UNIVERSITY OF CALIFORNIA,
IRVINE

Electron Microscopy of Dynamic Soft Matter and Nanoscale Materials: Integrating Liquid-Phase Transmission Electron Microscopy, Physics-Based Simulation, and Deep Learning

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Materials Science and Engineering

by

Zhaoxu Li

Dissertation Committee:
Professor Joseph P. Patterson, Chair
Professor Zhibin Guan
Associate Professor Stacy Copp

2026

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I am deeply grateful to my advisor, Professor Joseph P. Patterson, for his mentorship, scientific guidance, and support throughout my doctoral research. I also thank Professor Zhibin Guan and Associate Professor Stacy Copp for their service on my dissertation committee and their thoughtful feedback.

I thank the Patterson Group and the UCI Center for Complex and Active Materials, especially the IRG-2 community. Biweekly meetings, research presentations, and discussions provided valuable criticism and helped shape this work. I am especially grateful to Dr. Dipankar Barpuzary for synthesizing FFssFF and preparing the initial FFssFF solution used in the electrochemical LPTEM study described in Chapter 1.

This research was supported primarily by the National Science Foundation Materials Research Science and Engineering Center program through the UC Irvine Center for Complex and Active Materials (DMR-2011967). I acknowledge the Irvine Materials Research Institute and UCI Research Cyberinfrastructure Center for microscopy and high-performance-computing support.

Chapter 2 is adapted, with minor editorial modifications, from Y. Song, Z. Li, J. T. Mulvey, J. A. Freites, J. P. Patterson, and D. J. Tobias, “Cryogenic Electron Microscopy Informed Molecular Dynamics Simulations to Investigate the Disulfide Hydrogel Self-Assembly,” ChemPhysChem 2025, 26(11), e202401085. Used with permission of John Wiley & Sons - Books; permission conveyed through Copyright Clearance Center, Inc., License No. 1780538-1. The coauthors listed in this publication are Yuanming Song, Zhaoxu Li, Justin T. Mulvey, J. Alfredo Freites, Joseph P. Patterson, and Douglas J. Tobias. Yuanming Song and

Zhaoxu Li contributed equally; Zhaoxu Li performed the TEM simulations, and collaborators performed the molecular dynamics simulations.

During dissertation preparation, OpenAI ChatGPT (GPT-5.6 Pro) was used to improve the readability and consistency of the author's text and to assist with document formatting.

VITA

Zhaoxu Li

EDUCATION

Doctor of Philosophy in Materials Science and Engineering, University of California, Irvine, 2026

Master of Science in Materials Science and Engineering, University of California, Irvine, 2023

Bachelor of Medicine, Chongqing Medical University, 2019

FIELDS OF STUDY

Liquid-phase and cryogenic transmission electron microscopy; molecular dynamics and multislice TEM simulation; deep learning for microscopy; electron tomography.

SELECTED PUBLICATIONS AND MANUSCRIPTS

Song, Y.; Li, Z.; Mulvey, J. T.; Freites, J. A.; Patterson, J. P.; Tobias, D. J. Cryogenic Electron Microscopy Informed Molecular Dynamics Simulations to Investigate the Disulfide Hydrogel Self-Assembly. ChemPhysChem 2025, 26, e202401085.

Zeinalvand, L.; Barpuzary, D.; Ueon, T. H.; et al. Nanoscale Interfacial Organization Governs Maturation and Collapse in Passive versus Active Condensates. Journal of the American Chemical Society 2026, 148, 17967–17976.

Li, Z.; Sagar, D.; Abdelkarim, S.; et al. TSGNet: A Deep Convolutional Framework for Generating Missing Tilts and Restoring High-Quality Electron Tomography from Sparse Tilt Series. Manuscript submitted to Ultramicroscopy, 2026.

ABSTRACT OF THE DISSERTATION

Electron Microscopy of Dynamic Soft Matter and Nanoscale Materials: Integrating Liquid-Phase Transmission Electron Microscopy, Physics-Based Simulation, and Deep Learning

by

Zhaoxu Li

Doctor of Philosophy in Materials Science and Engineering

University of California, Irvine, 2026

Professor Joseph P. Patterson, Chair

Electron microscopy provides direct access to nanoscale structure, but quantitative interpretation is difficult when specimens are dynamic, weakly scattering, surrounded by liquid or vitreous ice, or sampled incompletely. This dissertation develops a traceable framework that combines liquid-phase transmission electron microscopy (LPTEM), cryogenic electron microscopy (cryo-EM), molecular dynamics, multislice image simulation, and deep learning. Across four studies, each conclusion is matched to directly measured image observables and to the reference information actually available.

First, electrochemical LPTEM was used to examine FFssFF-derived peptide coacervates through one three-cycle cyclic-voltammetry (CV) recording before electrochemical impedance spectroscopy (EIS) and two three-cycle recordings after EIS in one liquid-cell field. Initial domains were approximately 100–500 nm in projected diameter. Recording A showed no prominent resolved CV peak and minimal morphological change under the combined electrochemical and imaging conditions. The largest qualitative transition occurred across EIS: coacervate boundaries appeared sharper, while numerous smaller, lower-contrast aggregates appeared in the surrounding field. The impedance response had a predominantly capacitive appearance and lacked a resolved charge-transfer semicircle over the measured frequency range. During the subsequent CV recordings, small aggregates progressively diminished and the same four tracked domains lost projected area. Median endpoint area changes were +0.98%, −7.41%, and −20.67% in Recordings A, B, and C, respectively. These observations establish a time-

associated sequence of projected morphological change, but they do not isolate bias from electron exposure, identify a unique chemical mechanism, or equate projected contrast with mass or composition.

Second, cryo-EM measurements were used to constrain atomistic CSH–CSSC fiber models, which were converted into multislice TEM images for comparison with experiment in a common image domain. The experimentally measured individual-fiber diameter was 5.42 ± 1.55 nm, whereas the defocus-weighted simulated diameter was 4.78 ± 0.66 nm. Systematic simulations showed that defocus and water thickness altered apparent fiber width and contrast. Because the experimental diameter also informed the radial restraint, this agreement is interpreted as an image-domain consistency test rather than independent structural validation. The forward-modeling strategy was then extended to generate 500 physics-informed synthetic TEM image–reference pairs for low-contrast nanofiber segmentation. A U-Net trained exclusively on these data achieved Dice 0.9547, intersection over union 0.9134, and centerline Dice 0.9969 on the 78-image synthetic test set. In unlabeled experimental TEM images, its predictions followed many visually recognizable fiber-like trajectories, providing qualitative correspondence but not a quantitative estimate of experimental accuracy.

Finally, TSGNet was evaluated for generating intermediate projections in sparse simulated scanning transmission electron microscopy tilt series. Relative to linear interpolation, TSGNet improved the reported image-level mean-squared error, peak signal-to-noise ratio, and structural similarity. Insertion of the generated projections reduced both surface Chamfer distance and volumetric root-mean-square error relative to matched sparse baselines at every reported nominal tilt increment from 2° to 20° under both simulated configurations. These comparisons were made against a dense-series reconstruction reference and do not establish

experimental dose reduction or acquisition-time savings. Collectively, the studies show how synchronized in situ imaging, physically grounded forward simulation, and learning-based inference can extend electron microscopy toward traceable quantitative analysis while preserving the distinction between direct observations, model-based references, synthetic-domain evaluation, and experimental validation.

INTRODUCTION

Electron Microscopy of Dynamic Soft Matter and Nanoscale Materials

Electron microscopy is uniquely suited to interrogating materials whose decisive structural features occur at nanometer and sub-nanometer length scales. Conventional high-vacuum imaging, however, provides only a limited view of systems that assemble, reorganize, dissolve, or react in condensed environments. Liquid-phase transmission electron microscopy (LPTEM) extends direct imaging to specimens retained in liquid [1–3], whereas cryogenic electron microscopy (cryo-EM) preserves hydrated structures in vitrified water for structural characterization [4]. Together, these approaches create access to soft and dynamic materials that are difficult to study by dry-state microscopy alone.

The same experimental configurations that preserve a condensed environment also complicate image interpretation. In LPTEM, electron scattering from the liquid layer and silicon nitride windows reduces contrast, and local liquid thickness, defocus, drift, contamination, and detector noise can alter the apparent morphology of weakly scattering objects [1–3]. Electron irradiation of aqueous media further generates reactive radiolysis products that can modify polymers and other beam-sensitive materials [5]. In cryo-EM, ice thickness, contrast-transfer conditions, and limited signal similarly influence the width and visibility of nanoscale fibers [4]. An observed boundary is therefore not a direct copy of the material boundary; it is the output of specimen structure, electron scattering, instrument transfer, sampling, and noise.

These limitations are especially consequential for nonequilibrium soft materials. Biomolecular condensates and peptide-based coacervates can respond to chemical reactions, electric fields, and interfacial conditions by changing size, shape, and internal organization [6, 7]. A time-resolved image sequence can reveal when such changes occur, but temporal coincidence alone does not establish a unique electrochemical or molecular cause. Robust interpretation

therefore requires synchronization with the applied stimulus, tracking of the same objects, explicit operational definitions for image measurements, and a clear distinction between projected observables and physical quantities such as volume, mass, or composition.

The first experimental study in this dissertation applies this principle to FFssFF-derived peptide coacervates. Electrochemical LPTEM provides the spatial and temporal resolution needed to follow individual nanoscale domains during cyclic voltammetry, while the pre- and post-EIS sequence supports temporal comparison within the same field. The analysis emphasizes projected area and within-recording contrast because these quantities are directly defined in the recorded images. The study is based on four tracked domains in one liquid-cell field; repeated frames and cycles improve temporal sampling but do not constitute independent experimental replication. Interpretation is further bounded by the absence of calibrated dose, three-dimensional volume, and independent molecular-composition measurements.

From Molecular Structure to Image Formation

Microscopy images and atomistic simulations describe the same material at different representational levels. Molecular dynamics (MD) provides coordinates, interactions, and trajectories, whereas TEM records an intensity field after electron propagation through the specimen and transfer by the microscope. Directly comparing an atomistic rendering with a TEM micrograph can therefore be misleading. A physically meaningful comparison requires a forward model that converts atomic coordinates into a simulated image under specified accelerating voltage, specimen thickness, aberrations, defocus, detector response, and noise conditions.

The CSH-CSSC disulfide hydrogel system provides a useful test case. Cryo-EM established an approximately 5.4 nm individual-fiber diameter and revealed transient stacked structures during chemically driven assembly and disassembly [8]. Those measurements can

constrain the radial dimensions of atomistic fiber models, but satisfying a geometric restraint does not by itself show that a model reproduces experimental image features. In the second study of this dissertation, snapshots from cryo-EM-informed MD simulations are passed through multislice TEM calculations using abTEM [9, 10]. The simulated images can then be evaluated with the same transverse-intensity and full-width-at-half-maximum measurements used for the experimental micrographs.

This image-domain comparison serves two purposes. First, it tests whether selected atomistic structures reproduce the dark fiber core, underfocus rims, and apparent diameter observed experimentally. Second, it isolates the effects of imaging variables that cannot be separated easily in an experimental dataset. Varying defocus and water thickness demonstrates how the same structure can acquire different apparent widths and signal-to-noise ratios. Agreement is consequently interpreted as consistency between model and image under the tested conditions, not as proof of a unique molecular packing arrangement, particularly because the experimental diameter also informed the simulation restraint.

Forward image formation also creates a bridge from physical modeling to machine learning. Once a transformed atomic scene is available, one pathway can generate a TEM image from electron scattering while a separate geometric pathway generates a foreground reference from known fiber coordinates. The paired outputs share the same underlying structure but do not define the target by thresholding the simulated intensity. This separation is central to the synthetic-data strategy developed later in the dissertation.

Quantitative Image Analysis with Limited Ground Truth

Low-contrast nanofibers present a difficult segmentation problem. Their widths can approach the characteristic scale of image noise, weak segments may appear discontinuous, and

crossings, parallel overlaps, and bundles create ambiguous local boundaries. Pixel-level experimental annotation is therefore both labor intensive and intrinsically uncertain. Classical global or local thresholding can provide useful baselines, but a threshold selected from intensity alone may respond strongly to background gradients, thickness changes, or noise rather than to fiber morphology.

Supervised convolutional networks can learn contextual representations that are more robust than a single intensity rule, but their performance depends on the fidelity and scope of the training labels. U-Net architectures are well suited to dense semantic segmentation because encoder features provide context while decoder and skip pathways recover spatial detail [11]. For elongated structures, overlap metrics such as Dice and intersection over union should be supplemented with centerline-sensitive measures, because a prediction may preserve a fiber trajectory even when its boundary differs locally [12]. These metrics describe agreement with a defined reference; they do not remove uncertainty in how that reference was constructed.

The third study addresses the annotation problem with physics-informed synthetic data. Short hydrated CSH-CSSC fiber modules are assembled into longer, curved, and multi-fiber scenes, embedded in an aqueous volume, and converted into noisy TEM images by multislice simulation. Coordinate-derived masks provide reproducible foreground references for isolated fibers, crossings, overlaps, and bundles. A fixed train-validation-test split supports model development and retrospective synthetic-domain evaluation, while Otsu- and Sauvola-based pipelines provide frozen comparison methods.

Synthetic accuracy and experimental transfer are treated as separate questions. Quantitative metrics on held-out synthetic images measure agreement with the generator-defined target distribution. Unlabeled real-TEM images cannot provide the same accuracy estimate, so

their use is restricted to qualitative comparison, while tile-origin tests assess inference consistency rather than correctness. This distinction avoids using visual plausibility as a substitute for an experimental reference and identifies the additional work required for external validation, including annotated images, broader acquisition conditions, source-grouped synthetic splits, and scale-robust inference.

Recovering Three-Dimensional Information from Sparse Measurements

Electron tomography reconstructs a three-dimensional object from projections recorded across a range of specimen orientations [13]. Reconstruction quality depends on both image content and angular sampling. Dense sampling generally provides more complete information, but it also increases acquisition time and cumulative electron exposure. These costs are particularly restrictive for beam-sensitive materials and for in situ experiments in which the specimen may evolve, rotate, or leave the field of view during acquisition [14].

Sparse tilt series reduce the number of acquired projections but increase angular undersampling and reconstruction artifacts. Learning-based methods have been developed to restore missing information in tomographic data, including missing-wedge recovery and isotropic reconstruction [15, 16]. A generated projection, however, should not be evaluated only as a visually plausible image. Its value for tomography depends on whether insertion into the tilt series improves the reconstructed volume relative to the corresponding sparse acquisition.

The fourth study evaluates this downstream requirement with TSGNet, an encoder–decoder network that predicts an intermediate projection from neighboring simulated scanning transmission electron microscopy images. Image-level mean-squared error, peak signal-to-noise ratio, and structural similarity quantify agreement with held-out intermediate projections. The generated projections are then inserted into sparse tilt series, and the resulting reconstructions are

compared with a dense-series reconstruction reference using surface Chamfer distance and volumetric root-mean-square error.

This design separates two related but non-equivalent questions: whether a network predicts an image close to the missing projection and whether that prediction improves three-dimensional reconstruction relative to the corresponding sparse series. The simulated dataset provides exact intermediate projections for controlled testing across angular increments and imaging configurations. The dense-series reconstruction remains an algorithm-dependent reference, and a reconstruction-level linear-interpolation augmentation control was not included. Simulation-only performance therefore cannot establish general utility for experimental, dose-limited tomography.

Dissertation Scope and Organization

The central objective of this dissertation is to develop traceable links among nanoscale experiment, physical image formation, and computational inference. The four studies span different specimens and measurement problems, but they share a common methodological principle: the reported conclusion should be matched to the observable and to the reference information actually available. Experimental image changes are reported as projected morphology and contrast; atomistic models are tested in the image domain; segmentation is quantified against independently generated synthetic references; and generated tilt images are evaluated by their effect on three-dimensional reconstruction.

Chapter 1 examines electrochemically associated evolution of FFssFF-derived peptide coacervates by LPTEM. The same four domains are followed through one pre-EIS and two post-EIS cyclic-voltammetry recordings in a single field. Projected-area trajectories, an operational 5% area-loss threshold, fixed-reference contrast, and internal contrast redistribution are used to

characterize progressive image changes without assigning a unique chemical mechanism or treating repeated frames as independent replicates.

Chapter 2 connects cryo-EM measurements, collaborator-generated CSH–CSSC atomistic models, and multislice TEM simulation. The analysis measures the experimental fiber-diameter distribution, simulates selected restrained structures over a range of defocus values and water thicknesses, and compares simulated and experimental apparent widths as an image-domain consistency test. Chapter 3 extends the same forward-modeling logic to synthetic training data for low-contrast nanofiber segmentation. A U-Net is evaluated within a fixed synthetic distribution, compared with thresholding baselines, and applied without fine-tuning to unlabeled real TEM images for qualitative assessment.

Chapter 4 addresses sparse simulated electron tomography. TSGNet generates intermediate projections from neighboring tilts, and its performance is assessed at both the projection and reconstruction levels relative to a dense-series reconstruction reference. Chapter 5 integrates the four studies, distinguishes their evidence levels, and identifies the experimental and computational validation needed to extend the reported findings. Together, the chapters show how electron microscopy can be strengthened by explicit imaging physics and learning-based computation while keeping controls, independent replication, reference construction, and domain of validation visible.

References

1. de Jonge, N.; Ross, F. M. Electron Microscopy of Specimens in Liquid. Nature Nanotechnology 2011, 6, 695–704. https://doi.org/10.1038/nnano.2011.161.
2. Ross, F. M. Opportunities and Challenges in Liquid Cell Electron Microscopy. Science 2015, 350, aaa9886. https://doi.org/10.1126/science.aaa9886.
3. Sung, J.; Bae, Y.; Park, H.; Kang, S.; Choi, B. K.; Kim, J.; Park, J. Liquid-Phase Transmission Electron Microscopy for Reliable In Situ Imaging of Nanomaterials. Annual Review of Chemical and Biomolecular Engineering 2022, 13, 167–191. https://doi.org/10.1146/annurev-chembioeng-092120-034534.

4. Dubochet, J.; Adrian, M.; Chang, J.-J.; Homo, J.-C.; Lepault, J.; McDowall, A. W.; Schultz, P. Cryo-Electron Microscopy of Vitrified Specimens. Quarterly Reviews of Biophysics 1988, 21(2), 129–228. https://doi.org/10.1017/S0033583500004297.

5. Wu, H.; Sun, H.; Oerlemans, R. A. J. F.; et al. Understanding, Mimicking, and Mitigating Radiolytic Damage to Polymers in Liquid Phase Transmission Electron Microscopy. Advanced Materials 2024, 36, e2402987. https://doi.org/10.1002/adma.202402987.

6. Abbas, M.; Lipinski, W. P.; Nakashima, K. K.; Huck, W. T. S.; Spruijt, E. A Short Peptide Synthon for Liquid-Liquid Phase Separation. Nature Chemistry 2021, 13, 1046–1054. https://doi.org/10.1038/s41557-021-00788-x.

7. Gibson, W.; Mulvey, J. T.; Das, S.; et al. Observing the Dynamics of an Electrochemically Driven Active Material with Liquid Electron Microscopy. ACS Nano 2024, 18, 11898–11909. https://doi.org/10.1021/acsnano.4c01524.

8. Hurst, P. J.; Mulvey, J. T.; Bone, R. A.; et al. CryoEM Reveals the Complex Self-Assembly of a Chemically Driven Disulfide Hydrogel. Chemical Science 2024, 15, 1106–1116. https://doi.org/10.1039/D3SC05790A.

9. Song, Y.; Li, Z.; Mulvey, J. T.; Freites, J. A.; Patterson, J. P.; Tobias, D. J. Cryogenic Electron Microscopy Informed Molecular Dynamics Simulations to Investigate the Disulfide Hydrogel Self-Assembly. ChemPhysChem 2025, 26, e202401085. https://doi.org/10.1002/cphc.202401085.

10. Madsen, J.; Susi, T. The abTEM Code: Transmission Electron Microscopy from First Principles. Open Research Europe 2021, 1, 24. https://doi.org/10.12688/openreseurope.13015.2.

11. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In Medical Image Computing and Computer-Assisted Intervention - MICCAI 2015; Springer, 2015; pp 234–241. https://doi.org/10.1007/978-3-319-24574-4_28.

12. Shit, S.; Paetzold, J. C.; Sekuboyina, A.; et al. clDice - A Novel Topology-Preserving Loss Function for Tubular Structure Segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition; 2021; pp 16560–16569.

13. Midgley, P. A.; Dunin-Borkowski, R. E. Electron Tomography and Holography in Materials Science. Nature Materials 2009, 8, 271–280. https://doi.org/10.1038/nmat2406.

14. Koneti, S.; Roiban, L.; Dalmas, F.; et al. Fast Electron Tomography: Applications to Beam Sensitive Samples and In Situ TEM or Operando Environmental TEM Studies. Materials Characterization 2019, 151, 480–495. https://doi.org/10.1016/j.matchar.2019.02.009.

15. Ding, G.; Liu, Y.; Zhang, R.; Xin, H. L. A Joint Deep Learning Model to Recover Information and Reduce Artifacts in Missing-Wedge Sinograms for Electron Tomography and Beyond. Scientific Reports 2019, 9, 12803. https://doi.org/10.1038/s41598-019-49267-x.

16. Liu, Y.-T.; Zhang, H.; Wang, H.; et al. Isotropic Reconstruction for Electron Tomography with Deep Learning. Nature Communications 2022, 13, 6482. https://doi.org/10.1038/s41467-022-33957-8.

CHAPTER 1: Electrochemical Liquid-Phase TEM of FFssFF-Derived Peptide Coacervates

Introduction

Biomolecular condensates create compositionally enriched, membrane-free compartments whose structures and dynamics can respond to chemical and physical inputs. Short peptide coacervates provide a compact platform for connecting molecular interactions with mesoscale assembly behavior [8]. The disulfide-linked synthon FFssFF forms water-rich coacervates, and reduction of one FFssFF molecule produces two FFsH molecules [1]. Recent work on related passive and actively driven peptide condensates showed that nanoscale interfacial organization can govern their maturation and collapse [2].

Electric-field excitation can drive coacervate protocells away from equilibrium and alter their size, shape, and internal organization [7]. Electrochemical liquid-phase transmission electron microscopy (LPTEM) can record stimulus-associated morphology at nanometer length scales while a waveform is applied [3, 4]. Interpretation nevertheless requires caution because liquid thickness, window confinement, drift, phase contrast, and aqueous radiolysis can change apparent morphology; temporal association with an electrochemical step must therefore be distinguished from a calibrated causal or compositional assignment [5, 6].

This chapter follows the same four FFssFF-derived domains through one pre-EIS cyclic-voltammetry (CV) recording and two sequential post-EIS CV recordings in a single liquid-cell field. Projected area is used as the primary quantitative observable, and fixed-reference mean contrast and center–outer contrast provide secondary descriptions of image evolution. The study is descriptive: the four domains are repeated observations within one field, not four independent experimental replicates.

Methods

FFssFF was first dissolved in distilled water, for which the initial solution pH was approximately 5. Phosphate buffer at pH 7 was then added, and the sample was mixed gently by pipetting to form coacervates. The final buffer concentration was 200 mM, and the initial FFssFF concentration was 2.7 mM (2 mg/mL). Solid dithiothreitol (DTT) was added to a final concentration of 6 mM, the solution was mixed by gentle shaking, and the mixture was allowed to stand for approximately 2 h. DTT reduction cleaves the disulfide linkage so that one FFssFF molecule yields two FFsH molecules; the formulation is therefore reported in terms of the initial FFssFF concentration. Solid ferrocyanide, $[Fe(CN)_6]^{4-}$, was subsequently added to a final concentration of 50 mM and dissolved by gentle mixing.

Electrochemical LPTEM was performed using a DENSsolutions Stream liquid-cell holder equipped with a Stream Nano-Cell Liquid Biasing chip set containing Pt electrodes in a parallel-window configuration. The cell comprised two microfabricated chips, each containing an approximately 20 nm-thick silicon nitride (SiN) electron-transparent window. The chips were sealed with an ethylene propylene diene monomer (EPDM) O-ring, and the opposing windows were aligned to provide an unobstructed electron-beam path. The working, counter, and pseudo-reference electrodes were all Pt and were connected through the holder to a PalmSens4 potentiostat. All potentials are reported relative to the on-chip Pt pseudo-reference electrode.

Liquid was delivered using the DENSsolutions Liquid Supplying System. Polyether ether ketone (PEEK) tubing connected the supplying system to the holder valves, and internal silica tubing connected the holder valves to the liquid-biasing chip. The assembled holder was vacuum leak-tested before insertion into the microscope. Liquid flow was stopped before all still-image, EIS, and video acquisitions and was not restarted between the sequential electrochemical stages.

LPTEM experiments were conducted using a JEOL F2800 transmission electron microscope at the Irvine Materials Research Institute, University of California, Irvine. The microscope was operated at 200 kV, and images were acquired with a Gatan OneView camera using Gatan Microscopy Suite. A CL1 condenser aperture was used. The illumination intensity was set to the lowest practical setting (CL3 readout: FFFF) and was held constant throughout the electrochemical sequence. Before liquid was introduced, the aligned SiN windows were imaged under dry-cell conditions to verify window integrity. Liquid-phase images were recorded under slight underfocus to enhance phase contrast.

Still images were recorded at 4096 × 4096 pixels with 1 × 1 binning and a 1.0 s exposure. Image metadata gave a calibrated spatial sampling of 0.7626 nm pixel^{-1}. Video sequences were acquired at 1024 × 1024 pixels with 4 × 4 binning, 20 frames s^{-1}, and an exposure time of 0.05 s per frame, corresponding to 3.0504 nm pixel^{-1}. Calibrated beam current, illuminated area, electron flux, and cumulative dose were not recovered; no numerical electron dose is therefore reported.

The EIS measurement was acquired between Recordings A and B. It was performed at the open-circuit potential with a DC offset of 0 V relative to OCP and a 10 mV AC perturbation. The measurement contained 61 frequencies from 100 kHz to 1 Hz, corresponding to 12 points per decade. All measured points were retained, and no equivalent-circuit fit was applied because EIS served primarily as an electrochemical measurement within the imaging sequence.

Cyclic voltammetry was performed while TEM video was recorded under stopped-flow conditions. The PalmSens4 program used a 10 s equilibration period and began at 0.0 V. The first and second vertex potentials were −1.0 and +0.2 V, respectively, with a potential step of 1

mV, a scan rate of 50 mV s^{-1}, and three cycles per recording. The selected current range was 1 μA. Cell potential, working-electrode potential, and current were recorded.

Three video–electrochemistry recordings were acquired sequentially, with EIS performed between the first and second recordings (Figure 1.1). Recording A contained the CV block acquired before EIS, Recording B contained the first CV block acquired after EIS, and Recording C contained the second CV block acquired after EIS. Each recording contained three CV cycles. Video acquisition began before each CV block, and the measured delays from video start to CV start were 24, 36, and 24 s for Recordings A, B, and C, respectively.

Representative still images were acquired between electrochemical stages to document the field of view. Figure 1.6 compares the images obtained immediately before and after EIS. Because the TEM and potentiostat computers were not synchronized to a common clock, electrochemical–video alignment used the recorded relative offsets rather than absolute timestamps.

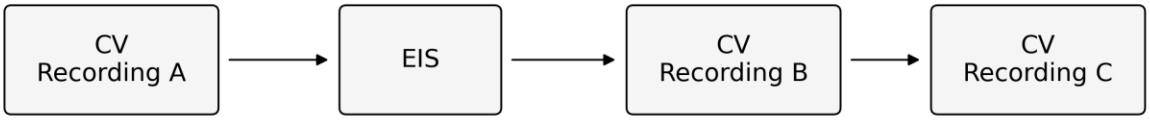

Figure 1.1. Electrochemical sequence used during LPTEM. A three-cycle CV block was recorded before EIS (Recording A), followed by EIS and two further three-cycle CV blocks (Recordings B and C).

Each analysis image was formed by averaging 20 source frames. Consecutive analysis windows advanced by 10 source frames, producing a 0.5 s analysis interval. Video time was aligned to the corresponding CV start using the recording-specific relative delay.

Four domains, designated O1–O4, were identified as the same physical objects across Recordings A–C and were followed within one liquid-cell field. The labels are ordered from the

smallest baseline projected area (O1) to the largest (O4). O4 exhibited a connected or bilobed outline in the post-EIS recordings and was treated as one object for the present analysis.

Projected area was defined as the number of pixels within a cleaned dynamic mask and therefore represents a two-dimensional image projection rather than volume or mass. Baseline values were the median of 21 pre-CV analysis windows, and endpoint values were the median of the final 11 windows. Area in square nanometers was calculated as $A = A_{px}p^2$, where A_{px} is the mask area in pixels and $p = 3.0504$ nm pixel^{-1}. Equivalent radius was calculated as $r_{eq} = \sqrt{(A_{px}/\pi)}p$. A radius-1 cross erosion and dilation provided a ±1-pixel boundary-sensitivity audit. The primary 5% area-loss latency was the first unsmoothed analysis-window center at or below 95% of baseline; observations that did not cross the threshold were treated as right-censored.

Raw image contrast was analyzed on the native stored gain-normalized detector scale without per-frame IQR normalization. For each O-label, a fixed reference, BA,pre, was defined as the median intensity of the 4–14-pixel local annulus across the 21 Recording A pre-CV windows and then held constant across Recordings A–C. Mean contrast was $M(t) = BA,pre - Imean(t)$, so positive values indicate that the domain was darker than the fixed reference. Because the common image background shifted between recordings, interpretation emphasized the within-recording endpoint-minus-pre-CV change rather than absolute contrast levels across Recordings A–C.

Spatial contrast redistribution was summarized by $DCO(t) = Iouter(t) - Icenter(t)$, where positive values indicate a darker center. The center and outer zones were defined by $r/R0 \leq 0.30$ and $0.65 \leq r/R0 \leq 0.90$, respectively. DCO was treated as an image-domain description of internal contrast distribution and not as a direct measurement of material density, concentration, or thickness.

Each triangular CV cycle was divided chronologically into three legs: approximately 0 to −1.0 V, −1.0 to +0.2 V, and +0.2 to 0 V. Potential dependence was assessed by comparing common-potential nodes across cycles and scan directions and by distinguishing repeatable within-cycle contrast from progressive drift. A response was considered potential-associated only when it was repeatable across cycles and scan directions rather than represented by a single coincident image event.

Results and Discussion

After liquid filling, approximately spherical, high-contrast coacervate domains were observed at multiple locations across the aligned window region (Figure 1.2). Their apparent projected diameters were approximately 100–500 nm, smaller than the micrometer-scale FFssFF droplets reported in benchtop optical microscopy [1]. The comparison is qualitative because solution composition, confinement, and imaging modality differed.

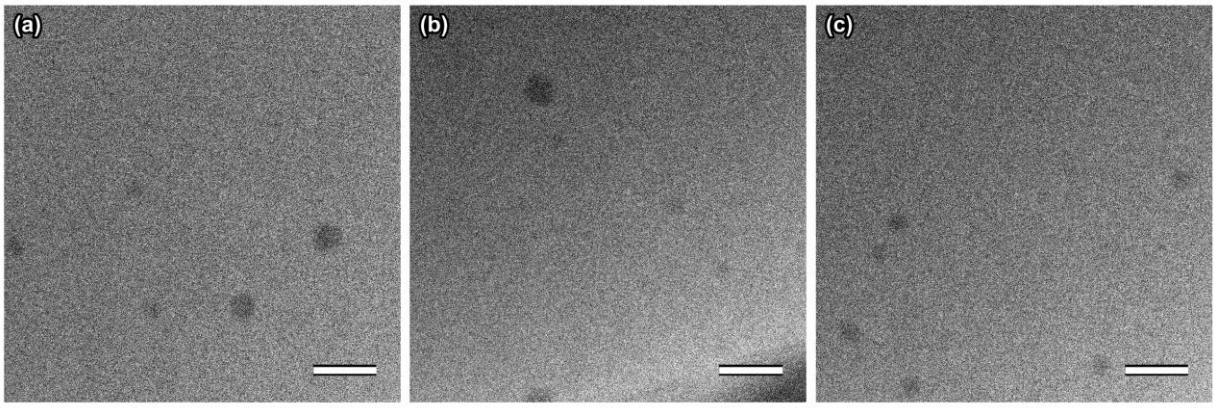

Figure 1.2. FFssFF-derived coacervate domains after liquid filling. Approximately spherical, high-contrast domains were observed at multiple locations across the aligned window area. The visually estimated diameter range was approximately 100–500 nm. Scale bars, 500 nm.

Representative time-averaged TEM frames from Recording A are shown in Figure 1.4. The panels sample the same field at matched potentials across the three CV cycles; each image is

a 20-frame average displayed with a common grayscale range to support direct morphological comparison.

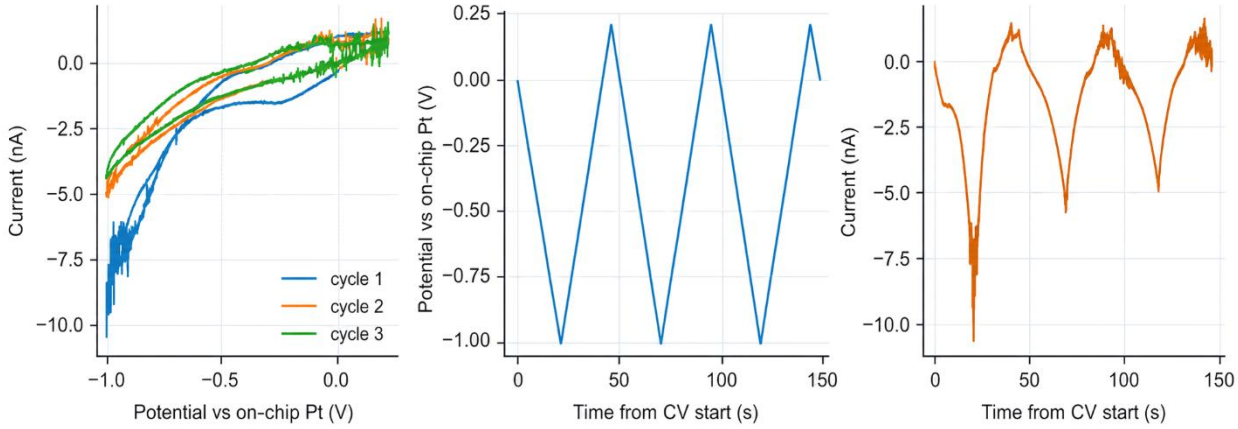

Figure 1.3. Electrochemical profile for Recording A. From left to right, the panels show the cyclic voltammograms (current versus potential) for cycles 1–3, the applied potential waveform as a function of time from the start of CV, and the corresponding current–time trace. Blue, orange, and green lines in the CV panel denote cycles 1, 2, and 3, respectively. Potentials are reported versus the on-chip Pt electrode.

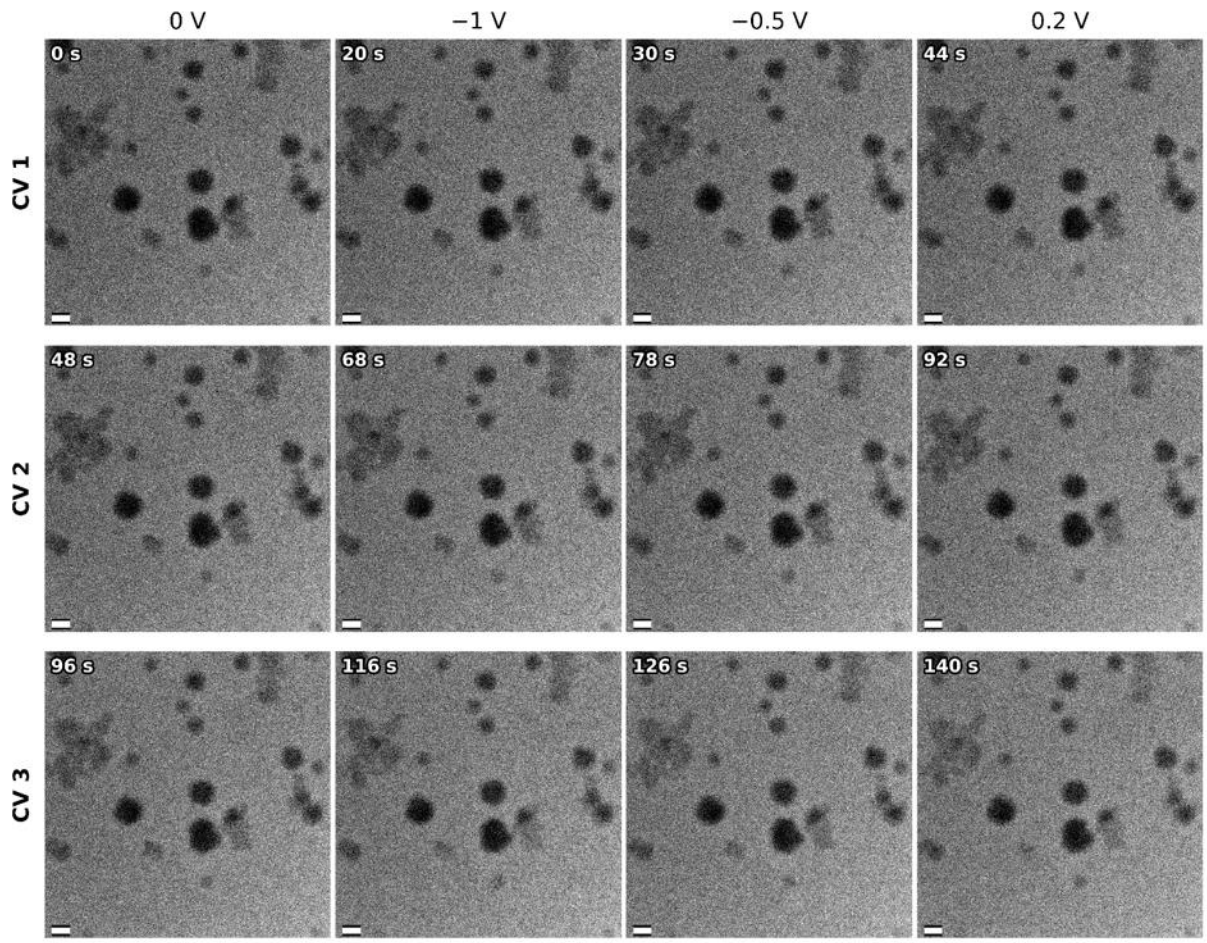

Figure 1.4. Representative liquid-phase TEM frames acquired during three consecutive cyclic voltammetry (CV) cycles in Recording A. Rows correspond to CV cycles 1–3, and columns show 0 V, −1 V, −0.5 V during the reverse sweep from −1 V toward 0 V, and 0.2 V. Timestamps indicate the calibrated elapsed time from the start of the first CV cycle. Each image is the average of 20 consecutive raw frames, with a common grayscale range applied across all panels. Scale bars, 200 nm.

Recording A was acquired during a three-cycle CV sequence. The CV curves are shown in Figure 1.3, and representative frames are shown in Figure 1.4. The CV curves did not show prominent resolved peaks, and morphological change was minimal. Recording A therefore showed no substantial resolved morphological response under the combined electrochemical and imaging conditions. Because separate bias-only and beam-only controls were not acquired, the contributions of applied potential and accumulated electron exposure cannot be distinguished.

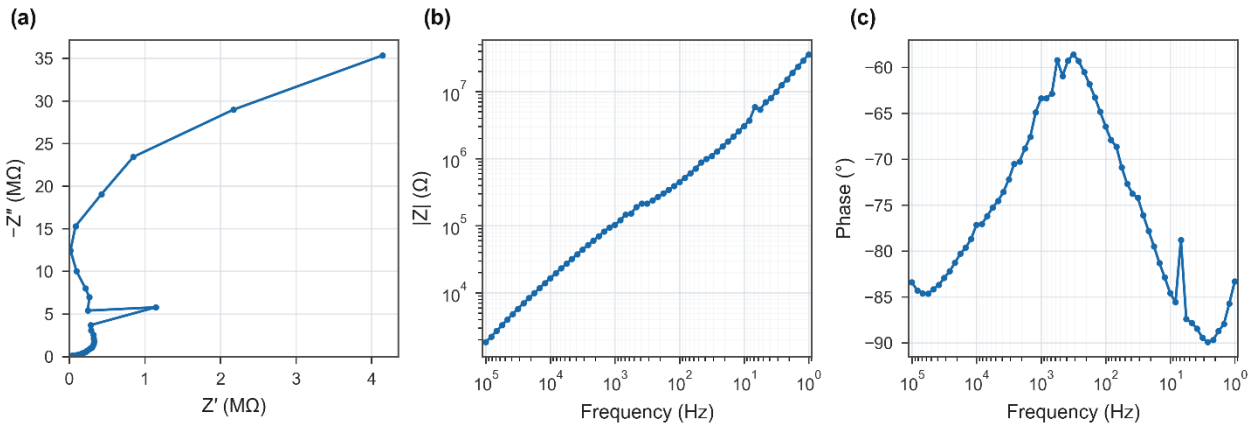

Figure 1.5. (a) Nyquist representation (Z′ versus −Z″), (b) impedance magnitude |Z|, and (c) phase angle. The sweep contains 61 logarithmically spaced frequencies from 100 kHz to 1 Hz; lines connect frequencies in acquisition order from high to low.

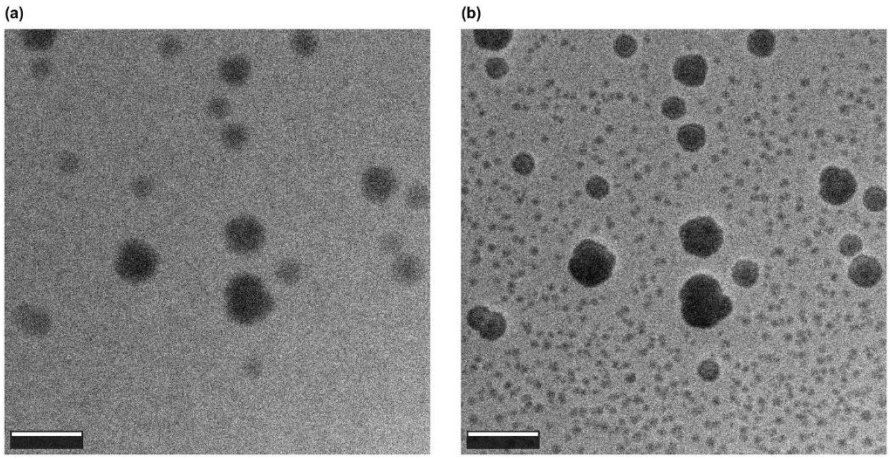

Figure 1.6. Representative still images acquired immediately (a) before and (b) after EIS. Scale bars, 500 nm.

The EIS data are shown in Figure 1.5. The impedance spectrum had a predominantly capacitive appearance and did not show a resolved charge-transfer semicircle over the measured frequency range; no equivalent-circuit parameters were extracted. Pronounced morphological differences were observed across the EIS interval (Figure 1.6). Before EIS, the coacervate boundaries appeared diffuse. After EIS, the boundaries appeared sharper, while numerous smaller, lower-contrast features appeared in the previously more uniform background. These changes could reflect electrochemical perturbation, elapsed beam exposure, redistribution of liquid or ions, or

changes in peptide organization. The present single-sequence experiment does not distinguish among these possibilities and does not establish a unique molecular mechanism.

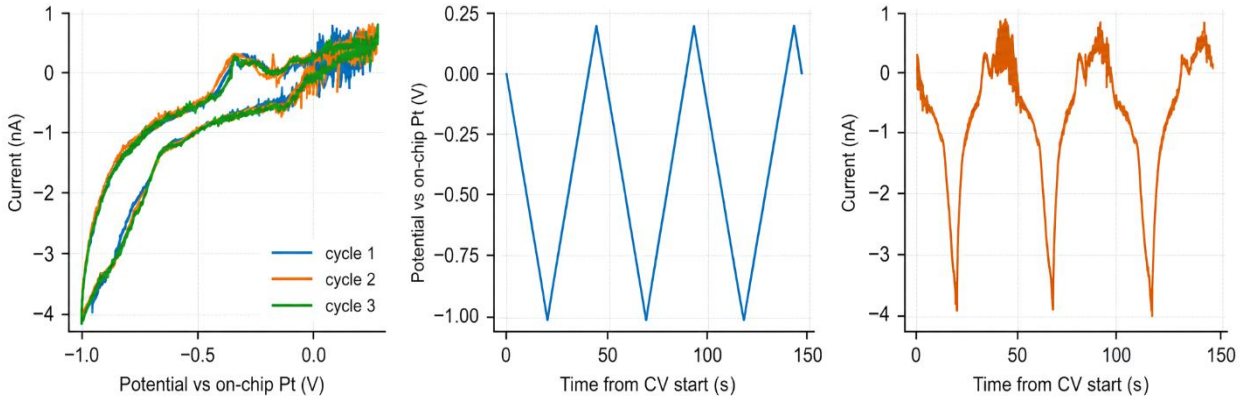

Figure 1.7. Electrochemical profile for Recording B. From left to right, the panels show the cyclic voltammograms (current versus potential) for cycles 1–3, the applied potential waveform as a function of time from the start of CV, and the corresponding current–time trace. Blue, orange, and green lines in the CV panel denote cycles 1, 2, and 3, respectively. Potentials are reported versus the on-chip Pt electrode.

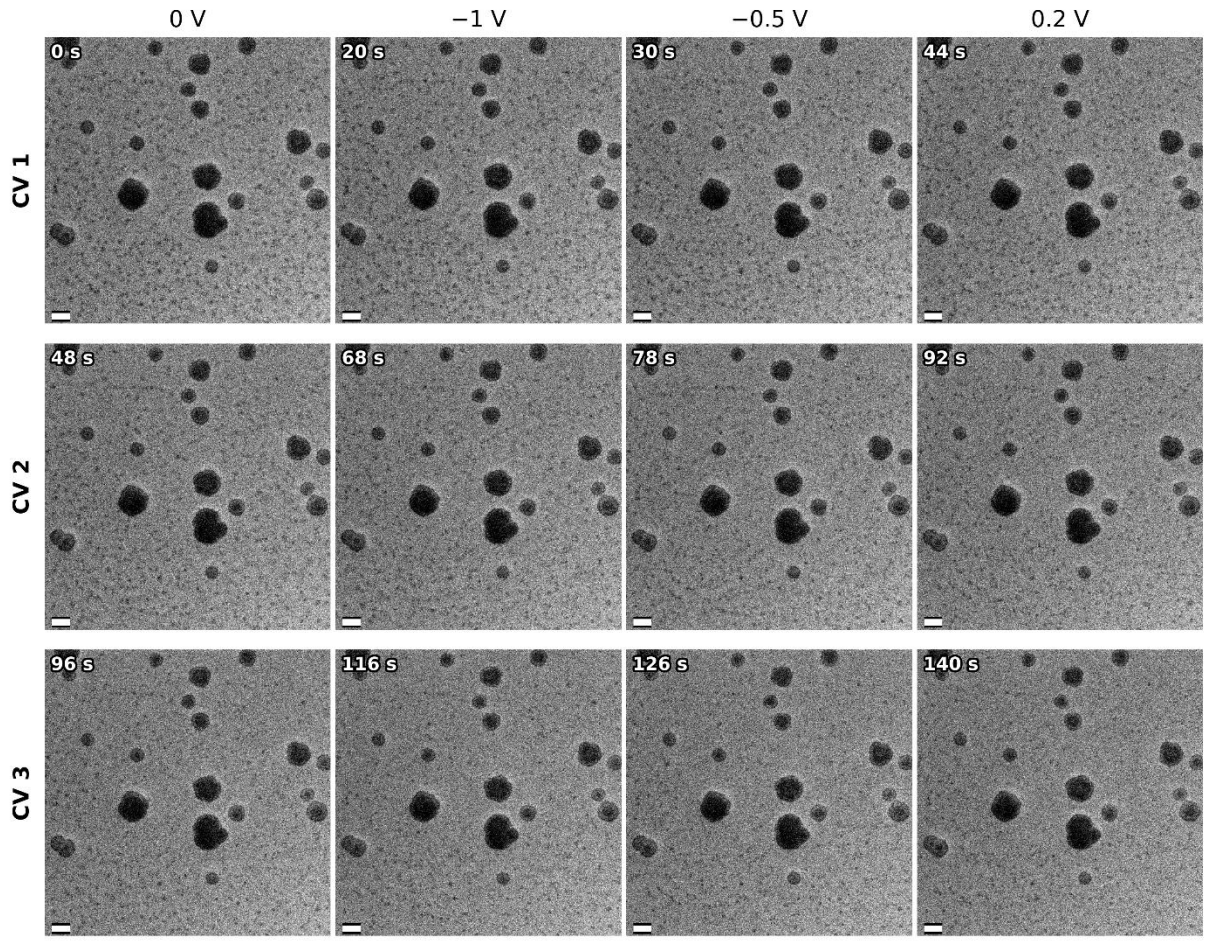

Figure 1.8. Representative liquid-phase TEM frames acquired during three consecutive cyclic voltammetry (CV) cycles in Recording B. Rows correspond to CV cycles 1–3, and columns show 0 V, −1 V, −0.5 V during the reverse sweep from −1 V toward 0 V, and 0.2 V. Timestamps indicate the calibrated elapsed time from the start of the first CV cycle. Each image is the average of 20 consecutive raw frames, with a common grayscale range applied across all panels. Scale bars, 200 nm.

Recording B was acquired during a three-cycle CV sequence. The CV curves are shown in Figure 1.7, and representative frames are shown in Figure 1.8. Cycle-to-cycle similarity was greater than in Recording A. The cathodic current increased continuously toward the −1.0 V switching potential, reaching approximately −4 to −4.5 nA, without forming a resolved reduction peak; this feature was therefore treated as a terminal cathodic-current envelope rather than assigned specifically to FFssFF reduction. A broad hysteretic response was also observed over

approximately −0.4 to +0.1 V. Its origin was not assigned because mediator redox chemistry, residual reductant, confined-cell transport, and electrode conditioning were not independently controlled. Morphological changes accompanied the electrochemical sequence: the tracked coacervates showed modest projected-area loss, while many smaller aggregates diminished or disappeared by the end of the third cycle. In this field, the smaller aggregates diminished earlier than the larger tracked domains, but this observation does not establish a general size-dependent reactivity law.

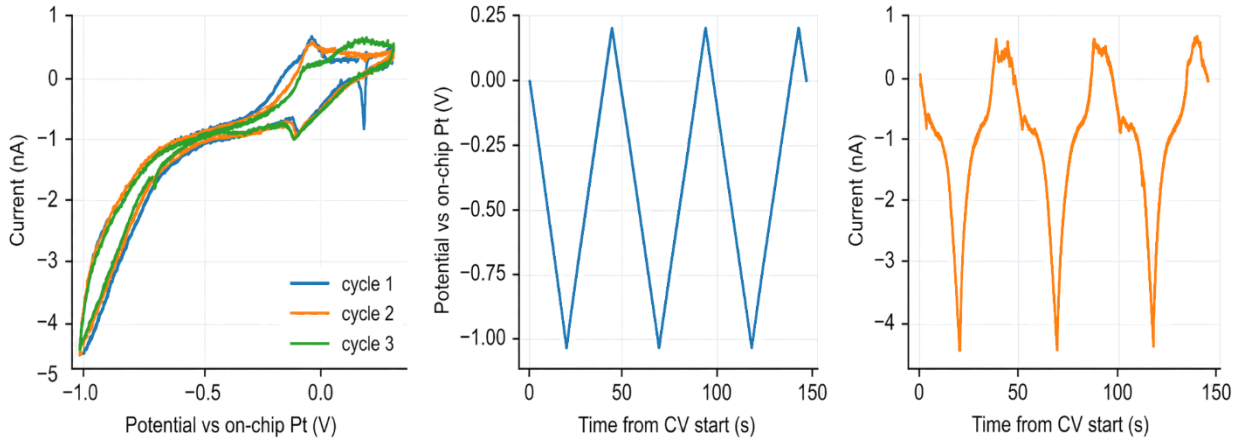

Figure 1.9. Electrochemical profile for Recording C. From left to right, the panels show the cyclic voltammograms (current versus potential) for cycles 1–3, the applied potential waveform as a function of time from the start of CV, and the corresponding current–time trace. Blue, orange, and green lines in the CV panel denote cycles 1, 2, and 3, respectively. Potentials are reported versus the on-chip Pt electrode.

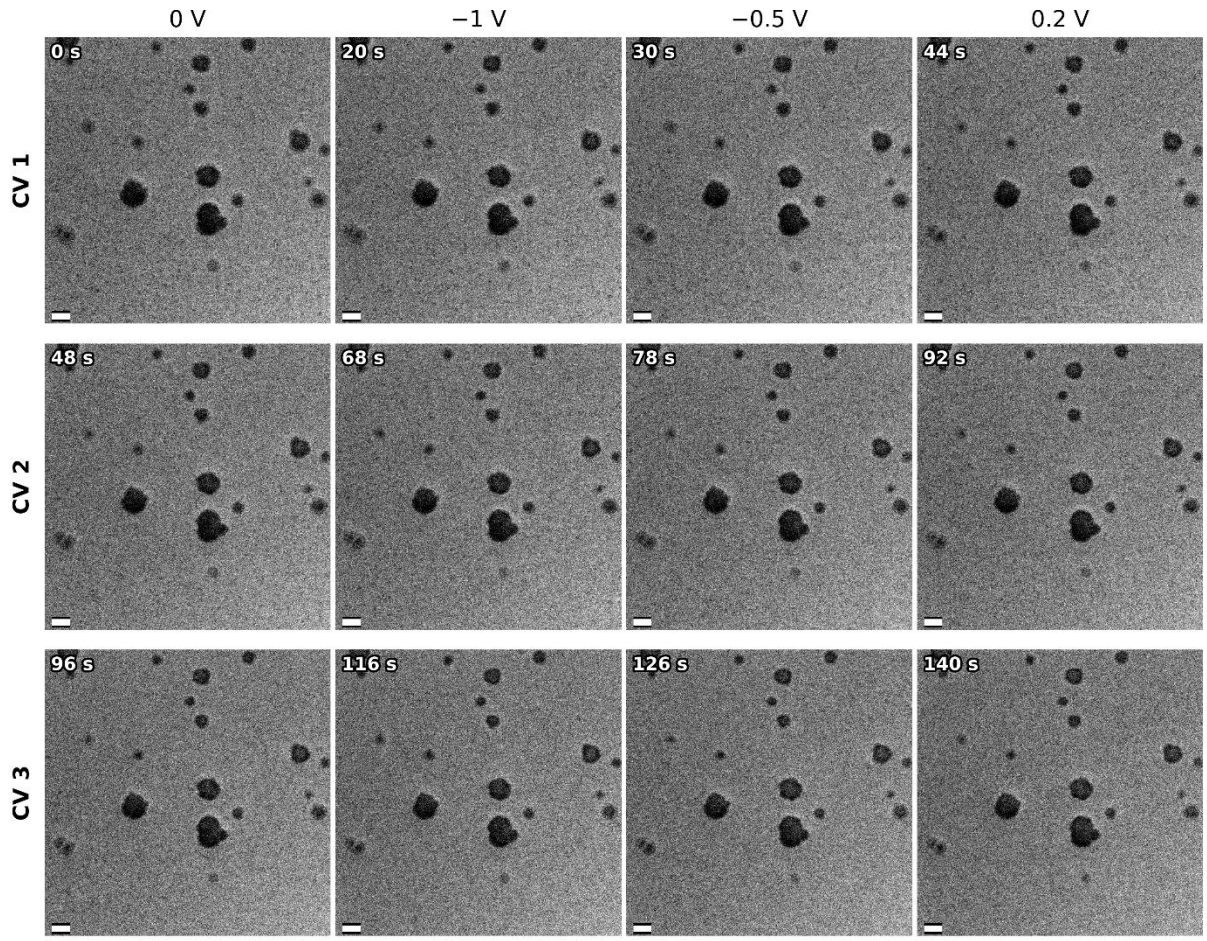

Figure 1.10. Representative liquid-phase TEM frames acquired during three consecutive cyclic voltammetry (CV) cycles in Recording C. Rows correspond to CV cycles 1–3, and columns show 0 V, −1 V, −0.5 V during the reverse sweep from −1 V toward 0 V, and 0.2 V. Timestamps indicate the calibrated elapsed time from the start of the first CV cycle. Each image is the average of 20 consecutive raw frames, with a common grayscale range applied across all panels. Scale bars, 200 nm.

Recording C was acquired during a three-cycle CV sequence. The CV curves are shown in Figure 1.9, and representative frames are shown in Figure 1.10. Recording C retained a similar cathodic-current magnitude and shape at negative potentials. Separation between the forward and reverse branches in the intermediate-potential region became more apparent, and the third cycle deviated more strongly from the preceding cycles, indicating continued evolution of the electrochemical response without identifying its source. Possible contributions include changes in solution composition, redistribution of mediator species within the confined cell, and continued

conditioning of the Pt electrode surface. Many smaller aggregates diminished or disappeared before the end of the second CV cycle, while the tracked coacervates showed greater projected-area loss than in Recording B. This size-associated pattern was observed in a single field and does not establish that smaller aggregates are intrinsically more reactive.

To quantify coacervate morphology, seeded local adaptive-intensity thresholding, connected-component selection, and morphological cleanup were used. Four well-segmented coacervates with different initial sizes were selected and labeled O1–O4 in ascending order of baseline projected area before Recording A. Pre-CV and endpoint segmentation overlays for O1–O4 across Recordings A–C are shown in Figure 1.11. This complete image set provides morphological context for the quantitative masks and avoids basing the comparison on a single representative domain. O4 retained a connected or bilobed outline in the post-EIS recordings and was analyzed as one object.

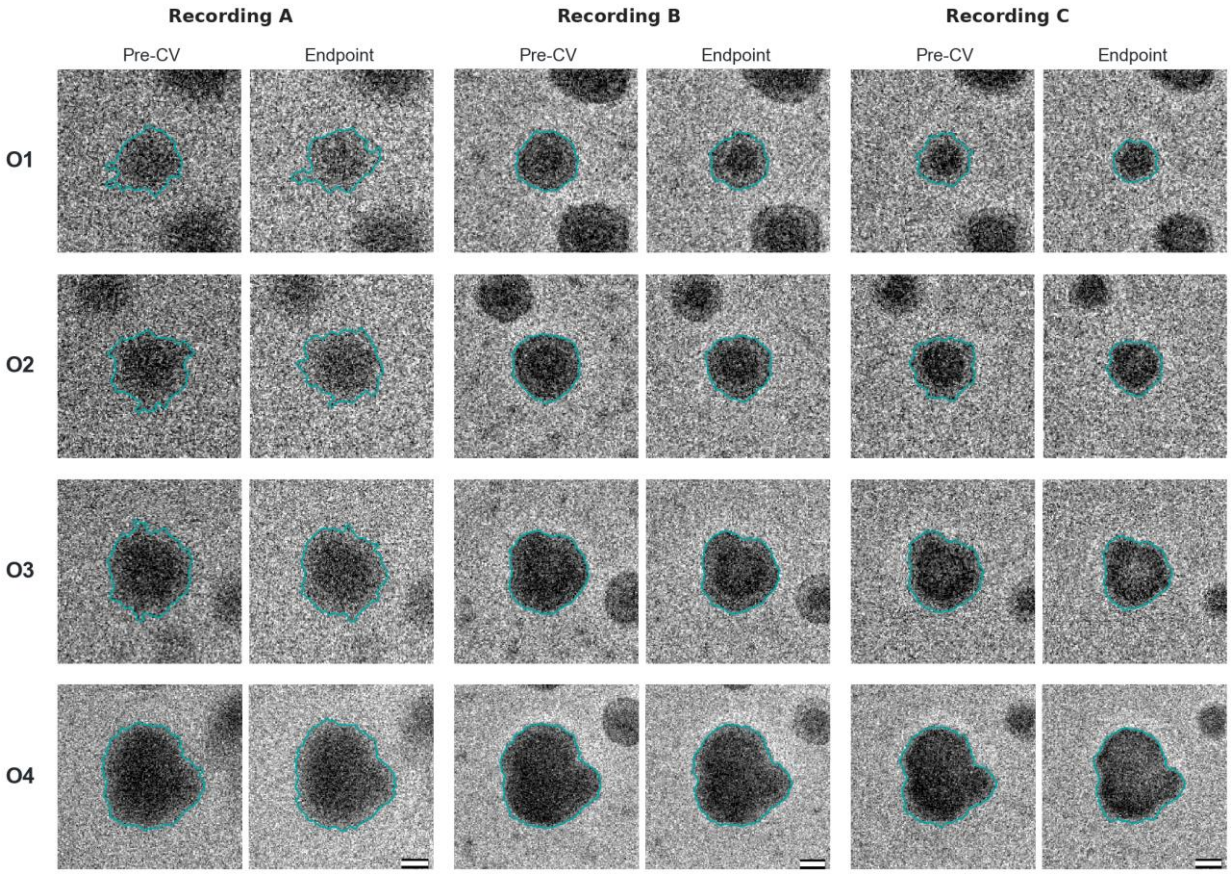

Figure 1.11. Pre-CV and endpoint morphology of O1–O4 across Recordings A–C. Rows show the same four domains, ordered from the smallest (O1) to largest (O4) baseline projected area. Each pair contains a pre-CV average and an endpoint average; cyan contours show the cleaned dynamic mask. A domain-specific grayscale range was held constant within each pre-CV/endpoint pair. Scale bars in the bottom row, 100 nm, apply to all panels.

Recording A fluctuated around its pre-CV projected-area baseline, whereas all four tracked domains lost area in Recordings B and C (Figure 1.12a). The median endpoint projected-area changes across O1–O4 were +0.98% in Recording A, −7.41% in Recording B, and −20.67% in Recording C (Table 1.1). The corresponding median absolute area changes were +512, −3,285, and −6,821 nm², and the median equivalent-radius changes were +0.67, −4.49, and −11.00 nm, respectively.

The individual endpoint area changes in Recording B were −13.85%, −8.88%, −5.94%, and −4.44% for O1–O4. In Recording C, the corresponding changes were −31.22%, −25.03%, −16.32%, and −7.78%. Within Recording C, the median area changes during the first, second,

and third cycles were −9.87%, −6.11%, and −2.25%, respectively, indicating that the largest net loss occurred during the first cycle and diminished over subsequent cycles.

Using the predefined unsmoothed 5% area-loss criterion, three of four Recording B domains crossed the operational threshold at 9.525, 58.025, and 101.025 s, while O4 remained right-censored at 143.025 s. All four Recording C domains crossed the threshold at 1.025, 6.025, 4.525, and 37.025 s for O1–O4, respectively. Under this operational definition, threshold crossing was earlier and more complete in Recording C. These times describe the selected image-analysis criterion and should not be interpreted as independently measured molecular-reaction onset times.

The monotonic O1–O4 ordering was stable under the ±1-pixel mask audit only for fractional area loss in Recordings B and C. It did not persist for absolute area loss, equivalent-radius loss, or radial recession rate (Figure 1.12b).

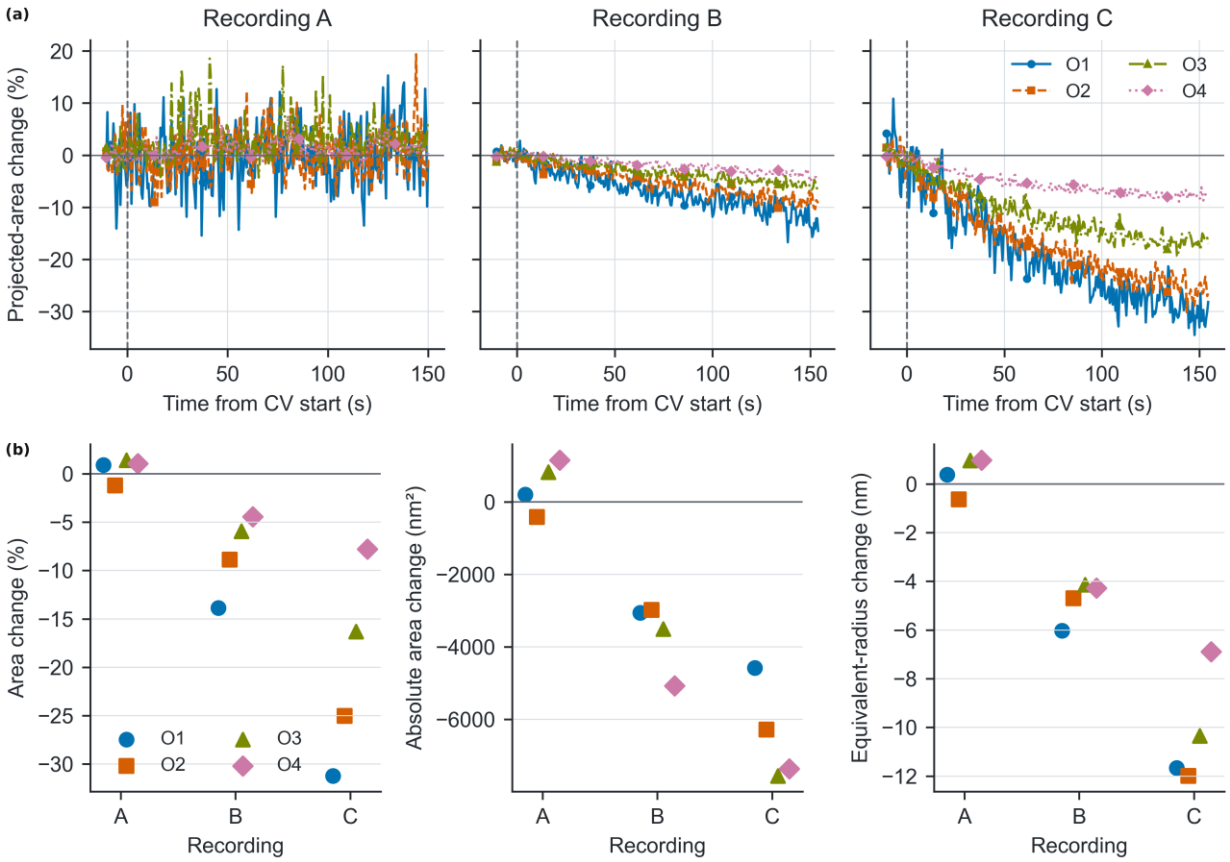

Figure 1.12. Projected-area trajectories and geometry-aware endpoint changes. (a) Projected-area change relative to each domain baseline for Recordings A–C. (b) Endpoint fractional area change, absolute area change, and equivalent-radius change for O1–O4.

Table 1.1. Median endpoint changes across O1–O4 in each recording.

Record.	ΔA (%)	ΔA (nm²)	Δr_eq (nm)	ΔM (a.u.)	ΔD_CO (a.u.)
A	+0.98	+512	+0.67	−19.13	−0.17
B	−7.41	−3,285	−4.49	−8.79	+9.49
C	−20.67	−6,821	−11.00	−5.15	−15.88

Note. Area, absolute-area, and equivalent-radius values compare the final 11 windows with the median of 21 pre-CV windows. Mean contrast is the within-recording endpoint-minus-pre-CV change using the fixed Recording A pre-CV local reference. $D_{CO} = I_{outer} - I_{center}$; a positive change indicates that the center became relatively darker. Values are descriptive medians across four repeated tracked domains in one liquid-cell field.

Fixed-reference raw contrast was interpreted by its within-recording change because the common image background shifted between recordings. In Recording B, mean contrast decreased for all four domains by −8.99, −8.09, −9.21, and −8.58 raw intensity units pixel^{-1} for

O1–O4, respectively (Figure 1.13a). In Recording C, O1 increased by +3.89, whereas O2–O4 changed by −0.99, −9.32, and −17.38 raw units pixel^{-1}. Recording C therefore combined strong area loss with heterogeneous internal contrast behavior rather than a uniform darkening or fading response.

The center-minus-outer difference showed a different spatial pattern (Figure 1.13b). DCO increased for all four domains in Recording B, with a median endpoint-minus-pre-CV change of +9.49 raw units pixel^{-1}, indicating that the centers became relatively darker than the internal outer zones. DCO decreased for all four domains in Recording C, with a median change of −15.88 raw units pixel^{-1}, indicating relative central contrast loss.

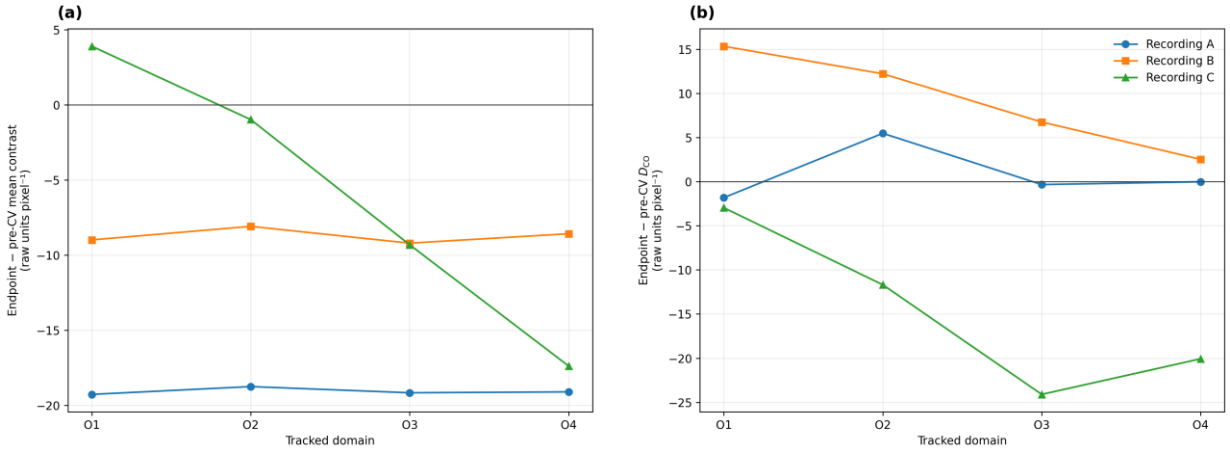

Figure 1.13. Within-recording raw-contrast responses. (a) Endpoint-minus-pre-CV mean-contrast change using the fixed Recording A pre-CV local-annulus reference. (b) Endpoint-minus-pre-CV change in DCO = Iouter − Icenter. Positive DCO indicates a darker center relative to the internal outer zone. Raw intensity units are native stored gain-normalized detector values and are not calibrated electron counts.

All four domains lost projected area in both post-EIS recordings, and the median loss was substantially larger in Recording C than in Recording B. No repeatable potential-locked morphological response was resolved across cycles and scan directions; the dominant pattern was progressive change over recording time. The observed area loss is compatible with partial dissolution, lateral collapse, fragmentation below the segmentation scale, or redistribution of

material outside the projected boundary. The two-dimensional images do not distinguish among these alternatives.

Conclusions

Electrochemical LPTEM resolved nanoscale FFssFF-derived coacervate domains and followed the same four objects through one pre-EIS and two post-EIS CV recordings in a single liquid-cell field. The largest qualitative image transition occurred across the EIS interval, and the two post-EIS recordings subsequently showed progressive projected-area loss. Median endpoint area changed by +0.98% in Recording A, −7.41% in Recording B, and −20.67% in Recording C. Under the predefined 5% operational criterion, threshold crossing occurred earlier in Recording C.

The apparent preference for greater percentage loss in smaller domains did not persist for absolute area or radius-based measures and therefore does not establish a general size law. Within-recording raw-contrast analysis showed different internal contrast responses in Recordings B and C, but these detector-scale measurements cannot be converted directly to mass, density, concentration, or composition. The chapter therefore provides a selected-case record of electrochemically associated projected morphology while leaving the relative contributions of applied potential, electron exposure, transport, and molecular chemistry unresolved.

References

1. Abbas, M.; Lipinski, W. P.; Nakashima, K. K.; Huck, W. T. S.; Spruijt, E. A Short Peptide Synthon for Liquid-Liquid Phase Separation. Nature Chemistry 2021, 13, 1046–1054. https://doi.org/10.1038/s41557-021-00788-x.
2. Zeinalvand, L.; Barpuzary, D.; Ueon, T. H.; Lai, W. B.; Song, Y.; Li, Z.; Dhaoui, R.; Idris, N. S.; Peddireddy, K. R.; Robertson-Anderson, R. M.; Tobias, D. J.; Guan, Z.; Green, J. R.; Patterson, J. P. Nanoscale Interfacial Organization Governs Maturation and Collapse in Passive versus Active Condensates. Journal of the American Chemical Society 2026, 148, 17967–17976. https://doi.org/10.1021/jacs.6c00997.

3. Gibson, W.; Mulvey, J. T.; Das, S.; Selmani, S.; Merham, J. G.; Rakowski, A. M.; Schwartz, E.; Hochbaum, A. I.; Guan, Z.; Green, J. R.; Patterson, J. P. Observing the Dynamics of an Electrochemically Driven Active Material with Liquid Electron Microscopy. ACS Nano 2024, 18, 11898–11909. https://doi.org/10.1021/acsnano.4c01524.

4. Mulvey, J. T.; Iyer, K. P.; Ortega, T.; Merham, J. G.; Pivak, Y.; Sun, H.; Hochbaum, A. I.; Patterson, J. P. Correlating Electrochemical Stimulus to Structural Change in Liquid Electron Microscopy Videos Using the Structural Dissimilarity Metric. Ultramicroscopy 2024, 257, 113894. https://doi.org/10.1016/j.ultramic.2023.113894.

5. Sung, J.; Bae, Y.; Park, H.; Kang, S.; Choi, B. K.; Kim, J.; Park, J. Liquid-Phase Transmission Electron Microscopy for Reliable In Situ Imaging of Nanomaterials. Annual Review of Chemical and Biomolecular Engineering 2022, 13, 167–191. https://doi.org/10.1146/annurev-chembioeng-092120-034534.

6. Wu, H.; Sun, H.; Oerlemans, R. A. J. F.; Li, S.; Shao, J.; Wang, J.; Joosten, R. R. M.; Lou, X.; Luo, Y.; Zheng, H.; Abdelmohsen, L. K. E. A.; Perez Garza, H. H.; van Hest, J. C. M.; Friedrich, H. Understanding, Mimicking, and Mitigating Radiolytic Damage to Polymers in Liquid Phase Transmission Electron Microscopy. Advanced Materials 2024, 36, e2402987. https://doi.org/10.1002/adma.202402987.

7. Yin, Y.; Niu, L.; Zhu, X.; Zhao, M.; Zhang, Z.; Mann, S.; Liang, D. Non-Equilibrium Behaviour in Coacervate-Based Protocells under Electric-Field-Induced Excitation. Nature Communications 2016, 7, 10658. https://doi.org/10.1038/ncomms10658.

8. Abbas, M.; Lipinski, W. P.; Wang, J.; Spruijt, E. Peptide-Based Coacervates as Biomimetic Protocells. Chemical Society Reviews 2021, 50, 3690–3705. https://doi.org/10.1039/D0CS00307G.

CHAPTER 2: Cryogenic Electron Microscopy Informed Molecular Dynamics Simulations to Investigate the Disulfide Hydrogel Self-Assembly

Publication Status and Author Contributions

Chapter 2 is adapted, with minor editorial modifications, from Song, Y.; Li, Z.; Mulvey, J. T.; Freites, J. A.; Patterson, J. P.; Tobias, D. J. "Cryogenic Electron Microscopy Informed Molecular Dynamics Simulations to Investigate the Disulfide Hydrogel Self-Assembly." ChemPhysChem 2025, 26(11), e202401085. https://doi.org/10.1002/cphc.202401085. Used with permission of John Wiley & Sons - Books; permission conveyed through Copyright Clearance Center, Inc., License No. 1780538-1.

For the work presented in this chapter, Zhaoxu Li performed the TEM simulations; the collaborators performed the MD simulations that provided the atomistic structural inputs. Accordingly, detailed MD analyses are condensed and retained only where needed to establish the provenance and suitability of those inputs.

Introduction

Molecular self-assembly serves as the foundation for creating functional nanostructures, offering customized properties vital for various applications [1]. Chemically fueled active self-assembly can generate transient, dynamically responsive materials maintained away from equilibrium [2, 6]. The CSH-CSSC system exhibits redox-driven transient dynamics and has been observed during electrochemical actuation [2, 3]. However, the detailed mechanisms underlying active self-assembly gelators such as CSH-CSSC remain incompletely understood, limiting systematic design and optimization of these materials.

In the redox-driven CSH-CSSC system, oxidation of CSH produces the disulfide hydrogelator CSSC, which forms fibrous assemblies [2]. Related N-acyl cystine derivatives are also known to form fibrous hydrogels stabilized by hydrogen bonding and aromatic interactions

[4, 5]. Through a redox reaction network governing hydrogelator formation and dissolution, this self-assembly system exhibits transient behavior that depends on reaction conditions [2]. The resulting fiber properties and self-assembly kinetics can be regulated spatially and temporally by electrical signals applied to microelectrodes [6].

Previous unrestrained MD simulations showed that both CSH and CSSC participate in aggregation and implicated π–π interactions and hydrogen bonding in stabilizing the assemblies [7]. Those results provide molecular context here, while this chapter focuses on the image-simulation bridge between atomistic fiber models and experimental cryo-EM.

Because those calculations addressed molecular aggregation rather than full experimentally sized fibers, they did not directly reproduce the complete fiber dimensions or experimental assembly timescales considered here [7, 10]. Cryo-EM and all-atom MD have been combined to resolve large protein assemblies, including the mature HIV-1 capsid [8]. Broader opportunities for integrating TEM observations with computational modeling are reviewed elsewhere [9].

Cryogenic transmission electron microscopy (cryo-EM) is a powerful method for visualizing the structural evolution of self-assembly systems in their native environment. Fibers formed by the CSH-CSSC system have been studied using time-resolved cryo-EM [10]. Hurst et al. highlighted the complexity of the synchronous process, wherein forward (assembly) and backward (disassembly) reactions occur simultaneously, contrasting with the sequential process, where these reactions happen in separate stages [10]. Cryo-EM revealed a transient, thermodynamically unstable stacked nanorod phase that emerged during the backward reaction and persisted for approximately six hours in the synchronous process [10].

The stacked-fiber phase provides an experimentally measured structural scale for individual fibers. Cryo-EM analysis reported a base diameter of approximately 5 nm and an apparent stacked-fiber diameter of approximately 4.6 nm [10]. The individual-fiber diameter distribution was therefore used to define the radial scale of restrained all-atom MD models. These restraints enabled generation of fiber-like atomistic inputs for image simulation, but direct simulation of the transient stacked phase and independent validation of its molecular packing were outside the scope of this study.

Although restraints derived from cryo-EM data have been applied previously in MD simulations, direct image-domain comparisons between MD-generated structures and cryo-EM observations remain limited. To bridge this gap, multislice TEM simulations were used to convert MD-derived structures into images that could be compared with experimental cryo-EM data. Imaging parameters were varied to determine whether the restrained structures reproduced principal fiber-edge features and to quantify how defocus and water thickness affected apparent fiber diameter.

This chapter focuses on the TEM-simulation component of the published study: generating multislice images from cryo-EM-constrained atomistic structures, varying defocus and water thickness, and quantitatively comparing simulated and experimental apparent fiber diameters. MD trajectories are summarized as the source of the atomistic inputs; detailed analyses of molecular packing and intermolecular interactions are outside the scope of this dissertation-focused adaptation.

Materials and Methods

Cryo-EM images were obtained from the dataset reported by Hurst et al. [10]. The segmentation masks, the 100-pixel (30.58 nm) profile-averaging procedure, and the FWHM

reanalysis of 770 images followed the workflow reported in the published source of this chapter [21].

After the diameter measurements were obtained, a Gaussian mixture model was fitted to characterize the fiber-diameter distribution in the cryo-EM images.

The atomistic snapshots used as TEM-simulation inputs were generated by collaborators with NAMD [13, 14]. CSH/CSSC parameters were generated with CGenFF [15, 16, 17], while CHARMM36/CHARMM36m terms were used where applicable [18, 19]; TIP3P water was used as solvent [20]. The exact topology, partial charges, penalty scores, and subsequent parameter refinement are documented in the published source and associated repository [21]. Each trajectory contained 1174 CSH equivalents in either pure CSSC or 70% CSSC systems, with the restraint conditions and trajectory lengths reported in Table 2.1. After initial spherical confinement, a cylindrical restraint with a 27 Å radius (5.4 nm diameter), derived from the experimental fiber diameter, was applied as illustrated in Scheme 2.1.

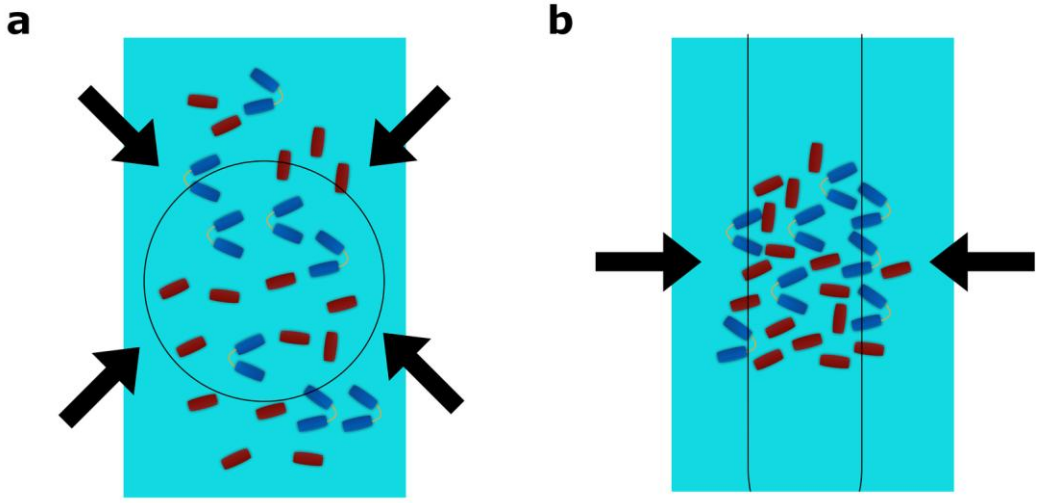

Scheme 2.1. Fiber restraint setup. The cyan box represents the solvent, with red and blue discs indicating CSH and CSSC molecules, respectively. Black circles and lines denote (a) spherical and (b) cylindrical restraints, respectively, while black arrows represent the direction of the force applied to any blue or red discs outside the black circle or line. Spherical restraints were applied only at the start of the simulations to concentrate all the CSH/CSSC molecules.

Table 2.1. Simulation trajectory composition and durations.

Conversion to CSSC (%)	Restrained species	CSH equivalents	Water molecules	Restrained duration (ns)	Unrestrained duration (ns)
70	All	1174	91,750	772.82	116.60
70	CSSC only	1174	91,750	479.66	116.60
100	All	1174	92,967	522.80	166.50
0	All	1174	99,182	95.86	70.46

Note. “All” denotes restraint of both CSH and CSSC molecules; “CSSC only” denotes restraint of CSSC molecules only.

Homogeneous water-only regions were randomly cropped from different snapshots of the pure-CSSC MD simulation in which all CSSC molecules were restrained and used to construct large water blocks. The blocks were then trimmed to the desired dimensions. For each selected snapshot, the CSH/CSSC region and its surrounding water were extracted and inserted at the center of a water block after overlapping water-block atoms within the extracted core region were deleted. This procedure generated structures with varying water thicknesses. The resulting structures were imported into abTEM [11] for TEM simulation. The accelerating voltage was set to 200 kV, and the frozen-phonon method was used to account for electron–phonon scattering. Multislice calculations used a slice thickness of 1 Å. The image sampling was set to 3.058 Å to match the cryo-EM images. The chromatic aberration coefficient, spherical aberration coefficient, and semi-angle cutoff in the contrast transfer function (CTF) were set to 2.1 mm, 2.0 mm, and 10.7 mrad, respectively, consistent with the JEOL 2100 TEM parameters. A range of defocus values was also applied in the CTF. After CTF application, image intensities from the frozen-phonon ensemble were averaged to generate each simulated TEM image.

Because the atomistic fibers were much shorter than the CSH–CSSC fibers in the cryo-EM images, simulated line profiles used a 10-pixel (3.058 nm) averaging width to reduce end effects, whereas experimental profiles used a 100-pixel (30.58 nm) width. The line-profile length was 100 pixels (30.58 nm) so that each profile included the fiber, surrounding rims, and adjacent water. The darkest point within the central 10 pixels was defined as the minimum intensity, and the average of the nearest local maxima on the two sides defined the maximum intensity. FWHM was then calculated to determine apparent fiber width (Figure 2.1). Because the averaging widths differed, the experimental and simulated measurements are compared at the level of aggregate apparent-width distributions rather than as identical point-by-point measurement operators.

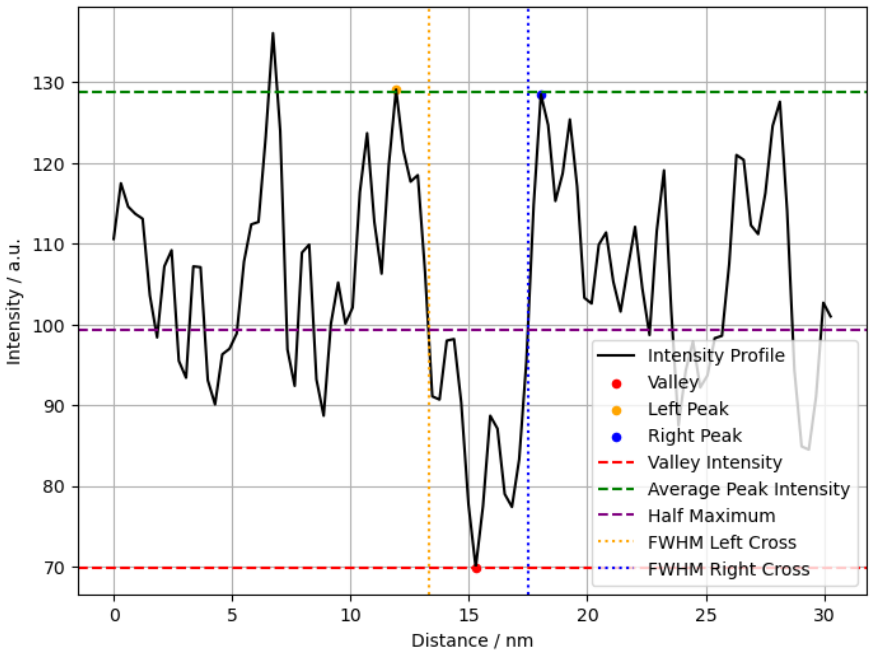

Figure 2.1. Illustration of the method used to calculate the full width at half maximum (FWHM), highlighting the key steps and parameters involved in the process.

Results

The reanalysis reported in the published source of this chapter included 770 cryo-EM images spanning 28 experimental conditions [21]. The reported segmentation method was used to identify fiber locations, and a custom MATLAB script calculated the full width at half maximum (FWHM) of the transverse intensity profile at each sampled location (Figure 2.1). Approximately 594 million spatial profile samples were combined into a histogram and fitted with a four-component Gaussian mixture model (Figure 2.2). These samples were spatially nested within fibers and images and were not treated as independent experimental replicates. The histogram contained two principal peaks: a peak near 5 nm associated with single-fiber measurements and a peak near 12 nm associated with profiles spanning adjacent fibers without clear separation. The component used to represent individual fibers had a mean of 5.42 nm and a standard deviation of 1.55 nm.

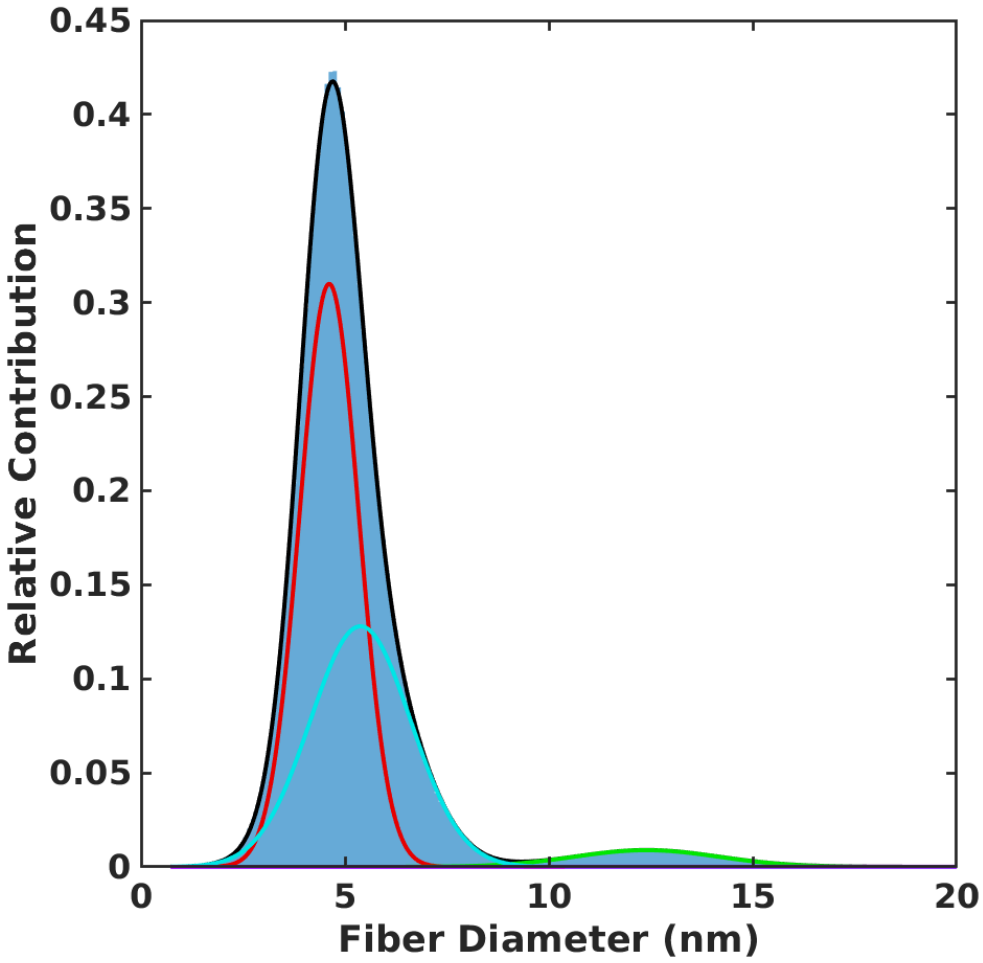

Figure 2.2. Fiber-diameter distribution and corresponding Gaussian-mixture fit obtained from 770 analyzed cryo-EM images.

TEM simulations used atomistic snapshots from three restrained systems, each containing 1174 CSH equivalents: pure CSSC with all CSSC molecules restrained, 70% CSSC with both CSH and CSSC molecules restrained, and 70% CSSC with only CSSC molecules restrained. The snapshot provenance and restraint conditions are summarized in Figure 2.3 and Table 2.1.

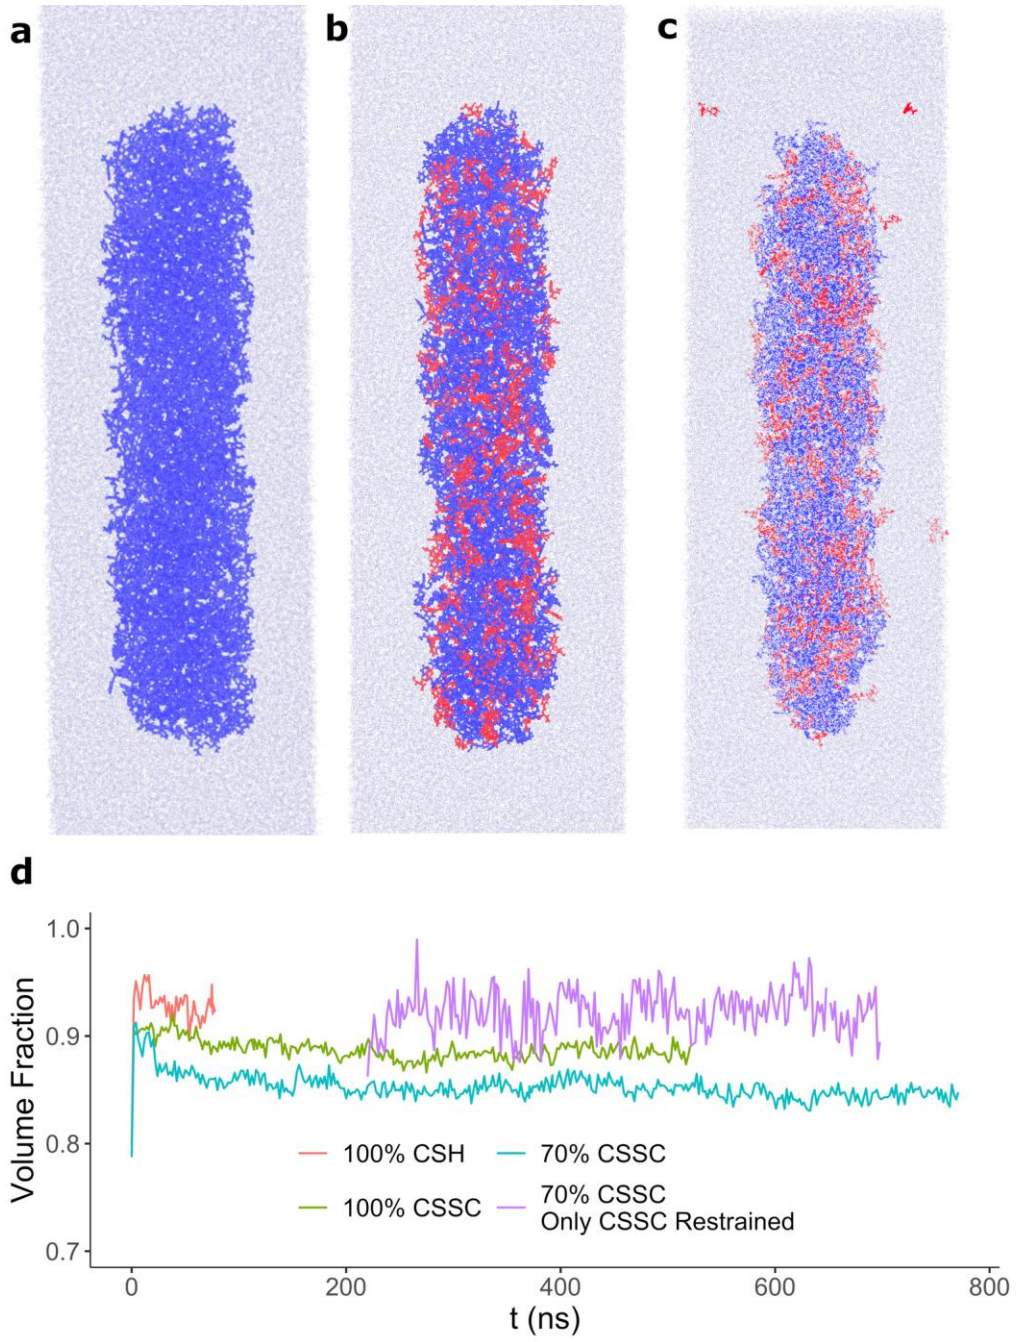

Figure 2.3. Restrained-fiber snapshots from collaborator-generated MD simulations used as atomistic inputs for TEM simulation: (a) pure CSSC with all CSSC molecules restrained, (b) 70% CSSC with both CSH and CSSC molecules restrained, and (c) 70% CSSC with only CSSC molecules restrained. (d) Volume-fraction trajectories for these systems and the pure-CSH control. Detailed MD setup and molecular analyses are reported in the published article [21].

Under the applied restraints, most molecules remained in stable fiber-like configurations, and the CSSC-containing systems retained elongated structures after restraint release over the

reported simulation intervals (Figures 2.3 and 2.4). The volume fractions equilibrated quickly and were used as a qualitative check of restraint stability before selecting snapshots for TEM simulation. Detailed analyses of density, conformation, hydrogen bonding, π–π interactions, and atomic contacts are reported in the published article and are not repeated here.

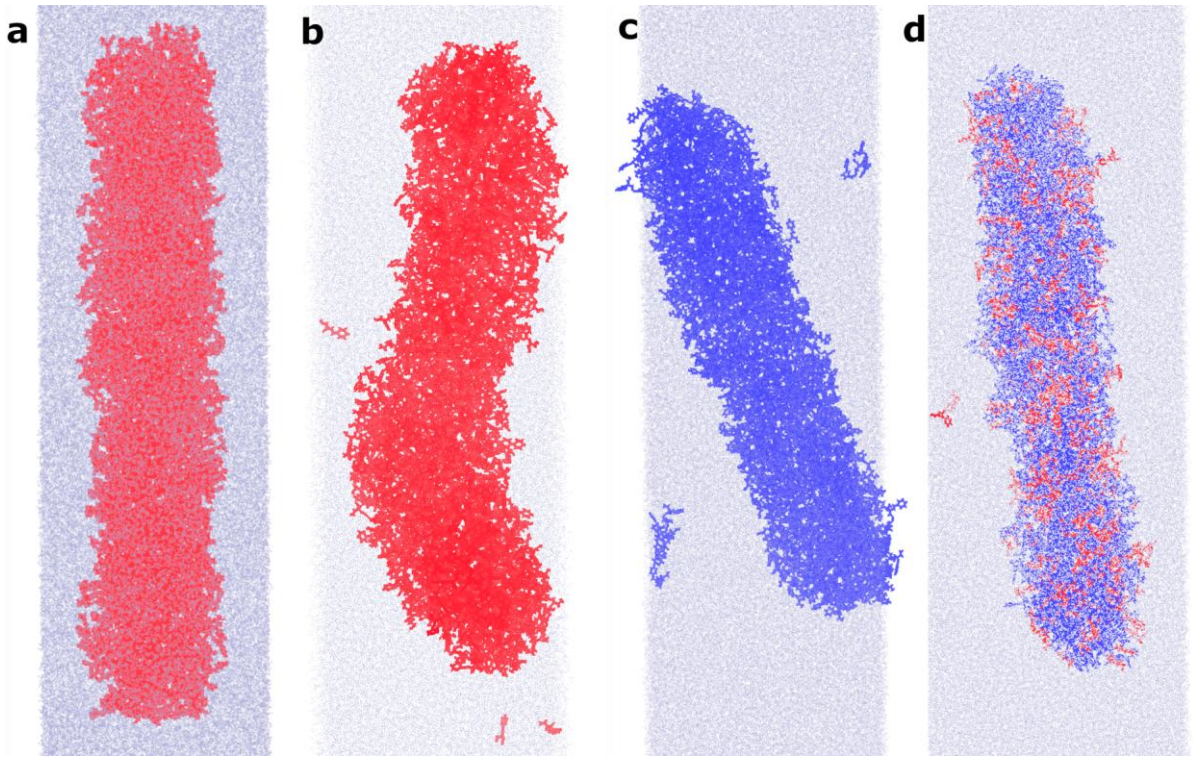

Figure 2.4. Snapshots of restrained systems: (a) pure CSH fiber under restraints, (b) pure CSH fiber 70 ns after restraint removal, (c) pure CSSC fiber 166 ns after restraint removal, and (d) mixed 70% CSSC fiber 116 ns after restraint removal. Retained here to document the stability screening used to select TEM-simulation inputs; detailed molecular interpretation is reported in the published article.

To compare the restrained MD structures with experimental cryo-EM observations, multislice TEM images were generated from the selected atomistic snapshots using abTEM [11]. Because specimen thickness can substantially change image appearance and signal-to-noise ratio [12], simulations were performed at study-specific defocus values from −4 to −18 µm in 0.5 µm increments and water thicknesses of 10, 20, 40, 60, 80, 100, 150, and 200 nm. Water-only regions from the fully restrained pure-CSSC simulation were used to expand the water volume surrounding each fiber.

The simulated TEM images of the pure-CSSC system with all CSSC molecules restrained, the 70% CSSC system with both CSH and CSSC molecules restrained, and the 70% CSSC system with only CSSC molecules restrained exhibited image features similar to those of fibers in the experimental cryo-EM images (Figure 2.5). Under the selected underfocus conditions, the simulated fibers exhibited a dark central feature bordered by bright edge fringes. The side rims appeared brighter than the end rims, and the surrounding water formed a noisy gray background.

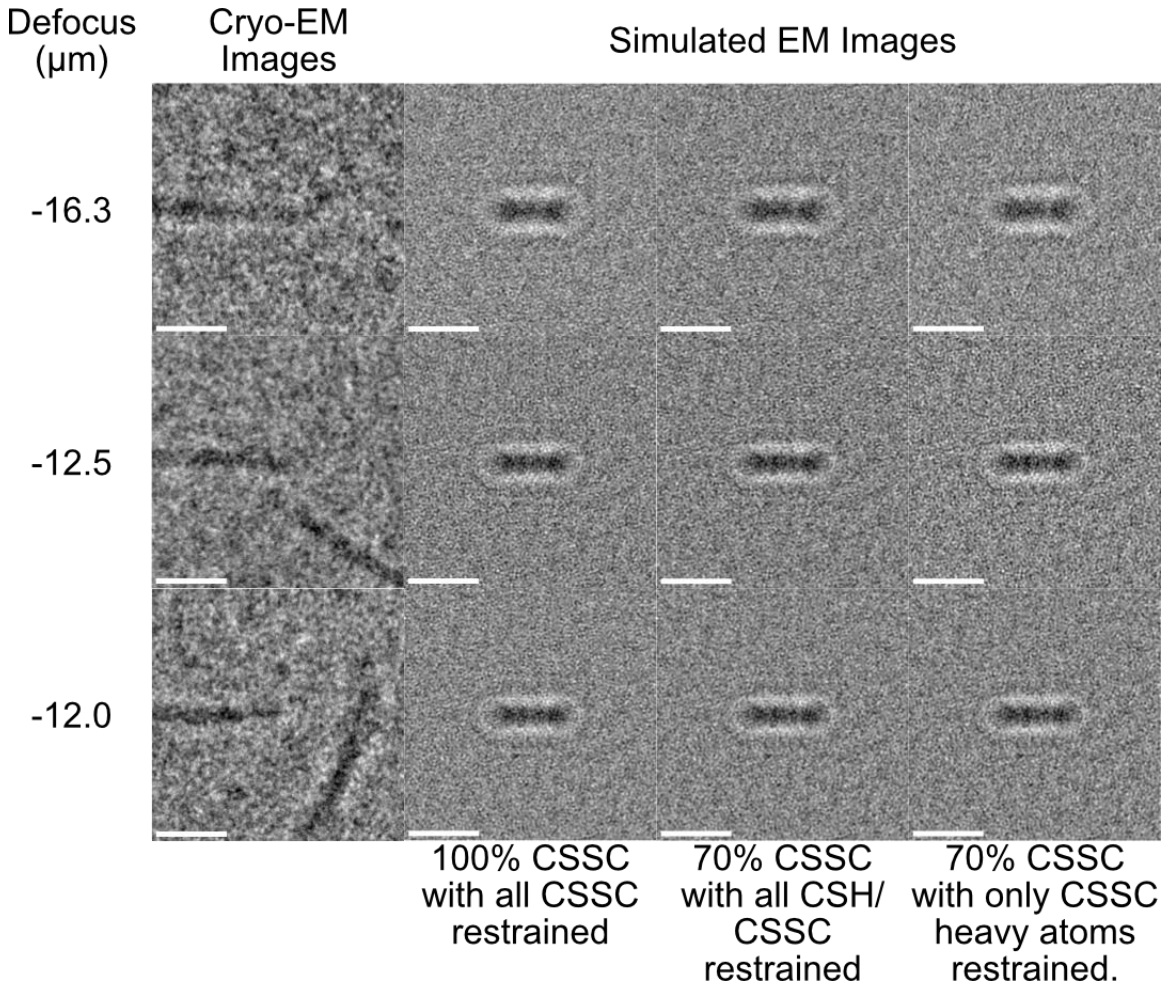

Figure 2.5. Comparison of fibers in experimental cryo-EM images with corresponding simulated images. The simulated images in the first row were generated at a defocus of −16.5 μm; those in the second and third rows were generated at the defocus values indicated for the corresponding experimental images. Scale bar: 20 nm.

Comparison across defocus values showed that defocus strongly influenced fiber appearance. Smaller underfocus magnitudes produced thinner apparent fibers and narrower rims, whereas larger underfocus magnitudes produced larger apparent fiber widths and rims (Figure 2.6a and Figures 2.7, 2.8, and 2.9). At lower water thicknesses, the fibers and their rims were distinctly visible. However, as the water thickness increased, the signal-to-noise ratio (SNR) progressively decreased, causing the fibers and rims to appear gray and blend more with the water background, making differentiation increasingly difficult (Figure 2.6b and Figures 2.10, 2.11, and 2.12).

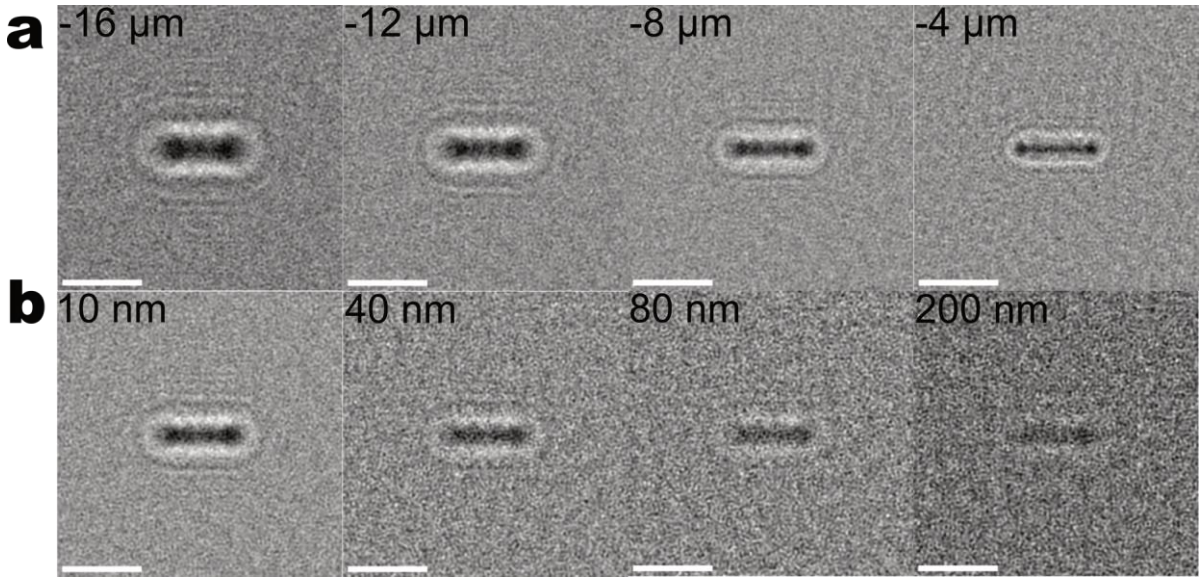

Figure 2.6. (a) Simulated cryo-TEM images generated from a single MD snapshot of 100% CSSC with all CSSC molecules restrained. Defocus values were set to −16, −12, −8, and −4 µm (from left to right) to illustrate the effect of defocus. (b) Simulated cryo-TEM images generated from the same MD snapshot at water thicknesses of 10, 40, 80, and 200 nm (from left to right) to illustrate the effect of water thickness. Scale bars: 20 nm.

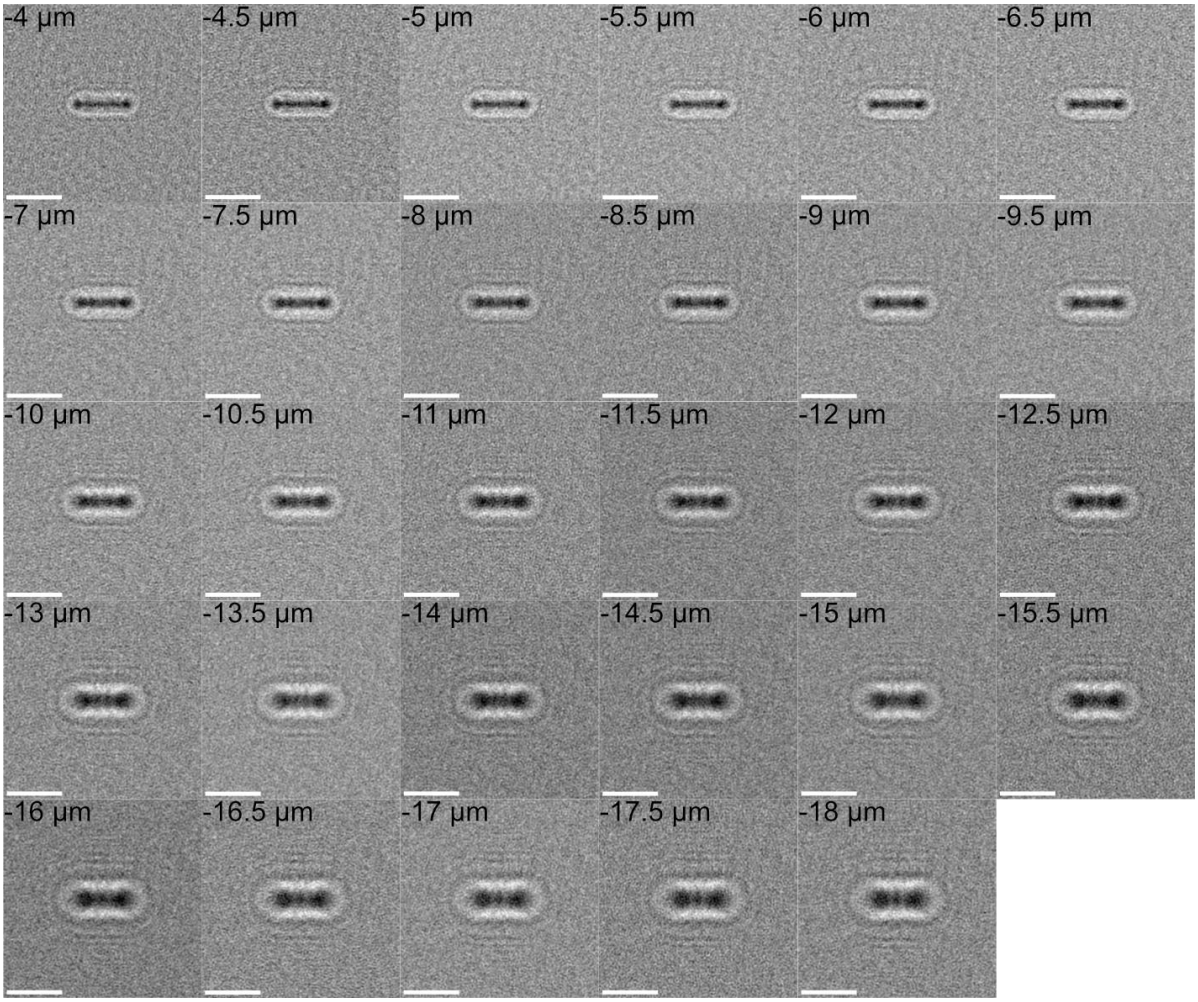

Figure 2.7. Simulated cryo-EM images at varying defocus values generated from a single snapshot of the MD simulation of 100% CSSC, with all CSSC molecules restrained. Scale bar: 20 nm.

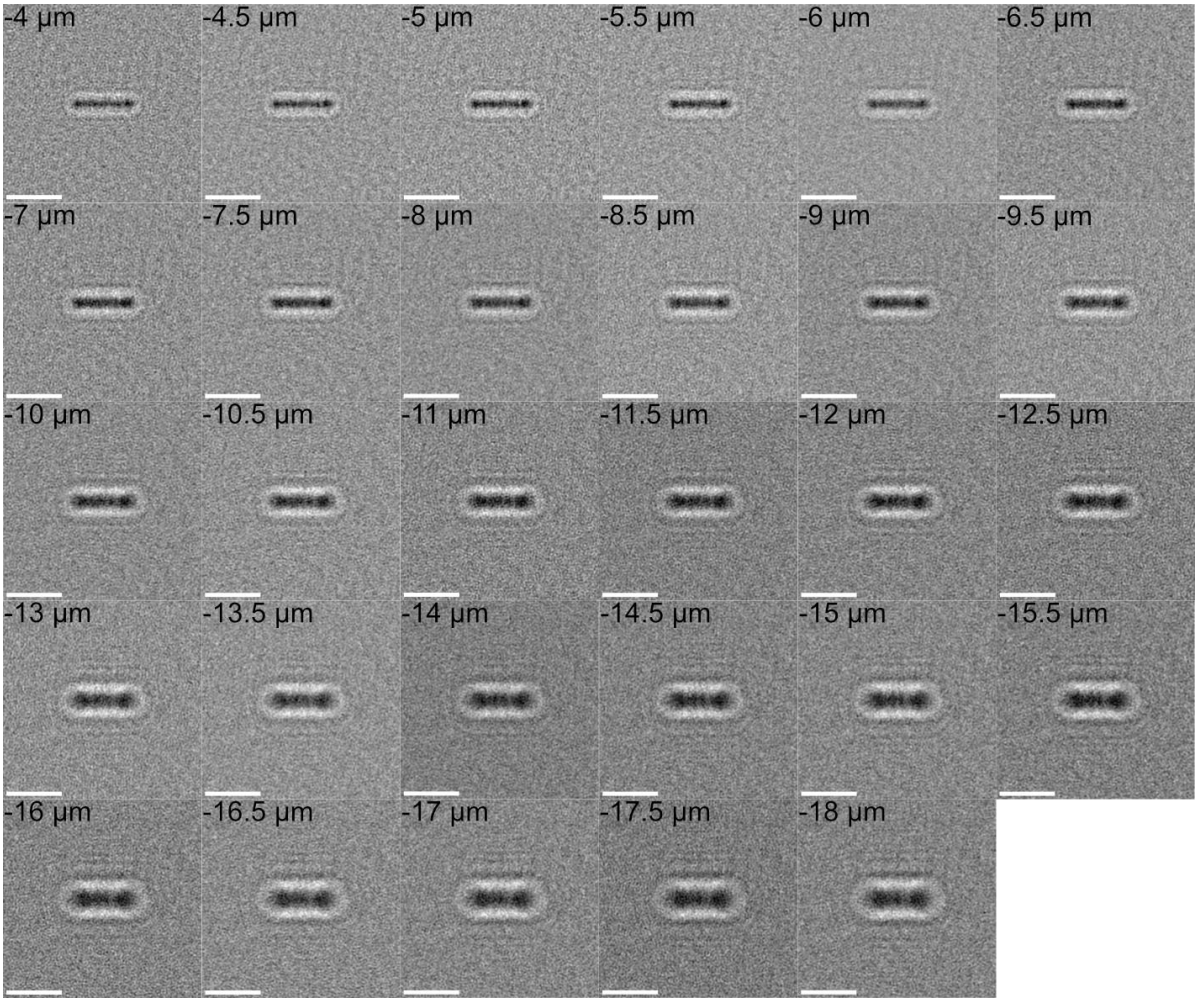

Figure 2.8. Simulated cryo-EM images at varying defocus values generated from a single snapshot of the MD simulation of 70% CSSC, with both CSH and CSSC molecules restrained. Scale bar: 20 nm.

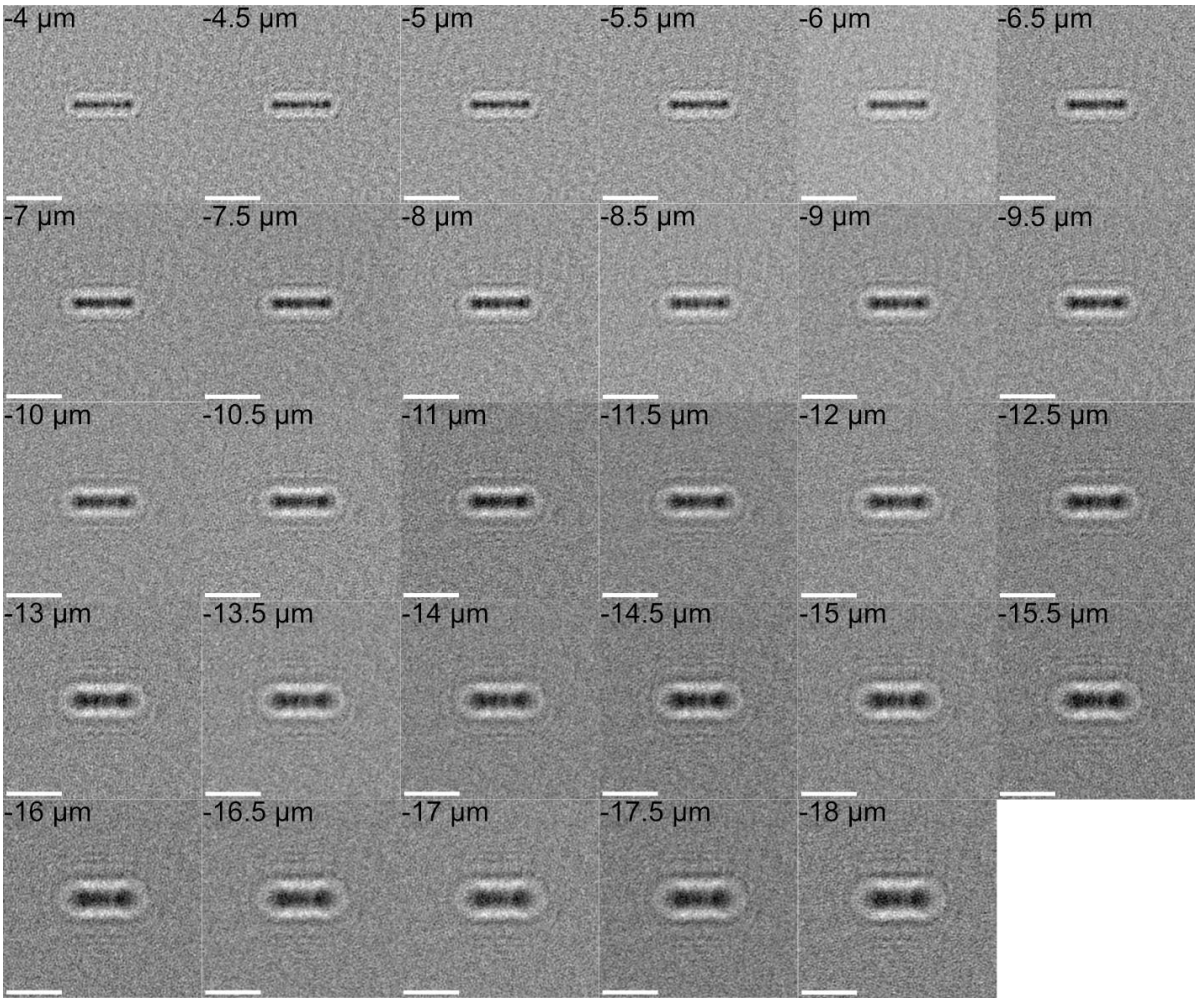

Figure 2.9. Simulated cryo-EM images at varying defocus values generated from a single snapshot of the MD simulation of 70% CSSC, with only CSSC molecules restrained. Scale bar: 20 nm.

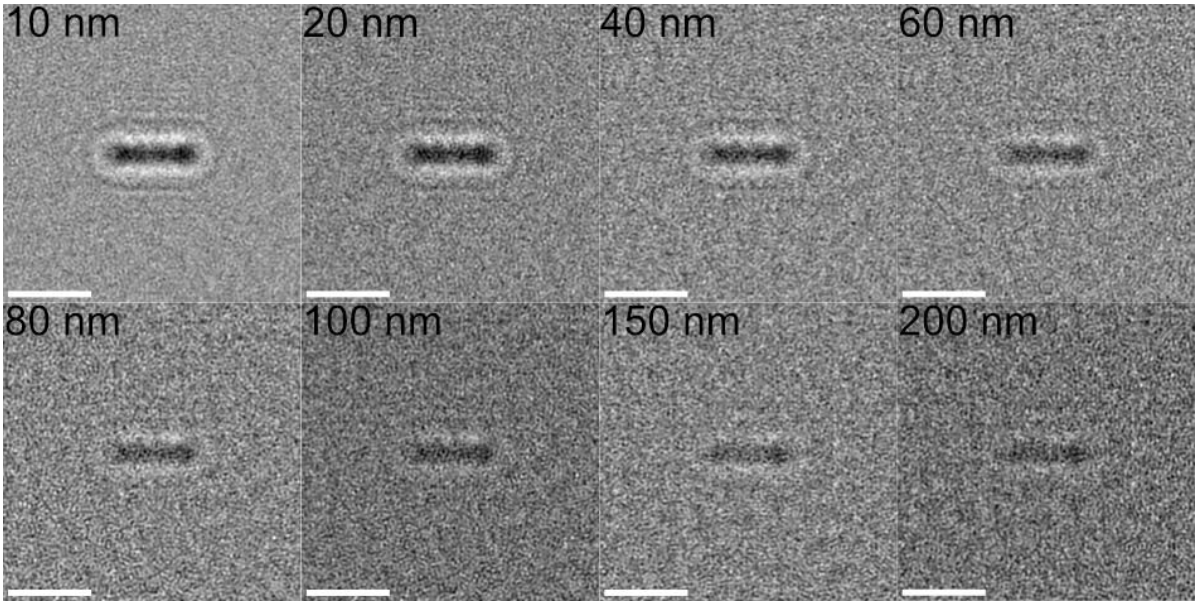

Figure 2.10. Simulated cryo-EM images at varying water thicknesses generated from a single snapshot of the MD simulation of 100% CSSC, with all CSSC molecules restrained. Scale bar: 20 nm.

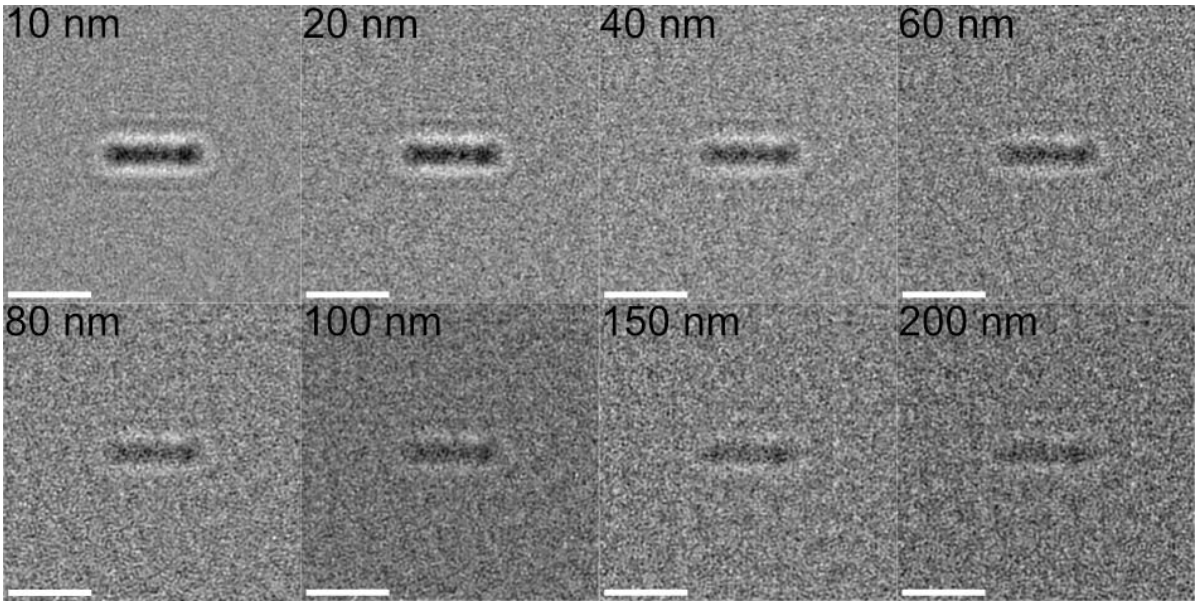

Figure 2.11. Simulated cryo-EM images at varying water thicknesses generated from a single snapshot of the MD simulation of 70% CSSC, with both CSH and CSSC molecules restrained. Scale bar: 20 nm.

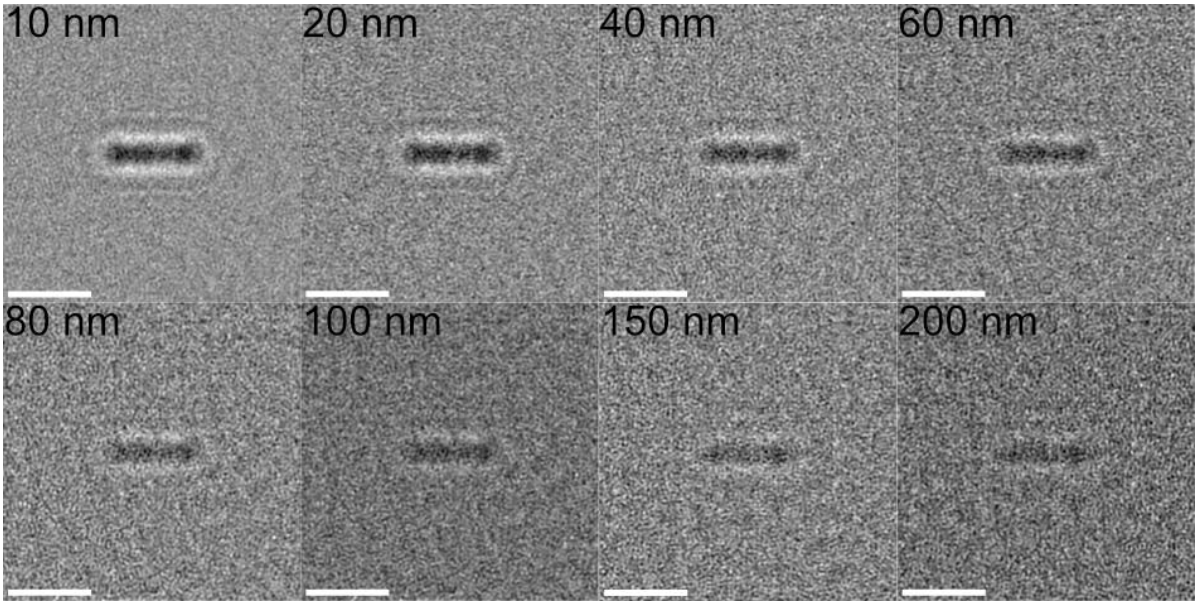

Figure 2.12. Simulated cryo-EM images at varying water thicknesses generated from a single snapshot of the MD simulation of 70% CSSC, with only CSSC molecules restrained. Scale bar: 20 nm.

To quantify these imaging-parameter effects, we measured vertical line profiles through the centers of fibers from the three simulated systems: pure CSSC with all CSSC molecules restrained, 70% CSSC with both CSH and CSSC molecules restrained, and 70% CSSC with only CSSC molecules restrained. Apparent diameters as functions of defocus for each water thickness are shown in Figure 2.13a–c. For each water thickness, apparent diameter increased approximately linearly with underfocus magnitude (Table 2.2). The diameter–defocus relationship became more linear as water thickness decreased from 200 to 10 nm, as indicated by the higher R^2 values. This trend was particularly pronounced below 100 nm.

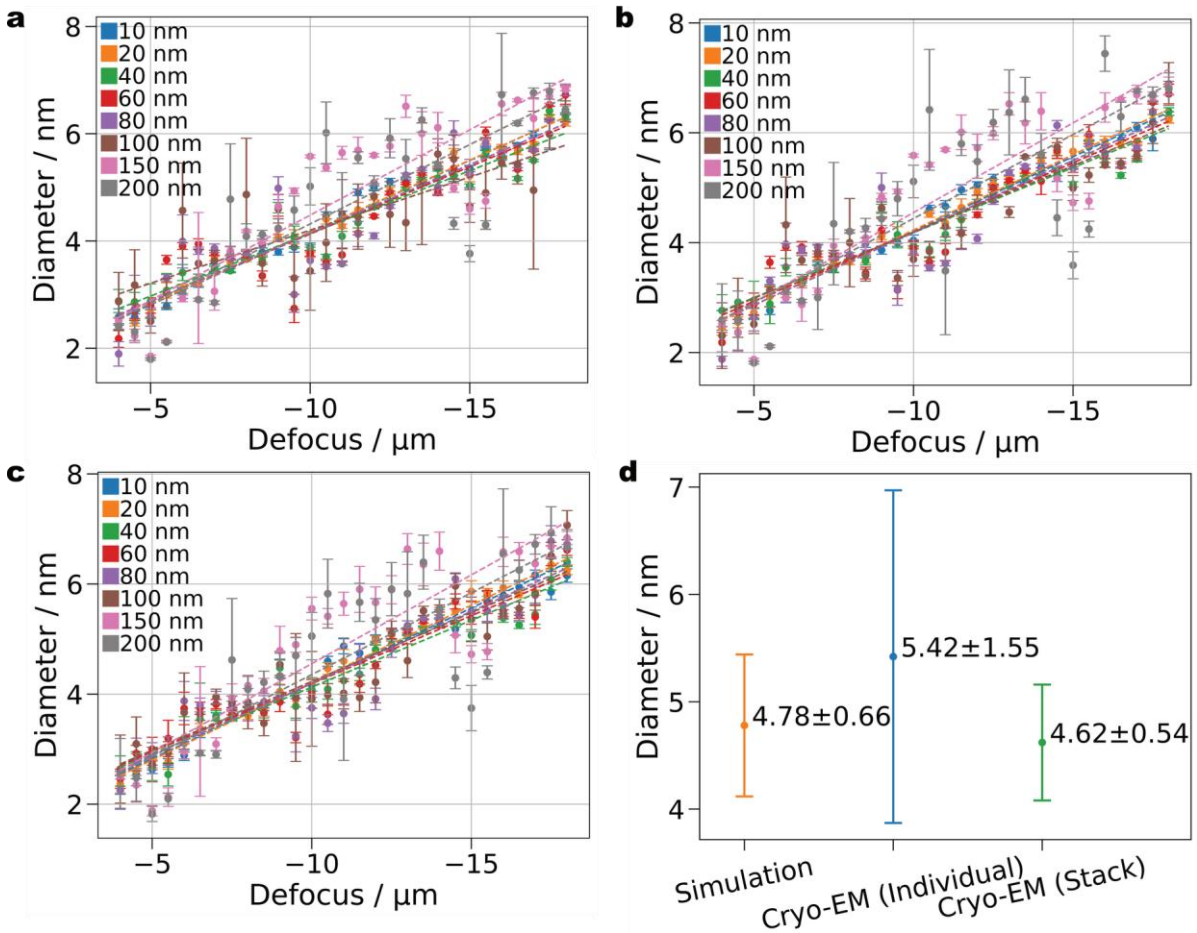

Figure 2.13. (a–c) Measured fiber diameters (n > 100) at different water thicknesses and defocus values, with corresponding linear regressions: (a) 100% CSSC with all CSSC molecules restrained, (b) 70% CSSC with both CSH and CSSC molecules restrained, and (c) 70% CSSC with only CSSC molecules restrained. (d) Average diameter and standard deviation of fibers in simulated images, individual fibers, and stacked fibers observed in cryo-EM.

Table 2.2. Linear-regression parameters.

Water Thickness (nm)	Pure CSSC			70% CSSC (CSH and CSSC Restrained)			70% CSSC (Only CSSC Restrained)		
	Slope	**Intercept**	**R-squared**	**Slope**	**Intercept**	**R-squared**	**Slope**	**Intercept**	**R-squared**
10	-0.269	1.465	0.977	-0.271	1.505	0.972	-0.273	1.460	0.969
20	-0.268	1.487	0.972	-0.274	1.484	0.974	-0.281	1.396	0.972
40	-0.234	1.803	0.919	-0.237	1.812	0.920	-0.243	1.699	0.912
60	-0.251	1.619	0.830	-0.252	1.669	0.856	-0.247	1.713	0.877
80	-0.255	1.595	0.786	-0.263	1.548	0.803	-0.262	1.577	0.825
100	-0.198	2.221	0.530	-0.241	1.778	0.776	-0.250	1.729	0.773
150	-0.318	1.300	0.763	-0.327	1.281	0.777	-0.322	1.329	0.749
200	-0.296	1.360	0.659	-0.306	1.381	0.595	-0.297	1.384	0.630

To further compare the experimental and simulated cryo-EM images, we calculated the mean and standard deviation across all simulated fibers. Before statistical analysis, the simulated defocus values were weighted according to the experimental defocus distribution from −7 to −18 μm. This range included 754 of the 770 analyzed cryo-EM images (97.9% of the dataset). The resulting simulated fiber diameter was 4.78 ± 0.66 nm, within the reported experimental variability for individual fibers (5.42 ± 1.55 nm) and fibers in the stacked phase (4.62 ± 0.54 nm) (Figure 2.13d and Figure 2.2). The standard deviation of individual-fiber diameters in the experimental cryo-EM images was larger than that of the simulated fibers. This difference likely reflects both measurement variability associated with defocus and water thickness and an intrinsic distribution of fiber diameters. By contrast, the standard deviation for fibers in the stacked phase was similar to that of the simulated fibers, consistent with the narrower apparent-width distribution of these structures [10].

Overall, the simulated cryo-EM images resembled the experimental images in their principal contrast features and apparent-diameter range. This agreement was partly expected because the experimental diameter was used to define the radial restraint, and the experimental images lacked high-resolution structural detail for a more discriminating comparison. The result is therefore useful as an image-domain consistency check and as a demonstration of how defocus and ice thickness can alter apparent fiber width and visibility, rather than as an independent validation of molecular packing or measurement accuracy.

Discussion

The principal result of this chapter is the image-domain comparison between experimental cryo-EM fibers and multislice TEM images generated from atomistic snapshots. The simulations reproduced the dark central fiber and bright underfocus rims observed

experimentally, and systematic variation of defocus and water thickness showed how both parameters changed apparent fiber width and contrast. After weighting the simulated defocus values to the experimental distribution, the simulated diameter remained within the variability of the experimentally measured individual-fiber diameter.

This workflow connects cryo-EM measurement, atomistic structural models, and multislice TEM simulation in a common image domain. The selected snapshots reproduced principal image-scale fiber features, and the simulations provided a controlled way to assess how imaging conditions affect apparent morphology. The approach is therefore most useful as an imaging and model-consistency test rather than as an independent determination of molecular packing.

The diameter agreement is not a fully independent validation because the experimentally measured fiber diameter of approximately 5.4 nm was used to define the 27 Å radial restraint (a 54 Å restrained diameter). In addition, the atomistic fibers were much shorter than the experimental fibers, and the experimental images did not contain sufficient high-resolution detail to discriminate among alternative molecular arrangements. The results therefore support image-scale consistency under the tested imaging conditions, but they do not establish a unique atomic packing model.

This chapter establishes a TEM-simulation workflow that connects atomistic structural models to experimentally observable cryo-EM contrast. By treating defocus and water thickness as explicit imaging variables and measuring simulated and experimental fibers in the same image domain, it provides a controlled framework for model testing while preserving clear limits on atomic-structure interpretation.

Conclusion

Cryo-EM diameter measurements constrained atomistic CSH/CSSC fiber models, and selected snapshots were converted into multislice TEM images. The simulated images reproduced the principal experimental image features, while defocus and water thickness altered apparent width and contrast. Because the experimental diameter informed the radial restraint and the profile operators differed between experimental and simulated images, the comparison supports image-scale consistency under the tested conditions but not a unique molecular packing model.

Data Availability Statement

The data supporting this study are available from the corresponding author upon reasonable request. The setup and analysis scripts, single-molecule graph representations used for the restrained simulations, and the CSH and CSSC parameter files are available in the public MD analysis repository.

Acknowledgments

Zhaoxu Li acknowledges Yuanming Song, Justin T. Mulvey, J. Alfredo Freites, Joseph P. Patterson, and Douglas J. Tobias, the other authors of the original article, for their contributions to the collaborative study. Yuanming Song and Zhaoxu Li contributed equally to the original article. The molecular dynamics simulations used to generate the atomistic input structures for this chapter were performed by collaborators. This research was supported by the National Science Foundation Materials Research Science and Engineering Center program through the UC Irvine Center for Complex and Active Materials (DMR-2011967). The molecular dynamics simulations were carried out on the HPC3 cluster administered by the UC Irvine Research Cyberinfrastructure Center.

Conflict of Interest

The authors declare no conflict of interest.

References

1. Pochan, D.; Scherman, O. Introduction: Molecular Self-Assembly. Chemical Reviews 2021, 121, 13699–13700. https://doi.org/10.1021/acs.chemrev.1c00884.
2. Ogden, W. A.; Guan, Z. Redox Chemical-Fueled Dissipative Self-Assembly of Active Materials. ChemSystemsChem 2020, 2, e1900030. https://doi.org/10.1002/syst.201900030.
3. Gibson, W.; Mulvey, J. T.; Das, S.; Selmani, S.; Merham, J. G.; Rakowski, A. M.; Schwartz, E.; Hochbaum, A. I.; Guan, Z.; Green, J. R.; Patterson, J. P. Observing the Dynamics of an Electrochemically Driven Active Material with Liquid Electron Microscopy. ACS Nano 2024, 18, 11898–11909. https://doi.org/10.1021/acsnano.4c01524.
4. Menger, F. M.; Yamasaki, Y.; Catlin, K. K.; Nishimi, T. X-Ray Structure of a Self-Assembled Gelating Fiber. Angewandte Chemie International Edition in English 1995, 34, 585–586. https://doi.org/10.1002/anie.199505851.
5. Menger, F. M.; Caran, K. L. Anatomy of a Gel. Amino Acid Derivatives That Rigidify Water at Submillimolar Concentrations. Journal of the American Chemical Society 2000, 122, 11679–11691. https://doi.org/10.1021/ja0016811.
6. Selmani, S.; Schwartz, E.; Mulvey, J. T.; Wei, H.; Grosvirt-Dramen, A.; Gibson, W.; Hochbaum, A. I.; Patterson, J. P.; Ragan, R.; Guan, Z. Electrically Fueled Active Supramolecular Materials. Journal of the American Chemical Society 2022, 144, 7844–7851. https://doi.org/10.1021/jacs.2c01884.
7. Song, Y.; Selmani, S.; Freites, J. A.; Guan, Z.; Tobias, D. J. Multiscale Molecular Dynamics Simulations of an Active Self-Assembling Material. Journal of Physical Chemistry B 2024, 128, 1266–1274. https://doi.org/10.1021/acs.jpcb.3c06572.
8. Zhao, G.; Perilla, J. R.; Yufenyuy, E. L.; Meng, X.; Chen, B.; Ning, J.; Ahn, J.; Gronenborn, A. M.; Schulten, K.; Aiken, C.; Zhang, P. Mature HIV-1 Capsid Structure by Cryo-Electron Microscopy and All-Atom Molecular Dynamics. Nature 2013, 497, 643–646. https://doi.org/10.1038/nature12162.
9. Rizvi, A.; Mulvey, J. T.; Carpenter, B. P.; Talosig, R.; Patterson, J. P. A Close Look at Molecular Self-Assembly with the Transmission Electron Microscope. Chemical Reviews 2021, 121, 14232–14280. https://doi.org/10.1021/acs.chemrev.1c00189.
10. Hurst, P. J.; Mulvey, J. T.; Bone, R. A.; Selmani, S.; Hudson, R. F.; Guan, Z.; Green, J. R.; Patterson, J. P. CryoEM Reveals the Complex Self-Assembly of a Chemically Driven Disulfide Hydrogel. Chemical Science 2024, 15, 1106–1116. https://doi.org/10.1039/D3SC05790A.
11. Madsen, J.; Susi, T. The abTEM Code: Transmission Electron Microscopy from First Principles. Open Research Europe 2021, 1, 24. https://doi.org/10.12688/openreseurope.13015.2.
12. van de Put, M. W. P.; Patterson, J. P.; Bomans, P. H. H.; Wilson, N. R.; Friedrich, H.; van Benthem, R. A. T. M.; de With, G.; O'Reilly, R. K.; Sommerdijk, N. A. J. M. Graphene Oxide Single Sheets as Substrates for High Resolution CryoTEM. Soft Matter 2015, 11, 1265–1270. https://doi.org/10.1039/C4SM02587C.
13. Phillips, J. C.; Braun, R.; Wang, W.; Gumbart, J.; Tajkhorshid, E.; Villa, E.; Chipot, C.; Skeel, R. D.; Kalé, L.; Schulten, K. Scalable Molecular Dynamics with NAMD. Journal of Computational Chemistry 2005, 26, 1781–1802. https://doi.org/10.1002/jcc.20289.
14. Phillips, J. C.; Hardy, D. J.; Maia, J. D. C.; Stone, J. E.; Ribeiro, J. V.; Bernardi, R. C.; Buch, R.; Fiorin, G.; Hénin, J.; Jiang, W.; McGreevy, R.; Melo, M. C. R.; Radak, B. K.; Skeel, R. D.;

Singharoy, A.; Wang, Y.; Roux, B.; Aksimentiev, A.; Luthey-Schulten, Z.; Kalé, L. V.; Schulten, K.; Chipot, C.; Tajkhorshid, E. Scalable Molecular Dynamics on CPU and GPU Architectures with NAMD. Journal of Chemical Physics 2020, 153, 044130. https://doi.org/10.1063/5.0014475.

15. Vanommeslaeghe, K.; Hatcher, E.; Acharya, C.; Kundu, S.; Zhong, S.; Shim, J.; Darian, E.; Guvench, O.; Lopes, P.; Vorobyov, I.; MacKerell, A. D. CHARMM General Force Field: A Force Field for Drug-Like Molecules Compatible with the CHARMM All-Atom Additive Biological Force Fields. Journal of Computational Chemistry 2010, 31, 671–690. https://doi.org/10.1002/jcc.21367.
16. Vanommeslaeghe, K.; MacKerell, A. D. Automation of the CHARMM General Force Field (CGenFF) I: Bond Perception and Atom Typing. Journal of Chemical Information and Modeling 2012, 52, 3144–3154. https://doi.org/10.1021/ci300363c.
17. Vanommeslaeghe, K.; Raman, E. P.; MacKerell, A. D. Automation of the CHARMM General Force Field (CGenFF) II: Assignment of Bonded Parameters and Partial Atomic Charges. Journal of Chemical Information and Modeling 2012, 52, 3155–3168. https://doi.org/10.1021/ci3003649.
18. Huang, J.; MacKerell, A. D. CHARMM36 All-Atom Additive Protein Force Field: Validation Based on Comparison to NMR Data. Journal of Computational Chemistry 2013, 34, 2135–2145. https://doi.org/10.1002/jcc.23354.
19. Huang, J.; Rauscher, S.; Nawrocki, G.; Ran, T.; Feig, M.; de Groot, B. L.; Grubmüller, H.; MacKerell, A. D. CHARMM36m: An Improved Force Field for Folded and Intrinsically Disordered Proteins. Nature Methods 2017, 14, 71–73. https://doi.org/10.1038/nmeth.4067.
20. Jorgensen, W. L.; Chandrasekhar, J.; Madura, J. D.; Impey, R. W.; Klein, M. L. Comparison of Simple Potential Functions for Simulating Liquid Water. Journal of Chemical Physics 1983, 79, 926–935. https://doi.org/10.1063/1.445869.
21. Song, Y.; Li, Z.; Mulvey, J. T.; Freites, J. A.; Patterson, J. P.; Tobias, D. J. Cryogenic Electron Microscopy Informed Molecular Dynamics Simulations to Investigate the Disulfide Hydrogel Self-Assembly. ChemPhysChem 2025, 26(11), e202401085. https://doi.org/10.1002/cphc.202401085.

CHAPTER 3: Physics-Informed Synthetic Data for Foreground Semantic Segmentation of Low-Contrast Nanofibers in Liquid-Phase Transmission Electron Microscopy

Introduction

Liquid-phase transmission electron microscopy (LPTEM) enables direct imaging of nanoscale structures in liquid, but scattering by the liquid and cell windows, thickness variation, defocus, counting noise, drift, contamination, and beam-induced changes can obscure thin, low-contrast fibers [1, 2]. In the nanofiber task considered here, crossings, parallel overlaps, bundles, and interrupted segments introduce an additional operational ambiguity in defining pixel-level foreground.

To obtain traceable training pairs without deriving labels from image intensity, this work generates each synthetic TEM image and its reference mask from the same transformed atomic scene through independent pathways. Multislice simulation produces the input image, whereas geometric projection of fiber coordinates produces the binary foreground reference. Synthetic scenes vary fiber number, orientation, curvature, overlap, defocus, and counting noise.

The task is foreground semantic segmentation: each pixel is classified as fiber or background, and individual-fiber identity is not retained at crossings or within bundles. The chapter develops a molecular-to-image workflow, evaluates a U-Net on a fixed 500-image synthetic dataset against Otsu- and Sauvola-based pipelines, and examines transfer to unlabeled real TEM images. Quantitative accuracy is limited to the synthetic domain; experimental transfer is assessed qualitatively and through prediction-stability checks.

Methods

The workflow linked molecular modeling, long-fiber construction, aqueous scene assembly, multislice TEM simulation, coordinate-defined reference generation, dataset

partitioning, and U-Net development. The image and reference mask were generated independently from the same transformed atomic configuration, as summarized in Figure 3.1 and Table 3.1.

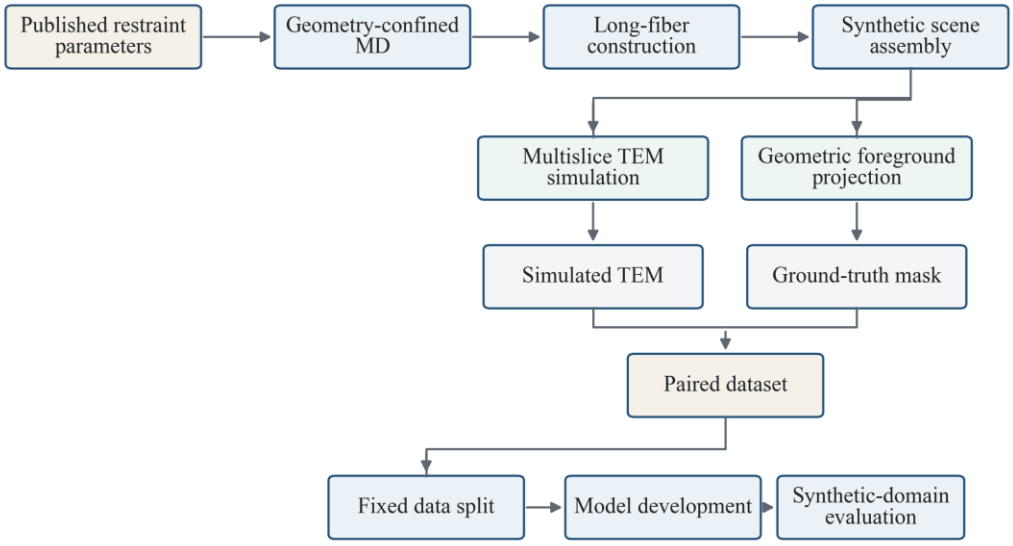

Figure 3.1. Physics-informed workflow for generating paired synthetic TEM images and coordinate-defined foreground references. Published restraint parameters were used in geometry-confined molecular dynamics to generate structural modules for long-fiber construction and multi-fiber scene assembly. Independent multislice TEM simulation and geometric foreground projection produced paired images and binary reference masks for fixed data partitioning, model development, and synthetic-domain evaluation.

Table 3.1. Overview of the physics-informed synthetic segmentation workflow.

Stage	Operation	Output
1	Periodic hydrated-fiber MD	Short periodic trajectory and late-stage frames
2	Long-fiber coordinate construction	Curved, locally hydrated fiber coordinates
3	Multi-fiber scene assembly and water embedding	Hydrated three-dimensional atomic scenes
4	TEM forward simulation	1024 × 1024 noisy synthetic TEM images
5	Heavy-atom projection and mask union	Foreground semantic reference masks independent of TEM intensity
6	Collection, audit, and fixed split	500 complete pairs; 348/74/78 fixed split
7	Four-stage U-Net development	47 completed adaptively selected training runs
8	Synthetic evaluation and real-TEM application	Selected U-Net and thresholding baselines evaluated on synthetic test data and applied qualitatively to real TEM images

The periodic molecular-dynamics source system contained 108 molecules of N-benzoyl-L-cysteine amide (CSH), 126 molecules of its disulfide dimer, N,N′-dibenzoyl-L-cystine diamide (CSSC), and 32,300 water molecules. The fiber was periodic along the x axis in an approximately $11 \times 10 \times 10$ nm³ simulation cell.

The cylindrical confinement geometry was adapted from the cryo-EM-informed molecular-dynamics model of Song et al. [3], which built on the experimental observations of Hurst et al. [4]. Cryo-EM measurements showed an individual-fiber diameter of approximately 5.4 nm; accordingly, a 2.7 nm radial confinement boundary was used. The corresponding flat-bottom restraint had a force constant of 836.8 kJ mol^{-1} nm^{-2} and acted radially on CSH and CSSC heavy atoms located outside the 2.7 nm cylindrical boundary, in the plane perpendicular to the fiber axis [3].

The transferred restraint geometry and its relationship to the periodic fiber axis are illustrated in Figure 3.2.

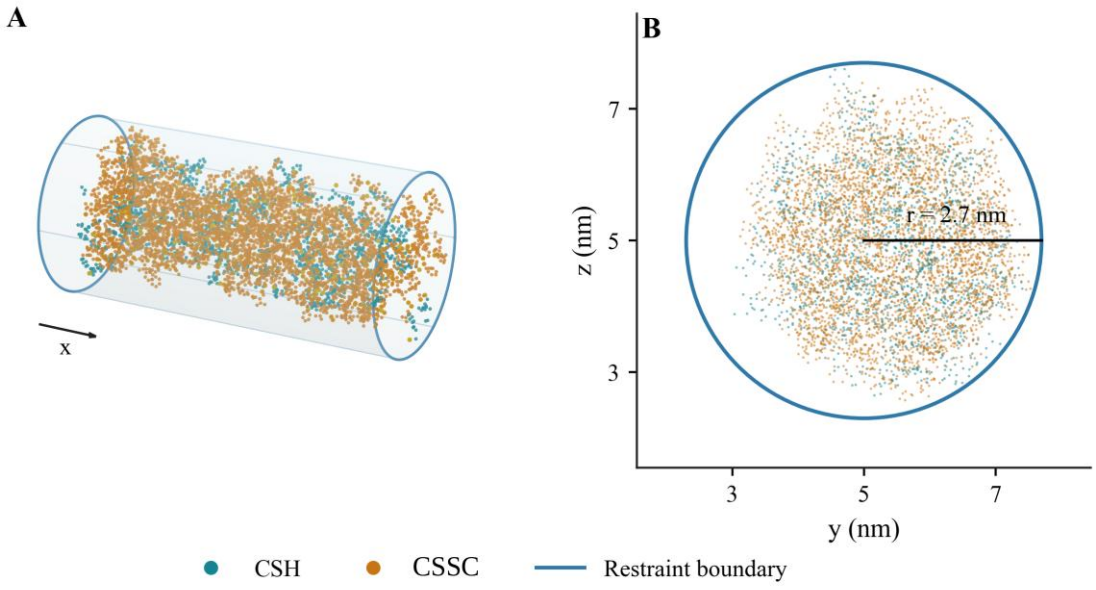

Figure 3.2. Geometry confinement used in periodic hydrated-fiber molecular dynamics. (A) Oblique view of the hydrated fiber aligned with the periodic x axis within the cylindrical flat-bottom restraint. (B) Transverse view of heavy atoms belonging to CSH and CSSC relative to the 2.7 nm restraint boundary. The restoring force acts radially toward the cylindrical boundary and is perpendicular to the fiber axis. The restraint geometry and parameters follow Song et al. [3] and were implemented in GROMACS.

The MD simulation was performed with GROMACS 2025.4 [7], using CHARMM/CGenFF parameters [6] and the CHARMM TIP3P water model [5]. The temperature, integration time step, particle-mesh Ewald settings, periodic-boundary conditions, equilibration schedule, and production duration were study-specific simulation inputs.

The simulation sequence consisted of energy minimization, a 50 ps constant-volume, constant-temperature equilibration, and segmented production runs to a cumulative duration of 10 ns. Eleven frames sampled at 500 ps intervals from the later portion of the trajectory supplied the structural modules used for long-fiber construction. These frames provided physically sampled local configurations, whereas the final 100–236.5 nm fibers were assembled geometrically from repeated modules rather than independently equilibrated as full-length fibers.

Representative late frames from this single completed trajectory are shown in Figure 3.3. The molecular composition and simulation conditions are summarized in Table 3.2.

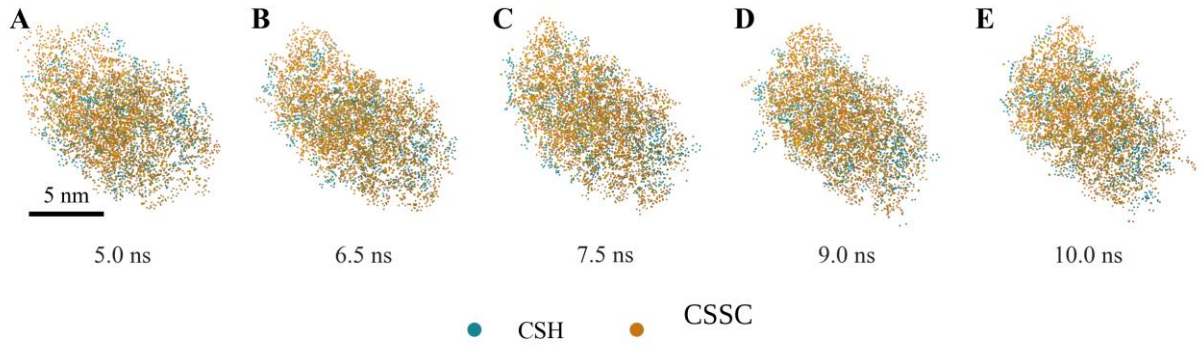

Figure 3.3. Molecular configurations sampled from the periodic hydrated-fiber molecular-dynamics trajectory. Panels A-E show configurations at 5.0, 6.5, 7.5, 9.0, and 10.0 ns, respectively, using a common projection, spatial scale, and sampling rule.

Table 3.2. Molecular system and simulation parameters for the periodic hydrated-fiber model.

Parameter	Value
Periodic box	Approximately 11 × 10 × 10 nm^3; fiber periodic along x
Composition	108 CSH molecules; 126 CSSC molecules; 32,300 water molecules
MD engine	GROMACS 2025.4
Force field / local water	CHARMM/CGenFF; CHARMM TIP3P
Temperature / time step	300 K / 1 fs
Electrostatics / boundaries	Particle-mesh Ewald / three-dimensional periodic boundaries
Fiber restraint	Flat-bottom cylinder with its axis along x; 2.7 nm radius; radial restoring force perpendicular to x; 836.8 kJ mol^{-1} nm^{-2}; applied to CSH and CSSC heavy atoms
Simulation sequence	Energy minimization; 50 ps NVT; segmented production to 10 ns
Frames used downstream	Eleven late frames sampled at 500 ps intervals

Long fibers were generated from the approximately 11 nm periodic modules by sampling the late MD frames and applying a randomized axial phase and rotation about the fiber axis to each module. The transformed modules were concatenated and then mapped to one of four centerline families: straight, circular-arc, sinusoidal, or random. The resulting fiber lengths were approximately 101.3, 133.6, 178.2, and 236.5 nm.

Local water molecules were not generated independently from the fiber coordinates. For each module, water molecules originating from the same MD frame and located within a 3.2 nm radial neighborhood of the solute were subjected to the same axial phase, rotation, and centerline mapping as the associated fiber coordinates. After transformation, clashes were removed using a 0.22 nm solute-heavy-atom-to-water-oxygen criterion and a 0.23 nm water-oxygen-to-water-oxygen criterion. This procedure preserved a local hydration environment tied to the sampled short MD frame while allowing the long fiber to adopt the target centerline geometry.

The recorded coordinate-construction operations from late MD modules through background-bath insertion are summarized in Figure 3.4.

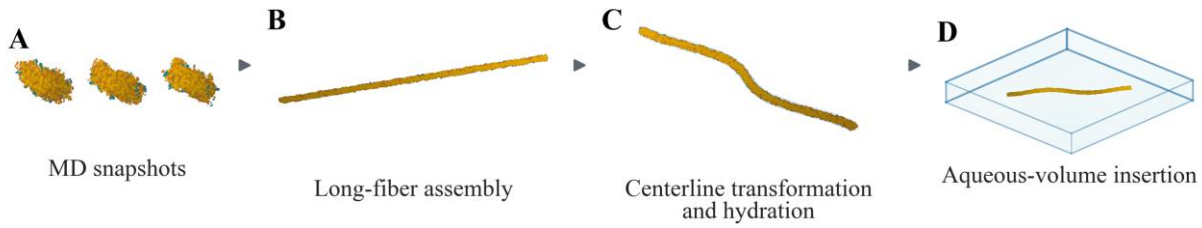

Figure 3.4. Assembly and hydration of a representative long-fiber structure. (A) Molecular configurations sampled from the geometry-confined molecular-dynamics trajectory. (B) Construction of a 236.5 nm fiber from time-, phase-, and rotation-mixed modules. (C) Centerline transformation of the assembled fiber together with its local hydration shell. (D) Insertion of the hydrated fiber into a 313.2 × 313.2 × 40 nm³ aqueous volume.

A representative long fiber and its transformed local hydration shell are shown in Figure 3.5.

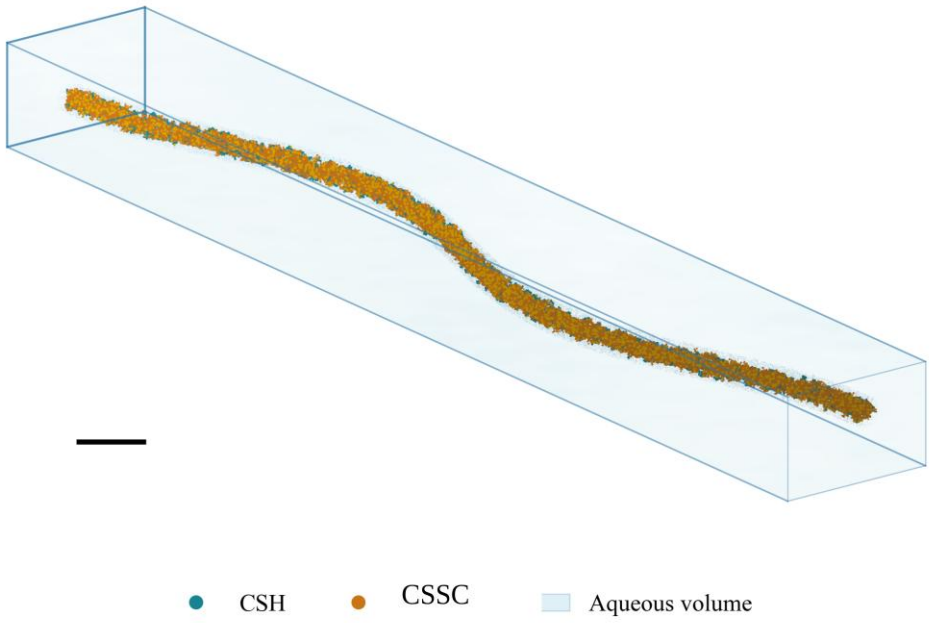

Figure 3.5. Three-dimensional representation of a representative 236.5 nm hydrated fiber comprising CSH and CSSC in the synthetic aqueous environment. Fiber heavy atoms and the transformed local hydration shell share a common projection within a continuous translucent aqueous volume. Scale bar, 20 nm.

Six scene classes were defined to span isolated fibers and progressively more complex projected arrangements: single fiber, two-fiber crossing, parallel overlap, loose bundle, crowded bundle, and mixed crossing/parallel. Their sampling priors were 0.08, 0.10, 0.18, 0.24, 0.24, and 0.16, respectively. Although the configuration allowed as many as ten fibers, the number of available hydrated sources limited a production scene to at most seven components.

Each component received an independent in-plane angle and a scene-conditioned two-dimensional position. Additional transformations included a z displacement sampled from −95 to +95 Å, an axial scale factor from 0.88 to 1.12, and a z-direction scale factor from 0.92 to 1.08. Across the 500 scenes, 2,321 transformed fiber components were used. The component-count distribution was 35 images with one fiber, 76 with two, 45 with three, 73 with four, 61 with five,

68 with six, and 142 with seven. The audited component-count distribution is summarized in Figure 3.6.

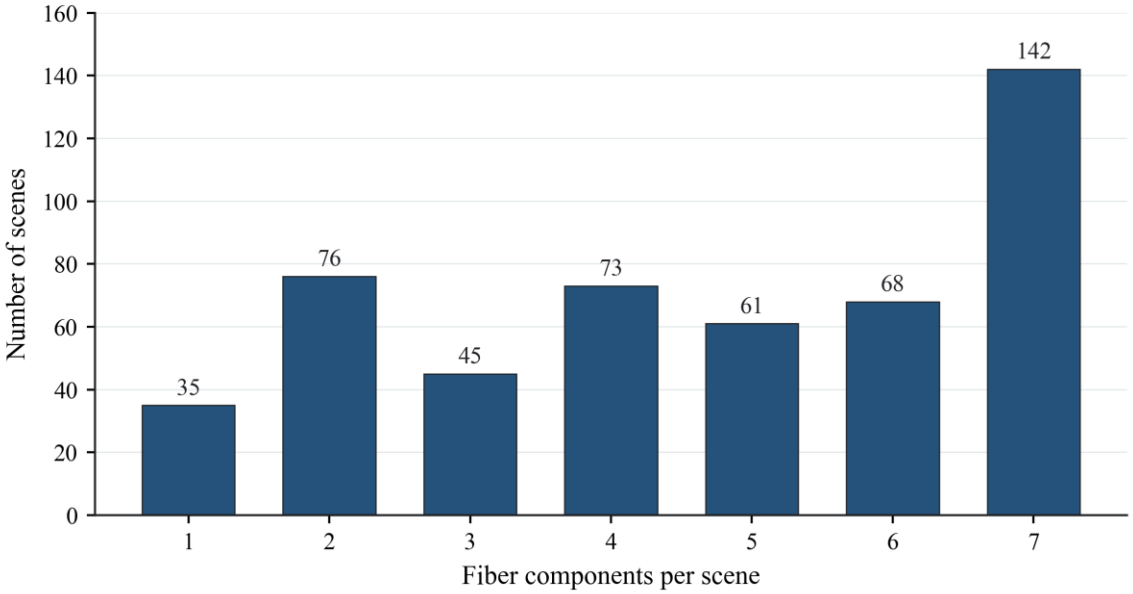

Figure 3.6. Distribution of fiber-component counts among the 500 synthetic scenes. Values above the bars indicate the number of scenes in each component-count category.

The multi-fiber scenes were embedded in a background-water reservoir assembled from a library of 400 equilibrated TIP4P/2005 water structures [8]. Each source box measured 150 × 150 × 150 Å³, was indexed in 15 × 15 × 5 Å³ spatial chunks, and contained approximately 1.1 × 10^5 water molecules.

Background water was placed in 40 Å slabs using randomized source boxes, source z origins, and fourfold rotations about the z axis, together with a 45° planar grid and ABC-type offsets in the xy plane. The complete atomic scene occupied 313.2 × 313.2 × 40 nm³. Water molecules within 2.6 Å of a solute heavy atom were removed, water oxygens closer than 2.25 Å were deduplicated, and the remaining coordinates were wrapped periodically in xy.

Two water models served distinct roles. The local hydration shell transformed with each fiber originated from the CHARMM TIP3P molecular-dynamics system, whereas the bulk

background bath originated from the TIP4P/2005 reservoir. In the three-dimensional overview figures, the surrounding bath is rendered as a continuous translucent aqueous volume, while the fiber and local hydration shell are shown from retained atomic coordinates.

Atomic scattering potentials were represented using the Lobato parameterization [9], and multislice propagation was implemented in abTEM [10]. Study-specific simulation settings were a 1024 × 1024 computational grid spanning 3132 Å in each lateral direction (3.05859375 Å per pixel), a 40 nm specimen thickness, a 1 Å slice thickness (400 slices), a 200 keV incident electron energy, complex64 numerical precision, and periodic boundaries in xy.

The contrast-transfer and detector configuration used spherical- and chromatic-aberration coefficients of 2 and 2.1 mm, respectively, an energy spread of 0.8 eV, five Gaussian focal-spread integration samples, a 10.7 mrad objective aperture, a 0.01 mrad aperture rolloff, and zero amplitude contrast. The dose was 5,000 e^{-} $Å^{-2}$. Detective quantum efficiency was set to 1, the modulation-transfer-function Gaussian width to 0 pixels, and readout noise to 0 electrons; Poisson electron-counting noise was nevertheless applied.

The underfocus magnitude was sampled independently for each image from candidate values between 5,000 and 18,000 nm in 10 nm increments; the realized dataset spanned 5,000–17,970 nm. The final synthetic PNG was generated by clipping the first and ninety-ninth intensity percentiles and linearly scaling the result to 8-bit grayscale. This encoding was separate from the per-image normalization used during model preprocessing. The complete synthetic image-formation and detector parameters are summarized in Table 3.3.

Table 3.3. Parameters for synthetic TEM image formation and detector modeling.

Parameter	Value
Image grid	1024 × 1024 pixels; 3.05859375 Å per pixel
Scene dimensions	313.2 × 313.2 × 40 nm³
Electron energy	200 keV
Multislice	1 Å slice thickness; 400 slices; periodic xy boundaries
Atomic potential	Lobato parameterization
Aberrations	C_s = 2 mm; C_c = 2.1 mm; energy spread = 0.8 eV
Aperture	10.7 mrad objective aperture; 0.01 mrad rolloff
Dose	5,000 e^- $Å^{-2}$
Detector	DQE = 1; MTF σ = 0 px; readout noise = 0 e^-; Poisson counting noise applied
Underfocus magnitude	Candidate range 5,000–18,000 nm in 10 nm increments; realized 5,000–17,970 nm
PNG encoding	First-to-ninety-ninth percentile clipping; 8-bit grayscale

The synthetic TEM image and reference mask originated from the same transformed atomic scene but were generated through separate operations. The image was simulated from multislice simulation, whereas the reference was produced from the projected fiber-solute heavy atoms; all water atoms were excluded.

Projected coordinates were converted to a binary fiber 2D image using a fixed geometric procedure designed to approximate the approximately 5.4 nm single-fiber diameter measured by cryo-EM [3, 4]. Each component was processed separately, and the component masks were combined by union.

The resulting target represented foreground semantic segmentation: each pixel encoded fiber or background, but the identity of individual fibers was not retained at crossings or within bundles. Because the boundary was defined computationally, it served as a reproducible synthetic reference rather than a claim that the same edge is directly visible in every experimental TEM image.

The final dataset comprised 500 complete samples. Each sample included a synthetic TEM image, a binary coordinate-defined foreground reference, an atomic-layout record, and an exit wave.

The dataset contained 132 crowded-bundle, 103 loose-bundle, 95 mixed crossing/parallel, 79 parallel-overlap, 56 two-fiber-crossing, and 35 single-fiber images. Images and reference masks were 1024 × 1024 pixels. Basic file checks verified complete image–reference pairing and binary masks containing both foreground and background. The foreground fraction ranged from 0.005854 to 0.061865 (median, 0.041008). Figure 3.7 shows the scene-category distribution, and Table 3.4 summarizes dataset composition and partitioning.

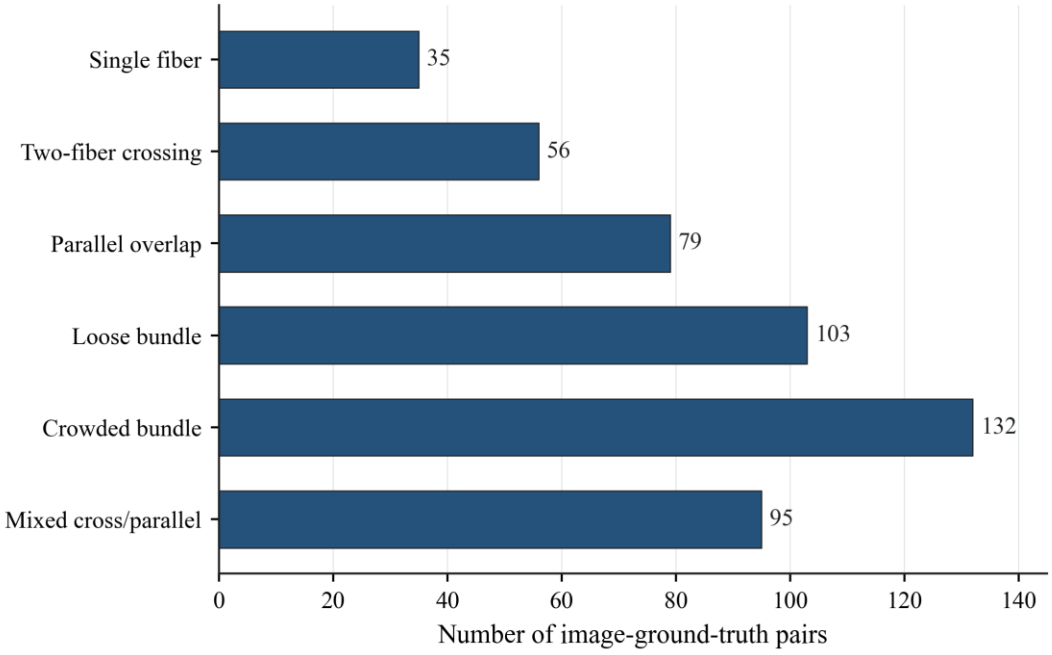

Figure 3.7. Distribution of the 500 synthetic scenes among six geometric categories. Values indicate the number of image–reference pairs in each category.

Table 3.4. Composition and partitioning of the synthetic dataset.

Item	Value
Total image–reference pairs	500
Image and reference-mask dimensions	1024 × 1024 pixels
Training / validation / test	348 / 74 / 78
Split method	Scenario-stratified; seed 42
Scene counts	Crowded 132; loose 103; mixed 95; parallel 79; two-cross 56; single 35
Fiber components	2,321 total; 1–7 per image
Foreground fraction	Minimum 0.005854; median 0.041008; maximum 0.061865

Using a fixed scenario-stratified split with random seed 42, the 500 pairs were divided into 348 training, 74 validation, and 78 synthetic test images. The validation set was used for model and baseline selection. The test set was not used for weight or parameter optimization; because its summaries were examined during chapter preparation, it is reported as a retrospective held-out evaluation rather than a prospective external benchmark.

Figure 3.8 shows one representative training example from each of the six scene categories.

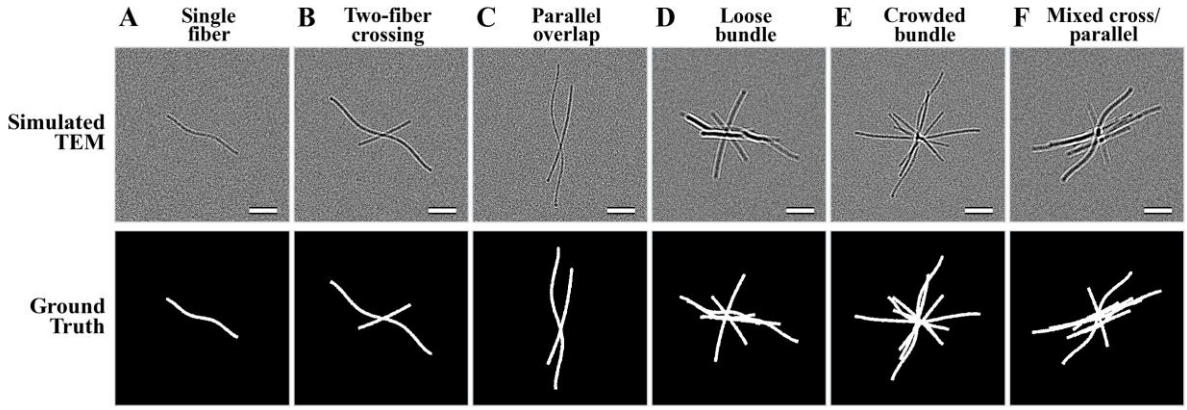

Figure 3.8. Representative training-set pairs from the six synthetic scene categories. Columns show single fiber, two-fiber crossing, parallel overlap, loose bundle, crowded bundle, and mixed cross/parallel; the upper row contains simulated TEM images and the lower row contains the corresponding coordinate-defined foreground reference masks. Scale bars, 50 nm.

Each TEM image was converted to single-channel float32. An 8-bit input was first divided by 255. For each image independently, the first and ninety-ninth percentiles were used for linear rescaling, followed by clipping to [0, 1]. Masks were converted to grayscale and binarized at values greater than 127. Training used full 1024 × 1024 images rather than patches. This choice avoided introducing artificial patch boundaries through long fibers, crossings, and bundles, although it constrained the batch size to one.

Geometric augmentation consisted of horizontal and vertical flips and random rotations of 0°, 90°, 180°, or 270°. These transformations were applied identically to the image and mask. Image-only intensity augmentation included a contrast factor sampled from 0.9 to 1.1, a gamma value sampled from 0.85 to 1.15, and Gaussian noise with $\sigma = 0.015$ applied with probability 0.5. Validation and test images were not augmented.

The model followed a U-Net encoder–decoder architecture with skip connections [11]. Study-specific modifications included GroupNorm [12], SiLU activation [13], bilinear

upsampling in the decoder [14], and optional contextual dilation [15]. The network accepted one grayscale channel and produced one logit channel. Each convolutional block contained two or three repetitions of a 3 × 3 convolution, GroupNorm, and SiLU. The encoder used max-pooling. The decoder used bilinear upsampling followed here by a 1 × 1 convolution, concatenation with the corresponding encoder feature map, and another convolutional block. Candidate networks used three or four downsampling levels, base channel counts of 16, 32, or 48, and optional contextual dilation in the bottleneck, deep encoder, or multibranch module. Dropout was not used. Sigmoid conversion was applied during loss calculation and evaluation rather than stored as the final model layer.

AdamW was used for all training runs [16]. The study-specific settings were an initial learning rate of 0.001, weight decay of 0.0001, batch size and gradient accumulation of one, and automatic mixed precision. Training was capped at 200 epochs. A ReduceLROnPlateau scheduler monitored validation loss with a factor of 0.5, patience of 10, and minimum learning rate of 10^{-6}. Early stopping occurred when the composite validation score failed to improve for 30 epochs. Training used CUDA acceleration.

Each epoch comprised 120 randomly ordered single-image updates rather than a complete pass through all 348 training images. Candidate models began from new random initializations, and configuration identifiers recorded inherited design choices rather than transferred weights.

Model development was conducted in four sequential stages comprising 47 runs: architecture screening, context and dilation screening, loss-function comparison, and seed replication for three finalist configurations.

The same 74-image validation split and candidate-ranking procedure were reused throughout the four stages. Table 3.5 and Figure 3.9 therefore document one adaptive model-selection process rather than independent estimates of generalization.

Table 3.5. Summary of the four-stage adaptive model-development sequence leading to the selected U-Net.

Stage	Screening role	Runs	ML model lineage composite Validation score
1	Architecture	12	0.959960
2	Context / dilation	12	0.960112
3	Loss function	14	0.960249
4	Three finalist configurations × three seeds	9	0.960146
Total	Sequential adaptive development	47	Selection record only

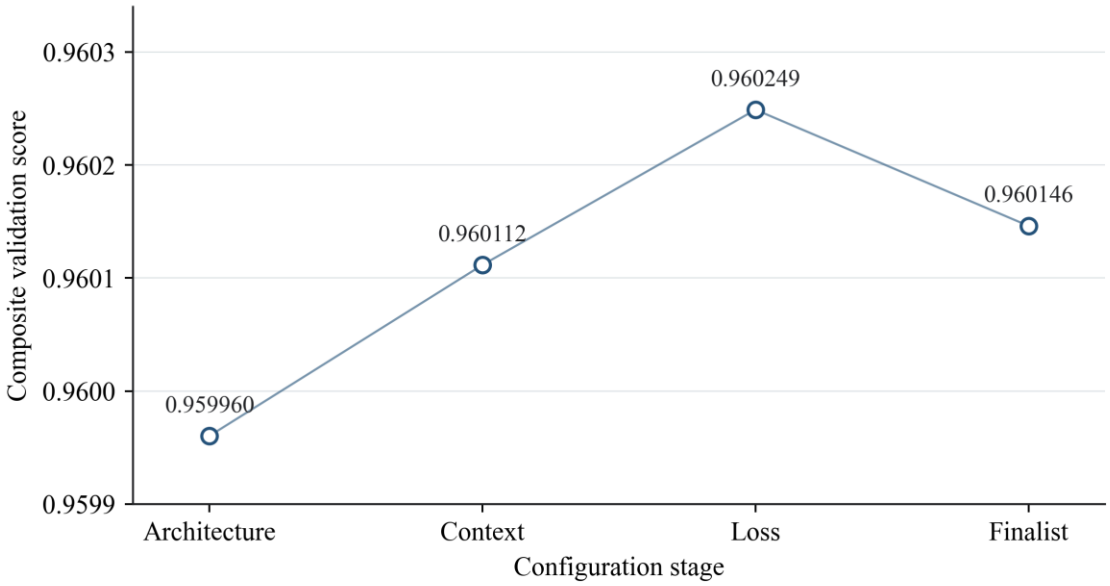

Figure 3.9. Composite validation scores during the four-stage internal model-development sequence. Values document candidate selection and are not independent test estimates.

Binary predictions were evaluated primarily with Dice similarity and were further characterized by intersection over union (IoU), precision, recall, and centerline Dice (clDice).

Let P denote predicted foreground, G the reference foreground, and S(·) the corresponding skeleton. The area-based metrics were:

$$Dice = 2|P \cap G| / (|P| + |G|)$$

$$IoU = |P \cap G| / |P \cup G|$$

$$Precision = |P \cap G| / |P|$$

$$Recall = |P \cap G| / |G|$$

For clDice, topological precision and sensitivity were Tprec = |S(P) ∩ G|/|S(P)| and Tsens = |S(G) ∩ P|/|S(G)|, respectively, and their harmonic mean defined centerline agreement [17]:

$$clDice = 2(Tprec)(Tsens) / (Tprec + Tsens)$$

For internal candidate ranking only, the metrics were combined as follows:

$$Composite = 0.25\ Dice + 0.25\ clDice + 0.25\ Recall + 0.15\ IoU + 0.10\ Precision$$

Dice was treated as the principal overlap metric when interpreting the final model; the composite score provided a consistent selection criterion across model and baseline tuning. Decision thresholds were scanned from 0.10 to 0.70 in increments of 0.05. Metrics were calculated per image and macro-averaged. Raw validation losses were not compared across different loss-function families because their numerical scales were not commensurate.

Checkpoint selection during training used a fixed 32-image subset of the validation set. Final checkpoint-level summaries were then calculated with the saved model across the full 74-image validation set and the 78-image synthetic test set at the selected threshold.

All model-specific results use one selected Stage 4 checkpoint. The network had four downsampling levels, two convolutions per block, 16 base channels, dilation 4 in the bottleneck, binary cross-entropy loss, and 1,811,729 trainable parameters. The saved epoch-119 checkpoint was applied at threshold 0.40 for validation, synthetic-test evaluation, and real-TEM inference.

Otsu global thresholding [18] and Sauvola local adaptive thresholding [19] were used as comparison baselines. The complete preprocessing, threshold offsets, smoothing scales, object filters, and hole-filling parameters were selected on the fixed 74-image validation split and were then frozen for the synthetic test set and real-TEM examples. No real-TEM parameter tuning was performed.

Using a reference fiber-width scale of $w = 18.4$ pixels, approximately consistent with the 5.4 nm cryo-EM diameter at the synthetic image scale, the Otsu-based pipeline used dark foreground, Gaussian smoothing $\sigma = 0.6w$, background estimation $\sigma = 2.0w$, zero threshold offset, minimum-object area $0.5w^2$, and maximum-hole area $1.0w^2$. The Sauvola-based pipeline used dark foreground, a 55-pixel window, $k = 0.90$, threshold offset 0.25, no minimum-object filter, and a maximum-hole area of 85 pixels.

Metrics were calculated per image and macro-averaged. Paired differences between methods were summarized with 10,000 image-level percentile-bootstrap resamples using seed 2026080817.

To examine application to experimental data, the selected U-Net was applied to 45 unlabeled 4096 × 4096 real-TEM images. The development cohort contained 27 images: 24 images from eight folders at 0.3058 nm pixel^{-1} and three images from one folder at 0.6116 nm pixel^{-1}. A frozen application cohort contained 18 additional images: 16 from the same eight primary-scale folder groups and two from the coarser-scale folder. The two pixel-size strata were analyzed separately.

All images were clipped at the first and ninety-ninth intensity percentiles and scaled to [0, 1]. U-Net inference used reflect padding, overlapping 1024 × 1024-pixel tiles at a stride of 512 pixels, and cosine-weighted probability fusion before thresholding at 0.40. The Otsu- and

Sauvola-based pipelines used the parameters selected on synthetic validation data. No illumination correction or real-TEM fine-tuning was applied.

Because pixel-level real-TEM annotations were unavailable, transfer was assessed qualitatively and supplemented by a U-Net tile-origin stability check. Four primary-scale development images spanning variation in fiber density, intersections, and contrast were selected for visualization with all three methods.

Results and Discussion

The production pipeline yielded 500 complete synthetic TEM image–reference pairs at 1024 × 1024 pixels. The dataset covered all six planned scene categories and contained one to seven fiber components per image.

Crowded and loose bundles together accounted for 235 images, whereas 35 images contained a single fiber. Aggregate metrics are therefore weighted toward complex multi-fiber scenes unless stratified by category. The coordinate-defined foreground occupied a median of 4.10% of image pixels, with a range from 0.585% to 6.19%. This foreground–background imbalance motivated the use of precision, recall, and centerline-based evaluation in addition to overlap metrics.

Across the four sequential development stages, 47 runs compared architecture, contextual dilation, loss formulation, and seed sensitivity. The stage-wise composite validation scores along the selected configuration lineage were 0.959960, 0.960112, 0.960249, and 0.960146.

The differences among these values were small. Because the same validation set was reused throughout adaptive screening, the sequence is interpreted as an internal selection record rather than four independent performance estimates.

All subsequent analyses use the selected U-Net checkpoint. Its full-validation composite score was 0.9626 at threshold 0.40.

On the 78-image synthetic test set, the selected U-Net achieved Dice 0.9547, IoU 0.9134, precision 0.9403, recall 0.9696, clDice 0.9969, and composite score 0.9613 at threshold 0.40. The complete metric summary is reported in Table 3.6.

Table 3.6. Descriptive performance of the selected U-Net on the fixed synthetic test set (n = 78; threshold = 0.40).

Metric	Mean
Composite score	0.9613
Dice	0.9547
IoU	0.9134
Precision	0.9403
Recall	0.9696
clDice	0.9969

These values demonstrate close agreement with the coordinate-defined foreground reference within the current synthetic distribution. They do not evaluate separation of individual fibers, generalization to an independent structural generator, or quantitative accuracy on real TEM images.

The synthetic test set was excluded from weight and baseline-parameter optimization. Because its summaries were reviewed during final chapter preparation, the results are interpreted as retrospective within-generator evaluations.

The selected U-Net outperformed both validation-optimized thresholding pipelines on every reported synthetic-test metric (Table 3.7). Dice was 0.9547 for the U-Net, 0.6389 for Otsu, and 0.3171 for Sauvola. The corresponding composite scores were 0.9613, 0.6763, and 0.3681.

Table 3.7. Image-macro performance of the selected U-Net and two validation-optimized thresholding baselines on the fixed synthetic test set (n = 78). The U-Net threshold was 0.40; baseline configurations were frozen after validation tuning.

Method	Dice	IoU	Precision	Recall	clDice	Composite score
ML model	0.95469	0.91340	0.94032	0.96962	0.99686	0.96134
Otsu	0.63889	0.51370	0.83569	0.65951	0.76432	0.67630
Sauvola	0.31707	0.19093	0.87710	0.19643	0.49358	0.36812

Relative to Otsu and Sauvola, the mean paired Dice differences were +0.3158 (95% percentile-bootstrap interval, +0.2595 to +0.3747) and +0.6376 (+0.6199 to +0.6546), respectively. Image-level bootstrap intervals favored the selected U-Net for every reported metric within this fixed synthetic test set. The intervals do not include training-seed, model-selection, source-trajectory, or generator uncertainty.

Illustrative method-specific semi-transparent mask overlays are shown in Figure 3.10.

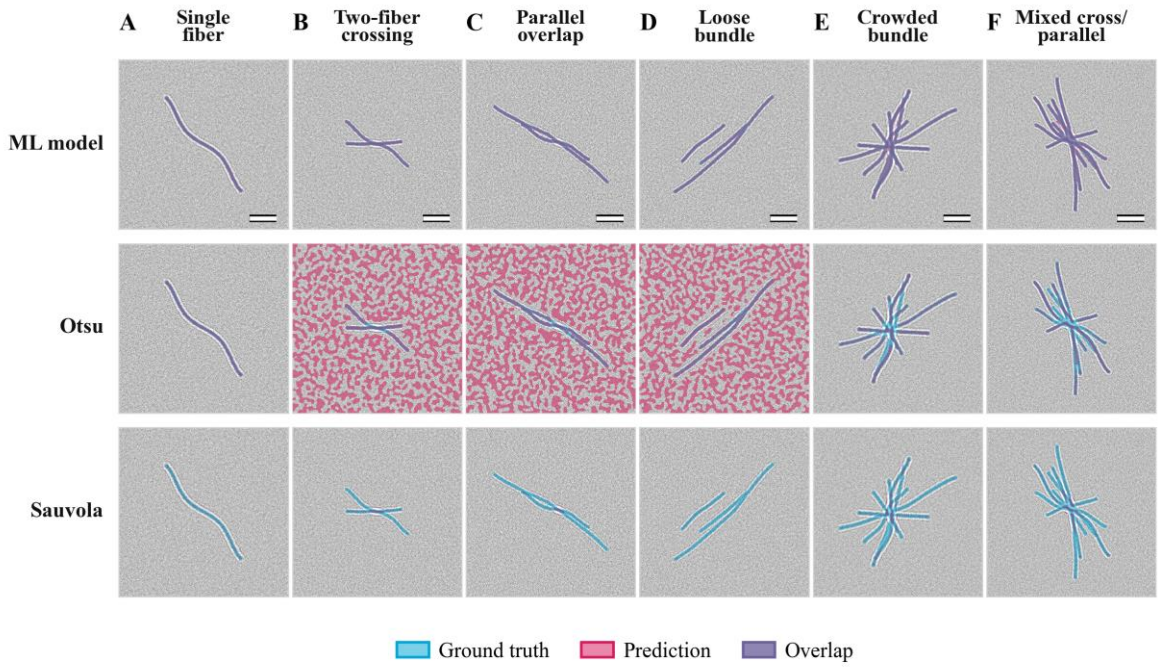

Figure 3.10. Semi-transparent mask-overlay comparison across the six synthetic fiber-scene categories. Columns show single fiber, two-fiber crossing, parallel overlap, loose bundle, crowded bundle, and mixed cross/parallel; rows show the ML model, Otsu, and Sauvola. Within each column, the simulated TEM field, coordinate-defined reference mask, field of view, intensity scaling, and overlay opacity are fixed; only the method prediction changes. Cyan marks the reference, magenta marks the prediction, and violet marks their pixelwise overlap. Scale bars, 50 nm.

For completeness, Table 3.8 and Figure 3.11 summarize the composite selection score on the validation and synthetic test sets. Validation was used for model or baseline selection, so the split-to-split differences are reported descriptively rather than interpreted as independent generalization gaps.

Table 3.8. Image-macro composite-score means and 95% image-bootstrap confidence intervals for the selected U-Net and two validation-optimized thresholding baselines on the 74-image validation split and retrospective 78-image synthetic test split. Values are mean [95% CI].

Method	Validation mean [95% CI] (n = 74)	Test mean [95% CI] (n = 78)
ML model	0.9626 [0.9613, 0.9637]	0.9613 [0.9598, 0.9628]
Otsu	0.7170 [0.6797, 0.7510]	0.6763 [0.6346, 0.7162]
Sauvola	0.3626 [0.3461, 0.3786]	0.3681 [0.3536, 0.3832]

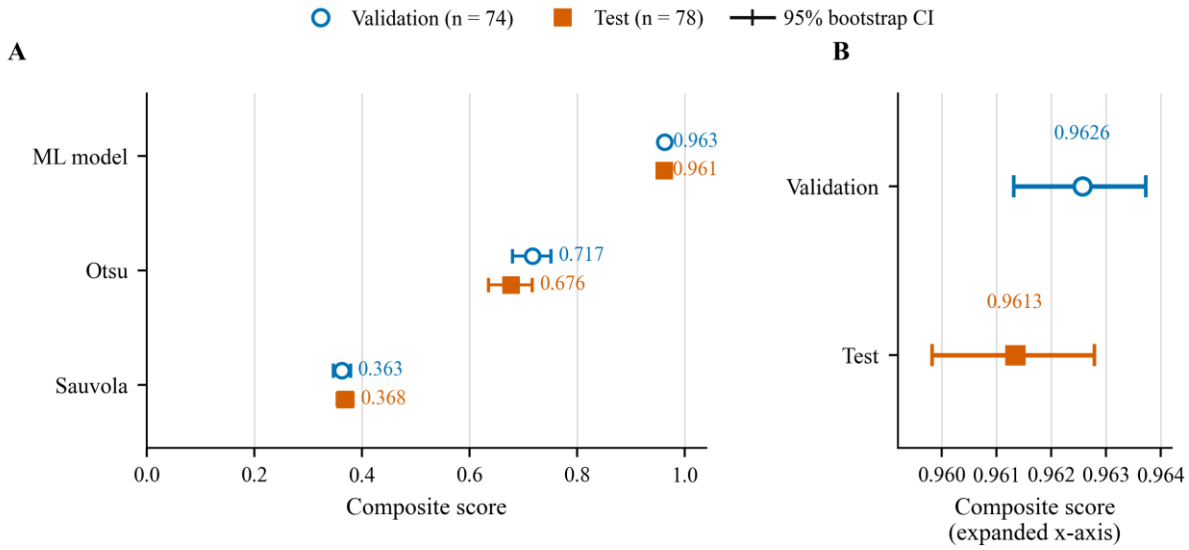

Figure 3.11. Composite-score means and 95% image-bootstrap confidence intervals. (A) Full-scale comparison of the ML model, Otsu, and Sauvola. (B) Expanded axis for the ML model showing the same Validation (n = 74) and retrospective Test (n = 78) estimates. Intervals use 10,000 image-level resamples, exclude training and selection variability, and do not support paired Validation-Test inference.

On 24 primary-scale development images, the overlapping-tile inference scheme reduced tile-origin probability sensitivity by 43.7% and border-band inconsistency by 43.0% relative to the original whole-image geometry. The frozen inference configuration was then applied once to

the 18-image application cohort. These measurements assess prediction consistency under changes in tiling and do not measure segmentation accuracy.

Across the representative real-TEM fields, the U-Net predictions followed many visually recognizable fiber-like features, including isolated segments, intersections, and dense networks (Figure 3.12). Otsu and Sauvola produced substantially different foreground extents. Because pixel-level annotations were unavailable, the comparison cannot quantify accuracy, sensitivity, specificity, or method ranking. The panels demonstrate qualitative image-level correspondence only.

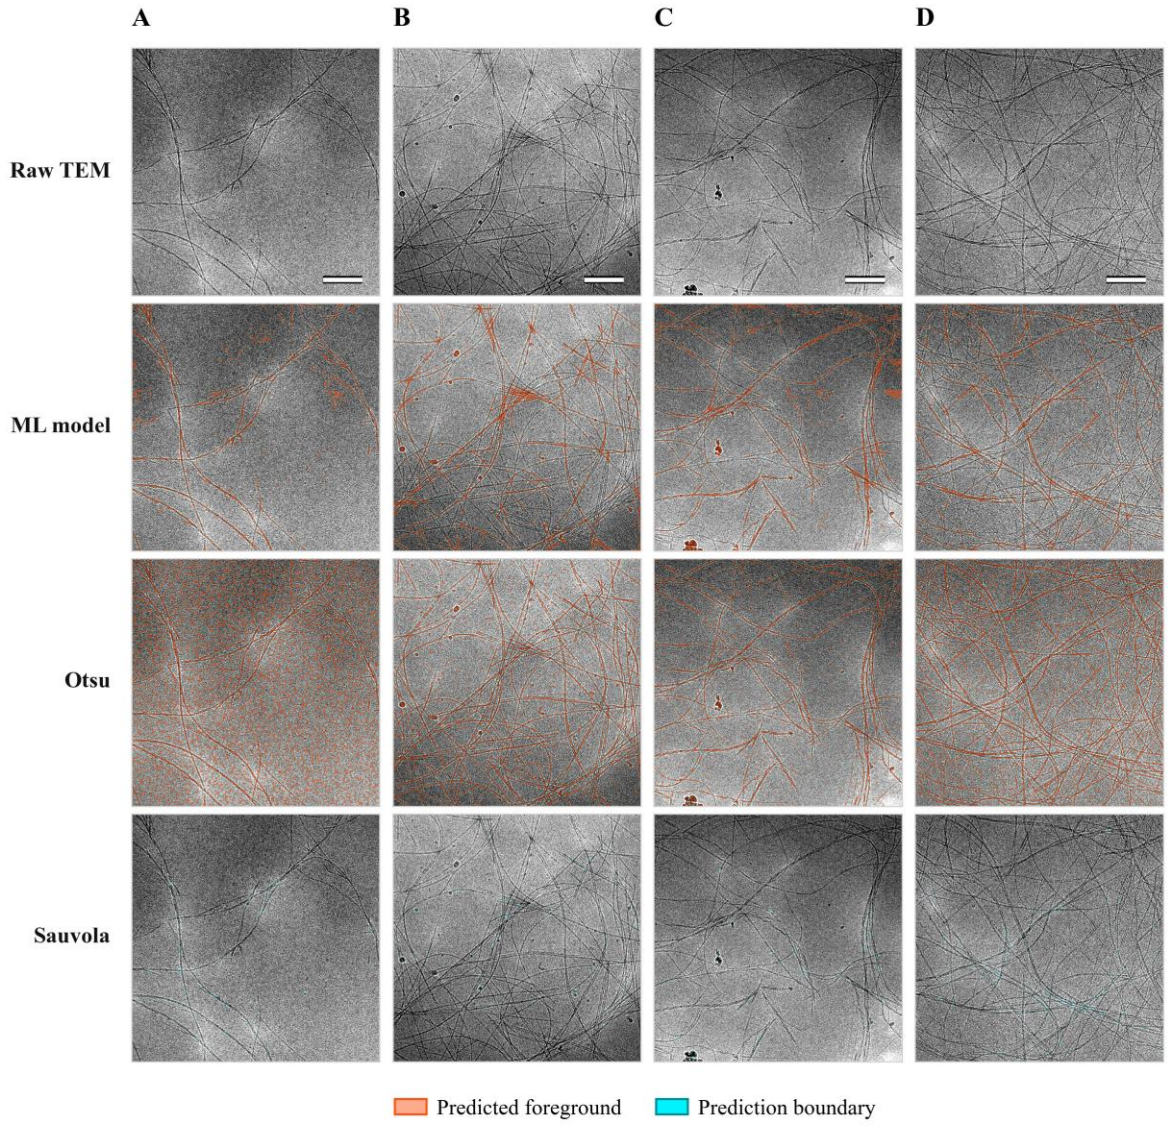

Figure 3.12. Qualitative comparison of whole-image foreground predictions on four representative primary-scale real-TEM fields. Columns A–D show the same fields; rows show raw TEM, the selected U-Net, Otsu, and Sauvola. Orange indicates predicted foreground and cyan delineates the prediction boundary. The U-Net used frozen overlapping-tile inference, and baseline parameters were transferred from synthetic validation without real-TEM tuning. The panels compare prediction morphology rather than quantitative accuracy. Scale bars, 200 nm.

The central result of this chapter is that a U-Net trained exclusively on synthetic, atomistically generated TEM image–reference pairs achieved high agreement on held-out images from the same generator-defined distribution. The workflow linked a cryo-EM-informed molecular model to long-fiber construction, aqueous scene assembly, multislice image

formation, and coordinate-defined foreground references. Application to unlabeled real TEM images produced qualitatively plausible fiber-like predictions but did not provide an experimental accuracy estimate.

The high clDice value indicates that the selected U-Net preserved the centerline geometry encoded by the synthetic reference, while Dice and IoU quantify agreement in projected foreground area. These metrics apply to the union foreground target and do not determine whether touching or crossing fibers were assigned separate identities.

Coordinate-derived references remove annotation ambiguity within the simulated task because the target is specified from atomic geometry rather than inferred from local TEM intensity. The exact boundary nevertheless depends on the fixed projection procedure and should be interpreted as a reproducible synthetic definition, not as a uniquely observable boundary in experimental TEM.

Within the current synthetic distribution, the selected U-Net agreed more closely with the reference masks than either thresholding pipeline. In representative real TEM images, its predictions traced many visually apparent fiber-like trajectories across fields with different densities and contrast levels. This correspondence supports further experimental evaluation but is not equivalent to validated synthetic-to-real transfer.

The four development stages reused the same validation set and therefore constitute one adaptive model-selection process rather than four independent tests. The synthetic test set was not used for optimization, but its results were reviewed before the final reporting choice, so the reported comparisons are retrospective.

Quantitative performance is confined to the current synthetic generator and its coordinate-defined foreground reference. The simulated images do not reproduce every source of

experimental variation, including the full range of specimen thickness, window scattering, contamination, drift, detector response, and beam-induced structural change.

The real-TEM images were unlabeled and clustered within acquisition folders. Consequently, Figure 3.12 cannot establish pixel-level accuracy or a general ranking among the U-Net, Otsu, and Sauvola. It shows only that the synthetic-only U-Net produced predictions that coincide with many visually recognizable fiber-like trajectories in the selected examples.

The task is foreground semantic segmentation and does not preserve individual-fiber identities at crossings or within bundles. In addition, the lower tile-origin stability observed at the coarser pixel size indicates sensitivity to scale mismatch. Experimental annotations and independently generated synthetic scenes will be needed to quantify transfer accuracy and robustness more rigorously.

Conclusions

This chapter established a physics-informed synthetic workflow for foreground semantic segmentation of low-contrast nanofibers. Cryo-EM-informed restraint geometry for the CSH–CSSC fiber system was implemented in GROMACS and combined with long-fiber construction, hydrated multi-fiber scene assembly, multislice TEM simulation, and coordinate-derived foreground references to generate 500 paired synthetic images and masks.

The selected U-Net was trained exclusively with synthetic data. On the 78-image synthetic test set, it achieved Dice 0.9547, IoU 0.9134, precision 0.9403, recall 0.9696, and clDice 0.9969 at threshold 0.40. Under the same synthetic evaluation, the validation-optimized Otsu- and Sauvola-based pipelines produced substantially lower Dice scores of 0.6389 and 0.3171, respectively.

Without experimental fine-tuning, the same U-Net produced predictions that followed many fiber-like structures in selected real TEM images and remained reasonably stable to tile origin at the matched pixel scale. These observations demonstrate qualitative plausibility and inference consistency, not quantitative experimental identification accuracy. External validation will require independently annotated images, source-grouped synthetic evaluation, broader acquisition conditions, and explicit assessment of scale robustness and individual-fiber instance separation.

References

1. de Jonge, N.; Ross, F. M. Electron Microscopy of Specimens in Liquid. Nature Nanotechnology 2011, 6, 695–704. https://doi.org/10.1038/nnano.2011.161.
2. Ross, F. M. Opportunities and Challenges in Liquid Cell Electron Microscopy. Science 2015, 350, aaa9886. https://doi.org/10.1126/science.aaa9886.
3. Song, Y.; Li, Z.; Mulvey, J. T.; Freites, J. A.; Patterson, J. P.; Tobias, D. J. Cryogenic Electron Microscopy Informed Molecular Dynamics Simulations to Investigate the Disulfide Hydrogel Self-Assembly. ChemPhysChem 2025, 26(11), e202401085. https://doi.org/10.1002/cphc.202401085.
4. Hurst, P. J.; Mulvey, J. T.; Bone, R. A.; Selmani, S.; Hudson, R. F.; Guan, Z.; Green, J. R.; Patterson, J. P. CryoEM Reveals the Complex Self-Assembly of a Chemically Driven Disulfide Hydrogel. Chemical Science 2024, 15(3), 1106–1116. https://doi.org/10.1039/D3SC05790A.
5. Jorgensen, W. L.; Chandrasekhar, J.; Madura, J. D.; Impey, R. W.; Klein, M. L. Comparison of Simple Potential Functions for Simulating Liquid Water. Journal of Chemical Physics 1983, 79, 926–935. https://doi.org/10.1063/1.445869.
6. Vanommeslaeghe, K.; Hatcher, E.; Acharya, C.; Kundu, S.; Zhong, S.; Shim, J.; Darian, E.; Guvench, O.; Lopes, P.; Vorobyov, I.; MacKerell, A. D., Jr. CHARMM General Force Field: A Force Field for Drug-Like Molecules Compatible with the CHARMM All-Atom Additive Biological Force Fields. Journal of Computational Chemistry 2010, 31, 671–690. https://doi.org/10.1002/jcc.21367.
7. Abraham, M. J.; Murtola, T.; Schulz, R.; Páll, S.; Smith, J. C.; Hess, B.; Lindahl, E. GROMACS: High Performance Molecular Simulations through Multi-Level Parallelism from Laptops to Supercomputers. SoftwareX 2015, 1–2, 19–25. https://doi.org/10.1016/j.softx.2015.06.001.
8. Abascal, J. L. F.; Vega, C. A General Purpose Model for the Condensed Phases of Water: TIP4P/2005. Journal of Chemical Physics 2005, 123, 234505. https://doi.org/10.1063/1.2121687.
9. Lobato, I.; Van Dyck, D. An Accurate Parameterization for Scattering Factors, Electron Densities and Electrostatic Potentials for Neutral Atoms That Obey All Physical Constraints. Acta Crystallographica Section A 2014, 70, 636–649. https://doi.org/10.1107/S205327331401643X.
10. Madsen, J.; Susi, T. The abTEM Code: Transmission Electron Microscopy from First Principles. Open Research Europe 2021, 1, 24. https://doi.org/10.12688/openreseurope.13015.2.
11. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. In Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015; Springer, 2015; pp 234–241. https://doi.org/10.1007/978-3-319-24574-4_28.

12. Wu, Y.; He, K. Group Normalization. In Computer Vision – ECCV 2018; Springer, 2018; pp 3–19. https://doi.org/10.1007/978-3-030-01261-8_1.
13. Elfwing, S.; Uchibe, E.; Doya, K. Sigmoid-Weighted Linear Units for Neural Network Function Approximation in Reinforcement Learning. Neural Networks 2018, 107, 3–11. https://doi.org/10.1016/j.neunet.2017.12.012.
14. Long, J.; Shelhamer, E.; Darrell, T. Fully Convolutional Networks for Semantic Segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2015; pp 3431–3440. https://doi.org/10.1109/CVPR.2015.7298965.
15. Yu, F.; Koltun, V. Multi-Scale Context Aggregation by Dilated Convolutions. In 4th International Conference on Learning Representations, ICLR 2016; 2016. https://arxiv.org/abs/1511.07122.
16. Loshchilov, I.; Hutter, F. Decoupled Weight Decay Regularization. In 7th International Conference on Learning Representations, ICLR 2019; 2019. https://openreview.net/forum?id=Bkg6RiCqY7.
17. Shit, S.; Paetzold, J. C.; Sekuboyina, A.; Ezhov, I.; Unger, A.; Zhylka, A.; Pluim, J. P. W.; Bauer, U.; Menze, B. H. clDice—A Novel Topology-Preserving Loss Function for Tubular Structure Segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition; 2021; pp 16560–16569.
18. Otsu, N. A Threshold Selection Method from Gray-Level Histograms. IEEE Transactions on Systems, Man, and Cybernetics 1979, 9(1), 62–66. https://doi.org/10.1109/TSMC.1979.4310076.
19. Sauvola, J.; Pietikäinen, M. Adaptive Document Image Binarization. Pattern Recognition 2000, 33(2), 225–236. https://doi.org/10.1016/S0031-3203(99)00055-2.

CHAPTER 4: TSGNet: Learned Intermediate-Projection Generation for Sparse Simulated Electron-Tomography Tilt Series

Manuscript Status and Author Contributions

An earlier version of this work, titled "Angular Sparsity Invariant Tilt Series Generation in Scanning/Transmission Electron Microscopy," was presented at the non-archival Machine Learning and the Physical Sciences workshop at NeurIPS 2025. At the time this chapter was prepared, the expanded manuscript underlying the present chapter had been submitted to Ultramicroscopy and had not been published as an archival journal article. The manuscript authors are Zhaoxu Li, Dikshant Sagar, Sherif Abdelkarim, Alexander M. Rakowski, Amin Tavakoli, Lingge Li, Jovany Merham, Pierre Baldi, and Joseph P. Patterson; Zhaoxu Li and Dikshant Sagar are identified as equal contributors.

The dissertation author's primary contributions to the work presented in this chapter were experimental design and three-dimensional reconstruction.

Introduction

Electron tomography (ET) reconstructs three-dimensional structures from two-dimensional electron microscopy projections acquired across specimen orientations. Reconstruction fidelity depends on image content and angular sampling: denser sampling generally improves reliability, whereas sparse sampling can introduce artifacts and loss of spatial detail. ET has been applied to heterogeneous catalysts, alloys, biomaterial interfaces, and nanoparticles [1–7].

Dense tilt-series acquisition is often impractical for beam-sensitive specimens and for time-resolved or in situ experiments in which particles evolve, rotate, or leave the field of view. These constraints motivate methods that recover more reconstruction information from fewer acquired projections [8, 9].

Deep convolutional networks have supported image restoration, frame interpolation, and related microscopy analysis, while learning-based tomography has addressed missing-wedge recovery, isotropic reconstruction, simultaneous denoising and missing-wedge reconstruction, and cryogenic ET tilt interpolation. However, the effect of generated intermediate projections on downstream surface and volume errors across multiple sparse angular-sampling conditions remains incompletely characterized [10–23].

This study introduces TSGNet, an encoder–decoder network that predicts an intermediate tilt image from two neighboring projections. Using simulated scanning transmission electron microscopy (STEM) tilt series under two imaging configurations labeled non-aberrated and aberrated in the source manuscript, the study evaluates image-level fidelity and whether generated projections improve three-dimensional reconstruction relative to the corresponding sparse Baseline series. Because the configurations also differ in convergence semi-angle, their comparison does not isolate a single aberration parameter.

Materials and Methods

The computational study comprised nanoparticle generation, dense STEM simulation, sparse input–target construction, TSGNet training and image-level evaluation, and reconstruction-based comparison of sparse and TSGNet-augmented tilt series under two simulated imaging conditions.

The training set contained 101 nanoparticles with icosahedral, decahedral, or octahedral shapes and elemental compositions of Pt, Ag, Au, or Cu. The test set contained 99 face-centered-cubic structures composed of Yb, Ir, Xe, or Al. Structures were generated with the Atomic Simulation Environment (ASE), and each structure was randomly rotated about the x, y, and z axes before image simulation [24].

High-angle annular dark-field STEM projections were simulated using PyQSTEM at an accelerating voltage of 200 kV and a detector collection semi-angle range of 54–175 mrad. The configuration labeled non-aberrated used zero defocus and a 35 mrad probe-convergence semi-angle. The configuration labeled aberrated used a spherical aberration of 1000 μm, zero defocus, and a 25 mrad convergence semi-angle [25]. These labels are retained for consistency with the source manuscript, but more than one optical parameter differed between the configurations.

For each structure and imaging condition, projections were simulated from −90° to +90° about the y axis at 1° increments, producing 181 images per tilt series. This yielded 18,281 training images and 17,919 test images per imaging condition.

Sparse stacks were produced from dense 512 × 512 × 181 tilt stacks by selecting frames at a fixed index gap g, where g ranged from 1 through 12, and an offset o, where o ranged from 0 through g−1. For a given gap and offset, retained frame indices satisfied $i \equiv o \pmod{g}$, enabling multiple sampling patterns from the same dense series. Each training example used two input projections, $f\alpha$ and $f\beta$, and the exact simulated intermediate projection, $f\gamma$, with $\theta\alpha < \theta\gamma < \theta\beta$; TSGNet predicted $\hat{f}\gamma$ from the two neighboring inputs.

TSGNet used a shared pyramid encoder to extract four levels of multi-scale features from two input frames. Four coarse-to-fine decoders combined these features and refined intermediate flow fields through residual blocks and transposed convolutions to generate an intermediate projection (Figure 4.1).

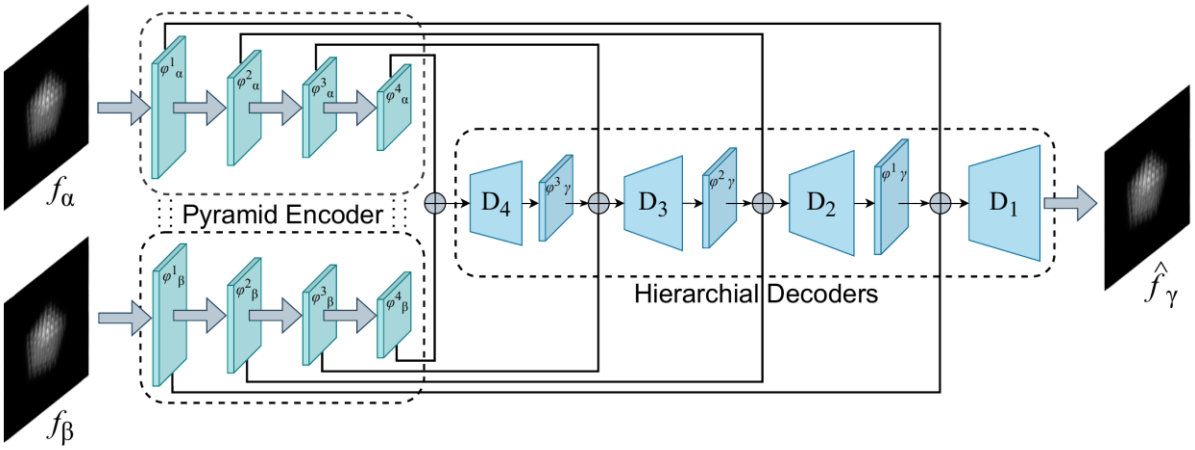

Figure 4.1. Overview of TSGNet for intermediate tilt-frame prediction. A shared pyramid encoder extracts multi-scale features from the two input frames, and four hierarchical decoders produce the predicted intermediate frame.

Training used 50 epochs on two NVIDIA A5000 GPUs with AdamW (weight decay 0.02), batch size 24, and a cosine-annealed learning rate from 10^{-4} to 10^{-5}. The objective combined a Charbonnier penalty ($\alpha = 0.5$, $\varepsilon = 10^{-3}$), a census loss on 7×7 patches, and a geometry-consistency loss across pyramid levels $k = 1–3$ with 3×3 patches; the geometry term was weighted by $\eta = 0.01$ [26–29].

TSGNet was compared with linear interpolation using mean-squared error (MSE), peak signal-to-noise ratio (PSNR), and structural similarity index (SSIM) between predicted and exact simulated intermediate projections. Lower MSE and higher PSNR and SSIM indicate closer image-level agreement; these metrics do not by themselves establish tomographic utility [30].

Reconstructions were performed in Python with TomoPy using the simultaneous iterative reconstruction technique (SIRT) for 100 iterations. The full 181-projection series served as the dense-series reconstruction reference. The Baseline retained projections at a specified nominal tilt increment. TSGNet-AUG used the same retained projections and inserted one predicted midpoint projection between adjacent acquired images, halving the effective angular increment.

The center of rotation was fixed at the image midpoint, and intensities were normalized independently for each projection [31–33].

Table 4.1. Tilt-sampling schedule for Baseline and TSGNet-AUG reconstructions.

Nominal increment (°)	TSGNet-AUG effective increment (°)	Retained Baseline projections
2	1	91
4	2	46
6	3	31
8	4	23
10	5	19
12	6	16
14	7	13
16	8	12
18	9	11
20	10	10

Reconstructed volumes were lightly smoothed with a Gaussian filter and segmented into object and background with global Otsu thresholding; surfaces were extracted with marching cubes. Reconstruction quality was evaluated against the dense-series reconstruction reference within the same imaging configuration using symmetric surface Chamfer distance and volumetric RMS error over the segmented object region. For each volume, the surface analysis used 1,000,000 uniformly sampled points, and Chamfer distances were reported in voxels [34–37].

Results

Across the reported test-set averages and evaluated frame gaps, TSGNet had lower MSE and higher PSNR and SSIM than linear interpolation under both simulated configurations. For

the aberrated configuration, TSGNet versus linear interpolation yielded MSE 1.51×10^{-3} versus 1.96×10^{-3}, PSNR 31.63 versus 30.48 dB, and SSIM 0.93 versus 0.90. For the non-aberrated configuration, the corresponding values were 1.87×10^{-3} versus 2.96×10^{-3}, 29.43 versus 28.57 dB, and 0.89 versus 0.86. Gap-dependent results are shown in Figure 4.2, and the reported dataset-level averages are listed in Table 4.2.

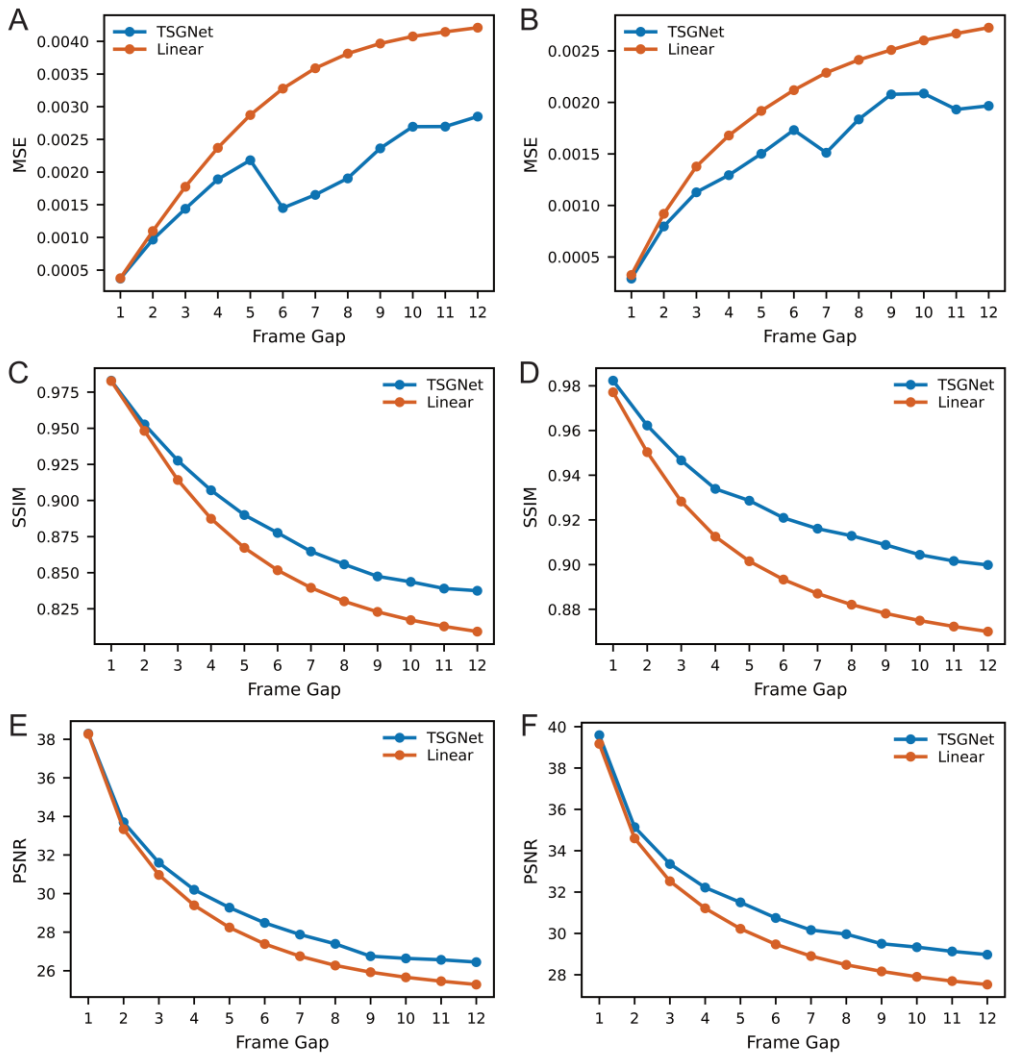

Figure 4.2. Quantitative comparison of TSGNet and linear interpolation across frame gaps 1–12. (A, C, E) Non-aberrated test set: MSE, SSIM, and PSNR, respectively. (B, D, F) Aberrated test set: MSE, SSIM, and PSNR, respectively.

Table 4.2. Reported intermediate-projection metrics for linear interpolation and TSGNet under the two simulated imaging configurations. Lower MSE and higher PSNR and SSIM indicate closer agreement with the exact simulated intermediate projections.

Imaging condition	Method	MSE	PSNR (dB)	SSIM
Aberrated	Linear	1.96×10^{-3}	30.48	0.90
Aberrated	TSGNet	1.51×10^{-3}	31.63	0.93
Non-aberrated	Linear	2.96×10^{-3}	28.57	0.86
Non-aberrated	TSGNet	1.87×10^{-3}	29.43	0.89

In the displayed example, both Baseline and TSGNet-AUG departed progressively from the dense-series reconstruction reference as the nominal tilt increment increased from 2° to 10°. Within each imaging configuration, TSGNet-AUG exhibited fewer visible external artifacts and retained a more compact reconstructed surface than the matched sparse Baseline. This observation is limited to the single simulated structure shown in Figure 4.3 and is not a population-level inference.

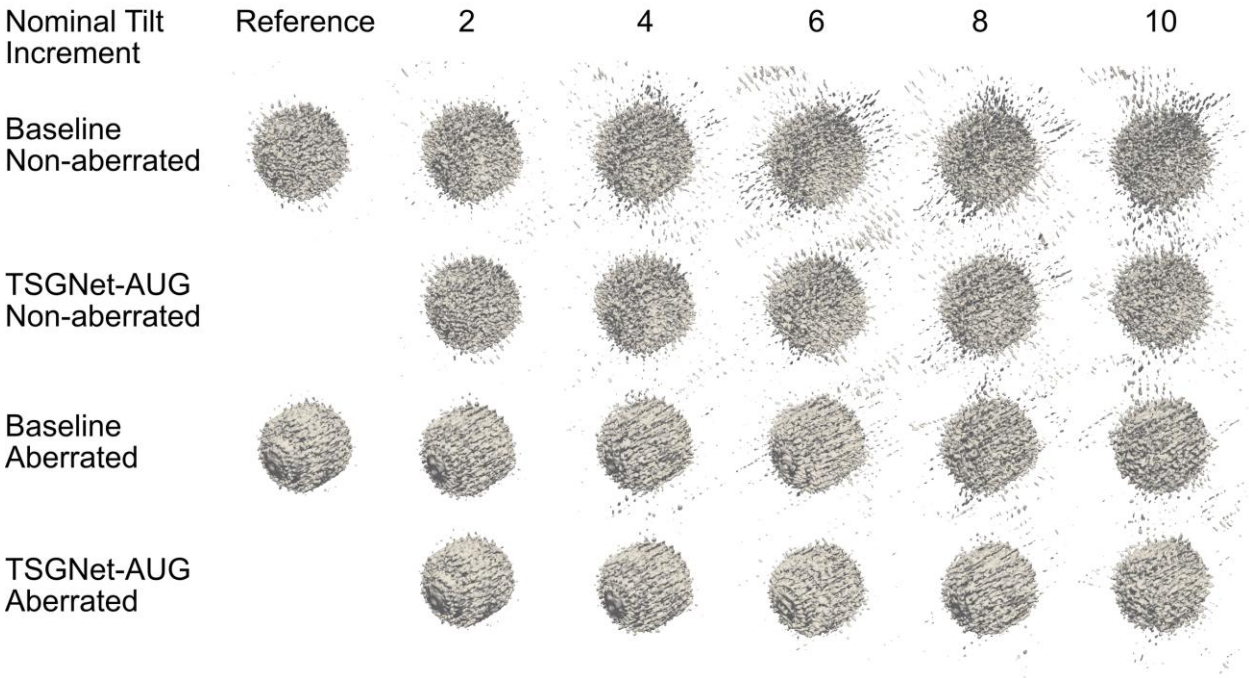

Figure 4.3. Example three-dimensional reconstructions of the same simulated structure. Columns show the dense-series reconstruction reference and nominal Baseline increments of 2°, 4°, 6°, 8°, and 10°; rows compare Baseline and TSGNet-AUG under the non-aberrated and aberrated configurations. Qualitative interpretation is limited to this displayed structure.

At every reported nominal increment from 2° through 20°, the tabulated improvements were positive for both Chamfer distance and RMS error under both simulated configurations,

indicating lower reported error for TSGNet-AUG than for the corresponding Baseline. The Baseline and TSGNet-AUG curves and the condition-specific improvement values are shown in Figures 4.4–4.6 and Tables 4.3–4.4.

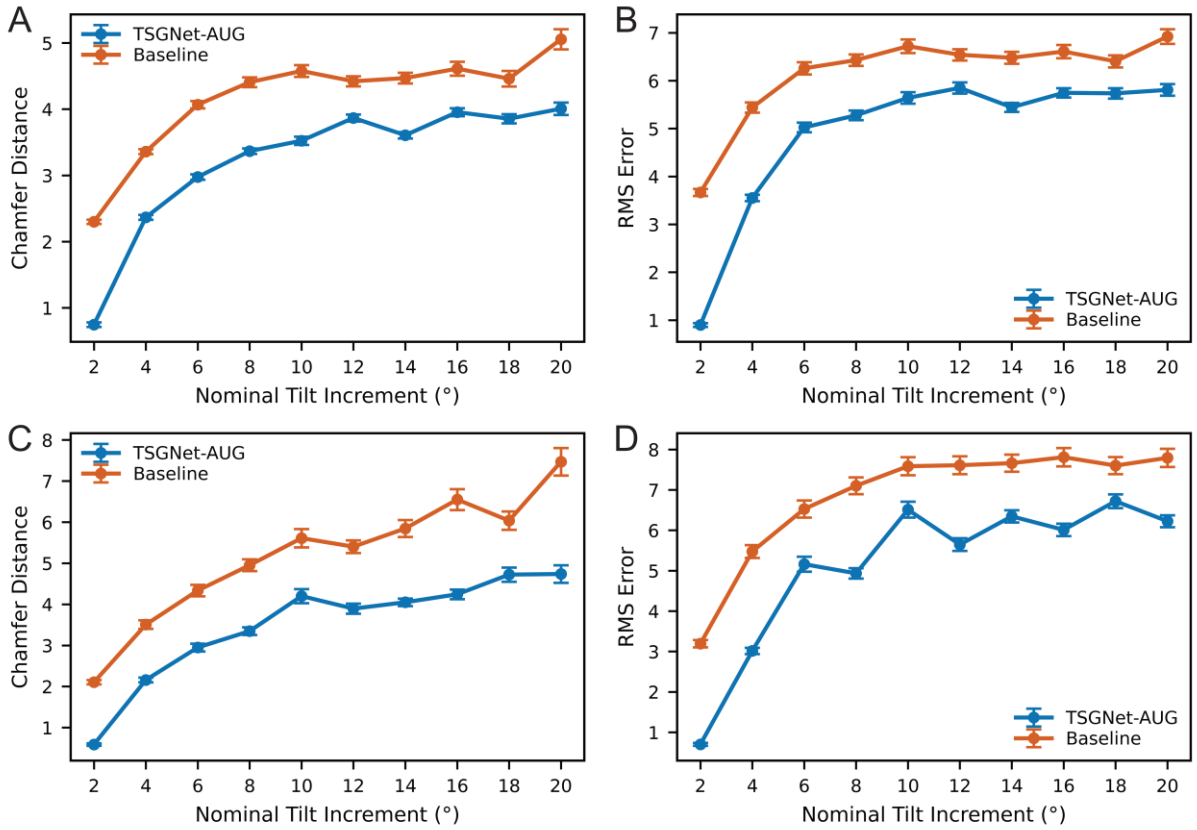

Figure 4.4. Chamfer distance and RMS error as functions of nominal tilt increment from 2° to 20° for Baseline and TSGNet-AUG. (A, B) Non-aberrated configuration: Chamfer distance and RMS error, respectively. (C, D) Aberrated configuration: Chamfer distance and RMS error, respectively. Error bars are reproduced from the source analysis; their statistical definition was not documented in the materials available for this dissertation.

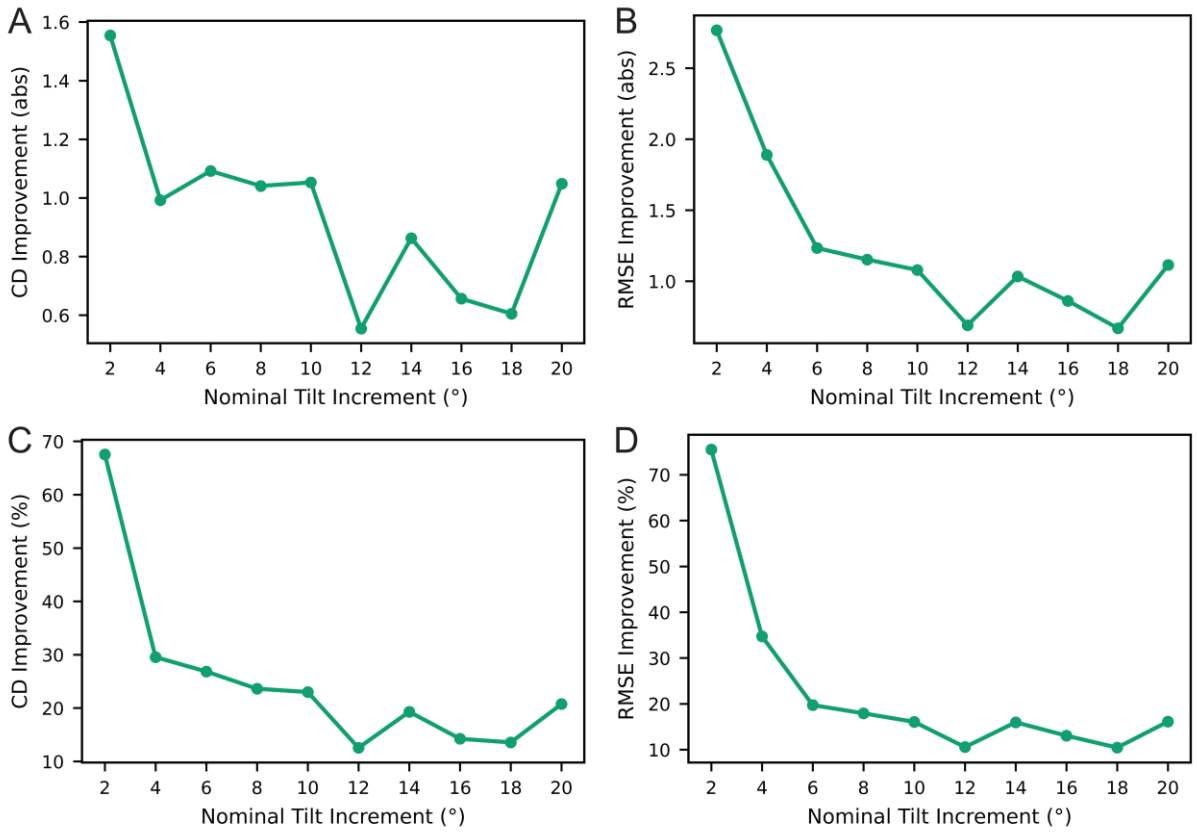

Figure 4.5. Absolute and percentage improvements in Chamfer distance and RMS error for TSGNet-AUG relative to Baseline under the non-aberrated configuration across nominal tilt increments from 2° to 20°. Positive values indicate lower error for TSGNet-AUG.

Table 4.3. Non-aberrated reconstruction-error improvements for TSGNet-AUG relative to Baseline. Positive values indicate lower reported error for TSGNet-AUG.

Nominal increment (°)	Chamfer improvement (absolute)	Chamfer improvement (%)	RMS improvement (absolute)	RMS improvement (%)
2	1.55	67.55	2.77	75.50
4	0.99	29.55	1.89	34.73
6	1.09	26.84	1.23	19.72
8	1.04	23.62	1.15	17.93
10	1.05	23.00	1.08	16.06
12	0.55	12.54	0.69	10.56
14	0.86	19.31	1.03	15.97
16	0.66	14.23	0.86	13.04
18	0.60	13.56	0.67	10.44
20	1.05	20.74	1.11	16.10

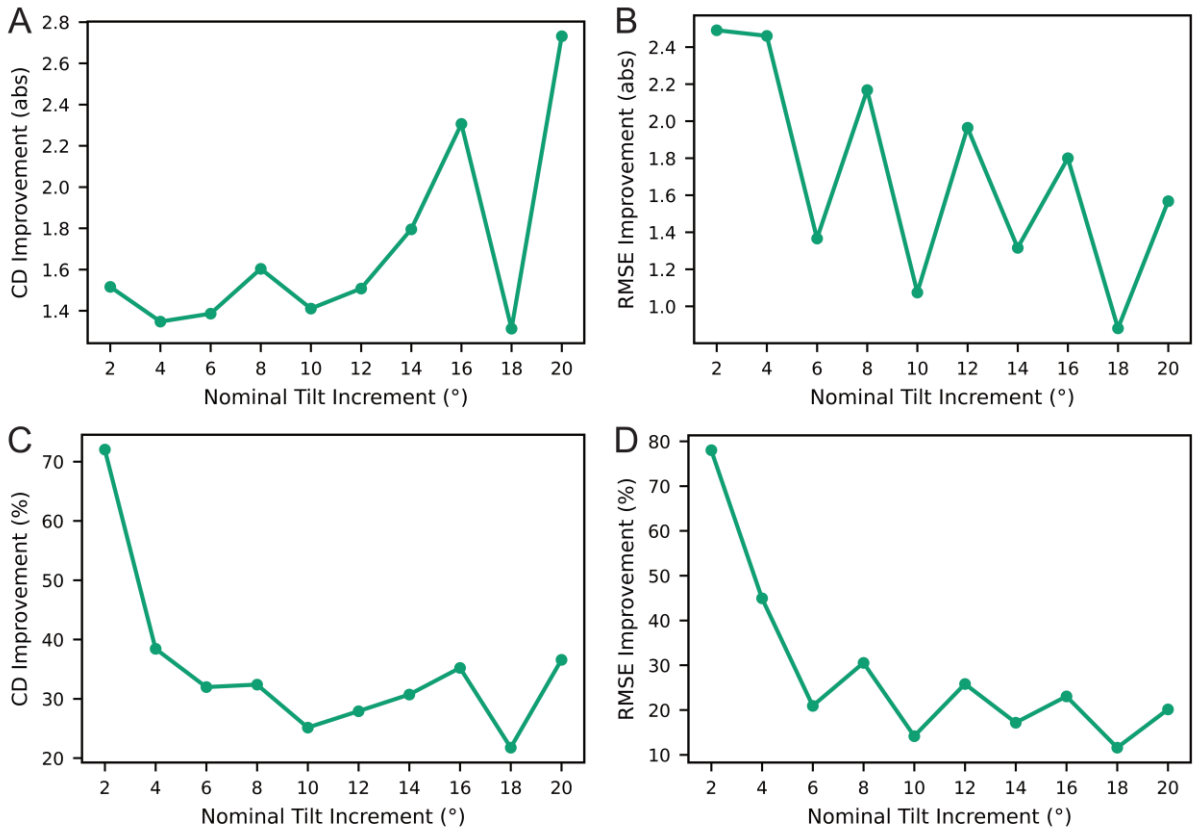

Figure 4.6. Absolute and percentage improvements in Chamfer distance and RMS error for TSGNet-AUG relative to Baseline under the aberrated configuration across nominal tilt increments from 2° to 20°. Positive values indicate lower error for TSGNet-AUG.

Table 4.4. Aberrated reconstruction-error improvements for TSGNet-AUG relative to Baseline. Positive values indicate lower reported error for TSGNet-AUG.

Nominal increment (°)	Chamfer improvement (absolute)	Chamfer improvement (%)	RMS improvement (absolute)	RMS improvement (%)
2	1.52	72.03	2.49	78.01
4	1.35	38.43	2.46	44.94
6	1.39	31.97	1.37	20.93
8	1.60	32.39	2.17	30.52
10	1.41	25.15	1.07	14.17
12	1.51	27.91	1.96	25.81
14	1.80	30.72	1.32	17.17
16	2.31	35.21	1.80	23.05
18	1.31	21.76	0.88	11.59
20	2.73	36.57	1.57	20.12

The largest percentage improvements in all four percentage series occurred at a nominal increment of 2°. At this setting, midpoint insertion restored an effective 1° spacing, matching the

nominal angular spacing of the dense series, so the maximum percentage improvement should not be interpreted as performance under the most severe sparsity. In the non-aberrated configuration, Chamfer distance and RMS error improved by 67.55% and 75.50%, respectively; in the aberrated configuration, the corresponding improvements were 72.03% and 78.01%.

Discussion

The image-level and reconstruction-level analyses support a limited conclusion: incorporating learned intermediate projections improved agreement with the dense-series reconstruction reference relative to using the sparse retained projections alone. Image-similarity metrics by themselves do not establish tomographic utility, so the reductions in downstream Chamfer distance and RMS error provide the more application-relevant evidence. The reconstruction analysis did not include a linear-interpolation augmentation control, however, and therefore does not isolate the learned model from the general effect of inserting interpolated projections.

Interpretation is further limited by the exclusive use of simulated data, simultaneous changes to multiple microscope parameters between configurations, the restricted range of structures and elements, and incomplete documentation of uncertainty across test structures. The dense-series SIRT volume is an algorithm-dependent reconstruction reference rather than the original atomic structure. Experimental generalization and complete reproducibility require a frozen correspondence among the manuscript version, data, model checkpoint, preprocessing, reconstruction code, metric implementation, and arrays underlying the reported figures and tables.

Because TSGNet-AUG inserts synthesized projections rather than additional acquired exposures, the approach may eventually support dose- or time-limited tomography; neither dose

nor acquisition time was measured in this study. Future evaluation should include experimental tilt series, a reconstruction-level linear-interpolation control, broader structural and imaging distributions, uncertainty quantification for generated projections, and explicit physical or tomographic consistency tests.

Conclusion

This chapter presented TSGNet for generating intermediate projections in sparse simulated STEM tilt series. TSGNet improved the reported projection-level metrics relative to linear interpolation, and TSGNet-AUG reduced Chamfer distance and RMS error relative to sparse Baseline reconstructions at every reported nominal increment from 2° to 20° under both simulated configurations. Within the reported reconstruction pipeline, the generated projections therefore improved agreement with the dense-series reconstruction reference. The study does not establish experimental accuracy, dose reduction, acquisition-time savings, or superiority over conventional interpolation at the reconstruction level.

Acknowledgments

This research was primarily supported by the National Science Foundation Materials Research Science and Engineering Center program through the UC Irvine Center for Complex and Active Materials (DMR-2011967).

The authors acknowledge the University of California, Irvine HPC support team and IMRI for high-performance-computing assistance; Colin Ophus, Ben Savitzky, and Steven Zeltmann for assistance with the EMD file format; and the Tomviz team for assistance with reconstruction-related aspects.

References

1. Midgley, P. A.; Dunin-Borkowski, R. E. Electron tomography and holography in materials science. Nature Materials 2009, 8(4), 271–280. https://doi.org/10.1038/nmat2406.
2. Friedrich, H.; de Jongh, P. E.; Verkleij, A. J.; de Jong, K. P. Electron Tomography for Heterogeneous Catalysts and Related Nanostructured Materials. Chemical Reviews 2009, 109(5), 1613–1629. https://doi.org/10.1021/cr800434t.
3. Zhou, X.; Luo, C.; Hashimoto, T.; Hughes, A. E.; Thompson, G. E. Study of localized corrosion in AA2024 aluminium alloy using electron tomography. Corrosion Science 2012, 58, 299–306. https://doi.org/10.1016/j.corsci.2012.02.001.
4. Grandfield, K.; Palmquist, A.; Engqvist, H. High-resolution three-dimensional probes of biomaterials and their interfaces. Philosophical Transactions of the Royal Society A 2012, 370(1963), 1337–1351. https://doi.org/10.1098/rsta.2011.0253.
5. Stewart, P. L. Cryo-electron microscopy and cryo-electron tomography of nanoparticles. WIREs Nanomedicine and Nanobiotechnology 2017, 9(2), e1417. https://doi.org/10.1002/wnan.1417.
6. Manassa, J.; Millsaps, W.; Schwartz, J.; Hovden, R. Optimal 3D chemical imaging with multimodal electron tomography. npj Computational Materials 2025, 11(1), 275. https://doi.org/10.1038/s41524-025-01750-y.
7. Hagen, W. J. H.; Wan, W.; Briggs, J. A. G. Implementation of a cryo-electron tomography tilt-scheme optimized for high resolution subtomogram averaging. Journal of Structural Biology 2017, 197(2), 191–198. https://doi.org/10.1016/j.jsb.2016.06.007.
8. Koneti, S.; Roiban, L.; Dalmas, F.; Langlois, C.; Gay, A.-S.; Cabiac, A.; Grenier, T.; Banjak, H.; Maxim, V.; Epicier, T. Fast electron tomography: Applications to beam sensitive samples and in situ TEM or operando environmental TEM studies. Materials Characterization 2019, 151, 480–495. https://doi.org/10.1016/j.matchar.2019.02.009.
9. Kang, S.; et al. Time-resolved Brownian tomography of single nanocrystals in liquid during oxidative etching. Nature Communications 2025, 16(1), 1158. https://doi.org/10.1038/s41467-025-56476-8.
10. Baldi, P. Deep Learning in Science. Cambridge University Press, 2021.
11. LeCun, Y.; Bengio, Y.; Hinton, G. Deep learning. Nature 2015, 521(7553), 436–444.
12. Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; Belongie, S. Feature pyramid networks for object detection. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition; 2017; pp 2117–2125.
13. Zhang, K.; Zuo, W.; Chen, Y.; Meng, D.; Zhang, L. Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. IEEE Transactions on Image Processing 2017, 26(7), 3142–3155.
14. Tian, C.; Xu, Y.; Zuo, W.; Zhang, B.; Fei, L.; Lin, C.-W. Coarse-to-fine CNN for image super-resolution. IEEE Transactions on Multimedia 2020, 23, 1489–1502.
15. Elharrouss, O.; Almaadeed, N.; Al-Maadeed, S.; Akbari, Y. Image inpainting: A review. Neural Processing Letters 2020, 51(2), 2007–2028.
16. Peleg, T.; Szekely, P.; Sabo, D.; Sendik, O. IM-Net for high resolution video frame interpolation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition; 2019; pp 2398–2407.
17. Lobato, I.; Friedrich, T.; Van Aert, S. Deep convolutional neural networks to restore single-shot electron microscopy images. npj Computational Materials 2024, 10(1), 10.
18. Horwath, J. P.; Zakharov, D. N.; Mégret, R.; Stach, E. A. Understanding important features of deep learning models for segmentation of high-resolution transmission electron microscopy images. npj Computational Materials 2020, 6(1), 108.

19. Su, Y.-T.; Lu, Y.; Chen, M.; Liu, A.-A. Spatiotemporal joint mitosis detection using CNN-LSTM network in time-lapse phase contrast microscopy images. IEEE Access 2017, 5, 18033–18041.
20. Ding, G.; Liu, Y.; Zhang, R.; Xin, H. L. A joint deep learning model to recover information and reduce artifacts in missing-wedge sinograms for electron tomography and beyond. Scientific Reports 2019, 9(1), 12803. https://doi.org/10.1038/s41598-019-49267-x.
21. Liu, Y.-T.; Zhang, H.; Wang, H.; Tao, C.-L.; Bi, G.-Q.; Zhou, Z. H. Isotropic reconstruction for electron tomography with deep learning. Nature Communications 2022, 13(1), 6482. https://doi.org/10.1038/s41467-022-33957-8.
22. Wiedemann, S.; Heckel, R. A deep learning method for simultaneous denoising and missing wedge reconstruction in cryogenic electron tomography. Nature Communications 2024, 15(1), 8255. https://doi.org/10.1038/s41467-024-51438-y.
23. Majtner, T.; Kreysing, J. P.; Tuijtel, M. W.; Cruz-León, S.; Liu, J.; Hummer, G.; Beck, M.; Turoňová, B. cryoTIGER: Deep-learning based tilt interpolation generator for enhanced reconstruction in cryo electron tomography. Communications Biology 2025, 8, 1443. https://doi.org/10.1038/s42003-025-08961-5.
24. Hjorth Larsen, A.; et al. The atomic simulation environment-a Python library for working with atoms. Journal of Physics: Condensed Matter 2017, 29(27), 273002. https://doi.org/10.1088/1361-648X/aa680e.
25. Madsen, J.; et al. A Deep Learning Approach to Identify Local Structures in Atomic-Resolution Transmission Electron Microscopy Images. Advanced Theory and Simulations 2018, 1(8), 1800037. https://doi.org/10.1002/adts.201800037.
26. Park, J.; Lee, C.; Kim, C.-S. Asymmetric bilateral motion estimation for video frame interpolation. In Proceedings of the IEEE/CVF International Conference on Computer Vision; 2021; pp 14539–14548.
27. Charbonnier, P.; Blanc-Feraud, L.; Aubert, G.; Barlaud, M. Two deterministic half-quadratic regularization algorithms for computed imaging. In Proceedings of the 1st International Conference on Image Processing; IEEE, 1994; Vol. 2, pp 168–172.
28. Meister, S.; Hur, J.; Roth, S. UnFlow: Unsupervised learning of optical flow with a bidirectional census loss. In Proceedings of the AAAI Conference on Artificial Intelligence; 2018; Vol. 32.
29. Loshchilov, I.; Hutter, F. Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101, 2017.
30. Wang, Z.; Bovik, A. C.; Sheikh, H. R.; Simoncelli, E. P. Image quality assessment: From error visibility to structural similarity. IEEE Transactions on Image Processing 2004, 13(4), 600–612.
31. Gürsoy, D.; De Carlo, F.; Xiao, X.; Jacobsen, C. TomoPy: A framework for the analysis of synchrotron tomographic data. Journal of Synchrotron Radiation 2014, 21(5), 1188–1193. https://doi.org/10.1107/S1600577514013939.
32. Gilbert, P. Iterative methods for the three-dimensional reconstruction of an object from projections. Journal of Theoretical Biology 1972, 36(1), 105–117. https://doi.org/10.1016/0022-5193(72)90180-4.
33. Pelt, D. M.; Palenstijn, W. J.; Batenburg, K. J. Integration of TomoPy and the ASTRA toolbox for advanced processing and reconstruction of tomographic synchrotron data. Journal of Synchrotron Radiation 2016, 23(3), 842–849.
34. Otsu, N. A threshold selection method from gray-level histograms. IEEE Transactions on Systems, Man, and Cybernetics 1979, 9(1), 62–66. https://doi.org/10.1109/TSMC.1979.4310076.
35. Barrow, H. G.; Tenenbaum, J. M.; Bolles, R. C.; Wolf, H. C. Parametric correspondence and Chamfer matching: Two new techniques for image matching. In Proceedings of IJCAI; 1977.

36. Borgefors, G. Hierarchical Chamfer matching: A parametric edge matching algorithm. IEEE Transactions on Pattern Analysis and Machine Intelligence 1988, No. 6, 849–865.
37. Fan, H.; Su, H.; Guibas, L. J. A point set generation network for 3D object reconstruction from a single image. In CVPR; 2017.

CHAPTER 5: General Conclusions and Outlook

Integrated Conclusions

This dissertation examined a common problem across several forms of electron microscopy: the recorded image is not identical to the underlying structure, and the strength of a scientific conclusion depends on how explicitly the image-formation pathway and the available reference are defined. The four studies addressed dynamic liquid-phase observations, atomistic-to-image forward simulation, synthetic-data segmentation, and sparse electron tomography. Although the specimens and computational tasks differed, each study was organized around the same principle: conclusions should remain at the level supported by the measurement and should not be extended beyond the available controls, reference data, or domain of validation.

Chapter 1 established a time-resolved record of projected morphological evolution for FFssFF-derived peptide coacervates in one electrochemical liquid-cell field. The largest qualitative image change occurred across the EIS interval, after which the same four tracked domains showed progressively greater projected-area loss in the two post-EIS CV recordings. The chapter therefore demonstrates that synchronized electrochemistry and LPTEM can localize when nanoscale changes occur and can quantify those changes using directly defined image observables. The experiment does not, however, separate the contributions of applied potential, elapsed electron exposure, interfacial charging, transport, or chemical conversion. Its principal contribution is a bounded, selected-case description of projected morphology rather than a population-level or mechanistic law.

Chapter 2 connected cryo-EM measurements and collaborator-generated atomistic fiber models through multislice TEM simulation. The resulting images reproduced the principal image-scale features of the experimental fibers, including a dark central feature, underfocus rims, and an apparent-width range consistent with the experimental measurements under the tested

imaging conditions. Systematic variation of defocus and water thickness further showed that apparent fiber width and visibility depend strongly on image formation. Because the experimental diameter was also used to define the radial restraint, the agreement is appropriately interpreted as an image-domain consistency test. It does not independently validate a unique molecular packing arrangement.

Chapter 3 extended this forward-modeling logic to supervised segmentation. Synthetic TEM images and coordinate-defined foreground references were generated independently from the same transformed atomic scenes, allowing a U-Net to be trained without deriving labels from local image intensity. The selected network achieved high agreement with the reference masks on the fixed synthetic test set and exceeded the two thresholding baselines within that generator-defined task. Application to unlabeled experimental TEM images produced predictions that followed many visually recognizable fiber-like trajectories, but those examples provide qualitative correspondence rather than a quantitative estimate of experimental accuracy. The quantitative conclusion is therefore restricted to held-out scenes generated by the same underlying structural and image-simulation framework.

Chapter 4 evaluated TSGNet at both the projection and reconstruction levels using simulated STEM tilt series. The network produced intermediate projections with better reported image-similarity metrics than linear interpolation, and insertion of those predictions reduced Chamfer distance and volumetric RMS error relative to the matched sparse baseline for every reported nominal increment from 2° to 20° under both simulated configurations. This result shows that predicted projections can improve agreement with a dense-series reconstruction reference within the reported pipeline. It does not establish experimental dose reduction, acquisition-time savings, or generalization to measured tilt series, and it does not isolate learned

interpolation from the broader effect of adding interpolated projections at the reconstruction stage.

Cross-Chapter Synthesis

Taken together, the studies define a progression from observation to inference. In Chapter 1, the reference is the same object followed over time, and the primary observables are projected area and image contrast. In Chapter 2, atomistic coordinates are converted into simulated images so that model outputs and experimental data can be compared in a common image domain. In Chapter 3, the same atomic scene supplies both an image-formation pathway and an independent geometric pathway for constructing a synthetic reference. In Chapter 4, the value of a predicted image is tested by its effect on a downstream three-dimensional reconstruction rather than by visual plausibility alone.

This progression also clarifies that “ground truth” is not a single category. A tracked LPTEM boundary is an operational image measurement; a coordinate-defined segmentation mask is exact only within the constructed synthetic task; and a full-series tomographic reconstruction remains an algorithm-dependent reference rather than the original three-dimensional atomic structure. Forward simulation improves traceability because specimen coordinates, microscope parameters, detector assumptions, and reference construction can be recorded explicitly. It does not remove model dependence or guarantee that the synthetic distribution spans the experimental one.

The strongest shared contribution of the dissertation is therefore methodological rather than tied to one specimen. Quantitative microscopy becomes more defensible when the analysis states what was directly measured, what was generated by a model, what served as the comparison reference, and which sources of uncertainty were not included. This framework

prevents image contrast from being treated automatically as mass or composition, prevents agreement with a restraint-informed model from being presented as independent structural validation, and prevents synthetic or simulation-only performance from being described as established experimental accuracy.

Evidence Boundaries

The principal limitation of the electrochemical LPTEM study is experimental scope. Four tracked domains in one field provide repeated temporal measurements but not independent biological or materials replication. The absence of bias-only and beam-only controls prevents causal separation of electrochemical and irradiation effects. The two-dimensional masks also cannot distinguish dissolution from axial reorganization, fragmentation below the segmentation scale, or redistribution outside the projected boundary.

The principal limitation of the atomistic-to-image comparison is partial circularity. The cryo-EM diameter informed the restraint applied to the MD fibers, so agreement in apparent diameter cannot serve as an independent test of that dimension. The experimental and simulated profile measurements also used different averaging widths, and the experimental images did not contain sufficient structural detail to discriminate among alternative molecular arrangements. The supported conclusion is image-scale consistency under specified conditions.

For segmentation, the fixed test images were distinct files but were generated from the same trajectory family, construction rules, and forward model used to create the training distribution. The reported scores consequently measure within-generator generalization. Reuse of the same validation set throughout adaptive model development and later review of the test summaries further mean that the results should not be interpreted as a prospective external

benchmark. The unlabeled real-TEM examples provide neither pixel-level sensitivity and specificity nor an independent ranking of the U-Net and thresholding methods.

For tomography, the dense 181-projection SIRT reconstruction is a reference produced by the same reconstruction and segmentation pipeline, not the original atomic structure. The two simulated imaging configurations changed more than one optical parameter, so their comparison does not isolate spherical aberration. In addition, reconstruction-level linear interpolation was not evaluated alongside TSGNet-AUG, uncertainty across test structures was not fully documented in the available source materials, and the analysis did not measure dose or acquisition time.

Outlook

For electrochemical LPTEM, the highest-priority next step is independent replication with experimental controls that separate bias, beam exposure, elapsed time, and solution chemistry. Calibrated electron flux, repeated liquid cells, no-bias imaging, bias-only electrochemistry, and matched mediator or reductant controls would allow the temporal associations reported here to be tested as causal effects. Three-dimensional or correlative measurements would also be needed to connect projected-area loss to volume, mass, composition, or molecular conversion.

For atomistic-to-image modeling, future work should compare alternative structural models under the same image-formation conditions rather than evaluating only structures constrained by the measured diameter. Independent trajectories, longer equilibrated fibers, matched experimental and simulated profile operators, and additional image observables could increase discrimination among candidate structures. Uncertainty should be propagated from specimen geometry and imaging parameters through the simulated image and final width measurement.

For synthetic segmentation, the most informative validation would combine source-grouped synthetic splits with independently annotated experimental images. Holding out entire MD frames, trajectories, or scene-generation families would test whether performance depends on repeated structural modules. A real-TEM benchmark should document inter-annotator boundary variability and should include different pixel sizes, liquid thicknesses, microscopes, detectors, contamination levels, and fiber densities. Domain randomization, uncertainty maps, and scale-aware inference may then be evaluated against a reference that is independent of the synthetic generator.

For sparse tomography, experimental tilt series and a reconstruction-level linear-interpolation control are essential. A complete comparison should include the sparse baseline, conventional interpolation augmentation, TSGNet augmentation, and the dense acquired reference under identical reconstruction settings. Structure-level uncertainty, failure cases, and dependence on particle shape, composition, alignment error, noise, missing angular range, and reconstruction algorithm should be reported. Only after acquisition dose and time are measured can the approach be evaluated as a low-dose or accelerated tomography strategy.

Across all four areas, reproducibility will be strengthened by archiving exact data versions, analysis scripts, model checkpoints, simulation parameters, and the arrays underlying summary figures and tables. Versioned repositories are particularly important for chapters based on published or submitted manuscripts, because public code, manuscript revisions, and dissertation analyses may not otherwise correspond to the same configuration.

Final Perspective

Electron microscopy of dynamic and weakly scattering materials requires more than higher spatial resolution. It requires a traceable connection between specimen state, image

formation, computational processing, and the comparison reference used to support a conclusion. The studies in this dissertation show that synchronized in situ imaging, physics-based forward simulation, and deep learning can each extend the information extracted from microscopy data when their roles are defined explicitly.

The resulting framework is intentionally conservative: it distinguishes direct observation from model consistency, synthetic-domain accuracy from experimental accuracy, and improvement relative to a reconstruction reference from recovery of an unknown true structure. Within those boundaries, the work provides quantitative tools for following nanoscale dynamics, testing atomistic models in the image domain, constructing traceable segmentation references, and evaluating missing-projection predictions through their downstream tomographic effect. These contributions support a broader transition from qualitative microscopy images toward quantitative inferences whose assumptions and limitations remain visible.