3b). In addition to aligning all the tilt images, tomviz needs the tilt axis to be centered and parallel to the x-axis. Users can automatically rotate and translate the tilt axis with algorithms available in the tomography drop-down menu. First, apply the auto rotation alignment (#2 in Fig. B.3b) and then shift alignment (#3 in Fig. B.3b). The data is ready for tomographic reconstruction after all alignments are applied.

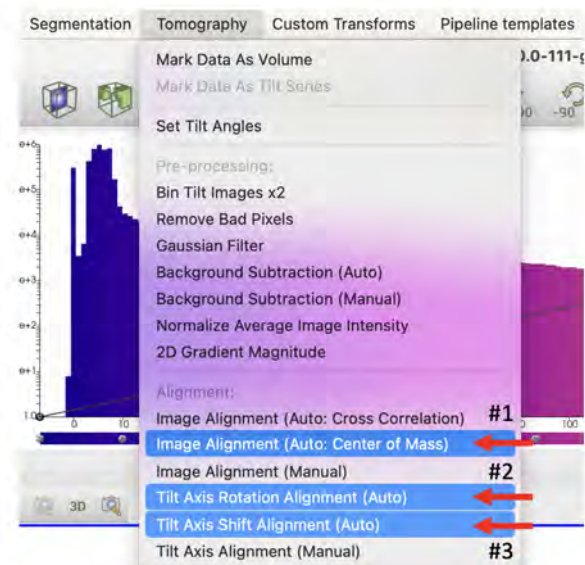


Figure B.3: **Align the Tilt Images.** The sample tilt series is slightly misaligned. We can automatically correct misalignments by centering the images with the Center of Mass method (#1), rotating the tilt axis (#2), and shifting its location (#3) so its center and parallel to x-axis.

(3) **Launch a live-tomographic reconstruction.** Several reconstruction algorithms are available with tomviz in the ‘Tomography’ drop-down menu. Run the SIRT algorithm for this live-reconstruction visualization demonstration (Fig. B.4).

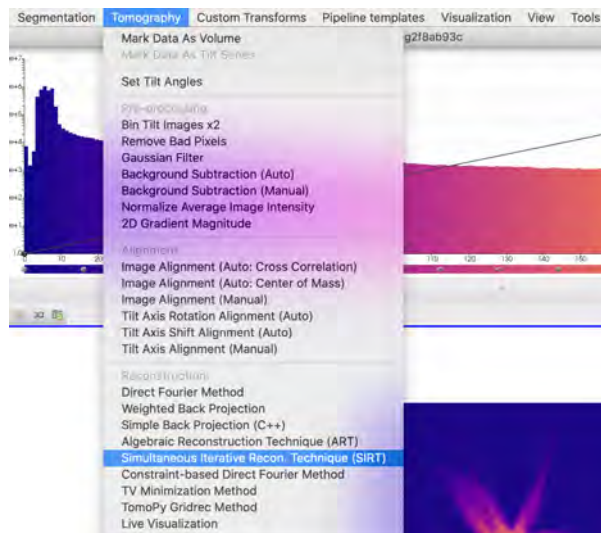


Figure B.4: **Select the reconstruction algorithm for live visualization.** The tomography menu contains many algorithms for electron tomography. The reconstruction sub-section contains all the iterative and direct algorithms for tomography. Here we will be selecting the Simultaneous Iterative Reconstruction Technique (SIRT).

SIRT is a fast and efficient iterative algorithm. The frequency of visual updates can be specified in the final input box. For this example, we recommend setting the percentage to 100% (Fig. B.5). Reducing the frequency of visual updates to 25-50% is useful when datasets are large or computational hardware is limited (e.g. anytime a reconstruction proceeds slower than 1 iteration per second). Tomviz also provides live-visualization capabilities for ART, WBP, or TVmin from the Tomography drop-down menu. Once SIRT parameters have been specified, press the blue ‘Ok’ button (Fig. B.5, bottom right hand corner) to begin.

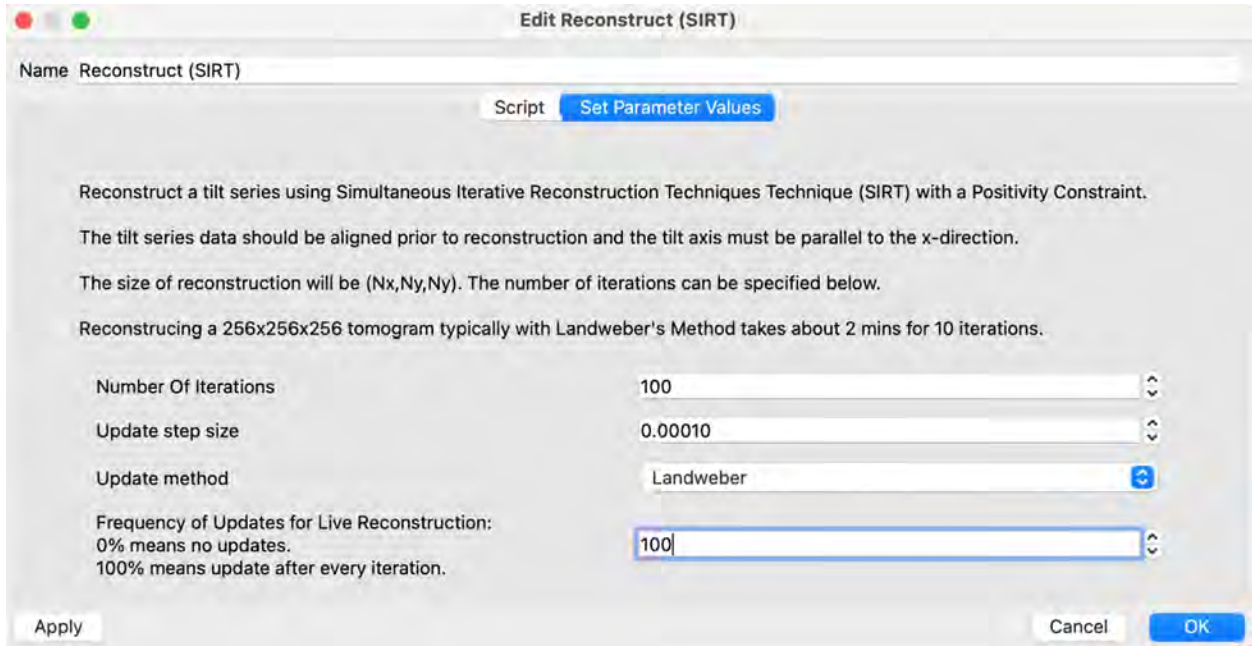


Figure B.5: Enter parameters for the SIRT algorithm

(4) **Visualize the live volumetric process.** Once the reconstruction begins, two additional elements in the ‘Pipelines’ labeled ‘Reconstruct (SIRT)’ and ‘Reconstruction’ will appear. Tomviz by default will continue visualizing the projection tilt images which can be identified from the lack of modules below the ‘Reconstruction’ dataset icon (Fig. B.6a). Delete the previous ‘Outline’ and ‘Slice’ modules (Fig. B.6b), or click the eye to make the previous displays invisible (Fig. B.6c), and reassign the modules to the reconstruction. You can visualize the 3D reconstruction by selecting the ‘Volume’ or ‘Slice’ from the visualization modules toolbar (Fig. B.7) after selecting the ‘Reconstruction’ dataset. These modules need to appear below the dataset icon (B.6b,c).

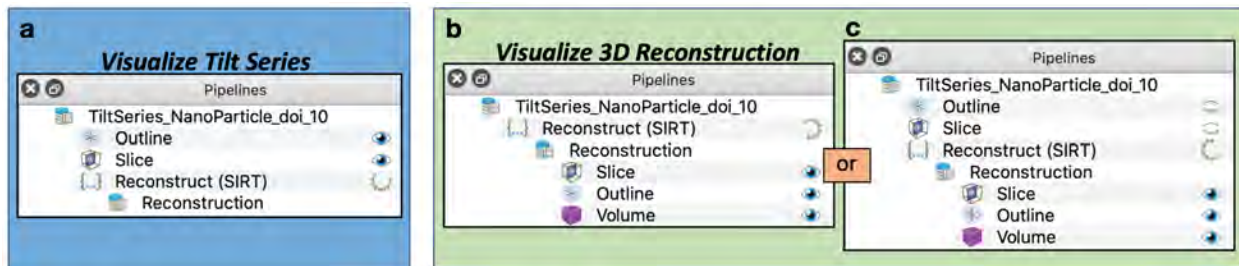


Figure B.6: **Visualizing Elements in the Data ‘Pipelines’**. Once a dataset has been loaded and live reconstruction begins, the pipeline can be populated with visualization modules of choice. **a**, By default, tomviz will continue displaying the input dataset after the reconstruction is initialized. **b-c**, To visualize updates for the live volumetric process, visualization modules should be present below the ‘Reconstruction’ dataset icon.

Volume rendering is an exceptional method of displaying volumetric objects. Each voxel in the volume is assigned both a color and an opacity based on its intensity. In tomviz, the color-opacity map can be adjusted interactively by dragging points on the line overlaying the histogram. The simplest method for examination of internal specimen structure is with individual 2D slices through the tomogram. Orthogonal slicing allows users to view slices through the data perpendicular to principal axes (x , y , z). Additional visualization modules such as a constant intensity contours can be selected for surface rendering.

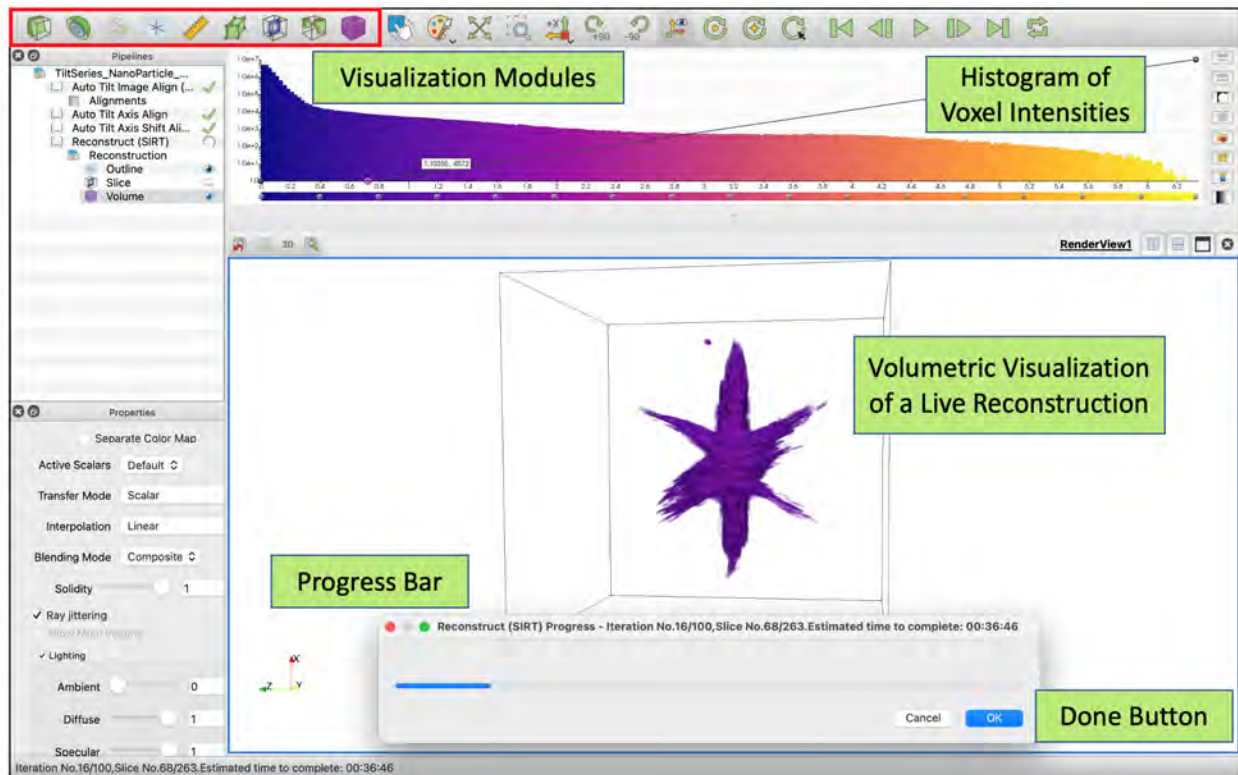


Figure B.7: **Visualizing a Live-SIRT reconstruction.** The tomviz graphical user interface for 3D visualization contains a variety of visualization tools for analyzing the 3D structure of specimens. 2D and 3D visualization modules can be selected from the top left menubar. Data transformations and visualizations are recorded in the Data Pipelines column on the left side for reproducible workflows. A histogram of voxel intensities is displayed on the top center where the black line represents the opacity map. Users can exit the reconstruction early by pressing the ‘Cancel’ or ‘Ok’ (Done) button.

B.2 Simulated Demonstration of Real-Time 3D Tomography

In this section we will simulate real-time tomography during tomographic experiments without the need for an electron microscope. Tomviz will monitor a local directory (folder) and automatically append new projections into the reconstruction process. To simulate an experiment, we will sequentially add pre-acquired images into the folder as though the data was being acquired. The steps in this tutorial closely correspond to an experimental scenario where tomviz has been installed on an electron microscope computer except the data has already been acquired.

(1) **Initialize real-time tomography.** To start a real-time reconstruction, load an empty dataset from the Sample Data drop down menu (Fig. B.8a) and press the blue ‘OK’ button (Fig. B.8b). A

volume comprised of elements with the value zero should appear. To start a real-time reconstruction, users can load any other dataset for this initial step. After a dataset is selected, select ‘Mark Data As Tilt Series’ from the Tomography drop down menu (Fig. B.9a). Here the tilt angles are irrelevant, continue by pressing the blue ‘Ok’ button and then ‘Initialize Real-Time Tomography’ (Fig. B.9b).




Figure B.8: Load an empty dataset to begin real-time tomography.

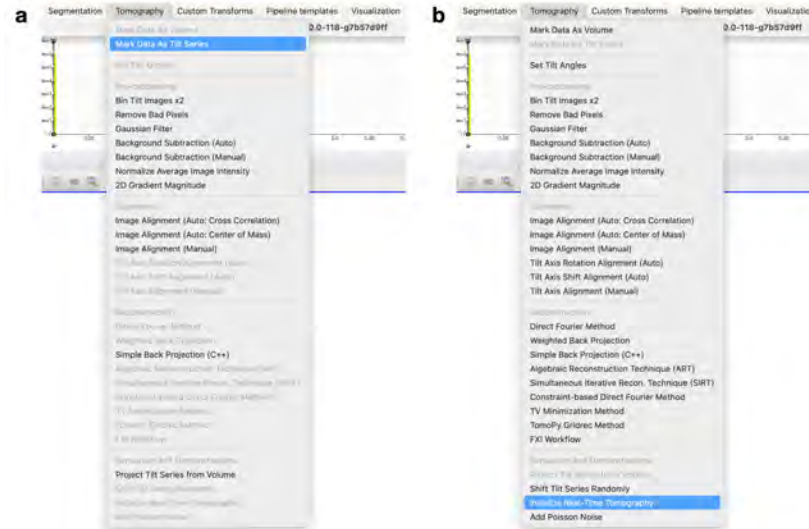


Figure B.9: Initialize a Real-Time Tomography Reconstruction.

(2) **Enter parameters to monitor the local directory where data will arrive.** Specify the directory that will be monitored for new projections (i.e. the folder where projections are saved). For this simulated experiment we will drag pre-acquired images into this folder. In a real tomographic experiment, the microscope acquisition software would automatically save projections images in this directory. For this demonstration we will define the target local directory as a folder on the Desktop: /Path/To/Desktop/real_time_tomo_demo (Fig. B.10).

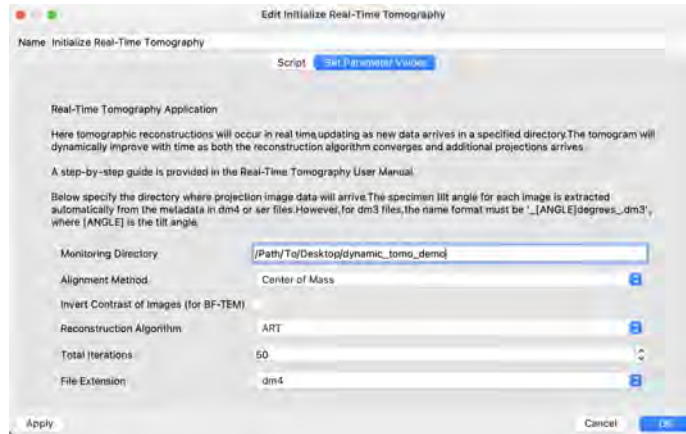


Figure B.10: GUI for real-time tomographic reconstruction.

(3) **Enter parameters for tomography reconstruction.** Reconstruction parameters such as the desired algorithm, number of iterations, and micrographs' file extension (dm4, dm3, ser) can be specified in this sub-menu. Tomviz can extract the tilt angle from the metadata embedded in dm4 and ser files. Unfortunately the metadata is not available in dm3 files, thus tomviz needs to parse the tilt angle from the file name in names in the following format: `*_[ANGLE]degrees_*.dm3`, where [ANGLE] is the tilt angle. The default parameters provided in the dialog box (Fig. B.10) will be sufficient. For this tutorial, use the tilt series provided as Supplementary Data Set 1. Download the dataset from the Supplemental, we will load the data into tomviz in the next step.

(4) **Run the real-time tomographic simulation.** Once ready, click the blue 'OK' button (Fig. B.10, bottom right hand corner). You should see a new element in the data pipelines column and a progress bar will appear (shown in Fig. B.7). The progress bar will initially display 'Initialize Real-Time Tomography Progress', this means the script is currently monitoring the folder `dynamic_tomo_demo` for new images. Start the simulation by dragging a few tilt-images into the target directory (Fig. B.11). Once projections are detected, the progress status message will update to 'Reconstructing Tilt Angles' and the reconstruction will begin.

Select the 'Reconstruction' dataset from the 'Pipelines' panel and either the volume or orthogonal slices from the visualization modules toolbar (Fig. B.7) to visualize the structure and observe the specimen evolution. New projections will be incorporated into the reconstruction process after the total number of iterations is complete (in this case 50). Copy 1-3 projections into the target folder every minute to simulate usual experimental acquisition speeds. We recommend copying data in chronological order, however, tomviz can still reconstruct the data if projections are passed in a random order. Once satisfied with the reconstruction, press 'Ok' (Fig. B.7) to exit the recon-

struction and continue with any desired post-processing.

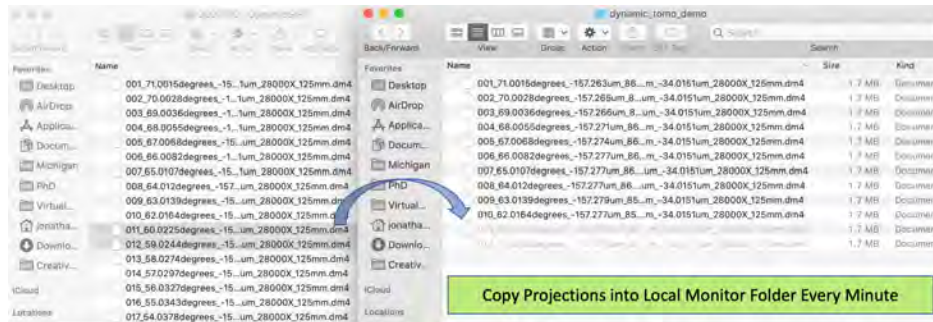


Figure B.11: **Copy individual projections into monitor directory.** Tomviz responds to changes in monitored folders once files are saved into the directory.

B.3 Tutorial: Live 3D Visualization During Tomography Experiments with Electron Microscopes

In this tutorial we will explain how to perform real-time tomographic experiments while projection images are collected on a scanning / transmission electron microscope (S/TEM). The process is nearly identical the previous section, however experimental real-time tomography requires that the microscope is aligned for tomography and the newly acquired data can be accessed by tomviz.

(1) **Align the microscope and prepare the tomography experiment.** We recommend using a high-tilt tomography holder to perform single-axis tomography. The microscope must be properly aligned before the acquisition begins to prevent any correctable distortions. Align the microscope as one normally would for high-resolution imaging to reduce beam aberrations (e.g. stigmatism or coma). Next, find the maximal allowable (positive and negative) tilt range such that the region of interest is visible. Achieving a $\pm 70^\circ$ tilt range or larger will provide best results.

Aligning the sample at eucentric height is essential to minimizing specimen drift and the need for any stage refinement during acquisition. Figure B.12a shows an object located below the eucentric height (\otimes) with the electron beam traveling from top to bottom. The object shifts and rotates when tilted around eucentric height, while proper alignment allows the specimen to simply rotate without any spatial translation (Fig. B.12b).




Figure B.12: **Specimen drift from eucentric height.** **a**, If the eucentric height for the goniometer is incorrect, tilting the specimen holder causes the specimen to sweep along an arc around the tilt axis. **b**, After correcting the stage position, the specimen drift is reduced significantly.

In order to bring the stage to eucentric height, one should measure the sample drift as the stage tilts from -50° to $+50^\circ$. Adjust the stage height by the drift's magnitude and iterate through this process until motion is minimized with minimal movement only due to the goniometer backlash. Choose a field of view which accommodates the specimen or region of interest across all tilts without the need to move the stage. Moreover, check that the detector camera length and gain does not clip image intensities at any one tilt. All calibrations and acquisition parameters for real-time tomography are now set properly and should not be changed during the tomography experiment. Move over to tomviz to set reconstruction parameters prior to starting the experiment.

(2) **Initialize real-time tomography.** To start a real-time reconstruction, load an empty dataset from the Sample Data drop down menu (Fig. B.8), Mark the Data as a Tilt Series (Fig. B.9a), and then select 'Initialize Real-Time Tomography' from the Tomography drop-down menu (Fig. B.9b).

(3) **Enter parameters for local directory monitoring.** Users can specify the directory that will be monitored for new projections (i.e. the folder where projections are saved). If tomviz is installed on the microscope computer, the acquisition software could automatically save projections images in this directory. However, this is not a necessary requirement as users can manually drag images into the folder or download it onto a personal machine (e.g. laptop). For this demonstration let's define the target local directory as a folder on the Desktop: `/Path/To/Desktop/real_time_tomo_demo` (Fig. B.10).

(4) **Enter parameters for tomography reconstruction.** Reconstruction parameters such as

the desired algorithm, number of iterations, and micrographs' file extension can be specified in this sub-menu. Tomviz will read the meta data to extract experimental parameters such as specimen alpha tilt. The default parameters provided in the dialog box (Fig. B.10) will be sufficient.

(5) **Start the real-time tomography experiment.** Once ready, click the blue 'OK' button (Fig. B.10, bottom right hand corner) and start collecting projection images. After collecting one or two tilt images, transfer the acquired micrographs into the target directory (Fig. B.11). The reconstruction will begin immediately after projections are detected inside the folder. Select either the volume or orthogonal slices from the visualization modules toolbar (Fig. B.7) to observe the specimen evolution. New projections will be incorporated into the reconstruction process after the total number of iterations is complete (in this case 50). Once satisfied with the reconstruction, press Ok (Fig. B.7) to exit the reconstruction and continue with any desired post-processing.

BIBLIOGRAPHY

- [1]
- [2] Wiebke Albrecht and Sara Bals. Fast electron tomography for nanomaterials. *J. Phys. Chem. C*, 124:27276–2786, 2020.
- [3] Leslie J Allen, SD Findlay, et al. Modelling the inelastic scattering of fast electrons. *Ultra-microscopy*, 151:11–22, 2015.
- [4] Anders Andersen. Algebraic reconstruction in ct from limited views. *IEEE Trans. Med. Imaging*, 8:50–55, 1989.
- [5] GR Anstis, DQ Cai, and DJH Cockayne. Limitations on the s-state approach to the interpretation of sub-angstrom resolution electron microscope images and microanalysis. *Ultra-microscopy*, 94(3-4):309–327, 2003.
- [6] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sciences*, 2:183–202, 2009.
- [7] James Bell, Muyuan Chen, Philip Baldwin, and Steven Ludtke. High resolution single particle refinement in eman2.1. *Methods*, 100:25–34, 2016.
- [8] H. Bender, C. Drijbrooms, and A. Radisic. Fib/sem structural analysis of through-silicon-vias. *AIP Conference Proceedings*, 1395:274–278, 2011.
- [9] Anuj Bhargava, Cindy Chen, Kapil Dhaka, Yuan Yao, Andrew Nelson, Kenneth Finkelstein, Christopher Pollock, Maytal Toroker, and Richard Robinson. Mn cations control electronic transport in spinel $\text{Co}_x\text{Mn}_{3-x}\text{O}_4$ nanoparticles. *Chem. Mater.*, 31:4228–4233, 2019.
- [10] Anuj Bhargava, Yuval Elbaz, Quynh Sam, Michelle A. Smeaton, Lena F. Kourkoutis, Maytal Caspary Toroker, and Richard D. Robinson. Enhanced li-ion diffusion and electrochemical performance in strained-manganese-iron oxide core-shell nanoparticles. *J. Chem. Phys.*, 155:144702, 2021.
- [11] Marouan Bouali and Said Ladjal. Toward optimal destriping of modis data using a unidirectional variational model. *IEEE Transactions on Geoscience and Remote Sensing*, 49:2924–2935, 2011.
- [12] Ronald Bracewell. Strip integration in radio astronomy. *Australian Journal of Physics*, 9:198–217, 1956.

- [13] Eric Brochu, Matthew Hoffman, and Nando de Freitas. Portfolio allocation for bayesian optimization. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*, pages 327–336. AUAI Press, 2011.
- [14] Hamish Brown, Jim Ciston, and Colin Ophus. Linear-scaling algorithm for rapid computation of inelastic transition of multiple electron scattering. *Phys. Rev. Research*, 1:033186, 2019.
- [15] Pierre Burdet, Z. Saghi, A.N. Filippin, A. Borrás, and P. Midgley. A novel 3d absorption correction method for quantitative edx-stem tomography. *Ultramicroscopy*, 160:118–129, 2016.
- [16] Jan-Willem Buurlage, Rob Bisseling, and Joost Batenburg. A geometric partitioning method for distributed tomographic reconstruction. *Parallel Computing*, 81:104–121, 2019.
- [17] Jan-Willem Buurlage, Federica Marone, Daniël Pelt, Willem Palenstijn, Marco Stamponi, Joost Batenburg, and Christian Schlepütz. Real-time reconstruction and visualization towards dynamic feedback control during time-resolved tomography experiments at tomcat. *Sci. Rep.*, 9:18379, 2019.
- [18] Thorsten Buzug. *From Photon Statistics to Modern Cone-Beam CT*. Springer, Berlin, 2008.
- [19] Vince Calhoun and Jing Sui. Multimodal fusion of brain imaging data: A key to finding the missing link(s) in complex mental illness. *Biol. Psychiatry: Cogn. Neurosci. and Neuroimaging*, 1:230–244, 2016.
- [20] Emmanuel Candes, Justin Romberg, and Terence Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inf. Theor.*, 52:489–509, 2006.
- [21] Emmanuel J. Candes and Michael B. Wakin. An introduction to compressive sampling. *IEEE Signal Processing Magazine*, 25:21–30, 2008.
- [22] Emmanuel Candès and Justin Romberg. Sparsity and incoherence in compressive sampling. *Inverse Problems*, 23:969, 2007.
- [23] Michael Cao, Zhen Chen, Yi Jiang, and Yimo Han. Automatic parameter selection for electron ptychography via bayesian optimization. *Scientific Reports*, 12:12284, 2022.
- [24] Yi Chang, Houzhang Fang, Luxin Yan, and Hai Liu. Robust destriping method with unidirectional total variation and framelet regularization. *Opt. Express*, 21:23307–23323, 2013.
- [25] Jinsong Chen, Yun Shao, Huadong Guo, Weiming Wang, and Boqin Zhu. Destriping cmodis data by power filtering. *IEEE Transactions on Geoscience and Remote Sensing*, 41:2119–2124, 2003.
- [26] Shu-wen Chen and Jean-Luc Pellequer. Destripe: frequency-based algorithm for removing stripe noises from afm images. *BMC Structural Biology*, 11:7, 2011.

- [27] Zhen Chen, Yi Jiang, Yu-Tsu Shao, Megan Holtz, Michal Odstrcil, Manuel Guizar-Scairos, Isabelle Hanke, Steffen Ganschow, Darrell Schlom, and David Muller. Electron ptychography achieves atomic-resolution limits set by lattice vibrations.
- [28] Lyndon Clarke, Ian Glendinning, and Rolf Hempel. The mpi message passing interface standard. *Programming Environments for Massively Parallel Distributed Systems*, pages 213–218, 1994.
- [29] Sean Collins and Paul A. Midgley. Progress and opportunities in eels and eds tomography. *Ultramicroscopy*, 180:133–141, 2017.
- [30] V.E. Cosslett. Radiation damage in the high resolution electron microscopy of biological materials: A review. *Journal of Microscopy*, 113:113–129, 1978.
- [31] A. Crewe, J. Wall, and J. Langmore. Visibility of single atoms. *Science*, 168:1338–1340, 1970.
- [32] Richard Crowther, Linda Amos, John Finch, Roiser. David, and Aaron Klug. Three dimensional reconstructions of spherical viruses by fourier synthesis from electron micrographs. *Nature*, 226:421–425, 1970.
- [33] Richard Anthony Crowther, D.J. DeRoiser, and Aaron Klug. The reconstruction of a three-dimensional structure from projections and its application to electron microscopy. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 317:319–340, 1970.
- [34] Imre Csiszár. Why least squares and maximum entropy? an axiomatic approach to inference for linear inverse problems. *Annals Stat.*, 19:2032–2066, 1991.
- [35] Paul Cueva, Robert Hovden, Julia Mundy, Huolin Xin, and David Muller. Data processing for atomic resolution electron energy loss spectroscopy. *Microscopy and Microanalysis*, 18:667–675, 2012.
- [36] Leonardo Dagum and Ramesh Menon. Openmp: An industry-standard api for shared-memory programming. *IEEE Comput. Sci. Eng.*, 5(1):46–55, 1998.
- [37] A. D’Alfonso, B. Freitag, D. Klenov, and L. Allen. Atomic-resolution chemical mapping using energy-dispersive x-ray spectroscopy. *Phys. Rev. B*, 81:100101, 2010.
- [38] Zicaho Wendy Di, Sven Leyffer, and Stefan Wild. Optimization-based approach for joint x-ray fluorescence and transmission tomographic inversion. *SIAM J. Imaging Sci.*, 9:1–23, 2016.
- [39] Zichao Wendy Di, Si Chen, Young Pyo Hong, Chris Jacobsen, Sven Leyffer, and Stefan Wild. Joint reconstruction of x-ray fluorescence and transmission tomography. *Opt. Express*, 25:13107–13124, 2017.
- [40] David Donoho. Compressed sensing. *IEEE Trans. Info. Theor.*, 52:1289–1306, 2006.

- [41] Kenneth Downing and Robert Glaeser. Restoration of weak phase-contrast images recorded with a high degree of defocus: The twin image problem associated with ctf correction. *Ultramicroscopy*, 108:921–928, 2008.
- [42] Marco F. Duarte and Yonina C. Eldar. Structured compressed sensing: From theory to applications. *IEEE Transactions on Signal Processing*, 59:4053–4085, 2011.
- [43] C Dwyer. The role of symmetry in the theory of inelastic high-energy electron scattering and its application to atomic-resolution core-loss imaging. *Ultramicroscopy*, 151:68–77, 2015.
- [44] Ray Egerton. Formulae for light-element microanalysis by electron energy-loss spectrometry. *Ultramicroscopy*, 3:243–251, 1978.
- [45] Ray Egerton and Shangcong Cheng. Characterization of an analytical electron microscope with a nio test specimen. *Ultramicroscopy*, 55:43–54, 1994.
- [46] R.F. Egerton. K-shell ionization cross-sections for use in microanalysis. *Ultramicroscopy*, 4:167–179, 1979.
- [47] R.F. Egerton. Formulae for light-element micro analysis by electron energy-loss spectrometry. *Ultramicroscopy*, 9:283–288, 1982.
- [48] R.F. Egerton. *Electron Energy-Loss Spectroscopy in the Electron Microscope*. Springer, 2011.
- [49] R.F. Egerton and Malac M. Eels in the tem. *Journal of Electron Spectroscopy and Related Phenomena*, 143:43–50, 2005.
- [50] Tommy Elfving, Per Christian Hansen, and Touraj Nikazad. Semi-convergence properties of kaczmarsz method. *Inverse Problems*, 30:1–16, 2014.
- [51] Tommy Elfving, Per Christian Hanser, and Nikazad Touraj. Semi-convergence and relaxation paramters for a class of sirt algorithms. *Electronic Transactions on Numerical Analysis*, 37:321–336, 2010.
- [52] Tommy Elfving, Nikazad Touraj, and Per Christian Hanser. Semi-convergence properties of kaczmarsz’s method. *Inverse Problems*, 30:055007, 2014.
- [53] Peter Ercius, Matthew Weyland, and David A. Muller. Three-dimensional imaging of nanovoids in copper interconnects using incoherent bright field tomography. *Appl. Phys. Lett.*, 88:243116, 2006.
- [54] S. Findlay, N. Shibata, H. Sawada, E. Okunishi, Y. Kondo, and Y. Ikuhara. Dynamics of annular bright field imaging in scanning transmission electron microscopy. *Ultramicroscopy*, 110:903–923, 2010.
- [55] F.L. Gadallah, F. Csillag, and J.M. Smith. Destriping multisensor imagery with moment matching. *International Journal of Remote Sensing*, 21:2505–2511, 2000.

- [56] Silvia Gazzola, Per Christian Hansen, and James Nagy. Ir tools: a matlab package of iterative regularization methods and large-scale test problems. *Numer. Algor.*, 81:773–811, 2018.
- [57] Peter Gilbert. Iterative methods for the three-dimensional reconstruction of an object from projections. *J. Theor. Biol.*, 36:105–117, 1972.
- [58] Jiseok Gim, Alden Koch, Laura M. Otter, Benjamin H. Savitzky, Sveinung Erland, Lara A. Estroff, Dorrit E. Jacob, and Robert Hovden. The mesoscale order of nacreous pearls. *P. Natl. Acad. Sci.*, 118:e2107477118, 2021.
- [59] Berit Goodge, Benjamin Geisler, Kyuho Lee, Motoki Osada, Bai Yang Wang, Danfeng Li, Harold Hwang, Pentcheva, Rossitza, and Lena Kourkoutis. Resolving the polar interface of infinite-layer nickelate thin films. *Nature Materials*, 22:466–473, 2023.
- [60] Bart Goris, Lakshmiarayana Polavarapu, Sara Bals, Gustaaf Van Tendeloo, and Luis Liz-Marzán. Monitoring galvanic replacement through three-dimensional morphological and chemical mapping. *Nano Lett.*, 14:3220–3226, 2014.
- [61] Matthew Guay, Wojciech Czaja, Maria Aronova, and Richard Leapman. Compressed sensing electron tomography for determining biological structure. *Scientific Reports*, 6:1–14, 2016.
- [62] Don-Hyung Ha, Andrew Caldwell, Matthew Ward, Shreyas Honrao, Kiran Matthew, Robert Hovden, Margaret Koker, David Muller, Richard Hennig, and Richard Robinson. Solid–solid phase transformations induced through cation exchange and strain in 2d heterostructured copper sulfide nanocrystals. *Nano Lett.*, 14:7090–7099, 2014.
- [63] Jordan A. Hachtel, Jingsong Huang, Ilja Popovs, Santa Jansone-Popova, Jong K. Keum, Jacek Jakowski, Tracy C. Lovejoy, Niklas Dellby, Ondrej L. Krivanek, and Juan Carlos Idrobo. Identification of site-specific isotopic labels by vibrational spectroscopy in the electron microscope. *Science*, 363:525–528, 2019.
- [64] David L Hall and James Llinas. An introduction to multisensor data fusion. *Proceedings of the IEEE*, 85(1):6–23, 1997.
- [65] Yoseb Han, Jaeduck Jang, Eunju Cha, Junho Lee, Hyungjin Chung, MyoungHo Jeong, Tae-Gon Kim, Byeong Gyu Chae, Hee Goo Kim, Shinae Jun, Sungwoo Hwang, Eunha Lee, and Jong Chul Ye. Deep learning stem-edx tomography of nanocrystals. *Nature Machine Intelligence*, 3:267–274, 2021.
- [66] Marcus D Hanwell, Kenneth M Martin, Aashish Chaudhary, and Lisa S Avila. The visualization toolkit (vtk): Rewriting the rendering code for modern graphics cards. *SoftwareX*, 1:9–12, 2015.
- [67] James L Hart, Andrew C Lang, Asher C Leff, Paolo Longo, Colin Trevor, Ray D Twesten, and Mitra L Taheri. Direct detection electron energy-loss spectroscopy: a method to push the limits of resolution and sensitivity. *Sci. Rep.*, 7(1):1–14, 2017.

- [68] P Hartel, H Rose, and C Dinges. Conditions and reasons for incoherent imaging in stem. *Ultramicroscopy*, 63(2):93–114, 1996.
- [69] Peter W. Hawkes. *The Electron Microscope as a Structure Projector*, pages 17–38. Springer, Boston, MA, 1992.
- [70] R. Hegerl and W. Hoppe. Influence of electron noise on three-dimensional image reconstruction. *Zeitschrift fur Naturforschung A*, 31:1717–1721, 1976.
- [71] L. Holzer, P.H. Gasser, A. Kaech, M. Wegmann, A. Zingg, R. Wepf, and B. Muench. Cryo-fib-nanotomography for quantitative analysis of particle structures in cement suspensions. *J. Microsc.*, 227:216–228, 2007.
- [72] W. Hoppe, J. Gassman, N. Hunsmann, H.J. Schramm, and M. Sturm. Three-dimensional reconstruction of individual negatively stained yeast fatty-acid synthetase molecules from tilt series in the electron microscope. *Hoppe Seylers Z Physiol Chem.*, 355:1483–1487, 1974.
- [73] W Hoppe and R Hegerl. Three-dimensional structure determination by electron microscopy (nonperiodic specimens). In *Computer processing of electron microscope images*, pages 127–185. Springer, 1980.
- [74] W. Hoppe, H.J. Schramm, M. Sturm, N. Hunsmann, and J. Gabmann. Three dimensional electron microscopy of individual biological objects - part iii. experimental results on yeast fatty acid synthetase. *Z. Naturforsch.*, 31:1380–1390, 1976.
- [75] Alain Hore and Djemel Ziou. Image quality metrics: Psnr and ssim. *2010 20th International Conference on Pattern Recognition*, pages 2366–2369, 2010.
- [76] Robert Hovden and David A Muller. Efficient elastic imaging of single atoms on ultrathin supports in a scanning transmission electron microscope. *Ultramicroscopy*, 123:59–65, 2012.
- [77] Robert Hovden, Huolin L Xin, and David A Muller. Channeling of a subangstrom electron beam in a crystal mapped to two-dimensional molecular orbitals. *Phys. Rev. B*, 86(19):195415, 2012.
- [78] Richard Huber, Georg Haberehner, Martin Holler, Gerald Kothleitner, and Kristian Bredies. Total generalized variation regularization for multi-modal electron tomography. *Nanoscale*, 11:5617–5632, 2019.
- [79] Martin Jacob, Julien Sorel, Rafael Bortolin Piinhiero, Frederic Mazen, Adeline Grenier, Thierry Epicier, and Zineb Saghi. Correlative stem-haadf and stem-edx tomography for the 3d morphological and chemical analysis of semiconductor devices. *Semicond. Sci. Technol.*, 36:035006, 2021.
- [80] Momin Jamil and Xin-She Yang. A literature survey of benchmark functions for global optimization problems. *Int. Journal of Mathematical Modelling and Numerical Optimisation*, 4:150–194, 2013.

- [81] Wenfeng Jiang, Zhi-Bei Qu, Prashant Kumar, Drew Vecchio, Yuefei Wang, Yu Ma, Joong Hwan Bahng, Kalil Bernardino, Weverson Gomes, Felipe Colombari, Asdrubal Lozada-Blanco, c Veksler, Emanuele Marino, Alex Simon, Christopher Murray, Sergio Muniz, Andre Moura, and Nicholas Kotov. Emergence of complexity in hierarchically organized chiral particles. *Science*, pages 642–648, 2020.
- [82] Yi Jiang, Zhen Chen, Pratiti Han, Yimo Deb, Hui Gao, Saien Xie, Prafull Purohit, Mark W. Tate, Jiwoong Park, Sol M. Gruner, Veit Elser, and David A. Muller. Electron ptychography of 2d materials to deep sub-ångström resolution. *Nature*, 559:343–349, 2018.
- [83] Yi Jiang, Elliot Padgett, Robert Hovden, and David A. Muller. Sampling limits for electron tomography with sparsity-exploiting reconstructions. *Ultramicroscopy*, 186:94–103, 2018.
- [84] Mockus Jonas. Application of bayesian approach to numerical methods of global and stochastic optimization. *J. Glob. Optim.*, 4:347–365, 1994.
- [85] Mundy Julia, Yasuyuki Hikita, Taekai Hidaka, Takeaki Yajima, Takuya Higuchi, Harold Hwang, David Muller, and Lena Kourkoutis. Visualizing the interfacial evolution from charge compensation to metallic screening across the manganite metal–insulator transition. *Nature Communications*, 5:3464, 2014.
- [86] Daniil Kazantsev, Jakob Jørgensen, Martin Andersen, William Lionheart, Peter Lee, and Philip Withers. Joint image reconstruction method with correlative multi-channel prior for x-ray spectral computed tomography. *Inverse Problems*, 34:064001, 2018.
- [87] Aaron Klug and Richard Anthony Crowther. Three-dimensional image reconstruction from the viewpoint of information theory. *Nature*, 238:435–440, 1972.
- [88] G. Kothleitner, M. Neish, N. Lugg, S. Findlay, W. Grogger, F. Hofer, and L. Alle. Quantitative elemental mapping at atomic resolution using x-ray spectroscopy. *Phys. Rev. Lett.*, 112:085501, 2014.
- [89] Paul Kotula, Dmitri Klenov, and Sebastian Harrach. Challenges to quantitative multivariate statistical analysis of atomic-resolution x-ray spectral. *Microscopy and Microanalysis*, 18:691–698, 2012.
- [90] Ondrej L Krivanek, Matthew F Chisholm, Valeria Nicolosi, Timothy J Pennycook, George J Corbin, Niklas Dellby, Matthew F Murfitt, Christopher S Own, Zoltan S Szilagy, Mark P Oxley, et al. Atom-by-atom structural and chemical analysis by annular dark-field electron microscopy. *Nature*, 464(7288):571–574, 2010.
- [91] Dana Lahat, Tülay Adali, and Christian Jutten. Multimodal data fusion: An overview of methods, challenges, and prospects. *Proceedings of the IEEE*, 103:1449–1477, 2015.
- [92] Wolfram Langheinrich and Heinz Beneking. The resolution of the inorganic electron beam resist lif(alf3). *Microelectron. Eng.*, 23:287–290, 1994.
- [93] Rowan Leary, Zineb Saghi, Paul Midgley, and Daniel Holland. Compressed sensing electron tomography. *Ultramicroscopy*, 131:70–91, 2013.

- [94] James M LeBeau, Scott D Findlay, Leslie J Allen, and Susanne Stemmer. Quantitative atomic resolution scanning transmission electron microscopy. *Phys. Rev. Lett.*, 100(20):206101, 2008.
- [95] K. Lepinay, F. Lorut, R. Pantel, and T. Epicer. Chemical 3d tomography of 28 nm high k metal gate transistor: Stem xeds experimental method and results. *Micron*, 47:43–49, 2013.
- [96] Barnaby DA Levin, Elliot Padgett, Chien-Chun Chen, MC Scott, Rui Xu, Wolfgang Theis, Yi Jiang, Yongsoo Yang, Colin Ophus, Haitao Zhang, et al. Nanomaterial datasets to advance tomography in scanning transmission electron microscopy. *Scientific data*, 3(1):1–11, 2016.
- [97] Meng-Yue Li, Wen-Duo Lu, Lei He, Ferdi Schüth, and An-Hui Lu. Tailoring the surface structure of silicon carbide support for copper catalyzed ethanol dehydrogenation. *Chem Cat Chem*, 11:481–487, 2019.
- [98] Xiao Liang, Yali Zang, Di Dong, Liwen Zhang, Mengjie Fang, Xin Yang, Alicia Arranz, Jorge Ripoll, Hui Hui, and Jie Tian. Stripe artifact elimination based on nonsubsampling contourlet transform for light sheet fluorescence microscopy. *J. Biomed Opt.*, 21:106005, 2016.
- [99] Ruoqian Lin, Enyuan Hu, Mingjie Liu, Yi Wang, Hao Cheng, Jinpeng Wu, Jin-Cheng Zheng, Qin Wu, Seongmin Bak, Xiao Tong, Rui Zhang, Wanli Yang, Kristin A. Persson, Xiqian Yu, Xiao-Qing Yang, and Huolin L. Xin. Anomalous metal segregation in lithium-rich material provides design rules for stable cathode in lithium-ion battery. *Nat. Commun.*, 10:1650, 2019.
- [100] Chi-Chun Liu, Elliot Franke, Yann Mignot, Ruilong Xie, Chun Wing Yeung, Jingyun Zhang, Chen Chi, Richard Farrell, Kafai Lai, Hsin-yu Tai, Nelson Felix, and Daniel Corliss. Directed self-assembly of block copolymers for 7 nanometre finfet technology and beyond. *Nature Electronics*, 1:562–569, 2018.
- [101] Li Liu, Yongxin Han, and Mingwu Jin. Fast alternating projection methods for constrained tomographic reconstruction. *PLoS One*, 12:1–23, 2017.
- [102] Yuan Lo, Chen-Ting Lio, Jihan Zhou, Arjun Rana, Charles Bevis, Guan Gui, Bjoern Enders, Kevin Canno, Young-Sang Yu, Richanrd Celestre, Nowrouzi Kasra, David Shapiro, Henry Kapteyn, Roger Falcone, Chris Bennett, Margaret Murnane, and Jianwei Miao. Multimodal x-ray and electron microscopy of the allende meteorite. *Sci. Adv.*, 5, 2019.
- [103] Michael Lustig, David Donoho, and John M. Pauly. Sparse mri: The application of compressed sensing for rapid mr imaging. *Magnetic Resonance in Medicine*, 58:1182–1195, 2007.
- [104] Yunhai Ma, Dali Cai, Yiru Li, Ning Wang, Usman Muhammad, Anna Carlsson, Dong Tang, Weizhong Qian, Yao Wang, Dangsheng Su, and Fei Wei. The influence of straight pore blockage on the selectivity of methanol to aromatics in nanosized zn/zsm-5: an atomic cs-corrected stem analysis study. *RSC Adv.*, 6:74797–74801.

- [105] Bruce F. McEwen, Michael Marko, Chyong-Ere Hsieh, and Carmen Mannella. Use of frozen-hydrated axonemes to assess imaging parameters and resolution limits in cryoelectron tomography. *J. Struct. Biol.*, 138:47–57, 2002.
- [106] G. McMullan, A. Faruqi, D. Clare, and R. Henderson. Comparison of optimal performance at 300keV of three direct electron detectors for use in low dose electron microscopy. *Ultramicroscopy*, 147:156–163, 2014.
- [107] Ohad Medalia, Igor Weber, Achilleas Frangakis, Daniela Nicastro, Günther Gerisch, and Wolfgang Baumeister. Macromolecular architecture in eukaryotic cells visualized by cryo-electron tomography. *Science*, 298:1209–1213, 2002.
- [108] Hamed Heidari Mezerji, Wouter Van den Broek, and Sara Bals. A practical method to determine the effective resolution in incoherent experimental electron tomography. *Ultramicroscopy*, 111:330–336, 2011.
- [109] Aaron Michelson, Brian Mineevich, Hamed Emamy, Xiaojing Huang, Yong Chu, Hanfei Yan, and Oleg Gang. Three-dimensional visualization of nanoparticle lattices and multimaterial frameworks. *Science*, 376:203–207, 2022.
- [110] Paul Midgley and Matthew Weyland. 3d electron microscopy in the physical sciences: The development of z-contrast and efem tomography. *Ultramicroscopy*, 96:413–431, 2003.
- [111] Paul A. Midgley, Matthew Weyland, John Meurig Thomas, and Brian F.G. Johnson. Z-contrast tomography: a technique in three-dimensional nanostructural analysis based on rutherford scattering. *Chem. Commun.*, pages 907–908, 2001.
- [112] Joseph Mindell and Nikolaus Grigorieff. Accurate determination of local defocus and specimen tilt in electron microscopy. *J. Struct. Biol.*, 142:334–347, 2003.
- [113] G. Mobus and B. Inkson. Three-dimensional reconstruction of buried nanoparticles by element-sensitive tomography based on inelastically scattered electrons. *Appl. Phys. Lett.*, 79:1369, 2001.
- [114] David Muller, Lena Kourkoutis, Matthew Murfitt, Jong Song, Harold Hwang, John Silcox, Niklas Dellby, and Ondrej Krivanek. Atomic-scale chemical imaging of composition and bonding by aberration-corrected microscopy. *Science*, 319:1073–1076, 2008.
- [115] Beat Munch, Pavel Trtik, Federica Marone, and Marco Stampanoni. Stripe and ring artifact removal with combined wavelet — fourier filtering. *Opt. Express*, 17:8567–8591, 2009.
- [116] A. Muray, M. Scheinfein, M. Isaacson, and I. Adesida. Radiolysis and resolution limits of inorganic halide resists. *Journal of Vacuum Science and Technology B: Microelectronics Processing and Phenomena*, 3:367–372, 1985.
- [117] B.K. Natarajan. Sparse approximate solutions to linear systems. *SIAM Journal on Computing*, 24:227–234.

- [118] Yu Nesterov. A method of solving a convex programming problem with convergence rate $O(1/k^2)$. *Sov. Math Dokl.*, 27:372–376, 1983.
- [119] Olivia Nicoletti, Francisco de la Peña, Rowan K. Leary, Daniel J. Holland, Caterina Ducati, and Paul A. Midgley. Three-dimensional imaging of localized surface plasmon resonances of metal nanoparticles. *Nature*, 502:80–84.
- [120] Michal Odstrčil, Andreas Menzel, and Manuel Guizar-Sicairos. Iterative least-squares solver for generalized maximum-likelihood ptychography. *Opt. Express*, 26:3108–3123, 2018.
- [121] Myoung Hwan Oh, Min Gee Cho, Dong Young Chung, Inchul Park, Youngwook Paul Kwon, Colin Ophus, Dokyoon Kim, Min Gyu Kim, Beomgyun Jeong, X. Wendy Gu, Jinwoung Jo, Ji Mun Yoo, Jaeyoung Hong, Sara McMains, Kisuk Kang, Yung-Eun Sung, A. Paul Alivisatos, and Taeghwan Hyeon. Design and synthesis of multigrain nanocrystals via geometric misfit strain. *Nature*, 577:359–363, 2020.
- [122] Elena Orlova and Helen Saibil. Structural analysis of macromolecular assemblies by electron microscopy. *Chem. Rev.*, 111:7710–7748, 2011.
- [123] Elliott Padgett, Robert Hovden, Jessica DaSilva, Barnaby Levin, John Grazul, Tobias Hanrath, and David Muller. A simple preparation method for full-range electron tomography of nanoparticles and fine powders. *Microscopy and Microanalysis*, 23:1150–1158, 2017.
- [124] M. Pan and P. Crozier. Quantitative imaging and diffraction of zeolites using a slow-scan ccd camera. *Ultramicroscopy*, 52:487–498, 1993.
- [125] Sanjaya Perera, Xiaoyue Ding, Anuj Bhargava, Robert Hovden, Andrew Nelson, Lena Kourkoutis, and Richard Robinson. Enhanced supercapacitor performance for equal co–mn stoichiometry in colloidal $\text{Co}_{3-x}\text{Mn}_x\text{O}_4$ nanoparticles, in additive-free electrodes. *Chem. Mater.*, 27:7861–7873, 2015.
- [126] M. Pfannmoller, H. Heidari, L. Nanson, O.R. Lozman, M. Chrapa, T. Offermans, G. Nisato, and S. Bals. Quantitative tomography of organic photovoltaic blends at the nanoscale. *Nano Lett.*, 15:6634–6642, 2015.
- [127] Hongliang Qi, Zijia Chen, and Linghong Zhou. Ct image reconstruction from sparse projections using adaptive tpv regularization. *Computational and Mathematical Methods in Medicine*, 2015:3544869, 2015.
- [128] Michael Radermacher. Three-dimensional reconstruction of single particles from random and nonrandom tilt series. *J. Electron. Microsc. Tech.*, 9:359–394, 1988.
- [129] J Radon. ber die bestimmung von funktionen durch ihre integralwerte lngs gewisser mannigfaltigkeiten. *Akad. Wiss.*, 69:262–277, 1917.
- [130] Preesan Rakwatin, Wataru Takeuchi, and Yoshifumi Yasuoka. Stripe noise reduction in modis data by combining histogram matching with facet filter. *IEEE Transactions on Geoscience and Remote Sensing*, 45:1844–1856, 2007.

- [131] Peter Rez. Cross-sections for energy loss spectroscopy. *Ultramicroscopy*, 9:283–288, 1982.
- [132] Alexis Rohou and Nikolaus Grigorieff. Ctffind4: Fast and accurate defocus estimation from electron micrographs. *J. Struct. Biol.*, 192:216–221, 2015.
- [133] David Roiser and Aaron Klug. A reconstruction of three dimensional structures from electron micrographs. *Nature*, 217:130–134, 1968.
- [134] David Rossouw and Robert Krakow. Blind source separation aided characterization of the γ strengthening phase in an advanced nickel-based superalloy by spectroscopic 4d electron microscopy. *Acta Materialia*, 107:229–238, 2016.
- [135] Steven Rozeveld, Wen-Sheng Lee, and Paolo Longo. Characterization of cobalt sulfide catalysts. *Microscopy and Microanalysis*, 26 (S2):1248–1250, 2020.
- [136] Leonid Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60:259–268, 1992.
- [137] Toby Sanders, Anne Gelb, Rodrigo Platte, Ilke Arslan, and Kai Landskron. Recovering fine details from under-resolved electron tomography data using higher order total variation l1 regularization. *Ultramicroscopy*, 174:97–105, 2017.
- [138] Benjamin Savitzky, Ismail Baggari, Alemayehu Admasu, Jaewook Kim, Sang-Wook Cheong, Robert Hovden, and Lena Kourkoutis. Bending and breaking of stripes in charge ordered manganite. *Nat. Commun.*, 8:1883, 2017.
- [139] Christopher Schankula, Christopher Anand, and Nabil Bassim. Plasma focused ion beam curtaining artifact correction by fourier-based linear optimization model. *Microscopy and Microanalysis*, 24:588–589, 2018.
- [140] Otto Scherzer. Über einige fehler von elektronenlinsen. *Zeitschrift für Physik*, 101:593–603, 1936.
- [141] Otto Scherzer. Sphaerische und chromatische korrektur von elektronenlinsen. *Optik*, 2:114–132, 1947.
- [142] P. Schlossmacher, D.O. Klenov, B. Freitag, and H.S. von Harrach. Enhanced detection sensitivity with a new windowless xeds system for aem based on silicon drift detector technology. *Microscopy Today*, 18:14–20, 2010.
- [143] Martin Schorb, Isabella Haberbosch, Hagen Wim, Yannick Schwab, and David Mastronarde. Software tools for automated transmission electron microscopy. *Nat. Methods*, 16:471–477, 2019.
- [144] Jonathan Schwartz, Zichao Wendy Di, Ayssa Fielitz, Don-Hyung Ha, Sanjaya Perera, Ismail El Baggari, Richard Robinson, Jeffery Fessler, Colin Ophus, Steve Rozeveld, and Robert Hovden. Imaging atomic-scale chemistry from fused multi-modal electron microscopy. *npj Comput. Mater.*, 8:16, 2022.

- [145] Jonathan Schwartz, Chris Harris, Jacob Pietryga, Huihuo Zheng, Prashant Kumar, Visheratina Anastasiia, Nicholas Kotov, Brianna Major, Patric Avery, Peter Ercius, Utkarsh Ayachit, Berk Geveci, David Muller, Alessandro Genova, Yi Jiang, Marcus Hanwell, and Robert Hovden. Real-time 3d analysis during tomographic experiments on tomviz. *Microscopy and Microanalysis*, 27(S1):2860–2862, 2021.
- [146] Jonathan Schwartz, Chris Harris, Jacob Pietryga, Huihuo Zheng, Prashant Kumar, Anastasiia Visheratina, Nicholas A. Kotov, Brianna Major, Patrick Avery, Peter Ercius, Utkarsh Ayachit, Berk Geveci, David A. Muller, Alessandro Genova, Yi Jiang, Marcus Hanwell, and Robert Hovden. Real-time 3d analysis during electron tomography using tomviz. *Nat. Comm.*, 2022:4458, 2022.
- [147] Jonathan Schwartz, Yi Jiang, Yongjie Wang, Anthony Aiello, Pallab Bhattacharya, Hui Yuan, Zetian Mi, Nabil Bassim, and Robert Hovden. Removing stripes, scratches, and curtaining with nonrecoverable compressed sensing. *Microscopy and Microanalysis*, 25:705–710, 2019.
- [148] Jonathan Schwartz, Huihuo Zheng, Marcus Hanwell, Yi Jiang, and Robert Hovden. Dynamic compressed sensing for real-time tomographic reconstruction. *Ultramicroscopy*, 219:113122, 2020.
- [149] Mary Scott, Chien-Chun Chen, Matthew Mecklenburg, Rui Xu, Peter Ercius, Ulrich Dahmen, Chris Regan, and Jianwei Miao. Electron tomography at 2.4-ångström resolution. *Nature*, 483:444–447, 2012.
- [150] MC Scott, Chien-Chun Chen, Matthew Mecklenburg, Chun Zhu, Rui Xu, Peter Ercius, Ulrich Dahmen, BC Regan, and Jianwei Miao. Electron tomography at 2.4-ångström resolution. *Nature*, 483(7390):444–447, 2012.
- [151] O. E. Shklyaev, M. J. Beck, M. Asta, M. J. Miksis, and P. W. Voorhees. Role of strain-dependent surface energies in Ge/Si(100) island formation. *Phys. Rev. Lett.*, 94:176102, 2005.
- [152] Emil Sidky, Kao Chien-Min, and Pan Xiaochuan. Accurate image-reconstruction from few-views and limited-angle data in divergent-beam ct. *J.X-Ray Sci. Technol.*, 14:119–139, 2006.
- [153] Emil Sidky, Yuval Duchin, and Xiaochuan Pan. A constrained, total-variation minimization algorithm for low-intensity x-ray ct. *Med. Phys.*, 38:S117–S125, 2011.
- [154] Emil Sidky and Xiaochuan Pan. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Phys. Med. Biol.*, 53:4777–4807, 2018.
- [155] Thomas Slater, Arne Janssen, Pedro H.C. Camargo, M. Grace Burke, Nestor J. Zaluzec, and Sarah J. Haigh. Stem-edx tomography of bimetallic nanoparticles: A methodological investigation. *Ultramicroscopy*, 162:61–73, 2016.
- [156] John Spence and J. Lynch. Stem microanalysis by transmission electron energy loss spectroscopy in crystals. *Ultramicroscopy*, 9:267–276, 1982.

- [157] Yingying Su, Marko Nykanen, Kristina Jahn, Renee Whan, Laurence Cantrill, Lilian L. Soon, Kyle Ratinac, and Filip Braet. Multi-dimensional correlative imaging of subcellular events: combining the strengths of light and electron microscopy. *Biophys. Rev.*, 2:121–135, 2010.
- [158] Mark W. Tate, Prafull Purohit, Darol Chamberlain, Kayla X. Nguyen, Robert Hovden, Celesta S. Chang, Pratiti Deb, Emrah Turgut, John T. Heron, Darrell G. Schlom, Daniel C. Ralph, Gregory D. Fuchs, Katherine S. Shanks, David A. Philipp, Hugh T. Muller, and Sol M. Gruner. High dynamic range pixel array detector for scanning transmission electron microscopy. *Microscopy and Microanalysis*, 22:237–249, 2016.
- [159] Dmitry Tegunov and Patrick Cramer. Real-time cryo electron microscopy data preprocessing with warp. *Nat. Methods*, 16:1146–1152, 2019.
- [160] John Meurig Thomas, Rowan Laery, Paul A. Midgley, and Daniel J. Holland. A new approach to the investigation of nanoparticles: Electron tomography with compressed sensing. *Journal of Colloid and Interface Science*, 392:7–14, 2013.
- [161] Jorge Torres and Sergio Infante. Wavelet analysis for the elimination of striping noise in satellite images. *Opt. Eng.*, 40:1309–1314, 2001.
- [162] Michael W. Vannier. Image reconstruction from projections implementation and applications. volume 32 in topics in applied physics. *Journal of Nuclear Medicine*, 21(10):1005–1005, 1980.
- [163] D. Vesely. Electron beam damage of amorphous synthetic polymers. *Ultramicroscopy*, 14:279–290, 1984.
- [164] J. Wall, Michael Isaacson, and John Langmore. The collection of scattered electrons in dark field electron microscopy. *Optik*, 39:359–374, 1974.
- [165] Shizhong Wang, Benjamin R. Jarrett, Susan M. Kauzlarich, and Angelique Y. Louie. Core/shell quantum dots with high relaxivity and photoluminescence for multimodality imaging. *J. Am. Chem. Soc.*, 129:1848–3856, 2007.
- [166] M. Watanabe and D.B. Williams. The quantitative analysis of thin specimens: a review of progress from the cliff-lorimer to the new ζ -factor methods. *Journal of Microscopy*, 221:89–109, 2006.
- [167] Rui Xu, Chien-Chun Chen, Li Wu, Mary C. Scott, Wolfgang Theis, Colin Ophus, Matthias Bartels, Yongsoo Yang, Hadi Ramezani-Dakhel, Michael R. Sawaya, Hendrik Heinz, Laurence D. Marks, Peter Ercius, and Jianwei Miao. Three-dimensional coordinates of individual atoms in materials revealed by electron tomography. *Nature Mater.*, 14:1099–1103, 2015.
- [168] Wei Xu, Fang Xu, Mel Jones, Bettina Keszthelyi, John Sedat, David Agard, and Klaus Mueller. High-performance iterative electron tomography reconstruction with long-object compensation using graphics processing units (gpu). *J. Struct. Biol.*, 171:142–153, 2010.

- [169] Reed Yalisove, Suk Hyun Sung, Peter Ercius, and Robert Hovden. Limits of three-dimensional resolution and dose for aberration-corrected electron tomography. *Phys. Rev. Appl.*, 15:014003, 2021.
- [170] Jiao Yan, Wenchun Feng, Ji-Young Kim, Jun Lu, Prashant Kumar, Zhengzhi Mu, Xiaochun Wu, Xiaoming Mao, and Nicholas Kotov. Self-assembly of chiral nanoparticles into semiconductor helices with tunable near-infrared optical activity. *Chem. of Mater.*, 32:476–488, 2019.
- [171] Yao Yang, Jihan Zhou, Yakun Yuan, Yakun Chang, Dilan Chang, Dennis Kim, Minh Pham, Arjun Rana, Xuezheng Tian, Yonggang Yao, Stanley Osher, Andreas Schmid, Liangbing Hu, Peter Ercius, and Jianwei Miao. Determining the three-dimensional atomic structure of an amorphous solid. *Nature*, 592:60–64, 2021.
- [172] Yao Yang, Jihan Zhou, Fan Zhu, Yakun Yuan, Dillan J Chang, Dennis S Kim, Minh Pham, Arjun Rana, Xuezheng Tian, Yonggang Yao, et al. Determining the three-dimensional atomic structure of an amorphous solid. *Nature*, 592(7852):60–64, 2021.
- [173] Yongsoo Yang, Chien-Chun Chen, Mary Scott, Colin Ophus, Rui Xu, Alan Pryor, Li Wu, Fan Sun, Wolfgang Theis, Jihan Zhou, Markus Eisenbach, Paul Kent, Renat Sabirianov, Hao Zeng, Peter Ercius, and Jianwei Miao. Deciphering chemical order/disorder and material properties at the single-atom level. *Nature*, 542:75–79, 2017.
- [174] S. Zaefferer, S.I. Wright, and D. Raabe. Three-dimensional orientation microscopy in a focused ion beam–scanning electron microscope: A new dimension of microstructure characterization. *Metall. Mater. Trans. A*, 39:374–389, 2008.
- [175] Rui Zeng, Yao Yang, Xinran Feng, Huiqi Li, Lauryn M. Gibbs, Francis J. Disalvo, and Hector D. Abruña. Nonprecious transition metal nitrides as efficient oxygen reduction electrocatalysts for alkaline fuel cells. *Science Advances*, 8:eabj1584, 2022.
- [176] Haitao Zhang, Don-Hyung Ha, Robert Hovden, Lena Kourkoutis, and Richard Robinson. Controlled synthesis of uniform cobalt phosphide hyperbranched nanocrystals using tri-n-octylphosphine oxide as a phosphorus source. *Nano Lett.*, 11:189–197, 2010.
- [177] Xiaomeng Zhang, Jing Wang, and Lei Xing. Ct metal artifact reduction in x-ray computed tomography by constrained optimization. *Med. Phys.*, 38:701–711, 2011.
- [178] Yichi Zhang, Daniel Apley, and Wei Chen. Bayesian optimization for materials design with mixed quantitative and qualitative variables. *Scientific Reports*, 10:4924, 2020.
- [179] Zezhong Zhang, Annick De Backer, Ivan Lobato, Sandra Van Aert, and Peter Nellist. Combining adf-edx scattering cross-sections for elemental quantification of nanostructures. *Micromscopy and Microanalysis*, 27(S1):600–602, 2021.