

# DIFFRACT: Diffusion-based Restoration via Adaptive Control and Thresholding for Diffraction Imaging

Nikolay Falaleev  
Fanis Technologies  
London, UK

n.falaleev@fanis.ai

Nikolai Orlov  
AMOLF  
Amsterdam, Netherlands

n.orlov@amolf.nl

## Abstract

This paper presents a novel approach for denoising Electron Backscatter Diffraction (EBSD) patterns using diffusion models. We propose a two-stage training process with a UNet-based architecture, incorporating an auxiliary regression head to predict the quality of the experimental pattern and assess the progress of the denoising process. The model uses an adaptive denoising strategy, which integrates quality prediction and feedback-driven iterative denoising process control. This adaptive feedback loop allows the model to adjust its schedule, providing fine control over the denoising process. Furthermore, our model can identify samples where no meaningful signal is present, thereby reducing the risk of hallucinations. We demonstrate the DIFFRACT - the successful application of diffusion models to EBSD pattern denoising using a custom-collected dataset of EBSD patterns, their corresponding Master Patterns, and quality values.

## 1. Introduction

EBSD (Electron Backscatter Diffraction) is an electron microscopy technique for studying crystallographic properties at the nanoscale. In this method, electrons are diffracted by crystallographic planes, producing patterns that reveal crystal orientation. Its high resolution provides detailed information on grains, even at the nanometer scale, offering insights into processes such as degradation mechanisms, electronic properties, and other material characteristics.

EBSD patterns are grayscale images with diffraction lines called Kikuchi lines, each representing a portion of the diffraction sphere. Similar to maritime navigation using constellations, Kikuchi patterns identify the crystallographic orientation of individual crystallites (Figure 1). This process, known as indexing, assigns crystallographic indexes to each Kikuchi line in the pattern.

The quality of EBSD results depends on several factors,

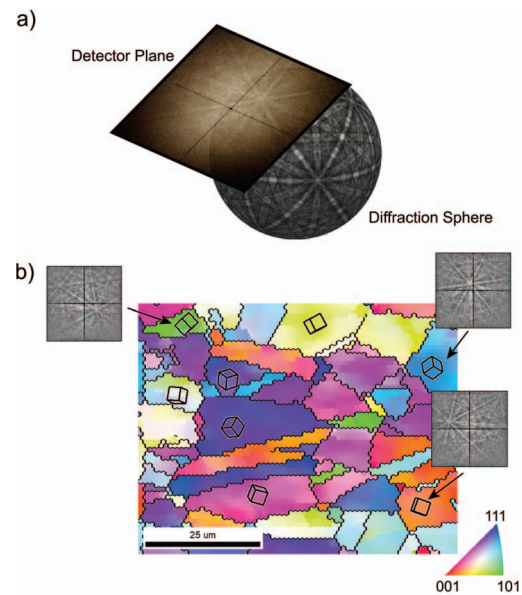


Figure 1. a) Pattern formation as a tangential projection of diffraction sphere on a detector plane; b) Inverse Polar Figure (IPF) map visualising the orientation of a grain by the color of the crystal face pointing perpendicular to the image plane. Each point of the IPF map is the result of EBSD pattern indexing.

including microscope settings (probe current, accelerating voltage, working distance, magnification), detector type (e.g., CCD or single-electron), and, most importantly, the sample itself (surface roughness, composition, crystallization degree, chemical and thermal stability) [6, 23]. Well-crystallized materials like steel or alumina produce clear Kikuchi patterns [15, 16], whereas beam-sensitive materials such as hybrid lead halide perovskites (promising for solar cell active layers) or metals in thin film form often yield weak, noisy patterns [14, 23].

Analyzing weak patterns is challenging, with three main approaches: Hough Indexing (HI) [12, 18], Dictionary Indexing (DI) [7], and the recently developed Spherical In-

dexing (SI) [19]. HI processes patterns in Hough space but depends on high-quality data, limiting its use for beam-sensitive materials. DI improves precision for noisy patterns by comparing them to simulated Master Patterns (Figure S6) but is computationally intensive. SI also uses a Master Pattern but matches via spherical functions rather than dot products, enabling faster analysis and higher precision than DI.

However, SI's effectiveness depends on preprocessing, which is subjective, variable, and often requires experience and, at times, luck. While SI can index very noisy patterns lacking visible Kikuchi lines, assessing indexing accuracy remains difficult. In this context, EBSD pattern restoration can enhance post-analysis in three key ways:

- reduce exposure time in situations where prolonged signal collection might damage samples, all while maintaining quality;
- make indexing more reliable, consistent, and less dependent on the preprocessing of patterns;
- enhance pattern processing to enable indexing of patterns that were previously unindexable.

Computer vision approaches offer a promising route for denoising EBSD patterns and extracting more information. Current methods include Autoencoders, GANs, CNNs, and diffusion-based techniques. While the first two have been applied to EBSD, their effectiveness depends on sufficiently high pattern quality to ensure stable results and avoid hallucinations. Diffusion learning provides an advantage with its step-by-step approach, but as shown in the Related Works section, even advanced methods can hallucinate when processing content-simple images like EBSD patterns.

To address these challenges and improve the reliability of EBSD pattern denoising, this paper aims to develop a control mechanism that limits the restoration process while preserving vital information in EBSD images. Our contributions include:

- We introduce a feedback-driven, adaptive denoising approach that integrates quality prediction to control the iterative process. To our knowledge, this is the first successful application of diffusion models to EBSD pattern denoising.
- Our model incorporates a mechanism for distinguishing between meaningful signal and noise, helping to flag potential hallucinations.

## 2. Related Works

Image denoising with Deep Learning has been actively developed for over a decade, with most methods based on convolutional architectures. NAFNet is notable as the first nonlinear-activation-free network and has been applied even to medical image denoising, reducing hallucination risk [5, 26, 29]. Another example is the Context Guided

Network for Semantic Segmentation (CGNet) [10, 24], which captures both local details and global context through a cascaded gaze mechanism.

Although strong baselines for image denoising, these methods have limitations in scientific imaging, especially microscopy, where long-distance pixel relationships are critical. For images such as EBSD patterns, approaches like GANs, transformers, and diffusion networks (discussed later) are often better suited.

Improving EBSD pattern (EBSP) quality with Deep Learning has been an active research topic in recent years. While diffraction image denoising often adapts techniques from real-world images, it differs due to distinct noise characteristics (less Gaussian, more Poisson) and the unique physics of image formation, involving reciprocal space and greater emphasis on fine details.

Since the pattern dataset forms the foundation for the entire indexing process, most studies first restore EBSD patterns with DL methods, then index the restored patterns and generate IPF maps.

Early EBSD denoising studies used UNet-based autoencoders [2], achieving good restoration quality. At the time, SI was not widely adopted, and many initial patterns were poorly indexed using HI. Today, SI can easily index patterns of that quality, so much of this earlier work focused on improving strain analysis, which demands higher-quality patterns.

Few studies address pattern restoration when images are mostly noise and nearly devoid of visible lines. A notable exception is [17], where patterns are almost lost but experienced users can still spot some Kikuchi lines. The scarcity of such research stems from the difficulty of verifying prediction accuracy and the challenges of collecting ground truth data.

The key advantage of diffusion models for image restoration is their stepwise handling of noisy images. Unlike classical models that map a noisy sample directly to a clean one, diffusion models progressively reduce noise over several steps, particularly valuable for scientific images shaped by physical sample properties. This gradual restoration occurs in image space, transitioning from noisy to clean, and was originally proposed as a general method for image restoration without training separate models for each task [4].

The diffusion models are represented by several types of approaches. The Denoising Diffusion Probabilistic Model introduced by [11] uses Markov chains to model the forward and reverse process. Step by step it applies noise during forward process. Then in backward process the model is learning noise distribution parameters such as mean  $\mu_\theta(x_t, t)$  and variance  $\beta_\theta(x_t, t)$  to restore image from noisy state back to original state.

Despite their general popularity and proved efficiency in denoising real-world imagery data, like photographs [27],

the diffusion models are still not widely used for scientific image and especially EBSD images denoising. The crucial difference in the [27] work is in method of sampling, which has helped stabilize the image restoration process. In previous works, every next image in the denoising pipeline is a combination of restored image from the previous state and normalized noise which is left after subtraction of restored image from the noisy one. In the [27] work, every next image is an image from the previous step and subtracted normalized difference between original noisy image and a restored image from the previous step.

As mentioned above, the main challenge for all denoising approaches is ensuring the stability of the image restoration process. In [8] was demonstrated, that the restored image is not a single fixed image but rather a variety of potential images, forming a data manifold. A promising approach suggested by [1] involves using an adaptive level of noise during each step of the reverse process, rather than applying a constant noise level. This adaptive noise level is estimated at each step based on the distance of the restored image  $x_t$  from the noiseless image  $x_0$ . To facilitate this, a separate network has been introduced to learn the level of noise correction, which is defined as follows:  $\hat{\sigma}_t := \sigma_t[1 + \hat{r}_t] \approx \text{dist}_\kappa(x_t)/\sqrt{n}$ .

Another approach to improving EBSD results, presented in the literature, involves processing EBSP using traditional imaging techniques, followed by enhancing the resulting Inverse Pole Figure (IPF) maps through Deep Learning. For example, in [3], IPF maps were restored using a diffusion-based approach. To stabilize the diffusion model, the authors employed Reduced-Order Generalized Spherical Harmonics (ROGSH), which incorporate additional information to improve restoration performance. In contrast, the present work focuses on restoring EBSD patterns directly, as this remains the most precise and reliable method for enhancing EBSD analysis.

### 3. Dataset

There are EBSD datasets available on the Internet (<https://zenodo.org/communities/ebbsd>). However, these datasets were not used in this work as our approach relies on Master Patterns as a ground truth and employs a patterns quality metric, which is not available in the public data. Therefore, we collected new datasets and generated corresponding pieces of Master Patterns using Spherical Indexing.

We utilized a scanning electron microscope ThermoFisher Verios 460, equipped with an EDAX Clarity detector utilizing a 4-chip Timepix detector offering noiseless digital integration and high detection efficiency [13] from Amsterdam Scientific Instruments (ASI). The data patterns were gathered using EDAX APEX software with beam currents ranging from 50 to 200 pA and accelerating voltages

5-20 kV. Pixel integration times varied from 30 to 250 ms, with different binning options applied (1, 2, 4, and 5 pixels). We analyzed multiple samples (Nickel (*Ni*), Ruthenium (*Ru*), Gold (*Au*), Silver (*Ag*)) to apply a diverse set of measurement conditions and to include various crystallographic variables, including different crystal lattice parameters and crystal systems (Cubic, Hexagonal). In total, the number of experimental patterns used for training was over 163000, with further 16500 patterns used for validation and over 6000 used as a hold-out test set.

### 3.1. Master Patterns

Master Patterns for each crystal structure were generated using EDAX OIM 9 software. This software was also utilized to process and index experimental and restored EBSD patterns using Spherical Indexing. The Master Patterns were used as ground truth for the models training and evaluation.

## 4. Methods

The data in this task differs from traditional image denoising because it includes both ideal images (Master Patterns) and experimental images with an associated 'confidence index' that quantifies the reliability of the experimental data. Additionally, within each grain, the computed Master Patterns are nearly identical, with only minor variations due to defects or crystal lattice bending. This allows experimental patterns within a grain to be considered noisy samples from the same distribution.


### 4.1. Proposed approach

Our approach combines a diffusion-based denoising model with an integrated quality assessment feedback mechanism. This dual-objective architecture allows the model to simultaneously denoise EBSD patterns while evaluating the denoisability and quality of the input patterns. Critically, our model includes a mechanism to identify and flag patterns that represent pure noise rather than noisy observations of actual crystal structures, thus preventing hallucination artifacts.

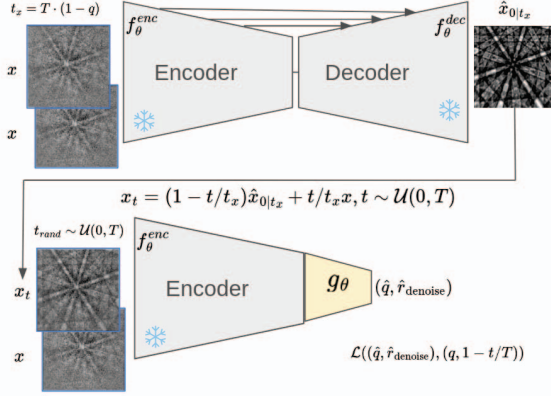
### 4.2. Training procedure

Model training was performed in two stages.

**Denoising model.** On the first stage we training the main model to perform image denoising. The pseudocode for this training procedure is presented in Algorithm 1, which is depicted in Figure 2a. During training, the network receives a noisy observation  $x_t$  as input and attempts to predict the corresponding clean image  $x_0$ . The denoising process is guided by a quality score  $q$  and a noise scale  $s_{\text{noise}}$ , with additional Gaussian noise injected at each diffusion step to simulate varying levels of corruption. The noise injection



(a) Stage 1: Training the denoising neural network. The example illustrates the process for  $t = 128$ .



(b) Stage 2: Training the auxiliary head to predict both the original image quality and denoising progress  $r_{denoise}$ . Instead of applying an external noise scheduler (as in Stage 1), the pretrained denoising network is used to provide partially denoised images. The function  $g_\theta$  is trained using a random value  $t_{rand}$  to decouple denoising stage prediction from step number embedding.

Figure 2. Proposed two-stage training procedure for image denoising and quality assessment.

process follows a diffusion model, where the noisy observation  $x_t$  is computed as a weighted combination of the clean image  $x_0$ , the noisy input  $x$ , and an additional noise term  $\omega_t$  drawn from a standard Gaussian distribution  $\mathcal{N}(0, I)$ , Eq. 1.

$$x_t = (1 - \alpha)x_0 + \alpha x + \beta\omega_t \quad (1)$$

The weight  $\alpha$  is determined by the current diffusion step  $t$ , relative to the total number of diffusion steps  $T$ , as  $\alpha = \frac{t}{t_x}$ , where  $t_x = T \cdot (1 - q)$ . The noise scale  $\beta$  is determined by the factor  $m \cdot \frac{s_{noise} \cdot t}{T}$ , where  $m$  is drawn from a Bernoulli distribution with probability  $p_{noise}$ , indicating whether additional noise should be applied, while  $s_{noise}$  controls the level of additional image degradation.

Notably, the algorithm allows the generation of images that can be more degraded than the observed image  $x$ , since for each experimental image the current step  $t$  can exceed  $t_x$ , leading to simulated further corruption of the image. This, together with the extra noise enables the model to recover clean images from increasingly corrupted data.

On each step, the network has access to the interpolated noisy image of the step  $x_t$  and is guided by the original experimental noisy image  $x$ , which gives the network more

context (both are stacked as a two channel image), and the current step  $t$ .

---

### Algorithm 1 Denoising Model Training

---

**Require:** Encoder-decoder neural network  $f_\theta$ , dataset  $\mathcal{D} = \{(x_0^{(i)}, x^{(i)}, q_i)\}_{i=1}^N$ , where  $x_0^{(i)}$  is the clean image (Master Pattern),  $x^{(i)}$  is the noisy observation, and  $q_i$  is the quality of  $x^{(i)}$ ; total number of diffusion steps  $T$ , noise probability  $p_{noise}$ , and noise scale  $s_{noise}$ .

- 1: **while** not converged **do**
  - 2:   Sample  $(x_0, x, q)$  from  $\mathcal{D}$
  - 3:   Calculate diffusion step of  $x$ :  $t_x = T \cdot (1 - q)$
  - 4:   Compute interpolation  $x_t = (1 - \alpha)x_0 + \alpha x$ , where  $t \sim \mathcal{U}(0, T)$ ,  $\alpha = \frac{t}{t_x}$ ,
  - 5:   Generate additional noise:  $\omega_t \sim \mathcal{N}(0, I)$
  - 6:   Add noise:  $x_t = x_t + \beta\omega_t$  where  $\beta = m \cdot \frac{s_{noise} \cdot t}{T}$ ,  $m \sim \text{Bern}(p_{noise})$ .
  - 7:   Normalize:  $x_t \leftarrow \frac{x_t - \mathbb{E}[x_t]}{\text{std}(x_t)}$
  - 8:   Predict the clean image:  $\hat{x}_0|t \leftarrow f_\theta(x_t, x, t)$
  - 9:   Update parameters of  $f_\theta$  to minimize loss:  $\mathcal{L}(\hat{x}_0|t, x_0)$
  - 10: **end while**
- 

**Quality Values Regression.** In the second stage of training (Algorithm 2, depicted in Figure 2b), we perform quality regression using the trained denoising network in a frozen state. Specifically, the bottleneck feature map, which captures an internal representation of the input tensor, is passed through an auxiliary head for the two-values regression task. The goal is to predict two values: the original image quality  $\hat{q}$ , and the progress of the denoising process  $\hat{r}_{denoise}$ .

The first predicted value,  $\hat{q}$ , represents the estimated quality of the original image at the current denoising step, corresponding to the same ground-truth value  $q$  from the dataset for all diffusion steps. Since the experimental noisy image  $x$  is part of the input tensor at each step, the information necessary for predicting this objective value is present in the input.  $q = 0$  for pure noise images.

The second predicted value,  $\hat{r}_{denoise}$ , estimates the progress of the denoising process. This value depends on the current diffusion step  $t$  and provides a measure of how far along the diffusion process is in recovering the clean image. More specifically,  $\hat{r}_{denoise}$  estimates how much noise has been removed by the current denoising step. The target value for  $\hat{r}_{denoise}$  is constructed using the following rule:

$$r_{denoise} = \begin{cases} 1 - \frac{t}{T}, & \text{if } x \text{ is not noise,} \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

When the image is not fully corrupted (i.e.,  $x$  is not pure noise), the value is inversely proportional to the diffusion step, with lower  $t$  values indicating a more advanced

stage of denoising;  $r_{\text{denoise}} = 0.0$  for pure noisy images. The sample  $x$  can represent pure noise in two cases: the dataset includes intentionally collected pure noise EBSD patterns, where noise collected as instrumental noise without a sample material installed, or synthetic noise samples generated during training by augmenting the data with pure noise drawn from a standard Gaussian distribution  $\mathcal{N}(0, I)$ .

The training of the auxiliary head involves direct estimation of the partially denoised image at diffusion step  $t$ :

$$x_t = \left(1 - \frac{t}{t_x}\right) f_\theta(x, x, t_x) + \frac{t}{t_x} x. \quad (3)$$

This component differs from the procedure in Algorithm 1, as the intermediate image state is generated using the actual denoising model rather than a linear interpolation scheduler, utilising the Master Patterns. Since the encoder’s input includes the diffusion step  $t$ , and since  $\hat{r}_{\text{denoise}}$  is a function of  $t$ , it is crucial to decouple the image generation process during training from the estimation of  $\hat{r}_{\text{denoise}}$ . This is done as follows: first,  $t \sim \mathcal{U}(0, T)$  is used to construct  $x_t$  and the target value for  $\hat{r}_{\text{denoise}}$ . Then during the forward pass, which updates the weights of  $g_\theta$ , a different random step  $t_{\text{rand}} \sim \mathcal{U}(0, T)$  is used as input to ensure that the auxiliary head relies on the actual visual features rather than the information from  $t$ -embedding, allowing the value to be used for denoising progress assessment.

As will be shown in the following sections, this two-values regression enables the model to estimate both the initial step at which the denoising process should begin and provide a feedback on the progress of the denoising process.

---

#### Algorithm 2 Quality Regression Training

---

**Require:** Frozen encoder( $f_\theta^{\text{enc}}$ )-decoder( $f_\theta^{\text{dec}}$ ) denoising neural network  $f_\theta$ , auxiliary quality prediction head  $g_\theta$ , dataset  $\mathcal{D}$ , total number of diffusion steps  $T$ .

- 1: **while** not converged **do**
  - 2:   Sample  $(x, q)$  from  $\mathcal{D}$
  - 3:   Calculate diffusion step of  $x$ :  $t_x = T \cdot (1 - q)$
  - 4:   Get partially denoised image:  $x_t = (1 - \alpha) f_\theta(x, x, t_x) + \alpha x$ , where  $\alpha = \frac{t}{t_x}, t \sim \mathcal{U}(0, T)$
  - 5:   Set the target  $r_{\text{denoise}}$  as per Eq. 2, using the same  $t$  as used for  $x_t$ .
  - 6:   Predict regression values:  $(\hat{q}, \hat{r}_{\text{denoise}}) \leftarrow g_\theta(f_\theta^{\text{enc}}(x_t, x, t_{\text{rand}}))$ , where  $t_{\text{rand}} \sim \mathcal{U}(0, T)$ .
  - 7:   Update parameters  $g_\theta$  to minimize loss:  $\mathcal{L}((\hat{q}, \hat{r}_{\text{denoise}}), (q, r_{\text{denoise}}))$
  - 8: **end while**
- 

### 4.3. Prediction

Conceptually, the prediction Algorithm 3, depicted in Figure 3, is similar to the improved sampling algorithm from [27], with two key novel features arising from the use of

the two regression values. First, we estimate the initial diffusion step based on a dry-run of the neural network (the input timestep is set to  $T$  during the run).

Second, in contrast to dynamic methods like [1], which adjust the amount of noise removed in each step based on a monotonous denoising process with predicted noise level, we adaptively select the next diffusion step during denoising based on the predicted denoising progress  $\hat{r}_{t_x}$ . This adaptive approach allows our algorithm to dynamically adjust the denoising steps: it can repeat a step if further denoising is required or skip steps if the progress indicates sufficient improvement. To prevent infinite loops, the total number of denoising steps is capped at  $T$ , and the predicted value  $\hat{r}_{t_x}$  provides a stopping criterion for the process, allowing it to exit at a desired denoising level  $R$  when required by the application.

Moreover, given the definition of  $\hat{r}_{t_x}$ , its value approaching zero can serve as a useful indicator to flag potential hallucinations or unreliable predictions. This is particularly valuable in ensuring the robustness of the denoising process and detecting situations where the model might be making implausible predictions. Usage of  $\hat{r}_{t_x}$  allows getting the best denoised image, based on the evaluation, rather than conventionally selected image from the last denoising step.

---

#### Algorithm 3 Denoising Prediction with Feedback

---

**Require:** Neural network  $E_\theta$ , consists of encoder( $f_\theta^{\text{enc}}$ )-decoder( $f_\theta^{\text{dec}}$ ) denoising neural network, auxiliary quality prediction head  $g_\theta$ , total number of diffusion steps  $T$ , desired final denoising  $R \in [0, 1]$

- 1:  $(\hat{q}_{\text{init}}, \hat{r}) \leftarrow g_\theta(f_\theta^{\text{enc}}(x, x, T))$
  - 2: Start diffusion step  $t_x = T \cdot (1 - \hat{q}_{\text{init}})$ ;  $x_t = x$
  - 3: Step counter  $n \leftarrow 0$
  - 4: **while**  $n < T$  and  $t_x > 0$  and  $\hat{r}_{t_x} < R$  **do**
  - 5:    $(\hat{x}_{0|t_x}, \hat{q}_{t_x}, \hat{r}_{t_x}) \leftarrow E_\theta(x_t, x, t_x)$
  - 6:    $\alpha_t = t_x/T$ ;  $\alpha_{t-1} = (t_x - 1)/T$
  - 7:    $\tilde{x}_t = (1 - \alpha_t)\hat{x}_{0|t_x} + \alpha_t x$
  - 8:    $\tilde{x}_{t-1} = (1 - \alpha_{t-1})\hat{x}_{0|t_x} + \alpha_{t-1} x$
  - 9:   Denoising:  $x_t = x_t - \tilde{x}_t + \tilde{x}_{t-1} = x_t + \frac{1}{T}(\hat{x}_{0|t_x} - x)$
  - 10:   Next step as predicted feedback:  $t_x = T \cdot (1 - \hat{r}_{t_x})$
  - 11:   Increment step counter  $n = n + 1$
  - 12: **end while**
  - 13: Final  $\hat{x}_0$  selected as  $x_t$  with the highest  $\hat{r}_{t_x}$ .
- 

Note that the proposed prediction schedule has no reliance on clean images (Master Patterns).

### 4.4. Architecture and Implementation details

All experiments used  $128 \times 128$  images and a UNet style CascadedGaze architecture [10], following the Real Image Denoising protocol from [10] with a network width of 60. The encoder contains [2, 2, 4, 6] CascadedGaze blocks, the middle layer has 10 NAF blocks [5], and the decoder




Figure 3. Diagram of the proposed denoising algorithm with feedback.

comprises [2, 2, 2, 2] blocks. The auxiliary head is composed of an InstanceNorm layer, followed by two convolutional blocks that each downsample the input by a factor of 2 and incorporate a skip connection. After global-average pooling, the resulting features are fed into a 3-layer MLP that employs batch normalization and a dropout with rate of 0.1. The hidden layer sizes are 256 and 128, and the final layer outputs two regression values, which are passed through a sigmoid activation. Diffusion-step values are encoded with 128-dimensional sine-cosine positional embeddings and then projected through a fully connected layer. The SiLU function is used as the non-linearity in this network.

We used the L1 loss function for both tasks, trained with the AdamW optimizer [20] with  $\beta_1 = 0.9, \beta_2 = 0.9$ , starting with a learning rate of  $10^{-3}$  and decays to  $10^{-7}$  using a cosine annealing scheduler over 96 epochs.

Two augmentations were applied during the first stage training: MixUp [28] with a probability of 0.25 and a maximum added pattern weight of 0.25, and 10% rate of Gaussian noise-only images paired with zero targets.  $s_{\text{noise}}$  and  $p_{\text{noise}}$  were set to 0.5. In Algorithms 2 and 3  $x_t$  values were normalized to range [0, 1].  $T = 384$  in all experiments.

The model was implemented in PyTorch 2.6 with mixed-precision training (bfloat16), and experiments ran on two Nvidia A10 GPUs with 24GB VRAM.

## 5. Results and discussion

In this section, we are comparing different training and inference strategies and evaluating the impact of our framework’s components on denoising performance.

All experiments use the same fixed data subsets, hyperparameters, and augmentations, with evaluation on a hold-out test set using Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM).

Since our model can predict a denoised image from any diffusion step, we first compare training strategies for the denoising component. As EBSD denoising has received little attention, we benchmark against representative scientific-image denoising methods (Table 1), measuring

One step process	PSNR ( $\uparrow$ )	SSIM ( $\uparrow$ )
Raw data	17.2198	0.5160
BM3D [9, 21]	18.5254	0.6310
NAFNet [5, 26]	24.9287	0.7991
CascadedGaze [10]	24.9309	0.7931
DIFFRACT	25.0333	0.8023
DIFFRACT with $t_x$ assessment	25.0337	0.8067

Table 1. Comparison of different procedures for direct patterns restoration. *Raw data* represents the test dataset without denoising. *BM3D*, *NAFNet* and *CascadedGaze* are reference training-free and DL methods. *Our method* is similar to the conventional linear algorithm from [27], expanded by incorporation of input data quality and linear extrapolation to more degraded images. *Our method with  $t_x$  assessment* follows one step prediction of  $\hat{x}_{0|t}$  by Algorithm 3 (lines 1-5).

Diffusion process	PSNR ( $\uparrow$ )	SSIM ( $\uparrow$ )
DDPM [11]	25.1763	0.8043
DDPM-NLC [1]	25.2081	0.8132
Baseline	25.0620	0.8001
+ start step: $t_x = T \cdot (1 - \hat{q})$	25.0905	0.8012
+ dynamic step: $t_{\text{next}} = T \cdot (1 - \hat{r})$	25.2816	0.8055
+ best step selection, based on $\hat{r}_{t_x}$	25.2825	0.8093

Table 2. The first two rows report the performance of the reference DDPM and DDPM-NLC models, both applied with deterministic sampling ( $\eta = 0$ ). *Baseline* represents the standard linear denoising approach similar to [27]. The remaining rows progressively add the components of our proposed denoising scheme with the final values demonstrating results of the complete Algorithm 3.

how well a clean image can be reconstructed directly from noisy input.

BM3D [9, 21] which is a widely used collaborative filtering non-local image denoising algorithm, performs far below learning-based methods, indicating that our dataset contains noise and structures beyond classical capabilities, despite BM3D’s success on synthetic noise and other imaging tasks [22, 25]. Since our method utilizes the same base architecture as CascadedGaze [10], which represents a state-of-the-art evolution of the widely used NAFNet convolutional architecture employed in scientific image denoising [5, 26], we trained the models with the original training protocol as a reference baseline. The results show that our training strategy outperforms the baseline approach, achieving an increase in both PSNR and SSIM metrics, confirming the value of quality-defined input generation and linear extrapolation to more degraded states. Predicted image quality for  $t_x$  selection benefits even direct denoising without diffusion.

Next, we investigate different inference-time diffusion-




Figure 4. Visual qualitative comparison of different EBSD patterns restoration methods.

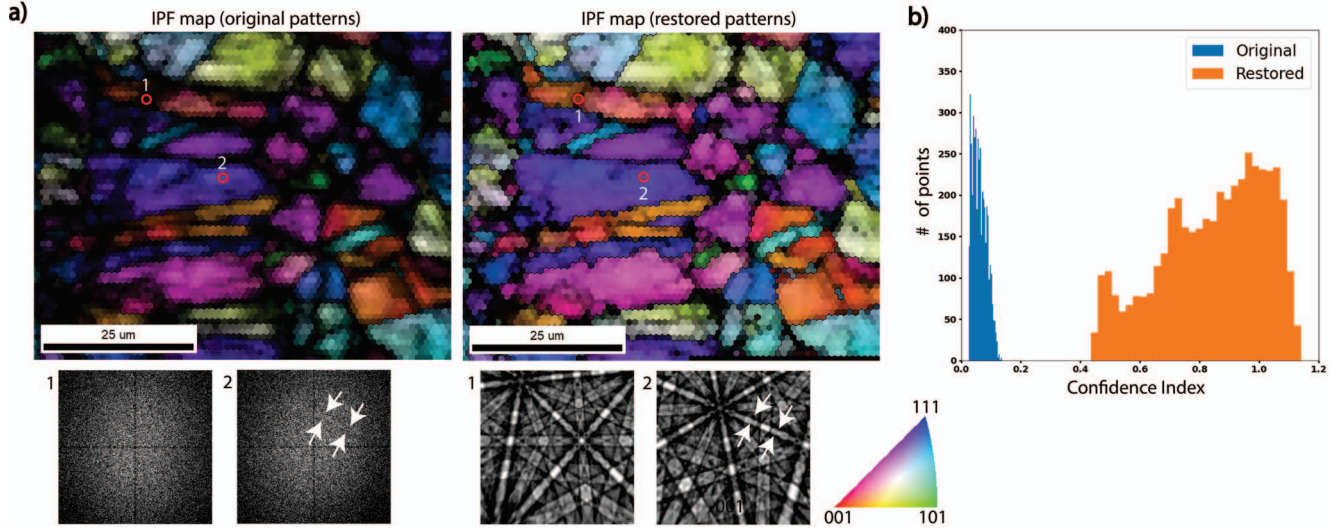


Figure 5. a) IPF maps of the original and restored datasets using DIFFRACT, with patterns shown before any image processing or neighboring-pattern averaging. b) Confidence Index (CI) values distribution for IPF map of the original (blue) and restored (orange) patterns.

based denoising strategies, analyzing how modifications to the step selection process impact image quality (Table 2). The results demonstrate that dynamically adjusting the diffusion starting step based on predicted quality improves the performance of iterative denoising as well. The approach of dynamically selecting the next diffusion step based on the assessment of the denoising progress  $\hat{r}$  proves beneficial and adds to the metric values. Selecting the final denoised pattern using the highest predicted  $\hat{r}$  demonstrates the effectiveness of adaptive denoising progress assessment. This result shows that optimal denoising can occur at intermediate steps, where further iterations may actually degrade quality. Adaptive stopping at the optimal step provides a clear advantage over fixed-step methods, emphasizing the importance of quality-based adaptive criteria for achieving optimal reconstruction. Examples of iterative denoising results are presented in Fig. 4 and visual comparison of different methods in Fig. S7.

We also compare our method with other diffusion-scheduling approaches (Table 2). For these experiments, we used the same architecture as our model

and the same data preparation, including artificial image degradation and augmentations, with the only difference being the scheduling. In particular, we benchmark against restoration with DDPM and DDPM-NLC (Algorithm 1 from [1] with  $\eta = 0$ ). The former is a base diffusion approach, whereas the latter is the closest to ours: its adaptive schedule varies the amount of noise at each step, whereas our approach selects the first and the next step number dynamically. Although the DDPM-NLC approach achieves slightly better SSIM value, it does not provide a straightforward way to stop denoising at a preferred image-quality level (the relative pattern quality should be preserved to maintain surface contrast differences like grain boundaries in the results after indexing) or to assess raw-image quality to reduce the risk of hallucination.

The observed results in Fig. S8 reveal a clear distinction between actual images and instrumental noise samples, based on the predicted values of  $\hat{q}$  and  $\hat{r}_{denoise}$ . For real EBSD patterns, the final predicted denoising progress values are significantly higher, indicating the model’s ability to recognize meaningful patterns. In contrast, for noise data,

the predicted values remain consistently low, reflecting the absence of structure in the input data. These findings suggest that the predicted values can serve as a criterion for differentiating between actual images and noise, offering a reliable approach for noise detection and filtering in practical applications.

Fig. 5 shows IPF maps, where each point's color and brightness result from EBSD pattern indexing with SI. In the "original" dataset, point 2 is brighter than point 1 due to a weak Kikuchi line in the EBSD pattern (indicated by white arrows) that increases the OIM9 software's quality parameter Confidence Index (CI) - confidence in crystal orientation assignment (point 1 does not relieve any visible lines). The same line, restored, is visible in the processed pattern. Original dataset patterns were enhanced using classical image processing (Adaptive Histogram Equalization, background subtraction, and neighboring-pattern averaging), whereas restored patterns were processed only by averaging neighboring patterns. Group of pixels with the same hue represent material grains - their size distribution and crystal orientation are the main information extracted from the EBSD analysis. Notably, the restored dataset preserved the surface-relief texture, reflecting the quality-assessment feedback mechanism that maintains differences between initially higher- and lower-quality patterns. This is consistent with the CI histogram (Fig. 5-b), which shows a wider CI distribution for the restored patterns. Thinner dark lines between grains in the "restored" map indicate better indexing quality and more precise grain-size estimation. Moreover, "restored" map reveal more smaller grains through improved indexing of restored patterns. At the end, our method preserves critical surface features and enhances grain boundary clarity, enabling more accurate grain-size estimation and revealing smaller grains, while requiring less preprocessing than classical approaches.

Our ablation study demonstrates that incorporating image quality assessment into the diffusion denoising process yields consistent improvements over baseline methods. The proposed quality-defined control and noisy pattern extrapolation strategies, combined with adaptive timestep selection, collectively improve denoising metrics. The limited performance of traditional methods like BM3D confirms that EBSD pattern complexity necessitates learning-based approaches, justifying the effectiveness of quality-aware training and adaptive denoising strategies for challenging real-world scenarios in scientific imaging.

## 6. Conclusion

We propose the DIFFRACT - a dynamic method for selecting diffusion steps based on predicted pattern quality,  $\hat{q}$ , and denoised quality,  $\hat{r}_{denoise}$ . This adaptive strategy: 1) ensures optimal termination of the denoising process; 2) integrates quality-defined control for generating noisy training

samples; 3) extrapolating experimental images to more degraded states. To the best of our knowledge, this is the first time such an approach has been applied to EBSD pattern restoration. Our approach significantly improves denoising performance compared to baseline methods. The predicted values of  $\hat{q}$  and  $\hat{r}_{denoise}$  also enable noise-only data detection and filtering. These contributions enhance the denoising process, improving reconstruction quality and robustness.

This method is proved useful in EBSD measurements, allowing for reduced exposure times while sacrificing some pattern quality, and subsequently restoring the information during post-processing. Furthermore, the suggested approach lessens the reliance on the experimentalist's selection of the image processing sequence and could be useful beyond EBSD context.

## 7. Acknowledgments

We would like to acknowledge the partial support of our computational resources through the Dutch national e-infrastructure with the support of the SURF Cooperative using grant no. EINF-9949.

Collecting EBSD patterns is performed at AMOLF, an NWO funded institute, and is part of the project "Achieving Semiconductor Stability From The Ground Up" (with project number 19459) which is financed by the Dutch Research Council (NWO), Gatan (EDAX), Amsterdam Scientific Instruments (ASI) and CL Solutions.

## References

- [1] Abulikemu Abuduweili, Chenyang Yuan, Changliu Liu, and Frank Permenter. Enhancing sample generation of diffusion models using noise level correction, 2025. 3, 5, 6, 7
- [2] Caleb E. Andrews, Maria Strantz, Nicholas P. Calta, Manyalibo J. Matthews, and Mitra L. Taheri. A denoising autoencoder for improved kikuchi pattern quality and indexing in electron backscatter diffraction. *Ultramicroscopy*, 10: 113810, 2023. 2
- [3] Michael O. Buzzy, Andreas E. Robertson, and Surya R. Kalidindi. Statistically conditioned polycrystal generation using denoising diffusion models. *Acta Materialia*, 267: 119746, 2024. 3
- [4] J. H. Rick Chang, Chun-Liang Li, Barnabas Poczos, B. V. K. Vijaya Kumar, and Aswin C. Sankaranarayanan. One network to solve them all — solving linear inverse problems using deep projection models. 2
- [5] Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *Computer Vision – ECCV 2022*, pages 17–33, Cham, 2022. Springer Nature Switzerland. 2, 5, 6
- [6] Y. Chen, J. Hjelen, S. Gireesh, and H. Roven. Optimization of ebsd parameters for ultra-fast characterization. *Journal of Microscopy*, 245:111–118, 2012. 1
- [7] Yu H Chen, Se Un Park, Dennis Wei, Greg Newstadt, Michael A Jackson, Jeff P Simmons, Marc De Graef, and Alfred O Hero. A dictionary approach to electron backscatter diffraction indexing. *Microscopy and Microanalysis*, 21(3):739–752, 2015. 1
- [8] Hyungjin Chung, Jeongsol Kim, Michael T. Mccann, Marc L. Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. 3
- [9] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE Transactions on Image Processing*, 16(8):2080–2095, 2007. 6
- [10] Amirhosein Ghasemabadi, Muhammad Kamran Janjua, Mohammad Salameh, Chunhua Zhou, and Di Niu Fengyu Sun. Cascadedgaze: Efficiency in global context extraction for image restoration. *Transactions on Machine Learning Research*, 2024. 2, 5, 6
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, pages 6840–6851. Curran Associates, Inc., 2020. 2, 6
- [12] P.V.C. Hough. Method and means for recognizing complex patterns. *U.S. Patent*, (3069654), 1962. 1
- [13] Jan Jakubek. Data processing and image reconstruction methods for pixel detectors. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 576(1):223–234, 2007. 3
- [14] Sarthak Jariwala, Hongyu Sun, Gede W.P. Adhyaksa, Andries Lof, Loreta A. Muscarella, Bruno Ehrl, Erik C. Garnett, and David S. Ginger. Local crystal misorientation influences non-radiative recombination in halide perovskites. *Joule*, 3:3048–3060, 2019. 1
- [15] Danka Katrakova and Frank Mücklich. Specimen preparation for electron backscatter diffraction (ebsd) - part i: Metals. *Prakt. Metallogr.*, 38:547–565, 2001. 1
- [16] Danka Katrakova and Frank Mücklich. Specimen preparation for electron backscatter diffraction (ebsd) - part ii: Ceramics. *Prakt. Metallogr.*, 39:644–662, 2002. 1
- [17] K. V. Mani Krishna, R. Madhavan, Mangesh V. Pantawane, Rajarshi Banerjee, and Narendra B. Dahotre. Machine learning based de-noising of electron back scatter patterns of various crystallographic metallic materials fabricated using laser directed energy deposition. *Ultramicroscopy*, 247:113703, 2023. 2
- [18] N.C. Krieger Lassen, D. Juul Jensen, and K. Conradsen. Image processing procedures for analysis of electron back scattering patterns. *Scanning Microscopy*, 6(1):115–121, 1992. 1
- [19] W.C. Lenthe, S. Sihgh, and M. De Graef. A spherical harmonic transform approach to the indexing of electron back-scattered diffraction patterns. *Ultramicroscopy*, 207:112841, 2007. 2
- [20] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization, 2019. 6
- [21] Ymir Mäkinen, Lucio Azzari, and Alessandro Foi. Collaborative filtering of correlated noise: Exact transform-domain variance for improved shrinkage and patch matching. *IEEE Transactions on Image Processing*, 29:8339–8354, 2020. 6
- [22] Rishant Pal, Shubham Kumar Gupta, Azeem Ahmad, Frank Melandsø, and Anowarul Habib. Block-matching and 3d filtering-based denoising of acoustic images obtained through point contact excitation and detection method. *Applied Acoustics*, 217:109843, 2024. 6
- [23] Hongyu Sun, Gede W.P. Adhyaksa, and Erik C. Garnett. The application of electron backscatter diffraction on halide perovskite materials. *Adv. Energy Mater.*, 10:2000364, 2020. 1
- [24] Tianyi Wu, Sheng Tang, Rui Zhang, and Yongdong Zhang. CGNet: A Light-weight Context Guided Network for Semantic Segmentation. *IEEE Transactions on Image Processing*, 2019. 2
- [25] Wentian Wang, Wei Du, and Zhuo Jia. Removing random noise of gpr data using joint bm3diam filtering. *Sensors*, 25(10), 2025. 6
- [26] Wei-jie Gu, Zheng Liu, Yun-jie Yang, Xuan-zhi Zhang, Liang-yu Chen, Fang-ning Wan, Xiao-hang Liu, Zhang-zhe Chen, Yun-yi Kong, and Bo Dai. A deep learning model, NAFNet, predicts adverse pathology and recurrence in prostate cancer using MRIs. *npj Precis. Onc.*, 7(7):134, 2023. 2, 6
- [27] Cheng Yang, Cong Wang, Lijing Liang, and Zhixun Su. Real-world image denoising via efficient diffusion model with controllable noise generation. *Journal of Electronic Imaging*, 33(4):043003, 2024. 2, 3, 5, 6
- [28] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *International Conference on Learning Representations*, 2018. 6
- [29] Zheng Liu, Wei jie Gu, Fang-ning Wan, Zhang-zhe Chen, Yun yi Kong, Xiao hang Liu, Ding wei Ye, and Bo Dai. Deep

learning network enhances imaging quality of low-b-value diffusion-weighted imaging and improves lesion detection in prostate cancer. *BMC Cancer*, 25(953), 2025. [2](#)