

Improving Multislice Electron Ptychography with a Generative Prior

Christian K. Belardi*, Chia-Hao Lee*, Yingheng Wang, Justin Lovelace,
Kilian Q. Weinberger, David A. Muller, Carla P. Gomes
Cornell University

{ckb73, c12696, david.a.muller}@cornell.edu, gomes@cs.cornell.edu

Abstract

Multislice electron ptychography (MEP) is an inverse imaging technique that computationally reconstructs the highest-resolution images of atomic crystal structures from diffraction patterns. Available algorithms often solve this inverse problem iteratively but are both time consuming and produce suboptimal solutions due to their ill-posed nature. We develop MEP-DIFFUSION, a diffusion model trained on a large database of crystal structures specifically for MEP to augment existing iterative solvers. MEP-DIFFUSION is easily integrated as a generative prior into existing reconstruction methods via Diffusion Posterior Sampling (DPS). We find that this hybrid approach greatly enhances the quality of the reconstructed 3D volumes, achieving a 90.50% improvement in SSIM over existing methods.

1. Introduction

Understanding the crystal structure of materials is pivotal for unraveling the fundamental interactions that govern the physical world. Electron microscopy has long been the cornerstone of atomic-scale imaging that provides unmatched spatial resolution. Among its advanced methodologies, *multislice electron ptychography* (MEP) has emerged as a powerful computational imaging technique, capable of achieving sub-angstrom resolutions by formulating the multiple electron scattering process as an inverse problem [5, 25, 56]. This capability allows researchers to reconstruct atomic structures from diffraction patterns with unprecedented clarity.

However, the current state-of-the-art in electron ptychography faces critical challenges. While effective at resolving structures in the lateral (i.e., height and width) dimensions, the depth resolution is often limited to around 2 nm, or 100 times worse than the lateral resolutions due to experimental constraints [5, 49]. Additionally, conventional iterative solvers are very time-consuming, and can take hours to days

*Equal Contribution

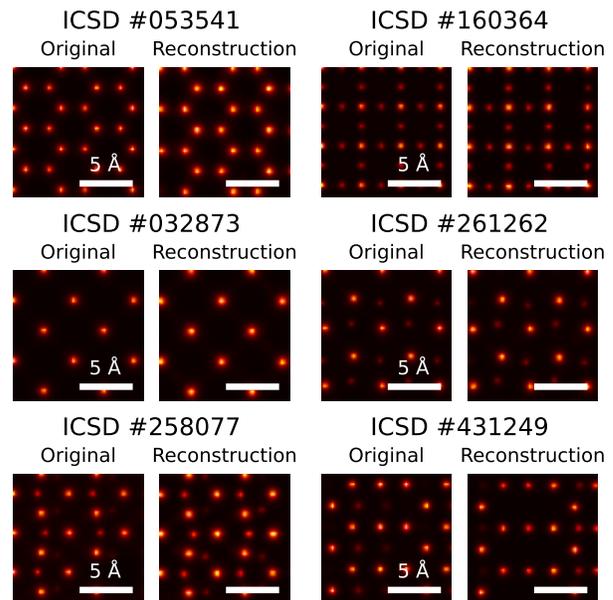


Figure 1. Side by side comparison of ground truth and MEP-DIFFUSION reconstruction for selected crystal structure. The corresponding diffraction patterns are shown in [Appendix H](#).

to return a solution, while the data acquisition only takes a few second [9, 34, 48, 54]. Lastly, these solvers can struggle with noisy measurements and produce suboptimal solutions due to the inherent ill-posed nature of the inverse problem.

To address these limitations, we develop a generative prior using diffusion probabilistic models [14, 43, 44] and demonstrate its utility when combined with existing iterative solver. Diffusion models have demonstrated remarkable capability in various inverse problems by leveraging learned priors to produce high-quality reconstructions [7, 11], and have been explored in Fourier ptychography [41, 51] and X-ray ptychography [2] for bulk specimens. In this work, we introduce MEP-DIFFUSION, which seamlessly integrates diffusion models into multislice electron ptychography. Our approach enables hybrid inverse problem solving, augmenting existing iterative solvers with

a data driven model that captures the manifold of possible crystal structures. Figure 1 shows that MEP-DIFFUSION produces high quality samples that match closely with the ground truth while remaining consistent with the physical measurements.

2. Preliminaries

Physical Forward Model In electron ptychography, we consider the measurement model in the following form:

$$\mathbf{y} = f(\mathbf{x}; \phi, \Omega)$$

where f represents the forward model, describing how the complex-valued electron probe ϕ interacts with the material \mathbf{x} at different probe positions Ω . The interaction produces a scattered wave function that propagates to a far-field detector, yielding real-valued diffraction patterns \mathbf{y} .

The materials are generally described by a complex-valued function $\mathbf{x} := \mathbf{O}(\mathbf{r}) \in \mathbb{C}^{D \times H \times W}$, where $\mathbf{r} = (x, y, z)$ represents spatial coordinates. Under the *strong phase approximation (SPA)* [23], $\mathbf{O}(\mathbf{r}) \approx \exp(i\sigma_e V(\mathbf{r}))$ where i is the imaginary unit, σ_e is the electron interaction parameter, $V(\mathbf{r})$ is the atomic scattering potential that captures the atom types and positions.

To simplify, we absorb the exponential function into the forward model, and redefine the crystal structure as $\mathbf{x} := \sigma_e V(\mathbf{r}) \in \mathbb{R}^{D \times H \times W}$. Because σ_e is just a scaling constant that does not depend on \mathbf{r} , \mathbf{x} contains all the structural information that electron ptychography aims to recover. By fixing both the probe ϕ and the probe positions Ω , we arrive at a simplified forward model:

$$\hat{\mathbf{y}} = f(\hat{\mathbf{x}}), \quad f: \mathbb{R}^{D \times H \times W} \rightarrow \mathbb{R}^{N_y \times N_x \times K_y \times K_x}$$

where N_y and N_x correspond to different probe positions, while K_y and K_x denote the number of detector pixels.

The goal of electron ptychography is to retrieve the crystal structure \mathbf{x} from diffraction patterns \mathbf{y} , which is commonly achieved by solving the following optimization problem:

$$\min_{\mathbf{x}} \|\hat{\mathbf{y}} - \mathbf{y}\|_2^2 + \gamma \|\mathbf{x}\|_1$$

where the first term ensures fidelity to the measured data, while the second term, which imposes sparsity in the solution and is controlled by the scalar weighting parameter γ is optional, but commonly added to promote discrete atomicity within crystal structures.

We refer to this formulation as the *electron ptychography reconstruction objective*. In prior work, this objective is typically minimized using first-order optimization methods, often requiring additional regularizations for stability.

Diffusion Probabilistic Models Given a dataset drawn from an unknown distribution $q(\mathbf{x})$, diffusion probabilistic

models (DPMs) [14, 20, 42, 46] learn a generative model $p_\theta(\mathbf{x})$ that approximates the data distribution $q(\mathbf{x})$. DPMs consist of a forward process and a generative process. The forward process defines a gradual transition between samples from the dataset distribution to samples from a Gaussian distribution. This introduces a series of increasingly noisy latent variables \mathbf{z}_t for timesteps $t \in [0, 1]$. For every timestep, the conditional $q_t(\mathbf{z}_t|\mathbf{x})$ is distributed such that

$$\mathbf{z}_t = \alpha_t \mathbf{x} + \sigma_t \epsilon, \quad \text{where } \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}).$$

Here, $\alpha_t, \sigma_t \in [0, 1]$ are smooth scalar functions of the timestep t which describe the *noise schedule*. The noise schedule is defined such that the process starts with the clean data, $\mathbf{z}_0 \approx \mathbf{x}$, and the final latent is approximately Gaussian, $q(\mathbf{z}_1) \approx \mathcal{N}(\mathbf{0}, \mathbf{I})$. The diffusion process can be equivalently expressed in terms of the *log signal-to-noise ratio* ($\log \text{SNR}$), $\lambda_t = \log(\alpha_t^2 / \sigma_t^2)$, which decreases monotonically with t [22].

The generative process inverts the forward process, defining a gradual transition from samples of Gaussian noise to samples from the data distribution. Given access to the *score function* $\nabla_{\mathbf{z}_t} \log q_t(\mathbf{z}_t)$, the forward process can be inverted exactly [46]. The key idea of diffusion models is to introduce a neural *score network* to approximate the score function $\mathbf{s}_\theta(\mathbf{z}_t, \lambda_t) \approx \nabla_{\mathbf{z}_t} \log q_t(\mathbf{z}_t)$. The score network is trained with a denoising score matching objective over noise scales and data samples $\mathbf{x} \sim q(\mathbf{x})$:

$$\mathcal{L}(\mathbf{x}) = \mathbb{E}_{\epsilon, t} [w(\lambda_t) \|\mathbf{s}_\theta(\mathbf{z}_t, \lambda_t) - \nabla_{\mathbf{z}_t} \log q_t(\mathbf{z}_t|\mathbf{x})\|_2^2],$$

where $\mathbf{z}_t = \alpha_t \mathbf{x} + \sigma_t \epsilon$.

In practice, the neural network can equivalently be parameterized as a noise (ϵ) prediction, data (\mathbf{x}) prediction, or a velocity ($\mathbf{v} := \alpha_t \epsilon - \sigma_t \mathbf{x}$) prediction network due to the following relationship

$$\nabla_{\mathbf{z}_t} \log q_t(\mathbf{z}_t|\mathbf{x}) = -\frac{\epsilon}{\sigma_t} = -\frac{\mathbf{z}_t - \alpha_t \mathbf{x}}{\sigma_t^2} = -\mathbf{z}_t - \frac{\alpha_t}{\sigma_t} \mathbf{v}.$$

These different parameterizations can be interpreted as imposing different weighting functions, $w(\lambda_t)$, over the denoising score matching objective and have been found to improve training stability and performance [14, 20]. We follow the best practices established in the image diffusion literature [20] and parameterize our score network as a velocity prediction model, $\mathbf{v}_\theta(\mathbf{z}_t, \lambda_t)$ [40].

Solving Inverse Problems with Diffusion Models During generation, we may want to impose some constraint over the generated data. For instance, we may have a partial measurement \mathbf{y} , and we want to draw samples \mathbf{x} consistent with our measurement \mathbf{y} . To draw such samples, it is sufficient to approximate the conditional score function $\nabla_{\mathbf{z}_t} \log p_t(\mathbf{z}_t|\mathbf{y})$. We can decompose the conditional score

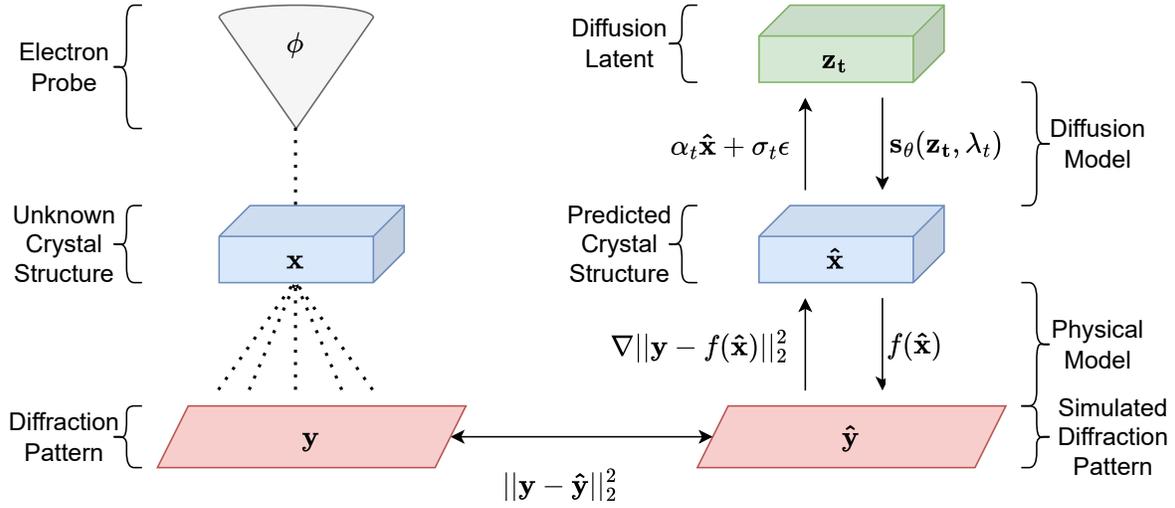


Figure 2. Overview connecting electron ptychography and diffusion posterior sampling. (Left) In the experimental setup, a focused electron probe is scanned across an unknown crystal structure on a 2D grid, recording a 2D diffraction pattern at each scan position. This results in a 4D dataset comprising a stack of 2D diffraction patterns over a 2D scan grid. (Right) MEP-DIFFUSION reconstructs the crystal structure by combining a diffusion model prior with a differentiable physical model. The physical model guides the diffusion sampling process to ensure consistency between the generated structure and the observed diffraction patterns.

function into the unconditional score and a likelihood term by Bayes' Rule:

$$\nabla_{\mathbf{z}_t} \log p_t(\mathbf{z}_t | \mathbf{y}) = \nabla_{\mathbf{z}_t} \log p_t(\mathbf{z}_t) + \nabla_{\mathbf{z}_t} \log p_t(\mathbf{y} | \mathbf{z}_t).$$

This expansion shows that an unconditional diffusion model, which estimates $\nabla_{\mathbf{z}_t} \log p_t(\mathbf{z}_t)$, can incorporate additional constraints during generation via the likelihood term $p_t(\mathbf{y} | \mathbf{z}_t)$. Critically, this is added during sampling and does not require re-training the diffusion network.

Diffusion Posterior Sampling For a variety of applications, we may have measurements for the *clean* data $p(\mathbf{y} | \mathbf{x})$, but obtaining the distribution over *noised* data $p_t(\mathbf{y} | \mathbf{z}_t)$ is infeasible. In such cases, we perform DPS [7] which uses the minimum mean squared estimate (MMSE) of the *clean* data provided by the diffusion model, $\mathbf{x}_\theta(\mathbf{z}_t, t)$, and the likelihood term over the *clean* data, $p(\mathbf{y} | \mathbf{x})$, to approximate the conditional score function:

$$\begin{aligned} \nabla_{\mathbf{z}_t} \log p_t(\mathbf{z}_t | \mathbf{y}) &= \nabla_{\mathbf{z}_t} \log p_t(\mathbf{z}_t) + \nabla_{\mathbf{z}_t} \log p_t(\mathbf{y} | \mathbf{z}_t) \\ &\approx \nabla_{\mathbf{z}_t} \log p_t(\mathbf{z}_t) + \nabla_{\mathbf{z}_t} \log p_t(\mathbf{y} | \mathbf{x}_\theta(\mathbf{z}_t, \lambda_t)). \end{aligned}$$

This enables the direct incorporation of clean data measurements into the sampling process. At each sampling step, regardless of the parametrization of the diffusion model, we can compute the MMSE in closed form.

3. Ptychographic Reconstruction with DPMs

In electron ptychography, a focused electron probe is scanned across a material with an unknown crystal structure, producing a series of diffraction patterns. Our objective is to reconstruct the crystal structure that gave rise to these observed patterns. To accomplish this, we train a diffusion model to learn the underlying distribution of possible crystal structures and combine it with a differentiable physical model to sample crystal structures from the diffusion model that are physically consistent with the observed diffraction patterns. An overview of this framework is provided in Figure 2.

DPMs for Crystal Structures We closely follow Kingma and Gao [20] in the design of our diffusion model, and parameterize the diffusion process in terms of the continuous log SNR, denoted λ_t . We use a variance-preserving forward process, with

$$\alpha_t^2 = \text{sigmoid}(\lambda_t), \quad \sigma_t^2 = \text{sigmoid}(-\lambda_t),$$

where $\lambda_t = \log(\alpha_t^2 / \sigma_t^2)$ is the log SNR.

We optimize our network for the “v”-prediction task proposed by Salimans and Ho [40] using a weighted mean squared error loss. Our training objective is therefore

$$\mathbb{E}_{\epsilon \sim \mathcal{N}(0, \mathbf{I}), t \sim \mathcal{U}(0, 1)} [w(\lambda_t) \|\mathbf{v}_\theta(\mathbf{z}_t, \lambda_t) - \mathbf{v}\|_2^2].$$

As shown in Figure 3, we observe that periodic crystal

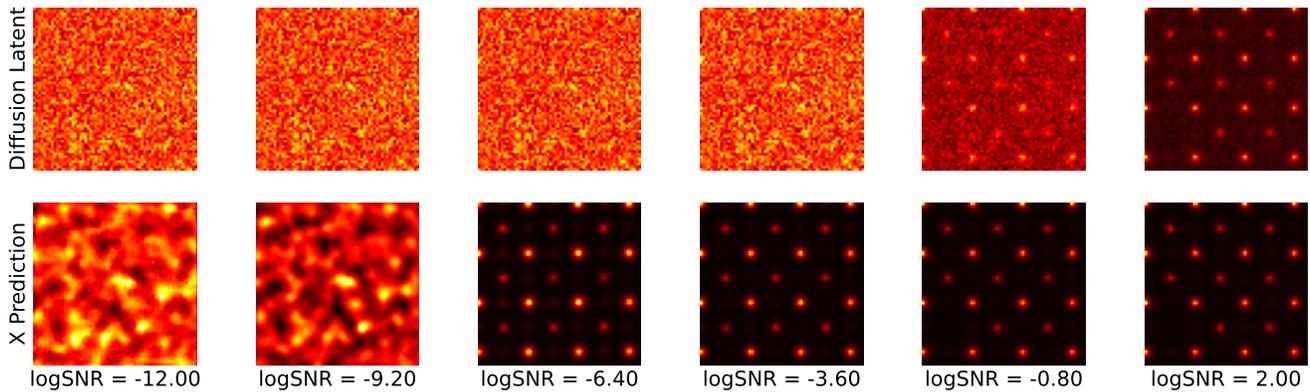


Figure 3. (Top) Visualization of the diffusion process across noise levels. (Bottom) Estimates of the original data by the trained diffusion network. We observe that periodic structure emerges very early in the diffusion process. Training at lower log SNR ranges compared to typical image diffusion models is critical for the model to learn a strong prior over atomic arrangements.

structure emerges very early in the diffusion process at relatively low log SNR. While diffusion models for natural images typically focus on moderate noise levels, we find it critical to concentrate training at higher noise levels where this periodic structure first emerges. See Appendix F for more discussion.

Following Karras et al. [19], we implement this using a Gaussian weighting in log SNR space: $w(\lambda_t) = \frac{1}{Z} \mathcal{N}(\lambda_t; -7, 3)$, where Z is a normalization constant ensuring the density integrates to 1. This formulation centers the weighting at $\lambda_t = -7$, emphasizing high-noise regions during training. We find that shifting the loss weighting towards higher noise levels is critical for generating globally coherent samples. This observation aligns with recent findings in high-resolution image diffusion tasks [4, 17, 18]. During training, we adopt the adaptive noise schedule from Kingma and Gao [20] to reduce variance.

Since the voxel intensities of the crystals are exponentially distributed, predicting \mathbf{x} as the zero tensor is a strong baseline. To discourage this, we rescale our data so the standard deviation is approximately 1, corresponding to a constant scaling factor of $c = 25$. This rescaling effectively up-weights the “x”-prediction component of the “v”-prediction task, for additional details see Appendix B.

Given the importance of capturing periodic structure at low log SNR, we find that traditional image diffusion schedules like the cosine schedule perform poorly. Instead, we use a linear log SNR sampling schedule that concentrates steps in the range $[-13, 4]$, where periodicity emerges. This enables our model to dedicate more computation to the noise levels where atomic arrangements first appear. For improved sampling efficiency, we adopt the second-order DPMSolver++ sampler [28], which outperforms first-order methods in our experiments.

DPS with the MEP Physical Model As shown in Chung et al. [7], we can compute the score of the posterior from the sum of the scores of the prior and likelihood. The score of the prior is given by our unconditional diffusion model, MEP-DIFFUSION. The score of the likelihood can be instantiated as the loss of a classifier, or in our case the loss of the physical model.

$$\nabla_{\mathbf{z}_t} \log p_t(\mathbf{z}_t | \mathbf{y}) = \underbrace{\nabla_{\mathbf{z}_t} \log p_t(\mathbf{z}_t)}_{\text{Diffusion Model}} + \underbrace{\nabla_{\mathbf{z}_t} \log p_t(\mathbf{y} | \mathbf{z}_t)}_{\text{Physical Model Gradient}}.$$

Practically, following Ho et al. [15], we can implement this guidance as an adjustment to the “x”-prediction:

$$\mathbf{x}'_{\theta}(\mathbf{z}_t, \lambda_t) = \mathbf{x}_{\theta}(\mathbf{z}_t, \lambda_t) - g(\lambda_t) \nabla_{\mathbf{z}_t} \|f(\mathbf{x}_{\theta}(\mathbf{z}_t, \lambda_t)) - \mathbf{y}\|_2^2$$

where $g(\lambda_t)$ is a guidance weight that controls the strength of the physical model contribution. While setting $g(\lambda_t) \propto \alpha_t$ yields $\mathbf{x}'_{\theta}(\mathbf{z}_t, \lambda_t) \approx \mathbb{E}_q[\mathbf{x} | \mathbf{z}_t, \mathbf{y}]$, we find that strong guidance towards the end of sampling leads to degenerate solutions with poor depth resolution (i.e., all the depth slices are nearly equivalent). We therefore anneal the guidance weight to smoothly reduce the influence of the physical model as sampling progresses, setting $g(\lambda_t) \propto \text{sigmoid}(4 - \lambda_t)^{1/2}$. We visualize our proposed guidance schedule in Figure 4.

Intuitively, the physical model guides the generation towards crystal structures consistent with the observed diffraction patterns while the diffusion model ensures atomic-scale details match the distribution of known crystal structures. This complementary approach is particularly effective for electron ptychography reconstruction — using physical guidance early in sampling establishes the correct periodic structure, while the learned prior later refines atomic-scale details that the physical model cannot resolve.

Table 1. Quantitative comparison of methods on instance-normalized reconstructions. We evaluate both full 3D reconstruction and depth-summed 2D projections using standard image quality metrics. Results show mean \pm standard error of the mean across our test set. DPS gives statistically significant improvements ($p < 1e - 4$) over baseline methods on all metrics.

METHOD	FULL MATERIAL RECONSTRUCTION		MATERIAL DEPTH SUM RECONSTRUCTION	
	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
LSQ-ML	27.4900 \pm 0.1455	0.3949 \pm 0.0090	27.9689 \pm 0.1541	0.4279 \pm 0.0096
PTYRAD[ADAM]	26.8527 \pm 0.1334	0.4629 \pm 0.0104	27.4223 \pm 0.1414	0.4896 \pm 0.0108
PTYRAD[L-BFGS]	31.4885 \pm 0.1501	0.6506 \pm 0.0090	33.5574 \pm 0.2003	0.7096 \pm 0.0095
MEP-DIFFUSION	33.7190 \pm 0.1243	0.7523 \pm 0.0053	41.5101 \pm 0.1721	0.8852 \pm 0.0041

This careful orchestration enables reconstruction of high-quality structures that are both physically consistent with measurements and structurally realistic at the atomic scale.

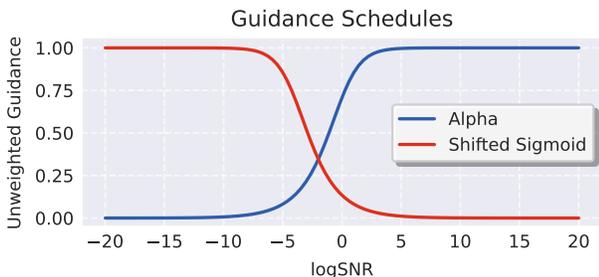


Figure 4. Illustration of guidance schedules: Alpha $g(\lambda_t) \propto \alpha_t$, and Shifted Sigmoid $g(\lambda_t) \propto \text{sigmoid}(4 - \lambda_t)^{1/2}$

3D UNet with Anisotropic Resampling We modify the classic UNet architecture to accommodate the volumetric nature of crystal structures, as shown in Figure 8 of Appendix D. The spatial sampling of our crystal data is anisotropic, with coarser sampling along the depth dimension (1.6 Å) compared to the lateral dimensions (0.2 Å). To handle this sampling anisotropy, we resample in two-stages. In the first stage, the UNet performs downsampling exclusively along height and width, reducing the feature spatial sampling from (1.6 Å, 0.2 Å, 0.2 Å) to (1.6 Å, 1.6 Å, 1.6 Å). In the second stage, downsampling is applied evenly across depth, height, and width. During upsampling, the process is reversed. First upsampling is applied uniformly across all dimensions and then only to the lateral dimensions, restoring the original anisotropic spatial sampling. Additional details are provided in Appendix D.

4. Experiments

Dataset Preparation We generated train, validation, and test datasets using the crystallographic information

files included in the Inorganic Crystal Structure Database (ICSD) [53]. ICSD is the largest database for experimentally identified crystal structures and contains more than 258,000 entries that describe the atom types and relative positions, which provides informative prior knowledge of existing crystals. We filter the ICSD dataset so only the crystal structures that have all three lattice vectors ($|\vec{a}|$, $|\vec{b}|$, $|\vec{c}|$) less than 20 Å are considered. After filtering we are left with \approx 211,000 unique materials each of which we render in three different orientations. We use abTEM [30], an electron microscopy simulation package, to render the crystal structures as 3-dimensional tensors. The final dataset contains \approx 633,000 examples.

We split off 5,000 materials (i.e. 15,000 examples) for our validation set, and another 100 materials (i.e. 300 examples) for our test set. For the 15,000 and 300 examples in our validation and test sets, respectively, we simulate diffraction patterns using abTEM with realistic detector noise equivalent to an electron dose of $10^6 e^-/\text{Å}^2$. These diffraction patterns will be used to evaluate the quality of different reconstruction methods. See Appendix A for more details.

Evaluation We evaluate all methods on all 300 pairs of (\mathbf{x}, \mathbf{y}) in our test set. That is, given a 4D diffraction dataset \mathbf{y} , recover the crystal structure \mathbf{x} . We quantitatively evaluate the reconstruction quality using two common distortion metrics: PSNR, and SSIM [50]. We set the sampling steps to 2000 if no additional notes are given. Since different crystal structures naturally exhibit varying intensity scales due to their atomic composition and density, and phase retrieval techniques like ptychography are generally only accurate up to an arbitrary constant phase offset [38], we apply instance normalization to each reconstruction so they range between [0,1] before computing metrics. This normalization focuses the evaluation on the relative spatial distribution of atomic densities rather than absolute intensity values, which can vary significantly between samples and weight differently without normalization. For our experiments, we perform paired bootstrap resampling [24] to

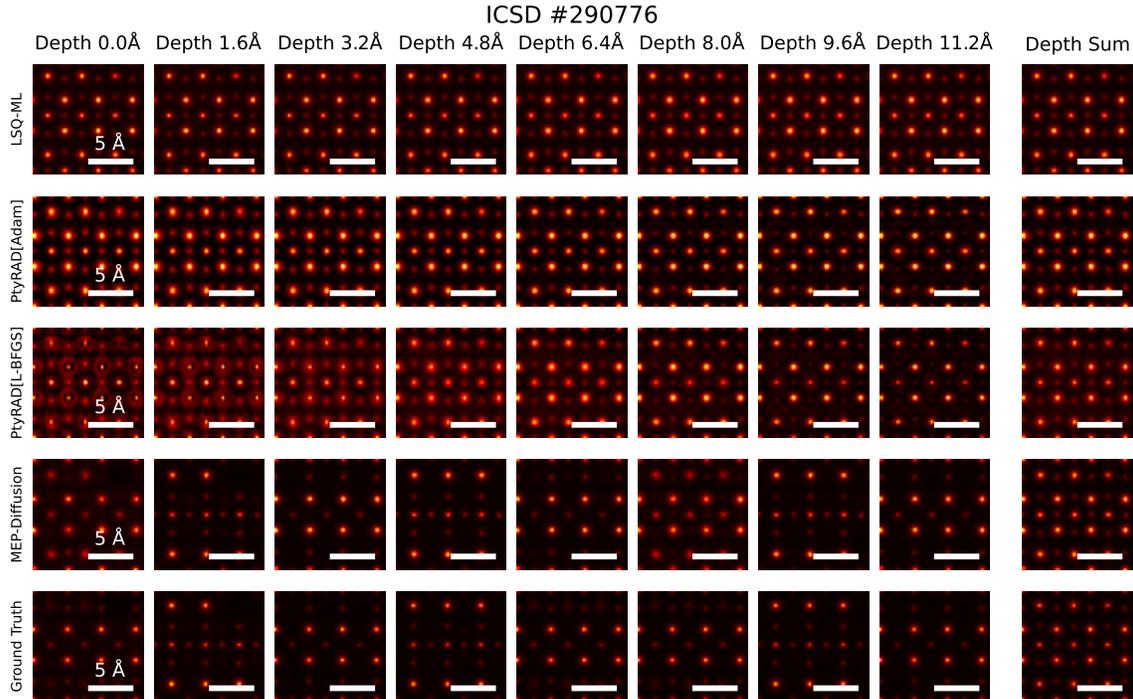


Figure 5. Qualitative comparison across depth slices of a crystal structure. Baseline methods struggle to recover depth variation, producing uniform reconstructions across slices. MEP-DIFFUSION successfully reconstructs the full crystal structure.

Table 2. Impact of our proposed Shifted Sigmoid guidance schedule. We evaluate both full 3D reconstruction (left) and depth-summed 2D projections (right) using standard image quality metrics. Results show mean \pm standard error of the mean across our test set. Our guidance schedule leads to statistically significant improvements ($p < 1e - 4$) over the baseline schedule on all metrics.

GUIDANCE SCHEDULE	FULL MATERIAL RECONSTRUCTION		MATERIAL DEPTH SUM RECONSTRUCTION	
	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
ALPHA	32.7677 ± 0.1126	0.6494 ± 0.0074	38.4307 ± 0.1517	0.7745 ± 0.0061
SHIFTED SIGMOID	33.7190 ± 0.1243	0.7523 ± 0.0053	41.5101 ± 0.1721	0.8852 ± 0.0041

quantify the statistical significance of our improvements. See [Appendix E](#) for training and sampling hyperparameters.

Baseline Methods We evaluate our approach against several baselines. Our primary comparison is with LSQ-ML [33], the current state-of-the-art algorithm for MEP. Additionally, we implement the same physical model using automatic differentiation to solve the *electron ptychography reconstruction objective* with either Adam or L-BFGS. This baseline is implemented using an open-source reconstruction package, PtyRAD [26], and is denoted as PtyRAD[Adam] or PtyRAD[L-BFGS] in our results [21, 27]. See [Appendix C](#) for more details.

Reconstruction Performance We report reconstruction results across methods in [Table 1](#). We visualize selected successful reconstructions in [Figure 1](#) and [Figure 5](#), as well as additional generations in [Appendix G](#). We observe that MEP-DIFFUSION quantitatively achieves more faithful reconstructions than all baselines across all metrics. The improvement over both baselines is statistically significant ($p < 1e - 4$) across all reconstruction metrics.

From our qualitative visualization, we observe that the baseline methods collapse to a uniform reconstruction across depth slices due to limitations of the physical model. MEP-DIFFUSION benefits from its learned prior which helps to recover distinct reconstructions across the depth di-

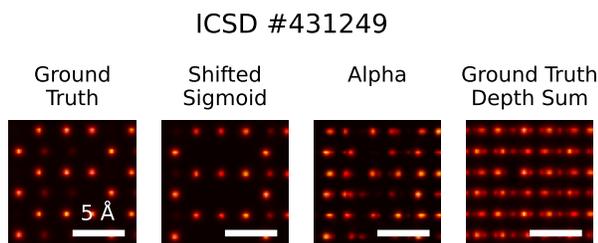


Figure 6. Qualitative comparison of our proposed Shifted Sigmoid guidance schedule and the baseline Alpha guidance schedule. Ground Truth, Shifted Sigmoid Guidance, and Alpha Guidance all display a single slice of the material, while Ground Truth Depth Sum shows the material summed over the depth dimension.

mension that are occasionally better aligned with the ground truth material. This demonstrates the utility of a generative prior to push updates from the physical model towards realistic crystal structures.

Impact of Guidance Schedule Our Shifted Sigmoid guidance schedule anneals the physical model influence to zero during generation, see Figure 4. This approach markedly differs from conventional guidance schedules in inverse problems, which typically increase guidance strength throughout generation. Without this carefully designed annealing, generation collapses along the depth dimension, exhibiting failure modes characteristic of direct optimization methods, see Figure 6. This collapse occurs because imprecise guidance from the physical model washes out the structure being formed by the diffusion model, forcing the generation toward degenerate solutions that satisfy physical constraints but are not realistic crystal structures. The Shifted Sigmoid schedule proves critical by gradually reducing physical guidance, enabling the diffusion model to resolve local atomic structure while still utilizing physical model guidance early in generation. The quantitative performance advantage persists across all metrics, see Table 2. This finding highlights a fundamental difference from typical inverse problems studied in machine learning: the guidance models available for real-world applications are often imperfect or inadequate, unlike the near-perfect models commonly studied.

Wall Clock Time Comparison We present the performance of MEP-DIFFUSION across different sampling steps in Figure 7. We observe that MEP-DIFFUSION demonstrates strong performance even with relatively few sampling steps. While L-BFGS converges quickly to a suboptimal solution and Adam requires longer optimization, MEP-DIFFUSION offers a flexible quality-to-time trade-off. With

just 100 sampling steps, MEP-DIFFUSION already outperforms both baselines in reconstruction quality. Increasing the number of steps yields further improvements, with diminishing returns beyond 1000 steps.

This flexibility is particularly valuable in practical applications, where different use cases may have different computational constraints. For rapid structure assessment, users can opt for fewer sampling steps while still obtaining better reconstructions than traditional methods. When reconstruction quality is paramount, additional sampling steps can be used to increase quality, albeit at increased computational cost. Importantly, this trade-off can be adjusted at inference time, without re-training the diffusion model.

5. Related Work

Inverse Problems Diffusion models have emerged as powerful tools for solving inverse problems through their learned priors, demonstrating effectiveness on natural image tasks like inpainting, super-resolution, and deblurring [7, 8, 29, 45], as well as diverse applications in medical imaging [47], motion synthesis [55], and physical simulation [10]. While most prior work focuses on simple forward models, our work shows how to effectively incorporate complex physical models from electron microscopy while learning priors over periodic structures.

Incorporating Physical Constraints Recent work has explored different approaches for incorporating physical constraints into deep generative models. Physics-informed neural networks (PINNs) incorporate general physical laws directly into the loss function during training [35–37]. Past work has also explored training the diffusion model to respect physical constraints [10]. In contrast, our work tackles physical inverse problems by incorporating guidance from a differentiable physical forward model using measurements directly at *inference* time.

Electron Ptychography Traditional ptychographic reconstruction relies on direct methods [1, 33, 39] or iterative algorithms [31, 33], with the latter achieving higher resolution at the cost of long reconstruction times [5, 32]. Recent deep learning approaches aim to improve efficiency by mapping diffraction measurements to reconstructions [3, 6, 12, 13, 16, 52], offering near-instant feedback but with reduced quality and flexibility. Our work integrates a generative prior from crystal structures with iterative electron ptychography, enhancing reconstruction quality while preserving the adaptability of traditional methods.

6. Limitations

MEP-DIFFUSION significantly improves reconstruction quality over existing methods but has several shortcomings.

Sweep Over Number of Sampling Schedule Steps

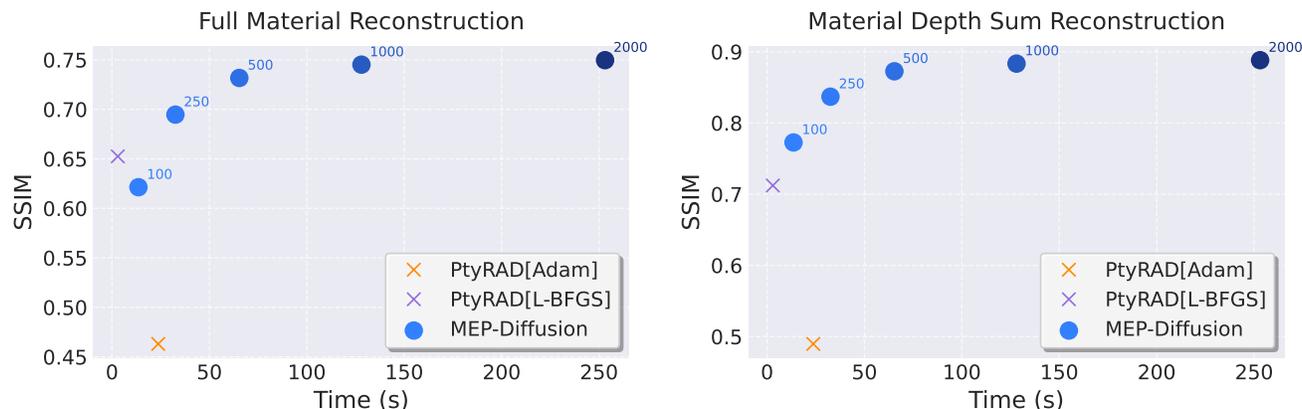


Figure 7. Trade-off between computational time and reconstruction quality. We evaluate MEP-DIFFUSION across different numbers of sampling steps (100-2000) and compare against PtyRAD[Adam] and PtyRAD[L-BFGS] baselines. All time experiments were run on a single NVIDIA H100 GPU. Even with just 100 steps, MEP-DIFFUSION achieves better reconstruction quality than baselines while maintaining reasonable computation time. Additional sampling steps provide further quality improvements at the cost of longer reconstruction times.

Errors range from minor issues, such as slice misordering and missing or extra atoms, to more significant ones, including incorrect local structures and resolution loss. Most errors involve atomic displacement along the depth dimension, even though the depth sum image remains accurate. These issues likely arise from limited depth information recorded in the 4D diffraction dataset [49], making true atomic structure recovery uncertain, especially within the depth dimension. Additionally, the generative prior may contribute to some errors. We have included more qualitative examples in [Appendix G](#) to illustrate these errors.

7. Future Directions

While our learning-based approach advances multislice electron ptychography (MEP), further improvements are needed. Specifically, enhancing robustness to weak guidance at low electron doses $10^3 e^-/\text{\AA}^2$ or lower is critical for dose-sensitive materials. Current methods struggle with depth reconstruction under these conditions. Additionally, developing principled methods for uncertainty quantification in atomic positions is crucial for experimental datasets, particularly when noise is significant. Another important direction is improving the quantitiveness of reconstructions, enabling accurate recovery of atomic species and local chemical compositions from phase signals. Addressing these challenges could improve MEP and provide valuable insights for both machine learning applications and scientific discovery. We will release our code and model checkpoints with a future manuscript that extends this work.

8. Conclusion

We present MEP-DIFFUSION, a diffusion model trained on crystal structures, which we combine with an existing physical forward model through DPS for multislice electron ptychography reconstruction. Our approach addresses limitations of current ptychographic reconstruction methods: traditional optimization methods are slow and struggle with depth ambiguity.

We find the shifted sigmoid guidance schedule is essential for guided sampling in this specific domain. Sampling with the typical guidance schedules leads to degenerate solutions with poor depth resolution. By gradually reducing the physical model's influence during sampling, MEP-DIFFUSION can recover realistic atomic-scale structures while maintaining consistency with measured diffraction patterns.

The combination achieves substantial quantitative improvements over existing methods, with statistically significant gains across all metrics. Most importantly, it successfully reconstructs distinct depth slices where baseline methods collapse to uniform solutions, addressing a key challenge in 3D ptychographic imaging. The method offers flexible quality-to-time trade-offs and performs well even with just 100 sampling steps. This work demonstrates that incorporating a generative prior over crystal structures into the reconstruction process can significantly improve ptychographic reconstruction quality.

Acknowledgments

CKB is supported by the National Science Foundation (NSF) through the NSF Research Traineeship (NRT) program under Grant No. 2345579. CHL is supported by the Eric and Wendy Schmidt AI in Science Postdoctoral Fellowship, a program of Schmidt Sciences, LLC. This work is also supported by the NSF through OAC-2118310; the National Institute of Food and Agriculture (USDA/NIFA); the Air Force Office of Scientific Research (AFOSR); the New York Presbyterian Hospital; arXiv; and the Schmidt AI2050 Senior Fellowship.

References

- [1] RHT Bates and JM Rodenburg. Sub-ångström transmission microscopy: a fourier transform algorithm for microdiffraction plane intensity information. *Ultramicroscopy*, 31(3): 303–307, 1989. 7
- [2] Refik Mert Cam, Junjing Deng, Rajkumar Kettimuthu, Mathew J Cherukara, and Tekin Bicer. Ptychographic image reconstruction from limited data via score-based diffusion models with physics-guidance. *arXiv preprint arXiv:2502.18767*, 2025. 1
- [3] Dillan J Chang, Colum M O’Leary, Cong Su, Daniel A Jacobs, Salman Kahn, Alex Zettl, Jim Ciston, Peter Ercius, and Jianwei Miao. Deep-learning electron diffractive imaging. *Physical review letters*, 130(1):016101, 2023. 7
- [4] Ting Chen. On the importance of noise scheduling for diffusion models. *arXiv preprint arXiv:2301.10972*, 2023. 4
- [5] Zhen Chen, Yi Jiang, Yu-Tsun Shao, Megan E Holtz, Michal Odstrčil, Manuel Guizar-Sicairos, Isabelle Hanke, Steffen Ganschow, Darrell G Schlom, and David A Muller. Electron ptychography achieves atomic-resolution limits set by lattice vibrations. *Science*, 372(6544):826–831, 2021. 1, 7
- [6] Mathew J Cherukara, Tao Zhou, Youssef Nashed, Pablo Enfedaque, Alex Hexemer, Ross J Harder, and Martin V Holt. Ai-enabled high-resolution scanning coherent diffraction imaging. *Applied Physics Letters*, 117(4), 2020. 7
- [7] Hyungjin Chung, Jeongsol Kim, Michael Thompson McCann, Marc Louis Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. In *The Eleventh International Conference on Learning Representations*, 2022. 1, 3, 4, 7
- [8] Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. In *Advances in Neural Information Processing Systems*, pages 25683–25696. Curran Associates, Inc., 2022. 7
- [9] Peter Ercius, Ian J Johnson, Philipp Pelz, Benjamin H Savitzky, Lauren Hughes, Hamish G Brown, Steven E Zeltmann, Shang-Lin Hsu, Cassio CS Pedroso, Bruce E Cohen, et al. The 4d camera: an 87 khz direct electron detector for scanning/transmission electron microscopy. *Microscopy and Microanalysis*, 30(5):903–912, 2024. 1
- [10] Xiantao Fan, Deepak Akhare, and Jian-Xun Wang. Neural differentiable modeling with diffusion-based super-resolution for two-dimensional spatiotemporal turbulence. *Computer Methods in Applied Mechanics and Engineering*, 433:117478, 2025. 7
- [11] Berthy T. Feng, Jamie Smith, Michael Rubinstein, Huiwen Chang, Katherine L. Bouman, and William T. Freeman. Score-based diffusion models as principled priors for inverse imaging. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10520–10531, 2023. 1
- [12] Thomas Friedrich, Chu-Ping Yu, Johan Verbeeck, and Sandra Van Aert. Phase object reconstruction for 4d-stem using deep learning. *Microscopy and Microanalysis*, 29(1):395–407, 2023. 7
- [13] Weijie Gan, Qiuchen Zhai, Michael Thompson McCann, Cristina Garcia Cardona, Ulugbek S Kamilov, and Brendt Wohlberg. Ptychodv: Vision transformer-based deep unrolling network for ptychographic image reconstruction. *IEEE Open Journal of Signal Processing*, 2024. 7
- [14] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, pages 6840–6851, 2020. 1, 2
- [15] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–8646, 2022. 4
- [16] Oliver Hoidn, Aashwin Ananda Mishra, and Apurva Mehta. Physics constrained unsupervised deep learning for rapid, high resolution scanning coherent diffraction reconstruction. *Scientific Reports*, 13(1):22789, 2023. 7
- [17] Emiel Hoogeboom, Jonathan Heek, and Tim Salimans. simple diffusion: End-to-end diffusion for high resolution images. In *International Conference on Machine Learning*, pages 13213–13232. PMLR, 2023. 4
- [18] Emiel Hoogeboom, Thomas Mensink, Jonathan Heek, Kay Lamerigts, Ruiqi Gao, and Tim Salimans. Simpler diffusion (sid2): 1.5 fid on imagenet512 with pixel-space diffusion. *arXiv preprint arXiv:2410.19324*, 2024. 4
- [19] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Advances in Neural Information Processing Systems*, 2022. 4
- [20] Diederik Kingma and Ruiqi Gao. Understanding diffusion objectives as the elbo with simple data augmentation. *Advances in Neural Information Processing Systems*, 36, 2024. 2, 3, 4
- [21] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [22] Diederik P Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. In *Advances in Neural Information Processing Systems*, 2021. 2
- [23] Earl J Kirkland. *Advanced Computing in Electron Microscopy*. Springer Science & Business Media, 2010. 2
- [24] Philipp Koehn. Statistical significance tests for machine translation evaluation. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain, 2004. Association for Computational Linguistics. 5

- [25] Harikrishnan Kp, Ruijuan Xu, Kinnary Patel, Kevin J Crust, Aarushi Khandelwal, Chenyu Zhang, Sergey Prosandeev, Hua Zhou, Yu-Tsun Shao, Laurent Bellaiche, et al. Electron ptychography reveals a ferroelectricity dominated by anion displacements. *Nature Materials*, pages 1–8, 2025. 1
- [26] Chia-Hao Lee, Steven E Zeltmann, Dasol Yoon, Desheng Ma, and David A Muller. Ptyrad: A high-performance and flexible ptychographic reconstruction framework with automatic differentiation. *Microscopy and Microanalysis*, 31(4):ozaf070, 2025. 6, 12
- [27] Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1):503–528, 1989. 6
- [28] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. In *Advances in Neural Information Processing Systems*, 2022. 4
- [29] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11461–11471, 2022. 7
- [30] Jacob Madsen and Toma Susi. The abtem code: transmission electron microscopy from first principles. *Open Research Europe*, 1, 2021. 5
- [31] Andrew M Maiden and John M Rodenburg. An improved ptychographical phase retrieval algorithm for diffractive imaging. *Ultramicroscopy*, 109(10):1256–1262, 2009. 7
- [32] Kayla X Nguyen, Yi Jiang, Chia-Hao Lee, Priti Kharel, Yue Zhang, Arend M van der Zande, and Pinshane Y Huang. Achieving sub-0.5-angstrom-resolution ptychography in an uncorrected electron microscope. *Science*, 383(6685):865–870, 2024. 7
- [33] Michal Odstrčil, Andreas Menzel, and Manuel Guizar-Sicairos. Iterative least-squares solver for generalized maximum-likelihood ptychography. *Optics express*, 26(3):3108–3123, 2018. 6, 7
- [34] Hugh T Philipp, Mark W Tate, Katherine S Shanks, Luigi Mele, Maurice Peemen, Pleun Dona, Reinout Hartong, Gerard van Veen, Yu-Tsun Shao, Zhen Chen, et al. Very-high dynamic range, 10,000 frames/second pixel array detector for electron microscopy. *Microscopy and Microanalysis*, 28(2):425–440, 2022. 1
- [35] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations. *arXiv preprint arXiv:1711.10561*, 2017. 7
- [36] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics informed deep learning (part ii): Data-driven discovery of nonlinear partial differential equations. *arXiv preprint arXiv:1711.10566*, 2017.
- [37] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. 7
- [38] John Rodenburg and Andrew Maiden. Ptychography. *Springer Handbook of Microscopy*, pages 819–904, 2019. 5
- [39] JM Rodenburg, BC McCallum, and PD Nellist. Experimental tests on double-resolution coherent imaging via stem. *Ultramicroscopy*, 48(3):304–314, 1993. 7
- [40] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*, 2022. 2, 3
- [41] Fahad Shamshad, Farwa Abbas, and Ali Ahmed. Deep ptych: Subsampled fourier ptychography using generative priors. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7720–7724. IEEE, 2019. 1
- [42] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015. 2
- [43] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 2256–2265, Lille, France, 2015. PMLR. 1
- [44] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. 1
- [45] Jiaming Song, Arash Vahdat, Morteza Mardani, and Jan Kautz. Pseudoinverse-guided diffusion models for inverse problems. In *International Conference on Learning Representations*, 2023. 7
- [46] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. 2
- [47] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. In *International Conference on Learning Representations*, 2022. 7
- [48] Mark W Tate, Prafull Purohit, Darol Chamberlain, Kayla X Nguyen, Robert Hovden, Celesta S Chang, Pratiti Deb, Emrah Turgut, John T Heron, Darrell G Schlom, et al. High dynamic range pixel array detector for scanning transmission electron microscopy. *Microscopy and Microanalysis*, 22(1):237–249, 2016. 1
- [49] EWC Terzoudis-Lumsden, TC Petersen, Hamish G Brown, PM Pelz, C Ophus, and Scott D Findlay. Resolution of virtual depth sectioning from four-dimensional scanning transmission electron microscopy. *Microscopy and Microanalysis*, 29(4):1409–1421, 2023. 1, 8
- [50] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 5
- [51] Kai Wu, An Pan, and Wei Gao. Fourier ptychographic reconstruction with denoising diffusion probabilistic models. *Optics & Laser Technology*, 176:111016, 2024. 1

- [52] Yudong Yao, Henry Chan, Subramanian Sankaranarayanan, Prasanna Balaprakash, Ross J Harder, and Mathew J Cherukara. Autophasenn: unsupervised physics-aware deep learning of 3d nanoscale bragg coherent diffraction imaging. *npj Computational Materials*, 8(1):124, 2022. [7](#)
- [53] Dejan Zagorac, H Müller, S Ruehl, J Zagorac, and Silke Rehme. Recent developments in the inorganic crystal structure database: theoretical crystal structure data and related features. *Journal of applied crystallography*, 52(5):918–925, 2019. [5](#)
- [54] P Zambon, S Bottinelli, R Schnyder, D Musarra, D Boye, A Dudina, N Lehmann, S De Carlo, M Rissi, C Schulze-Briese, et al. Kite: high frame rate, high count rate pixelated electron counting asic for 4d stem applications featuring high-z sensor. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1048:167888, 2023. [1](#)
- [55] Qinsheng Zhang, Jiaming Song, Xun Huang, Yongxin Chen, and Ming-Yu Liu. Diffcollage: Parallel generation of large content with diffusion models. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10188–10198. IEEE, 2023. [7](#)
- [56] Menglin Zhu, Michael Xu, Yu Yun, Liyan Wu, Or Shafir, Colin Gilgenbach, Lane W Martin, Ilya Grinberg, Jonathan E Spanier, and James M LeBeau. Insights into chemical and structural order at planar defects in pb2mgwo6 using multi-slice electron ptychography. *ACS nano*, 2025. [1](#)