

Diffusion-Driven Bayesian Inference for Robust Inverse Problem Solving in Live-Cell Microscopy

Kevin Sun

Introduction

Structured Illumination Microscopy (SIM) is an important tool for observing cellular structures such as the endoplasmic reticulum (ER). The ER forms a dynamic network of sheets and tubules that plays a significant role in protein folding, lipid synthesis, and calcium regulation. Therefore, understanding the ER morphology [1] is important for many diseases and biological processes. However, by reducing the phototoxicity and photobleaching [2], researchers usually get low-quality images because the images are affected by diffraction limited point spread function (PSF). Therefore, it became a problem to take clear images of the ER structure under fluorescence microscopy. The traditional imaging restoration method uses the Bayesian Inference framework with hand-crafted priors that are often too weak to capture complex cellular structure, usually requiring retraining [3]. Diffusion posterior sampling [4] (DPS) model provides an alternative by learning powerful data-driven image priors, so that measurement consistency can be enforced at test time without paired degraded data or retraining. For ER images degraded by a known PSF and additive Gaussian noise, the selected DPS configuration improves PSNR from 25.19 to 31.88 dB and achieves SSIM = 0.953, while also producing pixel-wise uncertainty maps that indicate regions of lower reconstruction confidence.

Methods

The DPS framework integrates a pretrained diffusion prior with Bayesian measurement consistency, enabling flexible inverse problem solving without retraining or paired supervision.




Figure 1. End-to-end DPS pipeline.

Step 1: The goal is to recover clean image $\mathbf{x}$ from degraded observation $\mathbf{y}$ through the forward model

$$\mathbf{y} = \mathcal{A}(\mathbf{x}) + \mathbf{n} \quad \mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma_y^2 \mathbf{I}), \quad (1)$$

where $\mathcal{A}$ is the degradation operator and $\mathbf{n}$ is additive Gaussian noise. By forming the Bayesian Inference framework, we use the equation:

$$p(\mathbf{x} | \mathbf{y}) \propto p(\mathbf{y} | \mathbf{x}) p(\mathbf{x}), \quad (2)$$

where $p(\mathbf{x})$ is the prior, which is the Denoising Diffusion Probabilistic Model (DDPM) and $p(\mathbf{y} | \mathbf{x})$ is the likelihood that can be calculated by the following equation:

$$p(\mathbf{x} | \mathbf{y}) \propto \exp(-\|\mathbf{y} - \mathcal{A}(\mathbf{x})\|^2 / 2\sigma_y^2) \cdot p(\mathbf{x}). \quad (3)$$

In order to get posterior, we only need to know the right side of Bayesian rule, which we know likelihood and now we need to find the prior.

Step 2: A DDPM [5] is trained on clean BioSR [3] ER images (128² patches, $T=1000$ timesteps) to learn the score function $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t) \approx -\epsilon_\theta(\mathbf{x}_t, t) / \sqrt{1 - \alpha_t}$ through the denoising objective. This requires **only unpaired clean data**—no supervised degradation-restoration pairs are needed.

By applying the DDPM prior, we can calculate **Diffusion Posterior Sampling** [4]. In order to sample from the posterior rather than the prior, each reverse diffusion step from unclean image to clean image is modified to include a measurement consistency gradient. At each step $t=T, \dots, 1$:

$$\mathbf{x}_{t-1} \leftarrow \underbrace{\mu_\theta(\mathbf{x}_t, t)}_{\text{DDPM step}} - \underbrace{\zeta \nabla_{\mathbf{x}_t} \|\mathbf{y} - \mathcal{A}(\hat{\mathbf{x}}_0)\|}_{\text{measurement consistency}} \quad (4)$$

where $\hat{\mathbf{x}}_0$ is the Tweedie estimate, an estimate of posterior mean without computing, and ζ is the hyperparameter, which is the step size of t . This formulation requires only a differentiable forward operator and the same trained prior could handle arbitrary degradations **without retraining**, which is one problem that traditional Bayesian restoration facing. Table 1 lists the four degradation operators used in this study.

Table 1. Degradation operators and their parameters.

Operator	Parameters
Gaussian Blur	$\sigma=1.5$, kernel $k=9$
Anisotropic Blur	$\sigma_x=3.0$, $\sigma_y=0.5$
Random Inpaint	10% blocks missing
Block Inpaint	24×24 centre mask

Results

We use hyperparameter sweep to identify optimal DPS operating points for PSF deblurring by applying four distinct degradation using in a single fixed model DDPM.

PSF Deblurring of Searching for Hyperparameter: 20 configurations (Denoising timesteps with $T \in \{100, 200, 500, 1000\}$ and step size $\zeta \in \{0.5, 1, 2, 3, 4\}$) are evaluated for Gaussian PSF deblurring with additive Gaussian noise ($\sigma_n=0.02$). Table 2 reports representative results showing the quality-efficiency trade-off:

Table 2. Representative hyperparameter configurations for PSF deblurring.

Config	T	ζ	PSNR
#1	500	4.0	32.40 dB
#2	1000	2.0	32.09 dB
#3	200	4.0	31.88 dB
#4	1000	4.0	31.82 dB
Baseline	—	—	25.19 dB

Config #3 was selected as the operating point: it achieves a comparable quality to the best configuration while requiring 2.5× fewer reverse steps. Figure 2 shows the visual quality of PSF deblurring with this configuration.




Figure 2. PSF deblurring results showing restoration of fine ER tubules.

Operator Agnosticism: The same trained model and fixed hyperparameters were applied to all four degradation operators without modification. Table 3 shows consistent performance improvements across fundamentally different forward models, demonstrating that the learned prior generalizes beyond specific degradation types:

Table 3. Performance across four degradation operators (Obs: observed, Rec: reconstructed).

Operator	Image A			Image B		
	Obs	Rec	SSIM	Obs	Rec	SSIM
Gaussian Blur	24.8	29.2	0.913	21.3	26.3	0.943
Random Inpaint	17.7	30.8	0.936	20.2	26.3	0.954
Block Inpaint	23.6	28.7	0.927	23.4	27.4	0.965
Aniso. Blur	24.3	27.5	0.880	20.1	23.9	0.902

Figure 3 compares visual restorations across operators. Even for anisotropic blur—a severely ill-posed problem—the method preserves biologically relevant structural features.




Figure 3. Restorations across four degradation operators.

Results (cont.): Uncertainty Quantification

DPS was run $N=16$ times with different random seeds ($T=200$, $\zeta=4.0$, $\sigma_n=0.1$). For representative ER fields of view, the main uncertainty statistics are summarized in Table 4:

$$\bar{\mathbf{x}} = \frac{1}{N} \sum_i \hat{\mathbf{x}}^{(i)}, \quad \mathbf{u} = \sqrt{\frac{1}{N} \sum_i (\hat{\mathbf{x}}^{(i)} - \bar{\mathbf{x}})^2} \quad (5)$$

Table 4. Summary uncertainty statistics.

Metric	Value
Posterior samples per field	16
PSNR after DPS	26.21–28.98 dB
SSIM after DPS	0.897–0.940
Mean / max pixel-wise std	0.021–0.032 / 0.112–0.134




Figure 4. Pixel-wise uncertainty maps.

Discussion

The results show that a single diffusion prior restores ER images across four degradation types without retraining. Config #3 achieves +6.69 dB of PSNR and gain with 2.5× fewer steps than the best configuration, indicating diminishing returns beyond $T=200$.

Traditional priors (TV [6], sparsity) favor smooth solutions that oversuppress thin and dynamic ER tubules. The learned diffusion prior performs better in preserving these biologically relevant features, maintaining SSIM above 0.88 across all operators.

Uncertainty variance concentrates at ER pores and junctions, high-frequency regions most affected by blur or missing data. These maps identify locations requiring cautious biological interpretation.

Limitations: (1) Gaussian PSF based on formulation is an approximation; experimental PSF based on true image could improve fidelity. (2) Step-size ζ requires per-operator tuning, meaning that we have to set this coefficient by ourself. (3) 16 posterior samples for uncertainty is expensive; accelerated samplers (DDIM) could help.

Conclusions

A single pretrained DDPM with DPS restores ER microscopy images across four operators with no paired data and no retraining.

PSF deblurring yields **PSNR** 25.19 → 31.88 dB (+6.69 dB) and **SSIM=0.953**, showing that the reconstructed image is close to the ground-truth and not overfitting. Based on sensitive analysis, **Pixel-wise uncertainty maps** localize ambiguous regions for cautious biological interpretation.

Applications: rapid, non-invasive ER imaging in drug screening and disease research.

Future work: 3-D ER volumes, spatially varying PSFs, and accelerated samplers.

References

- [1] J. Nixon-Abell, C. J. Obara, A. V. Weigel, D. Li, W. R. Legant, C. S. Xu, H. A. Pasolli, K. Harvey, H. F. Hess, E. Betzig, C. Blackstone, and J. Lippincott-Schwartz, "Increased spatiotemporal resolution reveals highly dynamic dense tubular matrices in the peripheral ER," *Science*, vol. 354, no. 6311, p. aaf3928, 2016.
- [2] P. P. Laisue, R. A. Alghamdi, P. Tomancak, E. G. Reynaud, and H. Shroff, "Assessing phototoxicity in live fluorescence imaging," *Nature Methods*, vol. 14, no. 7, pp. 657–661, 2017.
- [3] C. Qiao, D. Li, Y. Guo, C. Liu, T. Jiang, Q. Dai, and D. Li, "Evaluation and development of deep neural networks for image super-resolution in optical microscopy," *Nature Methods*, vol. 18, no. 2, pp. 194–202, 2021.
- [4] H. Chung, J. Kim, M. T. McCann, M. L. Klasky, and J. C. Ye, "Diffusion posterior sampling for general noisy inverse problems," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2023.
- [5] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 33, pp. 6840–6851, 2020.
- [6] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear Phenomena*, vol. 60, no. 1–4, pp. 259–268, 1992.