Score-Based Super-Resolution for Atomic-Scale MoS2 Imaging

Xinyuan Wang[1], *Mingli Huang*[2], *and Jingyuan Sun*[3]

1 Computational Imaging Group, Centrum Wiskunde & Informatica (CWI), Amsterdam, Science Park 123, 1098 XG Amsterdam, The Netherlands
2 School of Computer Science, University of Sheffield, Sheffield, United Kingdom
3 Department of Computer Science, EEMCS, Delft University of Technology, Echo Building, 2628 XE Delft, The Netherlands

Abstract

High-Angle Annular Dark-Field Scanning Transmission Electron Microscopy (HAADF-STEM) is a powerful tool for visualizing the atomic structure of 2D materials like molybdenum disulfide (MoS_2). However, obtaining high-resolution images with high-signal-to-noise ratio often requires high electron doses, which can induce beam damage and structural defects. In this work, we introduce Score-Based Super-Resolution (SBSR) for atomic-scale MoS_2 imaging, a conditional diffusion model for joint super-resolution and denoising of low-dose scanning transmission electron microscopy (STEM) images, tailored to atomic lattice reconstruction. We introduce a physics-informed degradation (PID) pipeline that simulates realistic electron microscopy noise and beam-induced blur. Our method demonstrates the ability to restore clear hexagonal lattice structures from noisy, low-resolution images, enabling atomic-scale reconstruction from low-dose STEM data and reducing reliance on high-dose imaging that risks beam-induced damage.

Introduction

Precise visualisation of atomic defects and lattice orientations is essential for the characterisation of nanomaterials, including MoS_2 nanowires. While HAADF-STEM provides atomic-resolution Z-contrast imaging of individual atomic columns, image acquisition is fundamentally constrained by the specimen damage threshold. As a result, low-dose imaging is often employed to mitigate beam-induced damage, but this comes at the cost of increased shot noise and reduced spatial resolution [1]. Traditional interpolation and filtering methods typically fail to recover high-frequency atomic details and instead introduce smoothing artefacts [2].

Recent advances in score-based generative models [3] have demonstrated strong performance in solving inverse problems by learning expressive data priors. In this work, we adapt these techniques to the microscopy domain and employ a conditional diffusion framework for atomic-scale super-resolution and denoising. Our approach reconstructs high-fidelity atomic maps from noisy, low-resolution STEM images using a curated dataset of high-quality MoS_2 scans.

Methodology

While [3] proposes an unsupervised inverse-problem solver based on an unconditional prior, scale ambiguity in microscopy data leads to instability in unconditional sampling. We therefore adopt a conditional training strategy, following the framework of Super-Resolution via Repeated Refinement (SR3) [4]. We employ a U-Net $\epsilon_\theta(x_t, t, y)$ with learnable parameter θ, to predict the injected noise ϵ at diffusion timestep t given the noisy high-resolution (HR) latent x_t and a low-resolution (LR) conditioning input y. The network input is formed by channel-wise concatenation of x_t and an upsampled version of y, obtained via bicubic interpolation to match the high-resolution spatial grid.

The training objective is:

$$\mathcal{L} = \mathbb{E}_{x_0,\epsilon,t}\left[\|\epsilon - \epsilon_\theta(x_t, t, y)\|_2^2\right].$$

This loss trains the network to predict the noise injected at each diffusion step while conditioning on the upsampled LR observation, thereby learning the conditional score of the HR data distribution. Full details of the diffusion process are provided in [3,4].

To train a robust super-resolution model, we employ the MoS_2_Nanowire dataset and select high-quality scans at a resolution of 1024×1024 pixels as ground-truth HR images, ensuring a consistent atomic scale across the training set. LR conditioning inputs are generated on the fly using a PID pipeline that models key aspects of the electron microscopy acquisition process:

1. **Anti-aliasing and downsampling:** HR patches are blurred with a Gaussian kernel to approximate the microscope point spread function, followed by spatial downsampling.
2. **Mixed noise injection:** A combination of signal-dependent Poisson noise (shot noise) and signal-independent Gaussian noise (readout noise) is applied to reflect the stochastic nature of electron detection.
3. **Intensity normalization:** Robust percentile normalisation is used to account for the high dynamic range of HAADF intensities.

Experimental Results and Discussion

The model was trained on 128×128 patches extracted from HR MoS_2 STEM images using a cosine noise schedule with T=1000 diffusion timesteps. The model effectively suppresses severe shot noise while enhancing atomic-column contrast, successfully recovering the characteristic hexagonal honeycomb lattice of MoS_2 without introducing spurious structures in vacuum regions, as shown in Figure 1. These results indicate that the model captures a physically meaningful atomic prior. Overall, the proposed SBSR framework, combined with PID modeling, enables robust atomic-scale reconstruction from low-dose STEM images and offers a promising direction for reducing dose requirements in beam-sensitive microscopy applications.

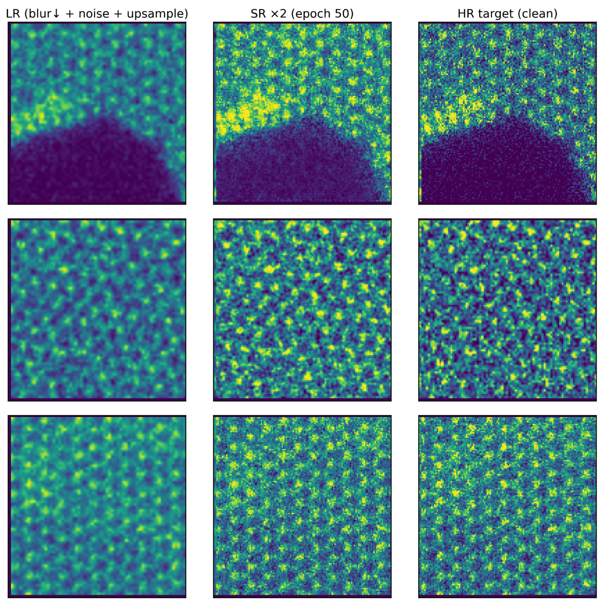

Figure 1: Comparison of LR inputs, SBSR super-resolution results (×2, epoch 50), and clean HR targets for MoS_2 STEM images.

Code: https://github.com/XinyuanWang283/microscopy_hackathon/tree/main

References

[1] Niklas Mevenkamp, Peter Binev, Wolfgang Dahmen, Paul M. Voyles, Andrew B. Yankovich and Benjamin Berkels. *Poisson noise removal from high-resolution STEM images based on periodic block matching and 3D filtering.* Advanced Structural and Chemical Imaging, 1(3):1–29, 2015.

[2] J. Han, Y. Zhang, Y. Kang, G. Lee and D. J. Smith. *Materials property mapping from atomic scale imaging via image enhancement and segmentation.* npj Computational Materials, 8(1):1–13, 2022.

[3] 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.

[4] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J. Fleet, and Mohammad Norouzi. *Image Super-Resolution via Iterative Refinement.* IEEE Transactions on Pattern Analysis and Machine Intelligence, 45(4):4713–4726, 2023.