

RESIRE-Net enables accurate 3D atomic structure determination from limited-angle electron tomography

Minh Pham

minhrose@ucla.edu

University of California, Los Angeles

Ada Morral

University of California, Los Angeles

Jianwei Miao

University of California, Los Angeles

Article

Keywords:

Posted Date: October 20th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11045178/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

RESIRE-Net enables accurate 3D atomic structure determination from limited-angle electron tomography

Minh Pham^{1,+,*}, Ada Morral^{1,+}, and Jianwei Miao^{1,*}

¹Department of Physics and Astronomy and California NanoSystems Institute, University of California, Los Angeles, CA 90095, USA

*minhrose@ucla.edu; j.miao@ucla.edu

+These authors contributed equally to this work

ABSTRACT

Determining three-dimensional (3D) atomic positions and chemical identities is essential for understanding defects, strain, and chemical order in materials. Atomic electron tomography provides this capability, but limited specimen tilting leaves a missing wedge of information that degrades depth resolution and atomic localization. We introduce RESIRE-Net, a self-supervised learning framework that combines a 3D U-Net with real-space iterative reconstruction (RESIRE). Using only the input dataset, RESIRE-Net generates training pairs from rotated volumes, tilt geometry, and the RESIRE reconstruction process, requiring neither full-angle ground truth nor external training data. In multislice-simulated metallic-glass and medium-entropy-alloy nanoparticles, RESIRE-Net improves the 3D positional precision from 21 to 14 pm and from 24 to 15 pm, respectively, while enhancing atom recovery and elemental classification. Applied to an experimental medium-entropy-alloy nanoparticle, RESIRE-Net sharpens atomic peaks along the missing-wedge direction while preserving twin boundaries and chemical short-range order. These results demonstrate that dataset-specific self-supervision can improve the reliability of 3D atomic structure determination.

Introduction

Three-dimensional atomic arrangements govern materials properties through local coordination, chemical order, strain, interfaces and defects. Electron tomography reconstructs 3D volumes from projections at different orientations, and atomic electron tomography (AET) extends this to the single-atom level by combining advanced electron microscopy with computational reconstruction¹⁻⁵. Because AET requires neither crystallinity nor averaging over nominally identical structures, it can directly resolve individual non-periodic materials, revealing dislocations and other defects⁶, short- and medium-range order in amorphous solids^{7,8}, local chemical order and lattice distortion in medium- and high-entropy alloys⁹, catalytically active sites in alloy nanoparticles¹⁰, and crystal nucleation and growth^{11,12}. These advances exemplify computational microscopy, where algorithms are integral to image formation and quantitative interpretation¹. Reliable AET therefore depends on the entire workflow, from projection denoising and alignment to reconstruction, atom tracing, elemental classification and coordinate refinement².

Tomographic reconstruction estimates a 3D object from a finite set of noisy projections. Filtered back projection is efficient when projections densely sample a nearly complete angular range¹³, but AET generally requires iterative reconstruction because tilt series contain relatively few projections over a restricted range. The oversampling principle led to equally sloped tomography^{14,15}, generalized Fourier iterative reconstruction (GENFIRE)¹⁶ and real-space iterative reconstruction (RESIRE)¹⁷. RESIRE minimizes the least-squares discrepancy between measured and calculated projections. Forward projections are computed using either the Fourier slice theorem with reciprocal-space oversampling or the Radon transform with subvoxel sampling, while back-projection drives iterative volume updates and refines the tilt geometry by minimizing the projection residual¹³. In numerical and experimental benchmarks under limited-angle conditions, RESIRE produced more accurate reconstructions and fewer missing-wedge artifacts than FBP, SIRT and GENFIRE¹⁶.

Restricted angular coverage remains a fundamental limitation of ET and AET. Mechanical clearance, specimen geometry, increased effective thickness at high tilt, drift and dose constraints typically limit acquisition to approximately $\pm(60^\circ-75^\circ)$. According to the Fourier slice theorem, each projection samples a plane in 3D reciprocal space¹³, so missing angles leave a continuous wedge-shaped region unsampled. This missing wedge differs from sparse angular sampling and makes the inverse problem ill-conditioned, producing anisotropic resolution, elongation and streaking along the least-sampled direction. At atomic resolution, these artifacts can corrupt atomic coordinates, elemental classification, defect identification and strain measurements.




Figure 1. RESIRE-Net concept and network architecture. **a.** The missing-wedge effect in limited-angle electron tomography is modeled by an operator $\mathbf{T}_\theta$ that maps a ground-truth object O^* to its RESIRE reconstruction O . **b.** RESIRE-Net employs a 3D U-Net with multiscale feature-extraction pathways, residual blocks, and skip connections. The network is designed to learn an approximation to the inverse operator $\mathbf{T}_\theta^{-1}$, enabling recovery of information degraded by the missing wedge and improving reconstruction quality.

Support and positivity constraints, total variation, compressed sensing and related regularizers can stabilize reconstruction^{18–20}, but their performance depends on the validity and weighting of imposed priors, which may smooth or bias genuine features. The unmeasured Fourier components therefore remain underdetermined by the projections alone.

Deep learning offers a complementary route by learning structural priors that are difficult to express analytically²¹. Supervised approaches have been applied to sinogram completion, artifact suppression and AET reconstruction^{22–24}, but their accuracy depends on training data that closely reproduce the specimen morphology, tilt geometry, noise statistics, scattering conditions and resolution. Mismatch between training and test conditions can therefore produce substantial domain shift when a network is transferred to a different material or acquisition setting. Self-supervised methods reduce this dependence by learning directly from the experimental data. Equivariant imaging exploits known object transformations without paired ground truth²⁵; IsoNet rotates subvolumes, adds a binary missing-wedge mask in Fourier space and iteratively trains a CNN on the resulting pairs²⁶; UsiNet performs unsupervised sinogram inpainting using ensembles of experimental nanoparticle projections without full-angle ground truth²⁷; and DeepDeWedge trains directly on 2D projections with a self-supervised objective for simultaneous denoising and missing-wedge reconstruction²⁸. Model-based methods have also combined neural network priors with iterative physics-based solvers^{29–31}. These studies highlight the importance of structural priors, imaging physics and projection fidelity, and motivate a complementary strategy that combines dataset-specific self-supervision with the experimental tilt geometry and the RESIRE reconstruction operator, without requiring an external training corpus.

Here we introduce RESIRE-Net, an operator-aware self-supervised framework for limited-angle electron tomography. Starting from a RESIRE reconstruction, the method rotates the 3D volume into multiple orientations, simulates projections at the experimental tilt angles and reconstructs each synthetic tilt series with RESIRE. These pairs train a 3D U-Net to map an additionally degraded reconstruction to its corresponding less-degraded target, and the network output is then used to regenerate the training pairs in an iterative refinement loop. The degradation operator therefore incorporates both the experimental angular sampling and the RESIRE reconstruction process. RESIRE-Net requires neither full-angle ground truth nor an external training set, and is formulated as a learned regularized estimator. For multislice-simulated nanoparticles, RESIRE-Net reduces the 3D atomic-position RMSD from 21 to 14 pm in a metallic glass and from 24 to 15 pm in a medium-entropy alloy, while improving atom recovery and elemental classification. Applied to an experimental medium-entropy-alloy nanoparticle, it sharpens atomic peaks along the missing-wedge direction while preserving the principal twin-boundary and chemical short-range-order features. These results demonstrate the potential of experiment-specific self-supervision for more reliable 3D atomic models of defects, strain and chemical order in materials.

Methods

Formulation

The goal of tomography is to reconstruct a three-dimensional image from a set of measured projections acquired at different tilt angles. A standard formulation is the least-squares problem

$$\min_O \varepsilon(O) = \frac{1}{2} \sum_{\theta} \|\Pi_{\theta}(O) - b_{\theta}\|_F^2 \quad (1)$$

where $\{b_{\theta}\}_{\theta}$ and $\{\theta\}_{\theta}$ denote the projections and their corresponding tilt angles. Π_{θ} is a linear projection operator with respect to the tilt angle θ . For simplicity, a single tilt axis is shown, though the formulation extends naturally to multiple axes




Figure 2. RESIRE-Net framework includes three stages: Stage 1: A RESIRE reconstruction O is rotated according to the symmetry group S_4 . Projections of each rotated volume are generated using the same tilt angles, and RESIRE is applied to obtain corresponding rotated reconstructions. Stage 2: Each rotated object O_i and its RESIRE reconstruction $\mathbf{T}_\theta(O_i)$ form a training pair for the U-Net. Stage 3: The trained network is applied to the original RESIRE reconstruction to produce a refined object, which is then fed back into Stage 1. This iterative process is repeated for 20-40 rounds.

using Euler angles. The l_2 metric is widely used due to its simplicity and its connection to classical reconstruction algorithms such as ART, IRT, and SIRT¹³.

Let O^* and O denote the ground-truth object and the corresponding least-squares solution, respectively. With sufficient angular sampling, the Nyquist-Shannon theorem guarantees uniqueness, and noise-free simulations yield $O = O^*$. In electron tomography, however, the tilt range is typically limited to $[-60^\circ, 60^\circ]$, making the inverse problem ill-posed. As a result, iterative reconstruction algorithms often converge to solutions that suffer from severe missing-wedge artifacts. Regularization methods, such as Tikhonov³², l_p norms, or total variation^{33,34}, can incorporate prior information, but they require manual parameter tuning and cannot fully recover missing information.

To formalize the effect of incomplete sampling, we define an operator $\mathbf{T}_{\theta, \text{RESIRE}}$ that maps a ground truth object O^* to its RESIRE reconstruction O obtained from a limited tilt series. Because RESIRE recovers substantially more missing-wedge information than classical methods^{17,35}, it provides a strong foundation for further enhancement using neural networks. Throughout this work, RESIRE is used as the conventional reconstruction operator. For simplicity, we denote $\mathbf{T}_{\theta, \text{RESIRE}}$ by $\mathbf{T}_\theta$ in the remainder of the paper. A visualization of $\mathbf{T}_\theta$ is shown in Fig. 1a.

The missing-wedge problem naturally leads to an inverse formulation, where the objective is to approximate $\mathbf{T}_\theta^{-1}$, or equivalently recover O^* from $\mathbf{T}_\theta(O^*)$. Since classical methods cannot reliably solve this inverse problem, we adopt a learning-based approach. Specifically, we introduce a neural network f_ϑ , parameterized by ϑ , designed to map $\mathbf{T}_\theta(O^*)$ back to O^* . The training objective is defined as follows:

$$\min_{\vartheta} \mathcal{L}(\vartheta) = L(f_\vartheta(\mathbf{T}_\theta(O^*)), O^*) \quad (2)$$

where common error metrics such as l_1 , l_2 , total variation, and perceptual loss can be used as the loss function L . To improve generalization, we incorporate measurement consistency by requiring that the projections of $f_\vartheta(\mathbf{T}_\theta(O))$ match those of O ²⁵. This yields an l_2 regularization term $\varepsilon \left\| \Pi_\theta(f_\vartheta(\mathbf{T}_\theta(O^*))) - \Pi_\theta(O^*) \right\|^2$.

Algorithm

Based on the formulation above, RESIRE-Net employs an iterative three-stage workflow consisting of (1) training-data generation, (2) network training, and (3) inference (object updating). These stages are iteratively repeated, progressively improving both the training dataset and the reconstruction quality. The details of each stage are described below.

Stage 1: Training-data generation

Supervised learning requires paired training data consisting of corrupted inputs and corresponding ground-truth outputs. In AET, however, reliable ground truth is generally unavailable because of the limited tilt range, missing-wedge effects, and experimental uncertainties. Although simulated datasets could provide training targets, they must accurately reproduce the experimental tilt geometry, noise characteristics, defocus conditions, and high-frequency image content. Such requirements make conventional supervised learning impractical for most AET applications and motivate the use of a self-supervised strategy.

To overcome this limitation, we generate training data through augmentation based on rotational symmetries. A reconstructed volume O is rotated by 90° about the X, Y, or Z axis, or by combinations of these operations, exploiting the cube symmetry group (isomorphic to S_4 ^{25,26,28}). This procedure produces 24 rotated variants O_i , from which transformed volumes $\mathbf{T}_\theta(O_i)$ are calculated to form training pairs $(O_i, \mathbf{T}_\theta(O_i))$. Four rotations are excluded because they preserve the same missing-wedge orientation as the original object. By exposing the network to multiple orientations of the same structure, the training set captures a broader range of artifact configurations.

The augmentation strategy can be extended further by introducing additional 45° Euler rotations about the x and y-axes. These operations increase the number of available training samples from 20 to 92 and provide a greater diversity of missing-wedge orientations. A larger and more diverse training set may improve robustness, particularly when the reconstructed volume contains anisotropic structural features. However, the expanded scheme also increases computational cost and training time. For simplicity and efficiency, we adopt the standard 20-rotation scheme throughout this work. Nevertheless, the 45° Euler rotation option remains available in our implementation and can be readily applied by users in their own experiments when additional augmentation is desired.

While $O = \mathbf{T}_\theta(O^*)$ is corrupted by the missing wedge in a single direction, the transformed volume $\mathbf{T}_\theta(O_i)$ contains an additional missing-wedge artifact oriented differently from that in O . The network is therefore trained to map a twice-corrupted volume, $\mathbf{T}_\theta(O_i)$, back to the corresponding once-corrupted volume, O_i . Through this process, the model learns to identify and compensate for wedge-related distortions without requiring ground-truth reconstructions. This self-supervised learning framework forms the basis of RESIRE-Net for missing-wedge recovery. For non-isolated objects, such as specimens containing high-tilt features or thin-film geometries, thin-film RESIRE or a three-dimensional mask can be incorporated during reconstruction to reduce interference from surrounding structures and improve reconstruction accuracy¹⁷.

Stage 2: Training

To learn how to recover O_i from $\mathbf{T}_\theta(O_i)$, we employ a 3D U-Net, a specialized convolutional neural network widely used for image deblurring, denoising, and segmentation^{36–38}. Its multiscale design is well suited for AET because it captures volumetric information across different spatial scales. The encoder-decoder structure extracts both local and global features, while skip connections preserve structural details that may otherwise be lost during downsampling³⁹. Residual learning improves optimization stability and fine-detail recovery⁴⁰, and adaptive skip connections further enhance reconstruction accuracy⁴¹. A visualization of the proposed 3D U-Net is shown in Fig. 1b.

The network architecture contains approximately 6.8 million trainable parameters. During training, volumes are randomly cropped to a size of $64 \times 64 \times 64$. Gaussian noise is added to the rotated volumes $\{O_i\}_{i=1}^{20}$ to suppress high-frequency artifacts and improve robustness²⁸. We choose the l_1 error metric to measure the discrepancy between $\mathbf{T}_\theta(O_i)$ and O_i due to its strong ability to preserve sharp features in missing-wedge-degraded images^{42,43}. The final loss function combines this fidelity term with the measurement-consistency regularizer:

$$\min_{\vartheta} \mathcal{L}(\vartheta) = \frac{1}{N} \sum_{i=1}^N \left\| f_{\vartheta}(\mathbf{T}_\theta(O_i)) - O_i \right\|_1 + \varepsilon \left\| \Pi_{\theta}(f_{\vartheta}(\mathbf{T}_\theta(O_i))) - \Pi_{\theta}(O_i) \right\|^2 \quad (3)$$

where $N = 20$ is the number of training pairs. The regularizer can be applied every 8 inner iterations to improve efficiency. Additional techniques can also be considered to improve deblurring and enhance sharp features, such as attention mechanisms⁴⁴, GAN-based⁴⁵ methods and perceptual losses derived from VGG⁴⁶ networks. We leave these methodologies open for readers to explore as part of their research.

Stage 3: Inference

After training, the network is applied to the original RESIRE reconstruction. Large volumes are partitioned into overlapping $64 \times 64 \times 64$ cubes, processed independently by the network, and subsequently blended to ensure smooth transitions between neighboring regions. The resulting output constitutes a refined reconstruction with reduced missing-wedge artifacts. This refined volume serves as an improved estimate of the object and provides the basis for subsequent refinement. By suppressing reconstruction artifacts while preserving structural details, the network generates a higher-quality input for the next iteration.

The refined object is then used to regenerate the training dataset for the next refinement round. Consequently, both the reconstructed object and the training samples improve progressively over successive iterations. In practice, the three stages

are typically repeated for 20–40 refinement rounds. As the training data become increasingly representative of the underlying structure, the quality of the reconstruction improves accordingly. A final RESIRE refinement step can optionally be performed to further enhance measurement consistency. The overall RESIRE-Net workflow is summarized in Algorithm 1 and Fig. 2.

Algorithm 1 RESIRE-Net

Input: RESIRE reconstruction O , tilt angles $\{\theta_i\}_{i=1}^n$, number of refinement K , number of iteration per refinement N .

Initialize: $O^0 = O$

// Outer refinement loop

for $k = 0, \dots, K - 1$ **do**

Stage 1: Training-data generation

for all rotation $i \in \hat{S}_4$ **do**

 Obtain rotation O_i^k from O^k .

 Obtain transformation $\mathbf{T}_\theta(O_i^k)$.

end for

Stage 2: Training (Inner refinement loop)

for $n = 1, \dots, N$ **do**

for all permuted $i \in \hat{S}_4$ **do**

$$\mathcal{L}(\vartheta) \leftarrow \left\| f_\vartheta(\mathbf{T}_\theta(O_i^k)) - O_i^k \right\|_1 + \varepsilon \left\| \Pi_\theta(f_\vartheta(\mathbf{T}_\theta(O_i^k))) - \Pi_\theta(O_i^k) \right\|^2$$

 Update ϑ via optimization algorithm such as AdaGrad, Adam, etc.

end for

end for

Stage 3: Inference, i.e. update $O^{k+1} = f_\vartheta(O^k)$

end for

Output: O^K

Atomic structure analysis

Atomic models were extracted from reconstructed volumes using previously published RESIRE-based atomic electron tomography (AET) analysis pipelines developed for the metallic glass dataset⁷. These workflows include polynomial tracing of local intensity maxima, k-means classification to distinguish atoms from noise and assign atomic species, limited manual correction, and refinement of atomic positions and species through forward-projection fitting. Together, these procedures provide a consistent framework for converting reconstructed volumes into quantitative atomic models suitable for structural analysis.

For all comparisons in this work, identical tracing, refinement, and analysis procedures were applied to both the RESIRE and RESIRE-Net reconstructions, with only the reconstruction algorithm replaced. Minor dataset-dependent adjustments were made to the initialization of the k-means classification procedure to improve stability for the sharper low-intensity distributions produced by RESIRE-Net; these modifications did not materially affect the final results. For simulated datasets, reconstructed atomic models were quantitatively compared with known ground-truth structures using nearest-neighbor matching and RMSD-based positional metrics.

Results

Metallic-glass simulation data

To validate RESIRE-Net, we applied it to a previously published simulated tilt series of a multi-component metallic glass nanoparticle, comprising Ni (Type 1), Pd (Type 2), and Pt (Type 3) atoms, with a known ground-truth atomic model⁷. The data were generated under experimentally realistic imaging conditions, including a limited angular range that produces a missing wedge of Fourier-space information, and projections were convolved with a Gaussian probe to model finite probe size and incoherent imaging effects (Supplementary Table 1). To evaluate RESIRE-Net, we replicated the original model-extraction pipeline, replacing only the reconstruction algorithm; all tracing, classification, refinement, and matching procedures were otherwise unchanged, enabling a direct comparison between the two reconstruction methods.

Reconstruction quality

Representative real-space slices and corresponding Fourier-space slices of the RESIRE reconstruction are shown in Fig. 3a–d and Supplementary Fig. 1. Atomic intensity maxima are well resolved in xy slices, but xz slices exhibit the characteristic




Figure 3. Demonstration of RESIRE-Net on previously-published multislice-simulated metallic glass nanoparticle data. **a, b,** Two 2.4 Å thick slices of the 3D RESIRE reconstruction in the xy (**a**) and xz (**b**) plane. Scale bar, 2 nm. **c, d,** Central reciprocal-space slices of the 3D Fourier transform of the RESIRE reconstruction in the $k_x k_y$ (**c**) and $k_x k_z$ (**d**) planes. **e–h,** Corresponding real-space slices (**e, f**) and reciprocal-space slices (**g, h**) of the RESIRE-Net reconstruction. **i,** Histograms of the integrated intensity of the subvolume around each potential atom following the polynomial tracing step. These histograms demonstrate the separation between atoms of each type, and non-atom noise, offered by each reconstruction method, which influences the global k-mean classification procedure used to remove noise from the model. **j,** Histograms of the 3D deviations of each model's atoms from their corresponding ground truth atoms. The peak, mean, and root-mean-square-deviation of the RESIRE histogram are 6 pm, 15 pm, and 21 pm, respectively. For the RESIRE-Net histogram we find those values to be 5 pm, 11 pm, and 14 pm. **k,** Spatial distribution of z-coordinate errors, $|\Delta z|$, for all ground-truth-matched atoms in each reconstruction, (top, RESIRE (17,987 atoms); bottom, RESIRE-Net (18,340 atoms)), shown in the xz plane.

Table 1. Reconstruction and model accuracy of RESIRE and RESIRE-Net models of a simulated metallic glass nanoparticle, evaluated against a ground truth model (18,356 atoms).

	RESIRE ^a	RESIRE-Net
# Model atoms	18,481	18,472
Refinement		
Error (R_1) ^b	4.76%	4.19%
Position accuracy		
3D RMSD (pm)	21	14
# Position matches to GT	17,987	18,340
% GT model matched	97.99%	99.91%
% Test model matched	97.32%	99.29%
Classification accuracy		
Type 1 accuracy (Ni)	94.87%	99.74%
Type 2 accuracy (Pd)	99.14%	99.83%
Type 3 accuracy (Pt)	100.00%	100.00%
Overall classification accuracy	97.35%	99.81%
# Total accurate matches	17,870	18,322

^a Model originally published in Ref.⁷.

^b The R_1 factor is defined as in equation 5 of Ref.⁴⁷.

elongation and blurring induced by the limited tilt range. In Fourier space, this appears as a wedge-shaped region of suppressed intensity and streaking along k_z (Fig. 3d). The corresponding RESIRE-Net slices are shown in Fig. 3e–h. The Fourier transforms of the RESIRE-Net reconstruction show reduced high-frequency artefacts and substantial filling of the missing-wedge region (Fig. 3h). Although xy slices remain visually similar between the two methods, xz slices from the RESIRE-Net reconstruction exhibit sharply localized atomic peaks with markedly reduced elongation, demonstrating improved recovery of depth (z) information, where missing-wedge effects are most pronounced. Quantitative analysis of the reconstructed atomic peak shapes further shows reduced broadening and directional anisotropy in the RESIRE-Net reconstruction for all atomic species (Supplementary Fig. 2).

Positional precision

Atomic coordinates extracted from each reconstruction were matched to the ground-truth model using nearest-neighbor assignment within a 1.5 Å cutoff, and displacement vectors between matched atoms were used to evaluate positional accuracy. Fig. 3j shows the distribution of three-dimensional atomic deviations relative to the ground truth, and Table 1 summarizes the resulting differences in accuracy between the two models (for more details, see Supplementary Table 2). For the RESIRE reconstruction, the peak, mean, and RMSD of the deviation distribution are 6 pm, 15 pm, and 21 pm, respectively, consistent with the original report. RESIRE-Net produces a noticeably narrower distribution, with corresponding values of 5 pm, 11 pm, and 14 pm, representing a 33% reduction in positional RMSD. Axis-resolved RMSDs indicate that this improvement is dominated by the depth (z) direction: the average z-deviation is reduced from 15 pm to 10 pm (Supplementary Table 2, Supplementary Fig. 3). Spatial maps of $|\Delta z|$ (Fig. 3k) show a systematic reduction in depth-direction error throughout the particle, demonstrating that the additional recovered Fourier information directly enhances atomic-scale depth localization.

Atom recovery and classification

RESIRE-Net also improves atom recovery and chemical classification. Although the extracted models contain nearly the same number of atoms, RESIRE-Net recovers 99.91% of the ground-truth atoms compared with 97.99% for RESIRE, while the number of spurious atoms decreases substantially, from 494 to 132. Integrated-intensity histograms after polynomial tracing (Fig. 3i) show clearer separation between atomic peaks and noise in the RESIRE-Net reconstruction, directly benefiting the k-means classification step (Supplementary Fig. 4); as a result, overall classification accuracy increases from 97.35% to 99.81%, with the largest gains for the weakest-scattering species (Type 1, Ni), whose accuracy rises from 94.87% to 99.74%. Notably, much of this improvement can be obtained without the manual adjustment step in the original model-extraction pipeline (Supplementary Table 3), which in turn reduces processing time and minimizes user-dependent variability. Together, these results demonstrate that, under experimentally realistic conditions and within an identical model-extraction pipeline, RESIRE-Net yields improved positional precision, atom recovery, and elemental classification compared to RESIRE.

Table 2. Reconstruction and model accuracy of RESIRE and RESIRE-Net models of a simulated MEA nanoparticle, evaluated against a ground truth model (8,141 atoms).

	RESIRE ^a	RESIRE-Net
# Model atoms	8,165	8,132
Refinement		
Error (R_1) ^b	5.28%	3.90%
Position accuracy		
3D RMSD (pm)	24	15
# Position matches to GT	7,677	8,101
% GT model matched	94.30%	99.51%
% Test model matched	94.02%	99.62%
Classification accuracy		
Type 1 accuracy (%)	77.03%	98.08%
Type 2 accuracy (%)	98.68%	99.17%
Type 3 accuracy (%)	100.00%	100.00%
Overall classification accuracy (%)	92.67%	99.08%
# Total accurate matches	7,544	8,066

^a Model originally published in Ref.⁹.

^b The R_1 factor is defined as in equation 5 of Ref.⁴⁷.

Medium-entropy alloy data

We next apply RESIRE-Net to a previously published medium-entropy alloy (MEA) nanoparticle dataset⁹. Unlike the metallic glass examined above, this material exhibits a predominantly FCC lattice, allowing reconstruction fidelity to be assessed using lattice-based metrics even in the absence of a ground-truth atomic model.

Simulation under experimental condition

Compared to the metallic glass dataset analyzed above, the original MEA study⁹ employed a lower-dose, reduced angular-range tilt series, resulting in a more challenging reconstruction problem. To assess the performance of RESIRE-Net under these conditions, we first considered a simulated dataset from the original study, generated using the same geometry and imaging parameters as the experiment (Supplementary Table 1), and replicated the originally published model-extraction pipeline⁹ in full, with only the reconstruction algorithm replaced by RESIRE-Net. As anticipated, the RESIRE reconstruction exhibits pronounced missing-wedge artefacts, manifested as elongation in xz slices and a wedge of suppressed intensity in Fourier space, whereas RESIRE-Net partially restores Fourier amplitudes within the unsampled k_z region and produces sharper, more localized atomic peaks (Supplementary Fig. 1). Results of the final model comparison are summarized in Table 2, Supplementary Table 4, and Supplementary Fig. 5: RESIRE-Net improves model fidelity, reducing the 3D RMSD from 24 pm to 15 pm (~37.5% lower positional error) while increasing ground-truth atom recovery (from 94.30% to 99.51%), lowering the refinement error (from 5.28% to 3.90%), and substantially enhancing classification accuracy (overall: from 92.67% to 99.08%; Type 1: from 77.03% to 98.08%). These results confirm that the method performs robustly under less ideal experimental conditions, with increasingly pronounced improvements over RESIRE as reconstruction difficulty increases.

Experimental reconstruction

We then apply RESIRE-Net to the experimental MEA tilt series (“MEA-2”) reported in Ref.⁹, whose original RESIRE reconstruction revealed an FCC lattice containing twin boundaries, lattice distortion, and heterogeneous chemical short-range order (CSRO). We analyze the RESIRE-Net reconstruction using the same model-extraction pipeline as in the original study to ensure methodological consistency. Consistent with the simulation results, the RESIRE-Net reconstruction exhibits reduced missing-wedge artifacts (Fig. 4c,d) and improved axial resolution in xz slices (Fig. 4a,b and Supplementary Fig. 1), and these qualitative improvements are reflected in the quantitative structural metrics summarized in Table 3. The pair distribution function (PDF) (Fig. 4e) shows a sharper first peak, and the first minimum shifts from 3.44 Å (RESIRE) to 3.41 Å, indicating closer agreement with an ideal FCC lattice. Local lattice distortion analysis further confirms this improvement: the mean deviation from a locally fitted FCC lattice decreases from 0.26 ± 0.12 Å in the RESIRE model to 0.23 ± 0.11 Å for RESIRE-Net (Fig. 4f).




Figure 4. Demonstration of RESIRE-Net on previously-published experimental medium entropy alloy (MEA) nanoparticle data. **a, b**, Two 2.5-Angstrom-thick slices through the xz plane of the experimental RESIRE (**a**) and RESIRE-Net (**b**) reconstructions. Scale bar, 2 nm. **c, d**, the FFTs corresponding to the real-space slices (**a, b**). **e**, Comparison of the normalized pair distribution function (PDF) of each reconstruction's model. The RESIRE-Net model had a slightly more pronounced first peak, corresponding to a more ideal FCC-structure, as expected. **f**, Distributions of the 3D atomic deviations of each model from an ideal FCC lattice, with the mean and standard deviation of the RESIRE model calculated as $0.26 \pm 0.12 \text{ \AA}$ and for the RESIRE-Net model calculated as $0.23 \pm 0.11 \text{ \AA}$. **g, j**, Comparison of the histograms of the average α_{NiPt} (**g**) and α_{PdPt} (**j**) values for each atomic layer along the $[111]$ direction, with dashed yellow lines indicating the positions of the two twin boundaries. **h, k**, 3D distribution of α_{NiPt} (**h**) and α_{PdPt} (**k**) in the RESIRE model. **i, l**, 3D distribution of α_{NiPt} (**i**) and α_{PdPt} (**l**) in the RESIRE-Net model. Scale bar, 1 nm. RESIRE results are reproduced from Ref.⁹ for comparison.

Table 3. Comparison of final model structure and composition for an experimental medium entropy alloy nanoparticle dataset.

	RESIRE ^a	RESIRE-Net
Final model		
# Type 1 atoms	2,522	2,698
# Type 2 atoms	3,247	3,397
# Type 3 atoms	2,580	2,569
# Total model atoms	8,349	8,664
Refinement		
Error (R_1) ^b	10.29%	9.97%
PDF statistics		
Location of 1st peak (Å)	2.76	2.75
Location of 1st minimum (Å)	3.44	3.41
Local fit to ideal FCC lattice		
Model size (excluding surface atoms)	6,850	7,137
Mean $\pm$ std (Å)	0.26 ± 0.12	0.23 ± 0.11

^a Model originally published as “MEA-2” in Ref.⁹.

^b The R_1 factor is defined as in equation 5 of Ref.⁴⁷.

CSRO, based on atomic positions and classifications, is preserved across reconstructions. Fig. 4g–i compares the α_{NiPt} ordering parameter obtained with RESIRE-Net and RESIRE, while Fig. 4j–l shows the corresponding comparison for α_{PdPt} ; the dominant features remain unchanged, with α_{NiPt} predominantly positive in the particle interior and α_{PdPt} largely negative, indicating a preference for Ni-Pt bonding and an avoidance of Pd-Pt bonding. The results of all six CSRO parameters are summarized in Supplementary Fig. 6.

Across both simulated and experimental datasets, RESIRE-Net improves reconstruction fidelity and atomic localization while remaining fully compatible with existing RESIRE-based analysis pipelines. By mitigating missing-wedge artifacts and producing sharper atomic intensity peaks, the method enhances positional precision, atom recovery, and chemical classification, thereby reducing the need for manual intervention in AET workflows and strengthening the reliability of downstream structural and chemical analysis.

Discussion

RESIRE-Net substantially improves missing-wedge recovery in electron tomography by integrating self-supervised learning with the RESIRE reconstruction algorithm in an operator-aware framework. Across both simulated and experimental datasets, the method consistently enhances depth resolution, suppresses missing-wedge artifacts, and improves atomic-scale accuracy. These improvements arise from explicitly modeling the missing-wedge effect as a reconstruction-dependent operator defined by the experimental tilt geometry and reconstruction algorithm. By learning an approximate inverse operator through rotational augmentation and iterative refinement, RESIRE-Net effectively recovers information that is otherwise difficult to reconstruct using conventional approaches alone.

In the metallic-glass simulation, RESIRE-Net reduces the three-dimensional positional RMSD by 33% relative to RESIRE, with the largest improvements occurring along the depth direction. The method also improves atom recovery, reduces spurious detections, and enhances chemical-species classification, demonstrating that the recovered structural information directly benefits downstream atomic-model extraction. Similar trends are observed in the medium-entropy alloy dataset, where RESIRE-Net sharpens atomic peaks, restores Fourier amplitudes within the missing-wedge region, and yields lattice metrics that more closely match ideal FCC behavior. Together, these results demonstrate the robustness of the framework across both simulated and experimental conditions.

Compared with existing self-supervised approaches, RESIRE-Net performs experiment-specific correction by explicitly accounting for the measured tilt geometry and reconstruction process. This coupling to the forward model provides a more physically grounded treatment of missing-wedge artifacts and enables recovery beyond what can be achieved using isotropy-based augmentation alone. Rather than relying solely on generic image priors, the framework incorporates reconstruction physics directly into the learning process, providing a flexible strategy that can be adapted to different acquisition conditions and reconstruction operators.

A limitation of the current framework is its reliance on rotational augmentation and the associated assumption of approximate isotropy. For objects with strong intrinsic anisotropy or limited symmetry, RESIRE-Net may partially interpret genuine structural anisotropy as missing-wedge artifacts, potentially leading to overcorrection. Experimental effects such as probe polarization, electron channeling, detector non-uniformity, noise, tilt-angle miscalibration, mechanical drift, and alignment errors may further limit reconstruction quality by introducing inaccuracies into the operator $\mathbf{T}_\theta$ ⁴⁸. Future work may explore generative priors, physics-informed constraints, and implicit neural representations to further improve reconstruction fidelity and robustness. Overall, RESIRE-Net provides a practical and operator-aware solution to missing-wedge artifacts in AET, enabling more accurate structural and chemical characterization at the atomic scale.

References

1. Miao, J. Computational microscopy with coherent diffractive imaging and ptychography. *Nature* **637**, 281–295 (2025).
2. Liao, Y. *et al.* Accurate determination of the 3d atomic structure of amorphous materials. *Nature* **649**, 1123–1129 (2026).
3. Midgley, P. A. & Dunin-Borkowski, R. E. Electron tomography and holography in materials science. *Nat. materials* **8**, 271–280 (2009).
4. Scott, M. C. *et al.* Electron tomography at 2.4-ångström resolution. *Nature* **483**, 444–447 (2012).
5. Miao, J., Ercius, P. & Billinge, S. J. L. Atomic electron tomography: 3D structures without crystals. *Science* **353**, aaf2157 (2016).
6. Chen, C.-C. *et al.* Three-dimensional imaging of dislocations in a nanoparticle at atomic resolution. *Nature* **496**, 74–77 (2013).
7. Yang, Y. *et al.* Determining the three-dimensional atomic structure of an amorphous solid. *Nature* **592**, 60–64 (2021).
8. Yuan, Y. *et al.* Three-dimensional atomic packing in amorphous solids with liquid-like structure. *Nat. Mater.* **21**, 95–102 (2022).
9. Moniri, S. *et al.* Three-dimensional atomic structure and local chemical order of medium-and high-entropy nanoalloys. *Nature* **624**, 564–569 (2023).
10. Yang, Y. *et al.* Atomic-scale identification of active sites of oxygen reduction nanocatalysts. *Nat. Catal.* **7**, 796–806 (2024).
11. Zhou, J. *et al.* Observing crystal nucleation in four dimensions using atomic electron tomography. *Nature* **570**, 500–503 (2019).
12. Yuan, Y. *et al.* Crystal nucleation and growth in high-entropy alloys revealed by atomic electron tomography. *Nat. Material* (in proceeding).
13. Kak, A. C. & Slaney, M. *Principles of Computerized Tomographic Imaging* (Society for Industrial and Applied Mathematics, 2001).
14. Miao, J., Charalambous, P., Kirz, J. & Sayre, D. Extending the methodology of X-ray crystallography to allow imaging of micrometre-sized non-crystalline specimens. *Nature* **400**, 342–344 (1999).
15. Miao, J., Förster, F. & Levi, O. Equally sloped tomography with oversampling reconstruction. *Phys. Rev. B* **72**, 52103 (2005).
16. Pryor, A. *et al.* GENFIRE: A generalized Fourier iterative reconstruction algorithm for high-resolution 3D imaging. *Sci. Reports* **7**, 1–12 (2017).
17. Pham, M., Yuan, Y., Rana, A., Osher, S. & Miao, J. Accurate real space iterative reconstruction (RESIRE) algorithm for tomography. *Sci. Reports* **13**, 5624 (2023).
18. Goris, B., Van den Broek, W., Batenburg, K. J., Mezerji, H. H. & Bals, S. Electron tomography based on a total variation minimization reconstruction technique. *Ultramicroscopy* **113**, 120–130 (2012).
19. Leary, R., Saghi, Z., Midgley, P. A. & Holland, D. J. Compressed sensing electron tomography. *Ultramicroscopy* **131**, 70–91 (2013).
20. Deng, Y. *et al.* ICON: 3D reconstruction with ‘missing-information’ restoration in biological electron tomography. *J. structural biology* **195**, 100–112 (2016).
21. Lee, J. & Yang, Y. Advancing atomic electron tomography with neural networks. *Appl. Microsc.* **55**, 7 (2025).
22. 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. *Sci. reports* **9**, 12803 (2019).

23. Wang, C., Ding, G., Liu, Y. & Xin, H. L. 0.7 Å resolution electron tomography enabled by deep-learning-aided information recovery. *Adv. Intell. Syst.* **2**, 2000152 (2020).
24. Lee, J., Jeong, C. & Yang, Y. Single-atom level determination of 3-dimensional surface atomic structure via neural network-assisted atomic electron tomography. *Nat. communications* **12**, 1962 (2021).
25. Chen, D., Tachella, J. & Davies, M. E. Equivariant imaging: Learning beyond the range space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4379–4388 (2021).
26. Liu, Y.-T. *et al.* Isotropic reconstruction for electron tomography with deep learning. *Nat. communications* **13**, 6482 (2022).
27. Yao, L., Lyu, Z., Li, J. & Chen, Q. No ground truth needed: unsupervised sinogram inpainting for nanoparticle electron tomography (UsiNet) to correct missing wedges. *npj Comput. Mater.* **10**, 28 (2024).
28. Wiedemann, S. & Heckel, R. A deep learning method for simultaneous denoising and missing wedge reconstruction in cryogenic electron tomography. *Nat. Commun.* **15**, 8255 (2024).
29. Zhao, C., Ge, M., Yang, X., Chu, Y. S. & Yan, H. Limited-angle X-ray nano-tomography with machine-learning enabled iterative reconstruction engine. *npj Comput. Mater.* **11**, 240 (2025).
30. Li, H. *et al.* Generative-model-based null-space iterative reconstruction for atomic electron tomography with sparse data. *Nat. Commun.* **17**, 7643 (2026).
31. Li, H. *et al.* A general data-driven framework for scalable electron tomography. *Natl. Sci. Rev.* **13**, nwag365 (2026).
32. Golub, G. H., Hansen, P. C. & O’Leary, D. P. Tikhonov regularization and total least squares. *SIAM journal on matrix analysis applications* **21**, 185–194 (1999).
33. Strong, D. & Chan, T. Edge-preserving and scale-dependent properties of total variation regularization. *Inverse problems* **19**, S165 (2003).
34. Chambolle, A. An algorithm for total variation minimization and applications. *J. Math. imaging vision* **20**, 89–97 (2004).
35. Pham, M., Lu, X., Rana, A., Osher, S. & Miao, J. Real space iterative reconstruction for vector tomography (RESIRE-V). *Sci. Reports* **14**, 9541 (2024).
36. Wolny, A. *et al.* Accurate and versatile 3D segmentation of plant tissues at cellular resolution. *Elife* **9**, e57613 (2020).
37. Gurrola-Ramos, J., Dalmau, O. & Alarcón, T. E. A residual dense u-net neural network for image denoising. *IEEE Access* **9**, 31742–31754 (2021).
38. Lian, Z. & Wang, H. An image deblurring method using improved U-Net model based on multilayer fusion and attention mechanism. *Sci. Reports* **13**, 21402 (2023).
39. Gao, H., Tao, X., Shen, X. & Jia, J. Dynamic scene deblurring with parameter selective sharing and nested skip connections. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3848–3856 (2019).
40. Dong, J., Pan, J., Yang, Z. & Tang, J. Multi-scale residual low-pass filter network for image deblurring. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12345–12354 (2023).
41. Peng, Y., Sonka, M. & Chen, D. Z. U-Net v2: Rethinking the skip connections of U-Net for medical image segmentation. *arXiv preprint arXiv:2311.17791* (2023).
42. Duan, L., Sun, S., Zhang, J. & Xu, Z. Deblurring turbulent images via maximizing L1 regularization. *Symmetry* **13**, 1414 (2021).
43. Eqtedaei, A. & Ahmadyfard, A. Blind image deblurring using both L0 and L1 regularization of Max-min prior. *Neurocomputing* **592**, 127727 (2024).
44. Cui, Y., Tao, Y., Ren, W. & Knoll, A. Dual-Domain attention for image deblurring. *Proc. AAAI Conf. on Artif. Intell.* **37**, 479–487 (2023).
45. Kupyn, O., Martyniuk, T., Wu, J. & Wang, Z. Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8878–8887 (2019).
46. Kupyn, O., Budzan, V., Mykhailych, M., Mishkin, D. & Matas, J. Deblurgan: Blind motion deblurring using conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8183–8192 (2018).
47. Yang, Y. *et al.* Deciphering chemical order/disorder and material properties at the single-atom level. *Nature* **542**, 75–79 (2017).

48. Iwashimizu, C., Haruta, M. & Kurata, H. Anisotropic atomic contrast by final-state selection in electron energy-loss spectroscopy. *Phys. Rev. B* **111**, 155147 (2025).

Acknowledgments

This work was supported by the US Department of Energy (DOE), Office of Science, Basic Energy Sciences (BES), Division of Materials Sciences and Engineering under grant number DE-SC0010378 on the electron tomography experiments as well as supported by STROBE: A National Science Foundation Science and Technology Center under grant number DMR 1548924 on the development of the RESIRE-Net. The ADF-STEM imaging with TEAM 0.5 was performed at the Molecular Foundry, Lawrence Berkeley National Laboratory, which is supported by the Office of Science, Office of Basic Energy Sciences of the U.S. DOE under Contract No. DE-AC02—05CH11231

Author contributions

J.M. directed the project; M.P. developed the RESIRE-Net algorithm with input from J.M. A.M. tested the algorithm and analyzed the results. All authors contributed to writing and reviewing the manuscript.

Data availability

MATLAB source codes of RESIRE-Net and simulated and experimental data are available at the GitHub repository <https://github.com/minhpham0309/RESIRE-Net>.

Competing interests

The authors declare no competing interests.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [Supplementary.pdf](#)