

<https://doi.org/10.1038/s44172-024-00331-z>

Microscopy image reconstruction with physics-informed denoising diffusion probabilistic model

Check for updates

Rui Li^{1,2,7}, Gabriel della Maggiora^{1,2,7}, Vardan Andriasyan³, Anthony Petkidis³, Artsemi Yushkevich⁴, Nikita Deshpande^{1,2}, Mikhail Kudryashev^{4,5} & Artur Yakimovich ^{1,2,6}

Light microscopy is a practical tool for advancing biomedical research and diagnostics, offering invaluable insights into the cellular and subcellular structures of living organisms. However, diffraction and optical imperfections actively hinder the attainment of high-quality images. In recent years, there has been a growing interest in applying deep learning techniques to overcome these challenges in light microscopy imaging. Nonetheless, the resulting reconstructions often suffer from undesirable artefacts and hallucinations. Here, we introduce a deep learning-based approach that incorporates the fundamental physics of light propagation in microscopy into the loss function. This model employs a conditioned diffusion model in a physics-informed architecture. To mitigate the issue of limited available data, we utilise synthetic datasets for training purposes. Our results demonstrate consistent enhancements in image quality and substantial reductions in artefacts when compared to state-of-the-art methods. The presented technique is intuitively accessible and allows obtaining higher quality microscopy images for biomedical studies.

Light microscopy (LM) is an important and accessible way to explore the hidden biomedical world of cells and sub-cellular structures. The decreasing cost of basic LM equipment is facilitating their accessibility in classrooms for educational purposes^{1–4} and in medical laboratories, particularly for applications such as cytometry¹. Moreover, it is hard to overstate the significant contributions made by advanced LM techniques such as fluorescence microscopy⁵, confocal microscopy⁶, or superresolution microscopy⁷ to a myriad of biomedical discoveries over the past century⁸. Yet, optical systems remain fundamentally limited owing to the principles of their design (Fig. 1a, b). The blur and imperfections of a point source (e.g. single molecule fluorescence) can be expressed mathematically as a point spread function (PSF)⁹. PSF describes the spread of light that occurs from scattering and diffraction as it passes through the optical components of the microscope. The advent of digital microscopy and image processing allowed us to attempt alleviating these limitations algorithmically^{10–15}. These methods include deconvolution methods^{9,16,17}, regularisation methods^{18–20} and bayesian methods²¹. However, the traditional methods fail to capture the complexity of the images. This leads to minuscule resolution improvement in reconstructions. Moreover, susceptibility to noise and artefacts during

image acquisition can compromise the performance of these methods, resulting in degraded quality during reconstruction.

The deep learning (DL) models, particularly convolutional neural networks (CNNs), have gained traction in image reconstruction tasks^{22–25}. CNNs learn complex mappings between images and their corresponding ground truth. For example, Xu et al. transformed a pseudo-inverse kernel for deconvolution into a CNN, capturing degradation characteristics²⁶. Ronneberger et al. incorporated deconvolution layers into the U-Net architecture²⁷. Generative adversarial networks have also been used for image restoration^{28,29}.

Another recent trend involves the utilisation of a type of likelihood-based generative models known as Denoising Probabilistic Diffusion Models (DDPMs)³⁰. They have shown promise in several tasks such as superresolution³¹, image colourisation, inpainting, uncropping, and JPEG restoration³². The diffusion models have desirable properties such as distribution coverage, stationary optimisation function, scalability, and training stability. They demonstrate promising results in terms of generating high-quality samples when compared to generative adversarial networks (GANs)³³.

¹Center for Advanced Systems Understanding (CASUS), Görlitz, Germany. ²Helmholtz-Zentrum Dresden-Rossendorf e. V. (HZDR), Dresden, Germany.

³Department of Molecular Life Sciences, University of Zurich, Zurich, Switzerland. ⁴In situ Structural Biology, Max Delbrück Center for Molecular Medicine in the Helmholtz Association, Berlin, Germany. ⁵Institute of Medical Physics and Biophysics, Charite-Universitätsmedizin, Berlin, Germany. ⁶Institute of Computer Science, University of Wrocław, Wrocław, Poland. ⁷These authors contributed equally: Rui Li, Gabriel della Maggiora. e-mail: a.yakimovich@hzdr.de

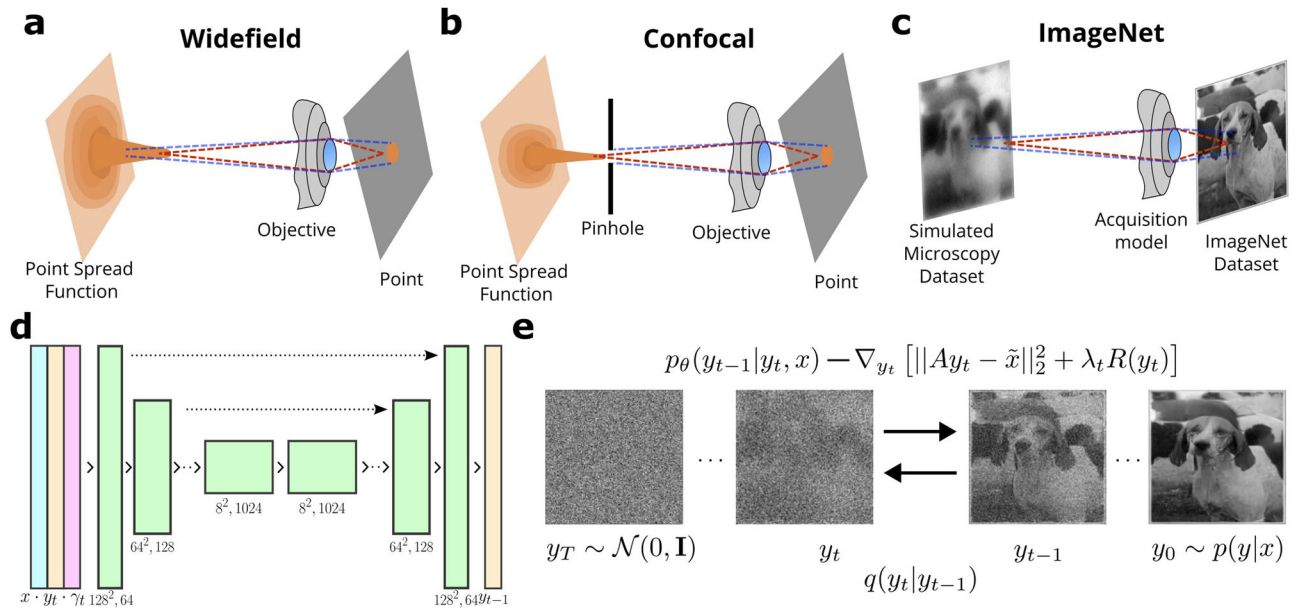


Fig. 1 | Proposed model with physics-informed probabilistic denoising diffusion for Microscopy image reconstruction. **a, b** simplified schematic depiction of widefield and confocal microscopy. **c** Schematic depiction of synthetic dataset generation using our image acquisition model. **d** Schematic depiction of our model’s architecture. Downsampling green blocks are convolution blocks with instance normalisation. Upsampling green blocks have an additional bilinear upsampling operator. Yellow blocks represent the diffusion state. Blue blocks represent the widefield image and red blocks represent the y function at timestep t (**e**) illustration

of the physics-informed denoising diffusion probabilistic model we propose. The diffusion process is carried out by distribution q , which progressively destroys the data structure until reaching a $N(0, I)$ distribution. The reverse process progressively recovers the structure of the data. This process is characterised by the learned distribution p , in which we incorporate the physics model to inform the recovery of the data (See Methods). The physics model is assumed to be linear, represented by matrix A , and can also incorporate a regulariser R , such as L_1 or L_2 .

However, most of these approaches ignore the great body of knowledge gathered on microscopy by optical physics. One common challenge in DL methods involves their inclination to produce structures not found in real images. This poses issues in fields that demand accurate reconstruction, such as medical diagnostic images and microscopy^{34,35}. To circumvent this, in other domains where the physics of the process is well understood, researchers have recently proposed an approach called physics-informed neural networks³⁶. In this approach, the prior knowledge regarding the physics models serves as a regularisation for DL models. While other approaches exist, incorporating physics by using the physical model as a regulariser is intuitive and widely used.

To overcome the shortcomings, we introduce a physics-informed (PI) diffusion model (Fig. 1d, e). In this model, the PI term is incorporated in the DDPM loss function (see Fig. 1e, Methods) using the technique shown in refs. 33,37. We show that this approach not only provides a simpler and more principled model but also produces more natural-looking results.

Results

Light microscopy image acquisition model

LM is an imaging technique in which an analogue-digital converter detects the electrical impulses generated by the light. As a result, image statistics can be well represented by a Poisson process. If several acquisitions are made and averaged, then according to the central limit theorem, the statistics of the image can be modelled by the Gaussian process. The mathematical model for optical systems assumes that the model is linear and time-invariant. Therefore the image acquisition model is described by the equation $I = \phi(h*x + b)$, where the image I , is the result of the convolution between object x and system PSF h with the background signal noise b added. The Poisson noise ϕ is applied afterwards over the true signal, given by the previous equation.

In LM, the diffraction pattern generated in an ideal optical system is the impulse response referred to as the PSF³. PSF of LM varies depending on the specifics of the technique employed, e.g. widefield and confocal LM (Fig. 1a, b). The diffraction pattern generated in an ideal optical system is the impulse

response referred to as the PSF. In fluorescence microscopy, usually, the illumination (excitation) and detection (emission) wavelength are not the same, so the most suitable model of the PSF³⁸ can be expressed as:

$$h(x, y, z) = |u_{\lambda_{ex}}(x, y, z)|^2 |u_{\lambda_{em}}(x, y, z)|^2, \tag{1}$$

where u_{λ} corresponds to the pupil function for a respective emission or excitation wavelength λ . This model is known as the Airy diffraction pattern³⁹. To add the effect of the pinhole used in confocal microscopy, we can convolve a disk function with the pupil of the emission sample. The disk function is usually modelled as:

$$T(x) = 1_{R^2 \leq x^2 + y^2}, \tag{2}$$

where R is the radius of the pinhole. Using this notion, we can rewrite the PSF as:

$$h(x, y, z) = |T(x, y) * u_{\lambda_{em}}(x, y, z)|^2 |u_{\lambda_{ex}}(x, y, z)|^2. \tag{3}$$

To model the pupil function u above, we employed the Arnison-Sheppard approach⁴⁰. This approach models the optical transfer function (OTF). The OTF is expressed as the autocorrelation of the pupil function u in the Fourier space. Thus, the u is the inverse Fourier transform of $C(\mathbf{K})$. Mathematically, in the k -space the OTF $C(\mathbf{K})$ can be denoted as below:

$$C(\mathbf{K}) = \iint \int Q(\mathbf{m} + \frac{1}{2}\mathbf{K}) \cdot Q^*(\mathbf{m} - \frac{1}{2}\mathbf{K}) d\mathbf{m}, \tag{4}$$

where Q corresponds to the complex vectorial pupil function⁴¹.

The complex vectorial pupil function is a complex-valued function $Q(\mathbf{m})$ of the position vector $\mathbf{m} = (k_x, k_y, k_z)$ within the aperture of an optical

system. This function can be expressed as:

$$Q(\mathbf{m}) = A(\mathbf{m})e^{i\phi(\mathbf{m})}, \quad (5)$$

where A is the amplitude transmission function with respect to the numerical aperture of the microscope and ϕ is the phase shift produced by aberrations and microscope imperfections. Finally, to obtain u the inverse Fourier transform is applied to the OTF.

Physics-informed denoising probabilistic diffusion models

Physics-informed methods arise from the need to incorporate physical knowledge into the construction or training of the model. In inverse problems, the objective is to determine $p(y|x)$ where y is the true object and x is some observation of the object that follows a possibly stochastic process $p(x|y)$. In image reconstruction problems such as MRI undersampled image reconstruction, deblurring or microscopy image reconstruction, the physical model is generally of the form $x = \phi(K*y)$ where ϕ is a function that applies noise according to a certain distribution. Particularly, model-based methods try to solve a problem of the form $\min_y \|Ay - x\|^2 + \lambda R(y)$. Where λ is the weight of the regulariser $R(y)$ and A is a matrix that represents the convolution operation with the PSF K .

To recover the true object from the observed variables. This formulation can be incorporated into the model-guided DDPM⁴² (See Methods for details on Denoising Diffusion Probabilistic Models and Conditioned Denoising Probabilistic Diffusion models, as well as Supplementary Materials) to obtain improvements over the conditioned case by using a physical prior. However, one cannot access K during inference in several cases. To remediate this, we suggest learning the shifted mean of a guided DDPM by the physical model. In this sense, we incorporate the gradient of the solution of the physical problem as a shift of the mean of the unconditioned case.

We describe the PI Model based on the framework of Guided Diffusion Models³³. We recall that the gradient of the data likelihood denoted as $g = \nabla_{y_t} \log p(x|y_t)$ guides the sampling through the chain. We approximate the transition $p(y_t|y_{t+1})p(x|y_t)$ with the expression:

$$p(y_t|y_{t+1})p(x|y_t) \approx \mathcal{N}(y_t; \mu + \sigma_t^2 g, \sigma_t^2 I) \quad (6)$$

For our specific problem of Microscopy image reconstruction, the conditioning variable x represents the measured data and follows the measurement model $x = Ay + n$, $n \sim N(n; \xi, \xi I)$, and A represents the convolution operation with the microscope PSF. Therefore, considering the posterior of the forward process $q(y_t|y_{t+1}, y_0)$, we express the target posterior as:

$$q(y_t|y_{t+1}, y_0, x) \approx q(y_t|y_{t+1}, y_0)p(x|y_t) \approx \mathcal{N}(y_t; \tilde{\mu} + \sigma_t^2 g, \sigma_t^2 I) \quad (7)$$

Here, $\tilde{\mu}$ is defined as $\tilde{\mu} = \frac{\sqrt{\gamma_t - \beta_t}}{1 - \gamma_t} y_0 + \frac{\sqrt{\alpha_t(1 - \gamma_t)}}{1 - \gamma_t} y_t$.

To train the transition $p_\theta(y_t|y_{t+1}, x)$, analogous to regular DDPMs, we define $p_\theta(y_t|y_{t+1}, x) = \mathcal{N}(y_t; \mu_\theta(y_t, x), \sigma_t^2 I)$, and optimise it using the following mean squared error loss:

$$L_{\text{mean}} := \mathbb{E} \left[\frac{1}{2\sigma_t^2} \|\tilde{\mu}(y_0, y_t) + \sigma_t^2 g - \mu_\theta(y_t, x)\|_2^2 \right] \quad (8)$$

Using the reparametrisation trick, we can rewrite the loss function as:

$$\mathbb{E} \left[\frac{1}{2\sigma_t^2} \left\| \frac{1}{\sqrt{\alpha_t}} \left(y_t(y_0, x, \epsilon) - \frac{\beta_t}{\sqrt{1 - \gamma_t}} \epsilon \right) + \sigma_t^2 g + \mu_\theta(y_t, x) \right\|_2^2 \right] \quad (9)$$

Here, we reparametrise the neural network prediction with $\epsilon_\theta(x, y_t)$, leading to:

$$\mathbb{E} \left[\frac{1}{2\sigma_t^2} \left\| \frac{1}{\sqrt{\alpha_t}} \left(y_t(y_0, x, \epsilon) - \frac{\beta_t}{\sqrt{1 - \gamma_t}} \epsilon \right) + \sigma_t^2 g - \frac{1}{\sqrt{\alpha_t}} \left(y_t(y_0, x, \epsilon) - \frac{\beta_t}{\sqrt{1 - \gamma_t}} \epsilon_\theta(x, y_t) \right) \right\|_2^2 \right] \quad (10)$$

To further simplify, we introduce a factorisation term $C = \frac{\beta_t^2}{2\sigma_t^2 \alpha_t (1 - \gamma_t)}$, resulting in:

$$L_{\text{pi-simple}} = \mathbb{E}_{\epsilon, t, (x, y_0)} \left[C \left\| \epsilon - \epsilon_\theta(\sqrt{\gamma_t} y_0 + \sqrt{1 - \gamma_t} \epsilon, x, y_t) + \frac{\sigma_t^2}{\sqrt{C}} g \right\|_2^2 \right] \quad (11)$$

We use a simplification of the loss function and set $C = 1$ to train our model.³⁰ One remaining challenge is defining the gradient term $\nabla_{y_t} \log p(x|y_t)$. In the context of linear inverse problems, particularly in linear inverse problems tasks, such as microscopy image super-resolution, that solve the following optimisation problem:

$$\min_y \|Ay - x\|_2^2 + \lambda R(y), \quad (12)$$

several approaches have been proposed to approximate $\nabla_{y_t} \log p(x|y_t)$. One line of work suggests an approximation of $\nabla_{y_t} \log p(x|y_t) \approx \frac{A^T(x - Ay_t)}{\sigma^2}$.⁴³ This approximation assumes white Gaussian noise in the measurement process, with variance σ^2 , and utilises the conversion of the convolution operation into its matrix representation A .

Another approach aims to approximate the gradient by evaluating the measurement model in the model estimation at timestep t .⁴² Specifically, they define that the distribution $p(x|y_t)$ can be approximated in the following manner:

$$p(x|y_t) \approx p(x|\hat{y}_0) \quad (13)$$

Where, $\hat{y}_0 = \mathbb{E}[y_0|y_t]$, and can be sampled with $\hat{y}_0 = \frac{y_t - \sqrt{1 - \gamma_t} \epsilon_\theta(y_t, x, y_t)}{\sqrt{\gamma_t}}$.

While this second approximation exhibits improved accuracy over the first one, it introduces a dependency on the model prediction, thereby generating a learning gradient through the function g . As our method seeks to avoid any learning signal through the function g , we find the second approximation unsuitable for our purposes. Introducing a dependency on model predictions may lead to deviations from the intended objective of learning the shifted mean of the forward process posterior.

We adopt the $\nabla_{y_t} \log p(x|y_t) \approx \frac{A^T(x - Ay_t)}{\sigma^2}$ approximation and adapt it to account for noise following a distribution $n \sim N(n; \xi, \xi I)$, where the measurement is given by $x = Ay + n$. Under this measurement model the approximation of $\nabla_{y_t} \log p(x|y_t)$ is

$$\nabla_{y_t} \log p(x|y_t) \approx \frac{A^T(x - \xi - Ay_t)}{\xi^2} \quad (14)$$

Since we train with synthetic data, we can access the noiseless version of the measurement, enabling us to replace $x - \xi$ with its noiseless counterpart. Using the same strategy described in⁴³, we employed annealing on the gradient of Eq. (14) with a variable ν_t that increases as $t \rightarrow 0$. Then the final gradient term is:

$$\nabla_{y_t} \log p(\tilde{x}|y_t) \approx \nu_t \nabla_{y_t} \|Ay_t - \tilde{x}\|_2^2 \quad (15)$$

We find in our experiments that the best results were obtained with $\nu_t = \eta \frac{\sqrt{\gamma_t}}{\|\nabla_{y_t} \log p(x|y_t)\|_2}$, where η is set to 10. Detailed algorithms of training

and sampling are provided in the Methods section in Algorithm 1 and Algorithm 2 respectively.

Training dataset

To ensure that a large dataset is available for training, we simulated the microscopy images (see Fig. 1c, Methods). For this, we have processed images derived from two large datasets ImageNet⁴⁴ and the biological image dataset for super-resolution microscopy (BioSR)²⁹ using our image acquisition model (IAM, see Methods, Supplement). Changes that such processing inflicts on the images are demonstrated in Figs. 1c and 2a. These images used for training bore a strong resemblance to images obtained using LM.

Model training and benchmark evaluation

Once our PI Denoising Diffusion Probabilistic Model (PI-DDPM, see Fig. 1d, e, Methods) was trained, we compared the PI-DDPM model's performance against other mainstream models, including U-Net, DDPM, and the Richardson-Lucy (RL) algorithm (Fig. 2b, c, Sup. Table 1). For this purpose, we trained U-Net and DDPM alongside PI-DDPM and included RL for comparison. Our visual assessment indicated that both DDPM and PI-DDPM exhibited fewer noise and processing artefacts than RL and U-Net. Notably, PI-DDPM excelled in preserving high-frequency details in both ImageNet and BioSR-derived images. To quantify this difference, we evaluated the performance with three metrics: peak single-to-noise ratio⁴⁵, multi-scale structural similarity index measure (MS-SSIM)⁴⁶, and normalised root mean square error (NRMSE). The results demonstrate the superiority of DDPM and PI-DDPM in image restoration (Table 1, BioSR, Sup. Table 1). Furthermore, applying the models to microscopy images from the BioSR test set further established the superiority of both DDPM and PI-DDPM over U-Net across all metrics (Table 1, BioSR). Notably, PI-DDPM improved over DDPM in this more realistic scenario. Additionally, we visually compared the reconstruction results on the BioSR dataset between DDPM and its PI counterpart proposed here. Crucially, DDPM hallucinated non-existent structures, while PI-DDPM remained close to the GT (Sup. Fig. S1a)

Next, to evaluate the model on another benchmark dataset we have employed widefield to structured illumination microscopy (W2S) benchmark⁴⁷. Results suggested that the performance of our PI-DDPM was superior over DDPM and U-Net in PSNR, MS-SSIM and NRMSE metrics (Table 1, W2S).

To expand our choice of metrics, we have computed Fourier ring correlation (FRC) and decorrelation-based resolution⁴⁸ on the individual subcellular structures of the BioSR dataset. These structures included clathrin-coated pits (CCP), endoplasmic reticulum (ER), microtubules (MT) and F-actin. Furthermore, we expanded our comparison to the GAN-based state-of-the-art specialist SR models including DFGAN²⁸ and DFCAN²⁹ (Sup. Fig. S1b). Results suggested that while the specialist DFGAN model often performed best, PI-DDPM performed second-best outperforming the baseline DDPM. Remarkably, when the FRC of high-resolution details was compared on MT and ER, PI-DDPM either outperformed the rest or performed comparably to DFGAN, outperforming the baseline DDPM by a margin.

Finally, we compared how DDPM and PI-DDPM perform in the presence of noise and with only limited data (Sup. Fig. S2). In this experiment, using the F-actin images from the BioSR dataset, we randomly sampled 100 images from the test set and simulated low-quality widefield acquisitions by introducing the increasing amount of controlled Poisson-Gaussian noise. Results suggested that PI-DDPM reconstructions preserved the triangular shape of the filamentous structures of F-actin more reliably than DDPM in high noise conditions (Sup. Fig. S2a, zoomed-in insets). Furthermore, PI-DDPM metrics deteriorated less with increased noise (Sup. Fig. S2b), suggesting higher robustness.

Model evaluation on non-benchmark superresolution microscopy

To assess the models' adaptability, we employed a published non-benchmark real-world microscopy dataset of Direct Stochastic Optical

Reconstruction Microscopy (dSTORM)⁴⁹. This dataset was previously unseen by our model and unprocessed. In this dSTORM dataset, authors provide widefield (low-resolution) and dSTORM (superresolved) microscopy paired images containing mid-zygotene nucleus immunostained for SYCP3 (red), DMC1 (green) and RAD51 (blue) proteins. PI-DDPM demonstrated exceptional performance on the dSTORM dataset, producing results closely aligned with the ground truth (Fig. 3a, b, Methods). As previously the quantitative evaluation was performed by computing metrics on images reconstructed from the widefield images and compared the actual dSTORM reconstruction as GT. The model surpassed both U-Net and DDPM by yielding sharper reconstructions and exhibiting fewer artefacts in low-signal scenarios. Remarkably, PI-DDPM produces visibly fewer artefacts in low signal (Fig. 3b, red channel, upper row) compared to DDPM. Additionally, PI-DDPM preserved the continuity of the filament structures better (Fig. 3b, red channel, lowest row) compared to DDPM. Interestingly, DDPM seems to overemphasise the green channel (DMC1) possibly due to the high signal-to-noise. However, all models struggled to capture low signal-to-noise punctae in RAD51 (blue channel). A thorough quantitative evaluation across different channels indicated that both DDPM and PI-DDPM consistently outperformed other models, with PI-DDPM having a slight edge in PSNR and NRMSE metrics and showing comparable results in MS-SSIM to DDPM (Table 1, dSTORM).

Model evaluation on real-world prospectively acquired microscopy

We further tested our model using a correlative widefield-confocal microscopy dataset of cell nuclei, which we acquired prospectively to test our model (Fig. 4)⁵⁰. In this dataset, both confocal and widefield stacks of the same region were taken by automated microscopy (see Methods). Since confocal microscopy (Fig. 4b) is known to have better resolution compared to widefield microscopy (Fig. 4a) it can serve as a guide on the correctness of image restoration in the absence of *bona fide* ground truth. Consistent with our previous observations, we noted that DDPM and PI-DDPM have shown significantly lower blur in the reconstructions (Fig. 4a, top row). Remarkably, comparing correlated fields of view (Fig. 4a,b, white asterisk) PI-DDPM reconstructions show more consistency with the confocal image than the conventional DDPM or U-Net irrespective of whether the input image comes from widefield or confocal microscopy. Furthermore, PI-DDPM showed more consistent output in the case of a low signal-to-noise ratio beyond the edge of the specimen (Fig. 4b, bottom).

Ablation study: the importance of regularisation and parameter η

To evaluate the importance of the regularisation term and the parameter η in our model, we conducted a series of experiments on a randomly selected subset of the BioSR test dataset (Fig. 5). This ablation study aimed to provide insights into the contributions and effects of these components on the overall performance of our proposed approach. Results are presented in Sup. Table S2 and S3, as well as Fig. 5.

One fundamental aspect under investigation is the impact of the regularisation term applied to the learning process. In the context of conditioned DDPMs (DDPM baseline), we follow the next regularisation formulation:

$$\nabla_{y_t} \log p(y_t) \approx \nabla_{y_t} \lambda_t R(y_t), \quad (16)$$

where $\nabla_{y_t} \log p(y_t)$ denotes the gradient of the log-likelihood with respect to the perturbed latent variable y_t , λ_t is the regularisation coefficient at time step t , and $R(y_t)$ represents the regularisation term.

In contrast, our proposed model introduces the learned gradient and the regularisation term as follows:

$$\nabla_{y_t} \log p(\tilde{x}|y_t) \approx \nabla_{y_t} [||Ay_t - \tilde{x}||_2^2 + \lambda_t R(y_t)], \quad (17)$$





Fig. 2 | Proposed model with physics-informed probabilistic denoising diffusion for Microscopy image reconstruction. **a** Examples of processing performed using Image Acquisition Model (IAM) on (left-to-right) Shepp-Logan phantom, ImageNet and BioSR images. Simulated examples. **b** Examples of input (Widefield) and reconstructed images using Richardson-Lucy (RL), U-Net, Denoising Diffusion Probabilistic

Models (DDPM) and physics-informed DDPM (PI-DDPM) in ImageNet and BioSR-derived images, respectively. Simulated dataset based on ImageNet dataset. **c** Denoising Diffusion Probabilistic Models (DDPM) and Physics-informed DDPM (PI-DDPM) in BioSR images. BioSR is a dataset of real widefield and SIM micrographs. GT stands for ground truth. MT stands for microtubules. Scale bar in micrographs is 2 μ m.

Fig. 3 | Model performance on unseen super-resolution dataset. a Image containing mid-zygote nucleus immunostained for SYCP3 (red), DMC1 (green) and RAD51 (blue) proteins from ref. 49. **b** Examples of input (Widefield) and reconstructed images using U-Net, Denoising Diffusion Probabilistic Models (DDPM) and physics-informed DDPM (PI-DDPM) in dSTORM images. The scale bar is 2 μm .



where A is the linear operator of the forward problem, \tilde{x} represents the measured image without noise, and the term $\|Ay_t - \tilde{x}\|_2^2$ captures the reconstruction error between the transformed latent variable Ay_t and the observation \tilde{x} . For our evaluation, we employed two common regularisers: L_1, L_2 . These regularisers emphasise common characteristics of images, such as sparsity and smoothness.

We also investigate the influence of the parameter η on the learning process (Sup. Fig. 1). The parameter η acts as a weight for the learned gradient, influencing how it contributes to optimisation. We observe as the parameter increases that the metrics improve especially in the case of NRMSE. However, beyond $\eta = 25$, the performance metrics start worsening. This effect is because the learned gradient is not perfect, as the method implicitly has to learn the PSF K . Furthermore, since we are also using an approximation of $\nabla_{y_t} \log p(\tilde{x}|y_t)$ which becomes worse as t approaches T , having a constant η will propagate errors further when obtaining the reconstruction.


Discussion

Despite the immense progress in microscopy, the ability to visualise the microscopic world remains limited due to hardware imperfections and physics boundaries⁸. While recent advances in DL and generative models promise to assist in overcoming these barriers, these models come with their own set of limitations. DDPMs have shown great promise for generative modelling^{30,51,52}.

In applications such as biomedical image restoration and superresolution microscopy, the primary focus is on preventing hallucinations, eliminating artefacts, and ensuring the accuracy of the produced structures.

In this work, we demonstrate that incorporating the optical physical prior into the DDPM model enhances stability and yields more realistic reconstruction results. We achieve this by introducing a novel Physics-Informed Diffusion Model (PI-DDPM), a framework designed to enhance Diffusion Probabilistic Models (DDPM) by incorporating physics-based terms. We demonstrate meaningful improvements over the traditional DDPM framework by integrating domain-specific physical knowledge. Our detailed contributions are as follows. Firstly, we propose a method to incorporate a physics-informed (PI) term directly into the DDPM framework, enabling the model to account for domain-specific physical constraints during generation. This formulation is applicable to scenarios where physical knowledge, such as point spread functions (PSFs), is otherwise unavailable. Secondly, we demonstrate that our model can effectively learn the forward operator during training, leveraging this knowledge to apply a PI term dynamically during inference. This allows the model to remain physics-aware even when key imaging parameters are unknown, improving interpretability and robustness. Finally, we introduce a regularisation strategy for the inference process that consistently enhances performance compared to traditional DDPM approaches. Our results show that adding the PI term and regularising the inference process leads to improvements in

Fig. 4 | Model performance on prospective correlative widefield-confocal microscopy. a Examples of widefield images of cell nuclei and their reconstructions using U-Net, Denoising Diffusion Probabilistic Models (DDPM) and physics-informed DDPM (PI-DDPM). **b** Examples of confocal images and their reconstructions. Top, centre and bottom refer to the Z-positions of the slices. The bottom lies beyond the edge of the specimen. The scale bar is 10 μm . Asterisk (*) marks correlated images of the same cell and focal plane.



model accuracy and consistency, especially in scenarios with limited data or high noise.

While the PI-DDPM we propose may not outperform specialist models like DFGAN²⁸, it significantly outperforms its baseline DDPM in various metrics and in producing outputs closer to GT. Using various benchmarks and real datasets, we demonstrate that introducing a physical prior lowers the hallucinations of generative models like DDPM. Our approach extends a recently introduced paradigm of PI neural networks³⁶ and extends the applicability of these methods to microscopy. Furthermore, since our model learns the distribution from the data, it estimates the mean and variance of the obtained reconstructions directly. This offers a convenient way to obtain confidence in the reconstructions, potentially facilitating broader adoption by the biomedical community.

Methods

Denoising diffusion probabilistic models

Unconditioned DDPMs start by setting a data distribution $y_0 \sim q(y_0)$ and a Markovian noising process q , which incrementally injects noise, resulting in noised samples y_1 to y_T . Each step of the noising process introduces Gaussian noise according to a variance schedule denoted by β_i :

$$q(y_t|y_{t-1}) = N(y_t; \sqrt{1 - \beta_t}y_{t-1}, \beta_t I) \quad (18)$$

It is unnecessary to apply q repeatedly to sample from $y_t \sim q(y_t|y_0)$. Instead, $q(y_t|y_0)$ can be represented as a Gaussian distribution. Defining

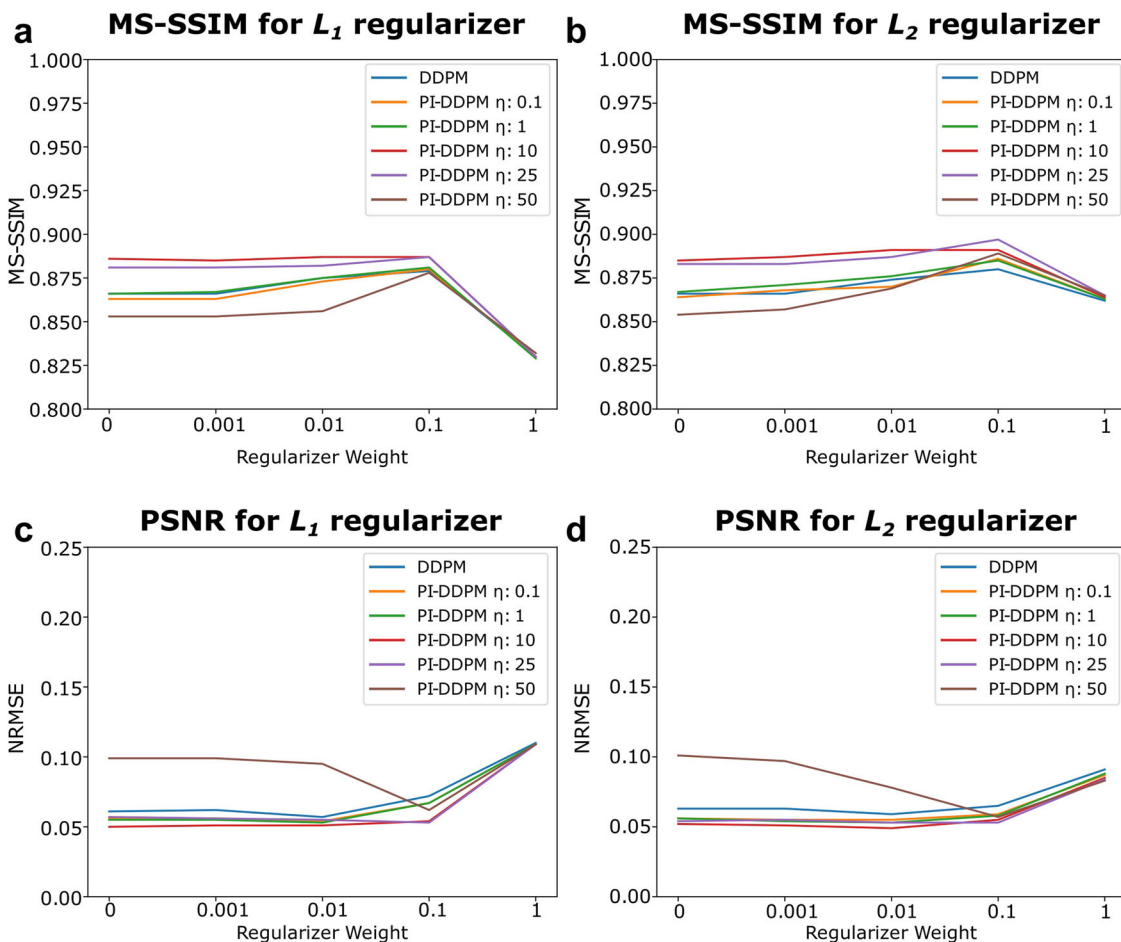


Fig. 5 | Influence of the parameter η , and different regularisers on the learning process. **a** Influence of decreasing η measured by Multi-Scale Structural Similarity (MS-SSIM) for L_1 regularisation. **b** Influence of decreasing η measured by

MS-SSIM for L_2 regularisation. **c** Influence of decreasing η measured by the peak signal-to-noise ratio (PSNR) for L_1 regularisation. **d** Influence of decreasing η measured by PSNR for L_2 regularisation.

$\alpha_t = 1 - \beta_t$ and $\gamma_t = \prod_{s=0}^t \alpha_s$, we have:

$$q(y_t|y_0) = N(y_t; \sqrt{\gamma_t}y_0, (1 - \gamma_t)I) \tag{19}$$

$$= \sqrt{\gamma_t}y_0 + \epsilon \sqrt{1 - \gamma_t}, \epsilon \sim N(0, I) \tag{20}$$

Here, $1 - \gamma_t$ indicates the noise variance at any given time step, which could also define the noise schedule instead of β_t .

Applying Bayes' theorem, the posterior $q(y_{t-1}|y_t, y_0)$ is also a Gaussian with mean $\tilde{\mu}_t(y_t, y_0)$ and variance $\tilde{\beta}_t$:

$$\tilde{\mu}_t(y_t, y_0) = \frac{\sqrt{\gamma_{t-1}}\beta_t}{1 - \gamma_t} y_0 + \frac{\sqrt{\alpha_t}(1 - \gamma_{t-1})}{1 - \gamma_t} y_t \tag{21}$$

$$\tilde{\beta}_t = \frac{1 - \gamma_{t-1}\beta_t}{1 - \gamma_t} \tag{22}$$

$$q(y_{t-1}|y_t, y_0) = N(y_{t-1}; \tilde{\mu}_t(y_t, y_0), \tilde{\beta}_t I) \tag{23}$$

To sample from the data distribution $q(y_0)$, one begins with $q(y_T)$ and sequentially samples reverse steps $q(y_{t-1}|y_t)$ until reaching y_0 . Given appropriate settings for β_t and T , $q(y_T)$ approximates an isotropic Gaussian, making the sampling of y_T straightforward. The subsequent step involves approximating $q(y_{t-1}|y_t)$ using a neural network to estimate its parameters:

mean μ_θ and a covariance matrix Σ_θ :

$$p_\theta(y_{t-1}|y_t) = N(y_{t-1}; \mu_\theta(y_t, t), \Sigma_\theta(y_t, t)) \tag{24}$$

To ensure $p(y_0)$ aligns with the true data distribution $q(y_0)$, one can optimise the variational lower-bound L_{vlb} for $p_\theta(y_0)$:

$$L_{vlb} = L_0 + L_1 + \dots + L_{T-1} + L_T$$

$$L_0 = -\log p_\theta(y_0|y_1)$$

$$L_{t-1} = D_{KL}(q(y_{t-1}|y_t, y_0) || p_\theta(y_{t-1}|y_t))$$

$$L_T = D_{KL}(q(y_T|y_0) || p(y_T))$$

However, an alternative objective yields better practical results. Instead of directly parameterizing $\mu_\theta(y_t, t)$ as a neural network, they trained a model $\epsilon_\theta(y_t, t)$ to predict ϵ from the aforementioned equation. The simplified objective is defined as:

$$L_{\text{simple}} = E_{t \sim [1, T], y_0 \sim q(y_0), \epsilon \sim N(0, I)} [||\epsilon - \epsilon_\theta(y_t, t)||^2] \tag{25}$$

For sampling, $\mu_\theta(y_t, t)$ can be derived from $\epsilon_\theta(y_t, t)$ via substitution:

$$\mu_\theta(y_t, t) = \frac{1}{\sqrt{\alpha_t}} \left(y_t - \frac{\beta_t}{\sqrt{1-\gamma_t}} \epsilon_\theta(y_t, t) \right) \quad (26)$$

and sample using the following recurrence:

$$y_{t-1} = \mu_\theta(y_t, t) + \sqrt{\Sigma_\theta} \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, I)$$

It is feasible to fix $\Sigma_\theta(y_t, t)$ to a constant value, such as $\beta_t I$ or $\tilde{\beta}_t I$, representing upper and lower bounds for the true reverse step variance, instead of learning the covariance³⁰.

Conditioned denoising probabilistic diffusion models

In the conditioned case, samples are drawn from an unknown distribution $p(y|x)$. This is referred to as distribution because conditioned image synthesis is, by nature, ill-posed. Specifically, there are many possible solutions y for any given input x . In the conditioned process, we want to learn an approximation of this distribution. To achieve that, it is possible to condition DDPMs in two ways.

The first way is to redefine the Markov chain. Given an image y and some corruption process $p(x|y)$, we want to learn $p(y|x)$. To achieve this, the diffusion Markov chain states are concatenated by the respective conditioning image x ³². Specifically, the distributions of the states of the Markov chain are generated by the diffusion process $q(y_t|y_{t-1})$ and concatenated to the conditioning image x . Where $y_t = \sqrt{\gamma_t} y_0 + \epsilon \sqrt{1-\gamma_t}$, $\epsilon \sim \mathcal{N}(\epsilon; 0, I)$. To learn the reverse process, a reverse Markov chain is established, where $p(y_T) = \mathcal{N}(y_T; 0, I)$:

$$p_\theta(y_{t-1}|y_t, x) = \mathcal{N}(y_{t-1}|\mu_\theta(x, y_t, \gamma_t), \sigma_t^2 I). \quad (27)$$

Using the same formulation as in regular DDPMs³², the model is trained to predict ϵ at each time step:

$$\min_{\theta} L_{\text{simple}} := \mathbb{E}_{t, (x, y_0), \epsilon} \|\epsilon - \epsilon_\theta(y_t, x, \gamma_t)\|_2^2. \quad (28)$$

Finally, to obtain y_0 , the same iterative denoising as in the regular DDPMs is applied.

$$y_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(y_t - \frac{\beta_t}{\sqrt{1-\gamma_t}} \epsilon_\theta(x, y_t, \gamma_t) \right) + \sqrt{\beta_t} \epsilon_t. \quad (29)$$

Model-guided denoising probabilistic diffusion models

The second way to obtain conditioned samples from a diffusion model is to condition an unconditioned reverse process³³. Given an unconditional reverse process $p_\theta(y_t|y_{t+1})$, to condition on label x we can factorise

$$p_{\theta, \phi}(y_t|y_{t+1}, x) = Z p_\theta(y_t|y_{t+1}) p_\phi(x|y_t), \quad (30)$$

where Z is a normalizing constant. This expression can be approximated as a perturbed Gaussian distribution. Since our unconditioned reverse process is a Gaussian, we have:

$$p_\theta(y_t|y_{t+1}) = \mathcal{N}(\mu, \sigma_t^2 I) \quad (31)$$

$$\log p_\theta(y_t|y_{t+1}) = -\frac{1}{2\sigma_t^2} (y_t - \mu)^T (y_t - \mu) + C. \quad (32)$$

Since, at infinity, the distribution of the reverse process tends to a delta distribution, then it is reasonable to approximate $p_\phi(x|y_t)$ by its Taylor

expansion around the mean.

$$\begin{aligned} \log p_\phi(x|y_t) &\approx \log p_\phi(x|y_t)|_{y_t=\mu} + (y_t - \mu) \nabla_{y_t} \\ \log p_\phi(x|y_t)|_{y_t=\mu} &= (y_t - \mu)g + C_1 \end{aligned} \quad (33)$$

where $g = \nabla_{y_t} \log p_\phi(x|y_t)$

Finally, by replacing and rearranging, we get:

$$p_\theta(y_t|y_{t+1}) p_\phi(x|y_t) \approx \mathcal{N}(\mu + \sigma_t^2 g, \sigma_t^2 I). \quad (34)$$

Thus, the reverse conditioned process approximates the unconditioned Gaussian transition with its mean shifted by $\sigma_t^2 g$.

Physics-informed denoising probabilistic diffusion models training

We trained our PI-DDPM employing the algorithm DDPM as outlined in Algorithm 1. For inference, we used the algorithm outline in Algorithm 2.

Algorithm 1. Physics-Informed Denoising Diffusion Model (PI-DDPM) Training

```

Input : Training data  $Y = \{y_j\}$ , Number of iterations  $N$ , Number of Timesteps  $T$ , Learning Rate  $\beta$ 
Output :  $\epsilon_\theta$ 
for  $k = 1$  to  $N$  do
    Generate PSF  $K$  according to random microscope parameters
     $\epsilon \sim \mathcal{N}(0, I)$ 
     $y_0 \sim Y$ 
     $t \sim U(0, T)$ 
     $\gamma \sim p(\gamma)$ 
     $y_t \leftarrow \sqrt{\gamma} y_0 + \sqrt{1-\gamma} \epsilon$ 
     $n \sim \mathcal{N}(\xi, \xi I)$ 
     $\tilde{x} \leftarrow K * y_0$ 
     $x \leftarrow K * y_0 + n$ 
     $g \leftarrow \nabla_{y_t} \|K * y_t - \tilde{x}\|_2^2$ 
     $\theta_k \leftarrow \theta_{k-1} - \omega \nabla_{\theta} L_{\text{pi-simple}}(\epsilon, \epsilon_\theta, x, \gamma, g)$ 

```

Algorithm 2. Physics-Informed Denoising Diffusion Model (PI-DDPM) Sampling

```

Input : Microscope data  $x$ , Number of Timesteps  $T$ , Regulariser  $R(y)$ 
Output : Predicted  $y_0$ 
 $y_T \sim \mathcal{N}(0, I)$ 
for  $t = T$  to  $0$  do
    if  $t = 1$  then
         $\epsilon \leftarrow 0$ 
    else
         $\epsilon \sim \mathcal{N}(0, I)$ 
     $y_{t-1} \leftarrow \frac{1}{\sqrt{\alpha_t}} \left( y_t - \frac{\beta_t}{\sqrt{1-\gamma_t}} \epsilon_\theta(x, y_t, \gamma_t) \right) + \sqrt{\beta_t} \epsilon + v_t g_\theta(y_t, x) + \lambda_t R(y_t)$ 

```

Metrics

To assess the performance of all the models in this study we used multi-scale structural similarity index measure (MS-SSIM)⁴⁶, normalised root mean square error NRMSE := $\frac{\sqrt{\text{MSE}}}{y_{\text{max}} - y_{\text{min}}}$, and peak single-to-noise ratio (PSNR). MS-SSIM metric is defined as:

$$\text{MS-SSIM}(x, y) := [I_M(x, y)]^{\alpha_M} \prod_{j=1}^M [c_j(x, y)]^{\beta_j} [s_j(x, y)]^{\gamma_j}, \quad (35)$$

where I_j , c_j , and s_j are the measures of luminance, contrast, and structure corresponding to scale j . We used five scales and $\alpha_j = \beta_j = \gamma_j = 1$ in accordance with the parameters reported in⁴⁶. PSNR was defined as: $\text{PSNR} := 20 \log_{10}(\text{MAX}_I) - 10 \log_{10}(\text{MSE})$, where MAX_I is the maximum pixel value of the image and MSE is the mean square error.

Training details

DDPM and PI-DDPM were trained in a cluster environment employing a single Nvidia A100 with 40GB of vRAM. The batch size was 16. Pre-training continued for 1 million iterations using an ImageNet-derived dataset and fine-tuned for 800,000 iterations on the mixed dataset. The U-Net model was trained in a cluster environment using a single Nvidia V100 GPU equipped with 32GB of vRAM. Pre-training continued for 800,000 iterations, and fine-tuning for 100,000.

Simulated dataset generation

To generate a simulated dataset for training our models, we utilised two primary sources of images: photographs from the ImageNet database⁴⁴ and structured illumination microscopy (SIM) images from the BioSR dataset²⁹. Our approach involved processing each image through a forward microscopy model, termed the Image Acquisition Model (see Methods and Fig. 1c). This model simulates the effects of different microscopes by convolving each image with a Point Spread Function (PSF) and then applying Poisson noise.

For the ImageNet-derived dataset, which comprises 1.2 million training images and 100,000 test images, we generated 30,000 unique PSFs. The generation of each PSF involved random sampling of physically plausible parameters for microscopy systems: numerical aperture (0.4 to 1.0), excitation wavelength (320 to 400 μm), emission wavelength (450 to 550 μm), and pinhole size (0.1 to 1000 μm). The focal plane was centered in the volume object with a refractive index of 1.33, corresponding to air. We then convolved the ImageNet images with these PSFs to create a training dataset that mimics microscopy images.

Additionally, we incorporated a simulated dataset derived from the publicly available BioSR dataset²⁹. BioSR contains approximately 20,000 SIM fluorescence images of various subcellular structures, such as clathrin-coated pits (CCPs), endoplasmic reticulum (ER), microtubules (MT), and F-actin, imaged at different fluorescence levels. For this part of the dataset, we took all high-resolution images from BioSR and processed them using the Image Acquisition Model to simulate low-resolution images. These simulated BioSR images, adhering to the original dataset's train-test split, were used for fine-tuning our models. The high-resolution BioSR images served as the ground truth during training, while the simulated low-resolution images were used as input. At test time, we employed BioSR's low-resolution (widefield) images as inputs.

This combined, processed dataset from ImageNet and BioSR, which includes both training and test images, was employed for training and fine-tuning various models, including U-Net²⁷, DDPM³⁰, and our proposed PI-DDPM model. The goal was to reconstruct high-resolution images from the simulated blurred micrographs.

Direct stochastic optical reconstruction microscopy dataset

To test how our method compares to the state-of-the-art single-molecule localisation microscopy (SMLM) we have employed a publicly available three-color Direct Stochastic Optical Reconstruction Microscopy (dSTORM) dataset⁴⁹. In this dataset authors provide a Widefield and SMLM reconstructed high-resolution image containing mid-zygotene nucleus immunostained for SYCP3 (red), DMC1 (green) and RAD51 (blue) proteins. Images in this dataset were acquired by Zeiss Elyra PS1 microscope using a 100x 1.46NA oil immersion objective. Further imaging details are provided by the authors in the following publication⁵³.

Prospective correlative widefield-confocal microscopy dataset

Finally, to test how our model would perform on a prospectively acquired dataset, we have obtained a correlated widefield-confocal microscopy dataset. For this, A549 lung carcinoma cell line cells were seeded in 96-well imaging plates a night prior to imaging, then fixed with 4% paraformaldehyde (Sigma) and stained for DNA with Hoechst 33342 fluorescent dye (Sigma). Cell culture was maintained similarly to the procedures described in⁵⁴. Next, stained cell nuclei were imaged using ImageXpress Confocal system (Molecular Devices) in either confocal or widefield mode employing Nikon 20X Plan Apo Lambda objective. To obtain 3D

information images in both modes were acquired as Z-stacks with 0.3 μm and 0.7 μm for confocal and widefield modes respectively. Confocal z-stack was Nyquist sampled. The excitation wavelength was 405 nm and the emission 452 nm. Using these settings we have obtained 72 individual stacks for both modalities, with each stack covering 2048 by 2048 pixels or 699 by 699 μm .

Widefield to SIM (W2S) dataset

Widefield to SIM (W2S)⁴⁷ is a collection obtained through conventional fluorescence widefield and SIM (Structured Illumination Microscopy) imaging techniques. This dataset encompasses 144,000 authentic fluorescence microscopy images, aggregating into a total of 360 distinct image sets. Each set consists of low-resolution (LR) widefield images exhibiting various levels of noise, an accompanying noise-free LR image, and a corresponding high-resolution (HR) SIM image of superior quality.

Source code

The source code for this work is available on GitHub (see Code availability). To train the model place your data generated by the dataset_generation script (if you are generating simulated data) or the STORM script if you are generating the respective dSTORM dataset. In the train_ddpm or train_unet script change the paths of the loading data to the ones that you generated. Next, choose a training modality, either widefield or confocal. Finally, run the script.

To test the model generate your testing dataset using the dataset_generation script. Change the paths corresponding to your data. Next, change the paths to the weights files that you wish to use. Finally, run the test script. Due to the size limitation of the submission for the prospective dataset containing correlative widefield-confocal fluorescence microscopy, we provide only a single stack as an example (data/teaser_c_w_test.npz).

Data availability

The correlated widefield confocal microscopy dataset acquired prospectively in this work is available under a Creative Commons license on the RODARE data repository⁵⁰. All other datasets used in this work are published and available from their respective sources.

Code availability

The source code for this work is available on GitHub under GNU General Public License v3.0 and can be accessed at⁵⁵ and <https://github.com/casus/pi-ddpm>. Additionally, the code for decorrelation-based resolution measurement can be found in⁴⁸ and <https://github.com/casus/pyres>. See Methods for details required to reproduce the results.

Received: 26 March 2024; Accepted: 6 December 2024;

Published online: 29 December 2024

References

- Knapper, J. et al. Fast, high-precision autofocus on a motorised microscope: Automating blood sample imaging on the openflexure microscope. *J. Microsc.* **285**, 29–39 (2022).
- Salido, J., Bueno, G., Ruiz-Santaquiteria, J. & Cristobal, G. A review on low-cost microscopes for open science. *Microsc. Res. Tech.* **85**, 3270–3283 (2022).
- Maia Chagas, A., Prieto-Godino, L. L., Arrenberg, A. B. & Baden, T. The€ 100 lab: A 3d-printable open-source platform for fluorescence microscopy, optogenetics, and accurate temperature control during behaviour of zebrafish, drosophila, and caenorhabditis elegans. *PLoS Biol.* **15**, e2002702 (2017).
- Aidukas, T., Eckert, R., Harvey, A. R., Waller, L. & Konda, P. C. Low-cost, sub-micron resolution, wide-field computational microscopy using open-source hardware. *Sci. Rep.* **9**, 7457 (2019).
- Conchello, J. & Lichtman, J. Fluorescence microscopy. *Nat. Methods* **2**, 910–919 (2005).

6. Minsky, M. Memoir on inventing the confocal scanning microscope. *Scanning* **10**, 128–138 (1988).
7. Schermelleh, L. et al. Super-resolution microscopy demystified. *Nat. cell Biol.* **21**, 72–84 (2019).
8. Nechiporuk-Zloy, V. *Principles of Light Microscopy: From Basic to Advanced* (Springer Nature, 2022).
9. Shaw, P. J. & Rawlins, D. J. The point-spread function of a confocal microscope: its measurement and use in deconvolution of 3-d data. *J. Microsc.* **163**, 151–165 (1991).
10. Richardson, W. H. Bayesian-based iterative method of image restoration. *J. Opt. Soc. Am.* **62**, 55–59 (1972).
11. Nasse, M. J. & Woehl, J. C. Realistic modeling of the illumination point spread function in confocal scanning optical microscopy. *J. Opt. Soc. Am. A* **27**, 295–302 (2010).
12. Lucy, L. B. An iterative technique for the rectification of observed distributions. *Astronomical J.* **79**, 745 (1974).
13. Dertinger, T., Colyer, R., Iyer, G., Weiss, S. & Enderlein, J. Fast, background-free, 3d super-resolution optical fluctuation imaging (sofi). *Proc. Natl Acad. Sci.* **106**, 22287–22292 (2009).
14. Salsman, J. & Dellaire, G. Super-resolution radial fluctuations (srrf) microscopy. In *Fluorescent Microscopy*, 225–251 (Springer, 2022).
15. Calisesi, G. et al. Compressed sensing in fluorescence microscopy. *Prog. Biophysics Mol. Biol.* **168**, 66–80 (2022).
16. Sarder, P. & Nehorai, A. Deconvolution methods for 3-D fluorescence microscopy images. *IEEE Signal Process. Mag.* **23**, 32–45 (2006).
17. McNally, J. G., Karpova, T., Cooper, J. & Conchello, J. A. Three-dimensional imaging by deconvolution microscopy. *Methods* **19**, 373–385 (San. Diego, Calif., 1999).
18. Lustig, M., Donoho, D. & Pauly, J. M. Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magn. Reson. Med.* **58**, 1182–1195 (2007).
19. Dutta, J., Ahn, S., Li, C., Cherry, S. R. & Leahy, R. M. Joint L1 and total variation regularization for fluorescence molecular tomography. *Phys. Med. Biol.* **57**, 1459 (2012).
20. Tang, J., Germain, R. N. & Cui, M. Superpenetration optical microscopy by iterative multiphoton adaptive compensation technique. *Proc. Natl Acad. Sci.* **109**, 8434–8439 (2012).
21. Yoo, S. et al. *Bayesian Approach for Automatic Joint Parameter Estimation in 3D Image Reconstruction from Multi-Focus Microscope* (2018).
22. Weigert, M. et al. Content-aware image restoration: pushing the limits of fluorescence microscopy. *Nat. methods* **15**, 1090–1097 (2018).
23. Buchholz, T.-O. et al. Content-aware image restoration for electron microscopy. *Methods cell Biol.* **152**, 277–289 (2019).
24. Noh, H., Hong, S. & Han, B. Learning deconvolution network for semantic segmentation. In *Proc. IEEE International Conference on Computer Vision*, 1520–1528 (2015).
25. Saguy, A. et al. This microtubule does not exist: Super-resolution microscopy image generation by a diffusion model. *bioRxiv* 2023–07 (2023).
26. Xu, L., Ren, J. S., Liu, C. & Jia, J. Deep convolutional neural network for image deconvolution. *Adv. Neural Inf. Process. Syst.* **27** (2014).
27. Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*, 234–241 (Springer, 2015).
28. Saqlain, A. S., Fang, F., Wang, L.-Y., Ahmad, T. & Abidin, Z. U. Dfgan: Image deblurring through fusing light-weight attention and gradient-based filters. In *2022 IEEE 2nd International Conference on Information Communication and Software Engineering (ICICSE)*, 110–114 (IEEE, 2022).
29. Qiao, C. et al. Evaluation and development of deep neural networks for image super-resolution in optical microscopy. *Nat. Methods* **18**, 194–202 (2021).
30. Ho, J., Jain, A. & Abbeel, P. Denoising Diffusion Probabilistic Models. In *Advances in Neural Information Processing Systems*, vol. 33, 6840–6851. <https://proceedings.neurips.cc/paper/2020/hash/4c5bfcfec8584af0d967f1ab10179ca4b-Abstract.html> (Curran Associates, Inc., 2020).
31. Li, H. et al. SRDiff: Single Image Super-Resolution with Diffusion Probabilistic Models. *arXiv:2104.14951 [cs]* <http://arxiv.org/abs/2104.14951>. ArXiv: 2104.14951 (2021).
32. Saharia, C. et al. Palette: Image-to-Image Diffusion Models. In *Special Interest Group on Computer Graphics*, 1–10 <https://doi.org/10.1145/3528233.3530757> (ACM, Vancouver BC Canada, 2022).
33. Dhariwal, P. & Nichol, A. Diffusion Models Beat GANs on Image Synthesis <http://arxiv.org/abs/2105.05233>. ArXiv:2105.05233 [cs, stat] (2021).
34. Bhadra, S., Kelkar, V. A., Brooks, F. J. & Anastasio, M. A. On Hallucinations in Tomographic Image Reconstruction. *IEEE Trans. Med. imaging* **40**, 3249–3260 (2021).
35. Uzunova, H., Wilms, M., Forkert, N. D., Handels, H. & Ehrhardt, J. A systematic comparison of generative models for medical images. *Int. J. Computer Assist. Radiol. Surg.* **17**, 1213–1224 (2022).
36. Raissi, M., Perdikaris, P. & Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput Phys* **378**, 686–707 (2019).
37. Ho, J. et al. Video Diffusion Models. *arXiv:2204.03458 [cs]* <http://arxiv.org/abs/2204.03458>. ArXiv: 2204.03458 (2022).
38. Pawley, J. *Handbook of biological confocal microscopy*, vol. 236 (Springer Science & Business Media, 2006).
39. Born, M. & Wolf, E. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light* (Cambridge University Press, 1999).
40. Arnison, M. R. & Sheppard, C. J. A 3d vectorial optical transfer function suitable for arbitrary pupil functions. *Opt. Commun.* **211**, 53–63 (2002).
41. Sheppard, C. & Larkin, K. Vectorial pupil functions and vectorial transfer functions. *Opt. - Int. J. Light Electron Opt.* **107**, 79–87 (1997).
42. Chung, H., Kim, J., Mccann, M. T., Klasky, M. L. & Ye, J. C. Diffusion posterior sampling for general noisy inverse problems 2209.14687 (2023).
43. Jalal, A. et al. Robust Compressed Sensing MRI with Deep Generative Priors <http://arxiv.org/abs/2108.01368>. ArXiv:2108.01368 [cs, math, stat] (2021).
44. Deng, J. et al. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255 (IEEE, 2009).
45. Elbadawy, O., El-Sakka, M. R. & Kamel, M. S. An information theoretic image-quality measure. In *Conference Proceedings. IEEE Canadian Conference on Electrical and Computer Engineering (Cat. No. 98TH8341)*, vol. 1, 169–172 (IEEE, 1998).
46. Wang, Z., Simoncelli, E. P. & Bovik, A. C. Multiscale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers*, 2003, vol. 2, 1398–1402 (IEEE, 2003).
47. Zhou, R. et al. W2s: Microscopy data with joint denoising and super-resolution for widefield to sim mapping 2003.05961 (2020).
48. Deshpande, N. & Yakimovich, A. Pyres: python-based decorrelation analysis for resolution estimation <https://doi.org/10.5281/zenodo.12699495> (2024).
49. Lieke Koornneef, W. B. & Slotman, J. dSTORM Imaging and Analysis of Recombination Foci in Mouse Spread Meiotic Nuclei. <https://www.ebi.ac.uk/biostudies/bioimages/studies/S-BIAD627> (2023).
50. Li, R. et al. Correlated widefield-confocal microscopy dataset <https://doi.org/10.14278/rodare.2668> (2024).

51. Nichol, A. Q. & Dhariwal, P. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, 8162–8171 (PMLR, 2021).
52. Lugmayr, A. et al. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11461–11471 (2022).
53. Koornneef, L. et al. Multi-color dSTORM microscopy in *hormad1*^{-/-} spermatocytes reveals alterations in meiotic recombination intermediates and synaptonemal complex structure. *PLoS Genet.* **18**, e1010046 (2022).
54. Yakimovich, A. et al. Plaque2.0—a high-throughput analysis framework to score virus-cell transmission and clonal cell expansion. *PLoS one* **10**, e0138760 (2015).
55. Li, R. et al. casus/pi-ddpm: Initial release <https://doi.org/10.5281/zenodo.10406506> (2023).

Acknowledgements

We thank Michael Hecht and his group for critical reading of this work. This work was partially funded by the Center for Advanced Systems Understanding (CASUS) which is financed by Germany's Federal Ministry of Education and Research (BMBF) and by the Saxon Ministry for Science, Culture, and Tourism (SMWK) with tax funds on the basis of the budget approved by the Saxon State Parliament. MK was supported by the Heisenberg award from the DFG (KU 3222/2-1), as well as funding from the Helmholtz Association. The authors acknowledge the financial support by the Federal Ministry of Education and Research of Germany and by Sächsische Staatsministerium für Wissenschaft, Kultur und Tourismus in the programme Center of Excellence for AI-research "Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig", project identification number: ScaDS.AI.

Author contributions

RL, GM, AYU, MK and AY conceived the idea and planned the computational experiments. RL and GM designed the mathematical models, wrote the program code and conducted computational experiments. ND implemented PyRes code. VA and AP prepared and acquired prospective experimental data. RL, GM, VA, AP, AYU, MK and AY wrote the manuscript.

Funding

Open Access funding enabled and organized by Projekt DEAL.

Competing interests

RL, GM, VA, AP, AYU, MK declare no competing interests. AY declares the following competing interest: role as an Editorial Board Member in Scientific Data.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s44172-024-00331-z>.

Correspondence and requests for materials should be addressed to Artur Yakimovich.

Peer review information *Communications Engineering* thanks Philip Wijesinghe and the other, anonymous, reviewer for their contribution to the peer review of this work. Primary Handling Editors: Or Perlman and Anastasiia Vasylichenkova.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024