

Complex-Valued Retrievals From Noisy Images Using Diffusion Models

Nadav Torem^{*†} Roi Ronen^{*†} Yoav Y. Schechner[†] Michael Elad[‡]

[†]Viterbi Faculty of Electrical and Computer Engineering & [‡]Department of Computer Science
Technion - Israel Institute of Technology, Haifa, Israel

{torem@campus, roironen@campus, yoav@ee, elad@cs}.technion.ac.il

Abstract

In diverse microscopy modalities, sensors measure only real-valued intensities. Additionally, the sensor readouts are affected by Poissonian-distributed photon noise. Traditional restoration algorithms typically aim to minimize the mean squared error (MSE) between the original and recovered images. This often leads to blurry outcomes with poor perceptual quality. Recently, deep diffusion models (DDMs) have proven to be highly capable of sampling images from the a-posteriori probability of the sought variables, resulting in visually pleasing high-quality images. These models have mostly been suggested for real-valued images suffering from Gaussian noise. In this study, we generalize annealed Langevin Dynamics, a type of DDM, to tackle the fundamental challenges in optical imaging of complex-valued objects (and real images) affected by Poisson noise. We apply our algorithm to various optical scenarios, such as Fourier Ptychography, Phase Retrieval, and Poisson denoising. Our algorithm is evaluated on simulations and biological empirical data.

1. Introduction

Light is a complex-valued field. During imaging, both the phase and intensity of the field change by the captured objects and the propagation within the optical system. Typically, the examined objects are assumed to have only real values. Recovering the phase of complex-valued objects is often crucial [1, 2]. However, image sensors can only measure real-valued non-negative intensities. This limitation gives rise to inverse problems such as phase-retrieval [3, 4, 5, 6, 7] ptychography [8, 9], Fourier ptychography [10, 11, 12, 13, 14], coded diffraction imaging [15] and holography. A common characteristic of all these problems is the *nonlinear* relationship between the measurements and the complex-valued unknowns.

Furthermore, the particle nature of light leads to Poissonian-distributed photon noise, in contrast to the commonly assumed Gaussian noise [16, 17, 18, 19, 20, 21, 22].




Figure 1: Poisson denoising. Results of BM3D + I + VST [23] and our algorithm on real data [18], obtained via two-photon microscopy. [Bottom] A zoom-in. Our recovery has high-perceptual quality (sharp) and low distortion.

Poisson noise complicates the recovery of phase-related problems as well as linear problems involving real-valued images. Therefore, addressing Poisson noise during restoration is of great importance, especially in scientific tasks, such as cell imaging (see Fig. 1).

Over the years, restoration algorithms have varied based on their objective. These algorithms, whether classical (NLM [24], KSVD [25], WNNM [26], and BM3D [27]) or based on deep-learning (TNRD [28], MALA [29], DnCNN [30], Noise2Void [31] and NLRN [27]), have been typically designed to minimize the mean squared error (MSE) between true and estimated images. These are MMSE estimators. Alternatively, other methods pursued maximum a-posteriori probability¹ (MAP) [25, 32].

The distortion-perception trade-off [33, 34] states that in image recovery, the resulting distortion (MSE score) is reciprocal to the resulting perceptual quality. Hence, there is a growing interest in recovery algorithms that yield high perceptual quality, rather than minimizing the MSE. Some prior works along this line use generative adversarial networks (GANs) for image recovery. GANs can be turned into samplers from the posterior distribution. However,

¹Interestingly, also works on these methods tend to report performance using MSE.

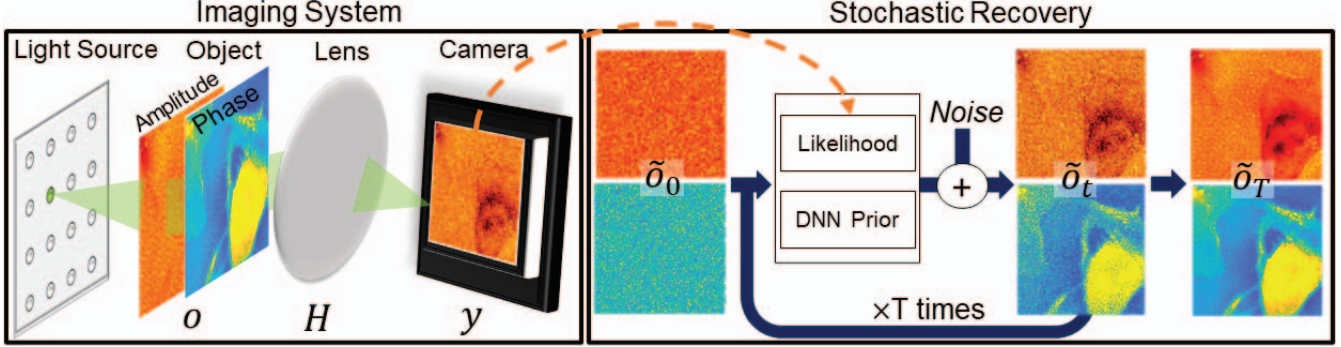


Figure 2: Imaging of a complex-valued objects. The imaging setup captures only real valued noisy intensities. Given the noisy measurements, our method stochastically recovers the objects amplitude and phase.

training difficulties and lack of solid mathematical guarantees hindered these techniques.

Newcomer generative models are deep diffusion models (DDMs). DDMs are a highly-performing class of iterative probabilistic generative models. DDMs have shown success in solving various inverse problems [35, 36, 37, 38, 39, 40, 41]. However, a DDM requires a likelihood term, which is often analytically intractable due to the imaging noise statistics. As a result, restoration algorithms that utilize DDMs have primarily concentrated on linear inverse problems affected by stationary Gaussian noise [37, 40, 39]. An exception is [42] which proposes DDMs for *nonlinear* inverse problems. It has two limitations: (i) Being restricted to real-valued images, and (ii) Requiring costly back-propagation through a DNN in each iteration. This lengthens computation time during recovery.

One type of DDM is *annealed Langevin Dynamics*. In this study, we generalize this DDM type, to tackle fundamental challenges in optical imaging: recovering complex-valued objects (and real images) affected by Poisson noise (see Fig. 2). To handle the untraceable likelihood term, we analytically develop a novel relaxed model, derived from the optical setting and the Poisson noise distribution. We show that the gradient of the prior can be expressed via an expectation, which is conveniently approximated by a deep neural network (DNN). We apply our algorithm² to various optical scenarios, such as Fourier ptychography, phase retrieval, and Poisson denoising. Our method samples from the posterior distribution, thus exposing uncertainties in the recovery, via the spread of the results. The uncertainty is valuable when imaging true objects. We show results on simulations and real empirical biological data .

2. Theoretical Background

2.1. Noise

Empirical light measurements are random, mainly due to the integer nature of photons and electric charges. This

²An initial description of this work appears in [43].

randomness is often regarded as photon noise. Dark noise, quantization noise, and photon noise are the dominant noise sources in common cameras [44]. Denote incorporation of noise into the expected noiseless signal \mathbb{I} by the operator \mathcal{N} . The noisy measurement is $y = \mathcal{N}(\mathbb{I})$.

Photon noise is modeled by a Poisson process [44]. A pixel is characterized by its full well capacity (FWC), that is, the maximal amount of charge that can be accumulated in the pixel before saturating. Let $x^e \in [0, \text{FWC}]$ be the expected (noiseless) value of a measured signal (photoelectrons). The corresponding noisy measurement is

$$y^e \sim \text{Poisson}(x^e). \quad (1)$$

Photon noise is fundamental. Its crucial nature of signal-dependency (1) is valid both for photon-hungry scenes, and even more so for well-lit scenes, as the noise variance is higher in brighter pixels.

Modeling the nature of noise, being signal-dependent, is important when solving scientific inverse problems in optical imaging. However, data analysis under Poisson noise is difficult to deal with in a closed form. Therefore, in many inverse problems, Poisson noise is approximated by a Gaussian distribution, $n^e \sim N(0, x^e)$. The variance still depends on the noiseless signal. Define normalized variables

$$x = \frac{x^e}{\text{FWC}} \in [0, 1], \quad y = \frac{y^e}{\text{FWC}}, \quad \sigma_0 = \frac{1}{\sqrt{\text{FWC}}}. \quad (2)$$

The Gaussian approximation of a measurement is

$$n^{\text{meas}} \sim N(0, \sigma_0^2 x). \quad (3)$$

The normalized approximation of Eq. (1) is

$$y \approx x + n^{\text{meas}}. \quad (4)$$

In Eqs. (1,4), the noise variance is linear in x , and the signal-to-noise-ratio (SNR) is the same. Hence, in this work, the parameter σ_0 controls the measurement noise level.

2.2. Imaging of Complex-valued Objects

Let \mathbf{r} represent the 2D spatial coordinate of an object. Let $a(\mathbf{r})$, $\phi(\mathbf{r})$ be its amplitude and phase, respectively, and $j = \sqrt{-1}$. An optically-thin object has transmittance $o(\mathbf{r}) = a(\mathbf{r}) \exp[j\phi(\mathbf{r})]$. Cameras only measure intensity. Therefore, optical imaging systems can often be formulated by an imaging model, which is nonlinear in o . Let ρ be the object irradiance. Let H be a linear transformation, which depends on the optical system. A noisy measurement y is modeled by

$$y = \mathcal{N}(\rho|H\{o(\mathbf{r})\}|^2). \quad (5)$$

For example, in the phase retrieval problem [45], H is equivalent to the Fourier transform \mathcal{F} .

Fourier ptychography uses a set of M linear transformations $H = \{H_m\}_{m=1}^M$. This method enables to expand the diffraction limit set by the numerical aperture of the optical system, thus increasing the optical resolution. Let \mathcal{F}^{-1} be the inverse Fourier transform. In the spatial frequency domain \mathbf{k} , the pupil function $P(\mathbf{k} - \mathbf{k}_m)$ is a 2D bandpass filter, centered at spatial frequency \mathbf{k}_m with a spatial bandwidth set by the system numerical aperture. Overall,

$$H_m\{o(\mathbf{r})\} = \mathcal{F}^{-1}\{P(\mathbf{k} - \mathbf{k}_m)\mathcal{F}\{o(\mathbf{r})\}\}. \quad (6)$$

Discretize $o(\mathbf{r})$ to q elements and the measurements $y(\mathbf{r})$ to $q^{\text{meas}} = Mq$ sampled pixels, creating vectors $\mathbf{o} \in \mathbb{C}^q$ and $\mathbf{y} \in \mathbb{R}^{q^{\text{meas}}}$. The discrete version of the linear transformation H is $\mathbf{H} \in \mathbb{C}^{q^{\text{meas}} \times q}$. The discrete versions of the Fourier transform, its inverse, and the pupil function are denoted by \mathbf{F} , \mathbf{F}^{-1} and \mathbf{P}_m , respectively. Then, $\mathbf{H}_m\mathbf{o} = \sqrt{\rho}\mathbf{F}^{-1}[\mathbf{P}_m \odot [\mathbf{F}\mathbf{o}]]$, where \odot is an element-wise multiplication. Assume that the dominant noise in Eq. (5) is photon noise. Following Sec. (2.1), during data analysis, we approximate this noise as Gaussian $\mathbf{n}^{\text{meas}} \sim N(0, \sigma_0^2|\mathbf{H}\mathbf{o}|^2)$. Then, Eq. (5) takes the form

$$\mathbf{y} = |\mathbf{H}\mathbf{o}|^2 + \mathbf{n}^{\text{meas}}. \quad (7)$$

2.3. Langevin Dynamics

Denosing is ill-posed: a noisy input may have multiple possible solutions. To address this, *Langevin Dynamics* can be used for denosing, converging stochastically to a single sharp possible solution. Here we detail it shortly.

Let $t \in [1, \dots, T]$ be an iteration index and $\alpha_t > 0$ be an appropriately chosen small function of t , having $\alpha_{t \rightarrow T} = 0$. To maximize the probability distribution $p(\mathbf{x}|\mathbf{y})$, the following iterations can be apply

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \alpha_t \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t|\mathbf{y}). \quad (8)$$

Eq. (8) may converge to an undesired local maximum. To address that, Langevin Dynamics [46] adds to Eq. (8) *synthetic* noise $\mathbf{n}_t^{\text{Langevin}} \sim N(0, \mathbf{I})$ at iteration t ,

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \alpha_t \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t|\mathbf{y}) + \sqrt{2\alpha_t} \mathbf{n}_t^{\text{Langevin}}. \quad (9)$$

Langevin Dynamics is a type of gradient descent. It relies on stochastic steps to converge to a probable solution.

In general, when solving an inverse problem, neither $p(\mathbf{x}_t|\mathbf{y})$ nor $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t|\mathbf{y})$ are known. Still, to use Eq. (9), there is a need to assess $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t|\mathbf{y})$. To address that, Ref. [47] extends Eq. (9) into *annealed* Langevin Dynamics. Hence, iterations would *not* perform Eq. (9), and would not attempt to update a vector \mathbf{x}_t whose probability cannot be computed. Instead, iterations would update a vector, that we denote $\tilde{\mathbf{x}}_t$, for which we can analytically assess the gradient $\nabla_{\tilde{\mathbf{x}}_t} \log p(\tilde{\mathbf{x}}_t|\mathbf{y})$. We thus need to properly define $\tilde{\mathbf{x}}_t$. Analogously to Eq. (9), annealed Langevin Dynamics has this iterative rule

$$\tilde{\mathbf{x}}_{t+1} = \tilde{\mathbf{x}}_t + \alpha_t \nabla_{\tilde{\mathbf{x}}_t} \log p(\tilde{\mathbf{x}}_t|\mathbf{y}) + \sqrt{2\alpha_t} \mathbf{n}_t^{\text{Langevin}}. \quad (10)$$

Any $\tilde{\mathbf{x}}_t$ differs from \mathbf{x} by an error $\mathbf{e}_t^{\text{annea}}$,

$$\tilde{\mathbf{x}}_t = \mathbf{x} + \mathbf{e}_t^{\text{annea}}. \quad (11)$$

We stress that this induced error $\mathbf{e}_t^{\text{annea}}$ is only a conceptual mathematical maneuver, and is not actually added during the iterative process. It is important that the probability distribution of $\mathbf{e}_t^{\text{annea}}$ would be known, despite not knowing the actual value of $\mathbf{e}_t^{\text{annea}}$. Moreover, the error is *designed* such that $\mathbf{e}_{t \rightarrow T}^{\text{annea}} \approx 0$.

Let us proceed with the description of annealed Langevin Dynamics Eq. (10). Using Bayes rule,

$$\nabla_{\tilde{\mathbf{x}}_t} \log p(\tilde{\mathbf{x}}_t|\mathbf{y}) = \nabla_{\tilde{\mathbf{x}}_t} [\log p(\mathbf{y}|\tilde{\mathbf{x}}_t) + \log p(\tilde{\mathbf{x}}_t)]. \quad (12)$$

There is a need to derive tractable expressions to the right-hand side of Eq. (12). The term $\nabla_{\tilde{\mathbf{x}}_t} \log p(\mathbf{y}|\tilde{\mathbf{x}}_t)$ in Eq. (12) can sometimes be derived analytically using the known distribution of $\mathbf{e}_t^{\text{annea}}$, as we show in Sections 3.1 and 3.2. To obtain $\nabla_{\tilde{\mathbf{x}}_t} \log p(\tilde{\mathbf{x}}_t)$, Refs. [40, 48] follow [49]: Before imaging takes place, they pre-train a DNN denoiser to approximate $\nabla_{\tilde{\mathbf{x}}_t} \log p(\tilde{\mathbf{x}}_t)$.

3. Langevin Dynamics for Optical Imaging

As described in Sections 2.1 and 2.2, real-world measurements have: (a) Poisson distributed noise; (b) Nonlinear transforms over complex-valued objects. There is a need to extend annealed Langevin Dynamics for both (a) and (b).

As discussed in Section 2.3, in general inverse problems $\nabla_{\mathbf{x}} \log p(\mathbf{x}|\mathbf{y})$ is unknown. Therefore, in annealed Langevin Dynamics, $\nabla_{\tilde{\mathbf{x}}_t} \log p(\tilde{\mathbf{x}}_t|\mathbf{y})$ is exploited, while using the connection

$$\lim_{t \rightarrow T} \nabla_{\tilde{\mathbf{x}}_t} \log p(\tilde{\mathbf{x}}_t|\mathbf{y}) \rightarrow \nabla_{\mathbf{x}} \log p(\mathbf{x}|\mathbf{y}). \quad (13)$$

In a similar fashion, to analytically handle signal dependency of noise, we model the noise as if the signal dependency is proportional to $\tilde{\mathbf{x}}_t$ rather than \mathbf{x} . This allows us

Algorithm 1: Langevin Dynamics for complex-valued objects. For real-valued images, $\tilde{\mathbf{o}}$ is replaced by $\tilde{\mathbf{x}}$.

```

Initialize  $\tilde{\mathbf{o}}_0$ , set a step size  $\epsilon > 0$ 
for  $t \leftarrow 1$  to  $T$  do
   $\alpha_t \leftarrow \epsilon(\sigma_t^2/\sigma_T^2)$ 
  Draw  $\mathbf{n}_t^{\text{Langevin}}$  from a Gaussian distribution
  if  $\tilde{\mathbf{o}}$  is complex-valued then
    From Eq. (39) compute gradient
     $\Delta_t \leftarrow \nabla_{\tilde{\mathbf{o}}_{t-1}} \log p(\mathbf{y}|\tilde{\mathbf{o}}_{t-1}) + h_{\hat{\Theta}}^{\text{complex}}(\tilde{\mathbf{o}}_{t-1}|\sigma_t)$ 
  else
    From Eq. (26) compute gradient
     $\Delta_t \leftarrow \nabla_{\tilde{\mathbf{o}}_{t-1}} \log p(\mathbf{y}|\tilde{\mathbf{o}}_{t-1}) + h_{\hat{\Theta}}^{\text{Poisson}}(\tilde{\mathbf{o}}_{t-1}|\sigma_t)$ 
  end
   $\tilde{\mathbf{o}}_t \leftarrow \tilde{\mathbf{o}}_{t-1} + \alpha_t \Delta_t + \sqrt{2\alpha_t} \mathbf{n}_t^{\text{Langevin}}$ 
end
 $\hat{\mathbf{o}} \leftarrow \tilde{\mathbf{o}}_T$ 

```

to analytically develop $\nabla_{\tilde{\mathbf{x}}_t} \log p(\mathbf{y}|\tilde{\mathbf{x}}_t)$ as described later on in Sections 3.1 and 3.3. From Eq. (11), notice that our approximation becomes accurate as $\tilde{\mathbf{x}}_t \rightarrow \mathbf{x}$. The iterative algorithm is described in Algorithm 1.

3.1. Denoising Poissonian Image Intensities

We now attempt to recover an image \mathbf{x} , given its noisy version \mathbf{y} . As discussed in Section 2.1, \mathbf{y} contains several types of noise. Indeed data is *generated* using a rich noise model: true Poissonian, compounded by graylevel quantization. For *data analysis*, however, we act as if \mathbf{y} is related by Eq. (4). The parameter σ_0 is known from the camera properties by Eq. (2). Here, all vector operations are *element-wise*. Analogously to [40, 48], we set a sequence of noise levels $\sigma_0 > \sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_T > \sigma_{T+1} = 0$. Let $\mathbf{e}_t^{\text{part}}$ be a statistically independent synthetic *partial* dummy error term, distributed as

$$\mathbf{e}_t^{\text{part}} \sim N[0, (\sigma_t^2 - \sigma_{t+1}^2)\mathbf{I}]. \quad (14)$$

A sum of zero-mean Gaussian variables is a Gaussian, whose variance is the sum of the variances. So,

$$\sum_{\tau=t_1}^{t_2} \mathbf{e}_\tau^{\text{part}} \sim N[0, (\sigma_{t_1}^2 - \sigma_{t_2}^2)\mathbf{I}]. \quad (15)$$

Using Eqs. (3) and (15), we design $\mathbf{e}_t^{\text{part}}$ such that

$$\mathbf{n}^{\text{meas}} = \sqrt{\mathbf{x}} \sum_{\tau=0}^T \mathbf{e}_\tau^{\text{part}}. \quad (16)$$

Moreover, recall that $\mathbf{e}_t^{\text{annea}}$ is a dummy error term. We set $\mathbf{e}_t^{\text{annea}}$ in this task as

$$\mathbf{e}_t^{\text{annea}} = \sum_{\tau=t}^T \mathbf{e}_\tau^{\text{part}}. \quad (17)$$

We do *not* set $\mathbf{e}_t^{\text{annea}}$ as Refs. [40, 48] because the noise in Refs. [40, 48] has stationary Gaussian distribution and here it is Poissonian (signal-dependent). We generalize Eq. (11) to Poisson noise, by defining an annealed variable

$$\tilde{\mathbf{x}}_t = \mathbf{x} + \sqrt{\mathbf{x}} \mathbf{e}_t^{\text{annea}} = \mathbf{x} + \sqrt{\mathbf{x}} \sum_{\tau=t}^T \mathbf{e}_\tau^{\text{part}}. \quad (18)$$

Eq. (18) allows us to derive $\nabla_{\tilde{\mathbf{x}}_t} \log p(\mathbf{y}|\tilde{\mathbf{x}}_t)$ analytically for Poisson denoising. Inserting Eqs. (16) and (18) in Eq. (4),

$$\mathbf{y} = \tilde{\mathbf{x}}_t - \sqrt{\mathbf{x}} \sum_{\tau=t}^T \mathbf{e}_\tau^{\text{part}} + \sqrt{\mathbf{x}} \sum_{\tau=0}^T \mathbf{e}_\tau^{\text{part}} = \tilde{\mathbf{x}}_t + \sqrt{\mathbf{x}} \sum_{\tau=0}^{t-1} \mathbf{e}_\tau^{\text{part}}. \quad (19)$$

Notice that $\tilde{\mathbf{x}}_t$ is statistically independent of the additional dummy error $\sum_{\tau=0}^{t-1} \mathbf{e}_\tau^{\text{part}}$. We apply the following steps. From Eq. (19),

$$p(\mathbf{y}|\tilde{\mathbf{x}}_t) = p(\mathbf{y} - \tilde{\mathbf{x}}_t|\tilde{\mathbf{x}}_t) = p\left(\left[\sqrt{\mathbf{x}} \sum_{\tau=0}^{t-1} \mathbf{e}_\tau^{\text{part}}\right]|\tilde{\mathbf{x}}_t\right). \quad (20)$$

Recall that \mathbf{x} in Eq. (20) is unknown. Therefore, to set an explicit analytical expression of Eq. (20), we approximate the remaining measurement noise as if the signal dependency is proportional to $\tilde{\mathbf{x}}_t$ rather³ than \mathbf{x} . Overall, Eq. (20) takes the form of

$$p(\mathbf{y}|\tilde{\mathbf{x}}_t) \approx p\left(\left[\sqrt{\tilde{\mathbf{x}}_t} \sum_{\tau=0}^{t-1} \mathbf{e}_\tau^{\text{part}}\right]|\tilde{\mathbf{x}}_t\right). \quad (21)$$

Finally, the expression $\sqrt{\tilde{\mathbf{x}}_t} \sum_{\tau=0}^{t-1} \mathbf{e}_\tau^{\text{part}}$ is both known and statistically tractable. This allows us to find an explicit expression of the probability distribution $p(\mathbf{y}|\tilde{\mathbf{x}}_t)$

$$p(\mathbf{y}|\tilde{\mathbf{x}}_t) \approx \frac{1}{\sqrt{2\pi(\sigma_0^2 - \sigma_t^2)\tilde{\mathbf{x}}_t}} \exp\left[-\frac{(\mathbf{y} - \tilde{\mathbf{x}}_t)^2}{2(\sigma_0^2 - \sigma_t^2)\tilde{\mathbf{x}}_t}\right]. \quad (22)$$

Computing the gradient of the logarithm of Eq. (22) yields

$$\nabla_{\tilde{\mathbf{x}}_t} \log p(\mathbf{y}|\tilde{\mathbf{x}}_t) = \frac{1}{2(\sigma_0^2 - \sigma_t^2)} \left(\frac{\mathbf{y}^2}{\tilde{\mathbf{x}}_t^2} - 1\right) - \frac{1}{2\tilde{\mathbf{x}}_t}. \quad (23)$$

In analogy to Section 2.3, we compute $\nabla_{\tilde{\mathbf{x}}_t} \log p(\tilde{\mathbf{x}}_t)$. The derivation is detailed in the supplementary material. Here we bring the result: $\nabla_{\tilde{\mathbf{x}}_t} \log p(\tilde{\mathbf{x}}_t)$ is given by,

$$\nabla_{\tilde{\mathbf{x}}_t} \log p(\tilde{\mathbf{x}}_t) = \mathbb{E}\left[\frac{\mathbf{x} - \tilde{\mathbf{x}}_t}{\sigma^2 \mathbf{x}}|\tilde{\mathbf{x}}_t\right]. \quad (24)$$

The expectation in Eq. (24) cannot be calculated analytically. The reason is that the probability distribution of \mathbf{x} is

³From Eq. (18), this approximation becomes accurate as $\tilde{\mathbf{x}}_t \rightarrow \mathbf{x}$.

not accessible. So, before imaging takes place, we pre-train a DNN $h_{\Theta}^{\text{Poisson}}$ having learnable parameters Θ to estimate $\nabla_{\tilde{\mathbf{x}}_t} \log p(\tilde{\mathbf{x}}_t)$, using synthetic pairs of corresponding clean and noisy images $(\mathbf{x}, \mathbf{x}')$,

$$\hat{\Theta} = \arg \min_{\Theta} \mathbb{E} \left[\left\| \frac{\mathbf{x} - \mathbf{x}'}{\sigma^2 \mathbf{x}} - h_{\Theta}^{\text{Poisson}}(\mathbf{x}'|\sigma) \right\|_2^2 \right]. \quad (25)$$

The expectation \mathbb{E} is approximated by an empirical average of all pairs in the training set of corresponding clean and noisy images $(\mathbf{x}, \mathbf{x}')$. Overall, Eqs. (23,24,25) yeild

$$\begin{aligned} \nabla_{\tilde{\mathbf{x}}_t} \log p(\tilde{\mathbf{x}}_t|\mathbf{y}) &= \nabla_{\tilde{\mathbf{x}}_t} \log p(\mathbf{y}|\tilde{\mathbf{x}}_t) + h_{\hat{\Theta}}^{\text{Poisson}}(\tilde{\mathbf{x}}_t|\sigma_t) = \\ &= \frac{1}{2(\sigma_0^2 - \sigma_t^2)} \left(\frac{\mathbf{y}^2}{\tilde{\mathbf{x}}_t^2} - 1 \right) - \frac{1}{2\tilde{\mathbf{x}}_t} + h_{\hat{\Theta}}^{\text{Poisson}}(\tilde{\mathbf{x}}_t|\sigma_t). \end{aligned} \quad (26)$$

3.2. Recovering Complex-Valued Objects

This section formulates annealed Langevin Dynamics for complex-valued, nonlinear imaging models. Let $\Re(\cdot)$, $\Im(\cdot)$ be real and imaginary arguments of a complex-valued number, respectively. Analogously to Eq. (9), for Langevin Dynamics, $\nabla_{\mathbf{o}_t} \log p(\mathbf{o}_t|\mathbf{y})$ is needed. However, $p(\mathbf{o}_t|\mathbf{y})$ is not known, thus analogously to Eq. (10), annealed Langevin Dynamics for a complex-valued object is

$$\tilde{\mathbf{o}}_{t+1} = \tilde{\mathbf{o}}_t + \alpha_t \nabla_{\tilde{\mathbf{o}}_t} \log p(\tilde{\mathbf{o}}_t|\mathbf{y}) + \sqrt{2\alpha_t} \mathbf{n}_t^{\text{Langevin}}. \quad (27)$$

Here, $\mathbf{n}^{\text{Langevin}}$ is a random complex-valued vector, where $\Re(\mathbf{n}^{\text{Langevin}})$, $\Im(\mathbf{n}^{\text{Langevin}}) \sim N(0, \mathbf{I})$ and independent. For simplicity, we present here the special case $\mathbf{H} = \mathbf{I}$ in Eq. (7). Thus, most derivations can be given in scalar formulation, that is, $y = |o|^2 + n^{\text{meas}}$. The generalization of the model to a general \mathbf{H} is detailed in the supplementary material. Analogously to Eq. (11), we set

$$\tilde{o}_t = o + e_t^{\text{annea}}. \quad (28)$$

Using Eq. (28) in Eq. (7) yields

$$y = |\tilde{o}_t|^2 - e_t^{\text{couple}} + |e_t^{\text{annea}}|^2 + n^{\text{meas}}, \quad (29)$$

where

$$|o_t|^2 = \Re(o_t)^2 + \Im(o_t)^2, \quad (30)$$

$$e_t^{\text{couple}} = 2[\Re(\tilde{o}_t)\Re(e_t^{\text{annea}}) + \Im(\tilde{o}_t)\Im(e_t^{\text{annea}})], \quad (31)$$

and $|e_t^{\text{annea}}|^2 = \Re(e_t^{\text{annea}})^2 + \Im(e_t^{\text{annea}})^2$. The non-linearity of Eq. (7) yields two terms in Eq. (29) that pose analytic challenges: a coupling term of e_t^{couple} , and a squared annealing error term $|e_t^{\text{annea}}|^2$. We handle them by introducing a *relaxed* model, creating annealed Langevin Dynamics for complex-valued objects, as

$$y = |\tilde{o}_t|^2 - e_t^{\text{couple}} + |e_t^{\text{annea}}|^2 + \tilde{n}_t^{\text{meas}} - \sigma_t^2 z_t. \quad (32)$$

We detail Eq. (32) in Section 3.3.

3.3. Relaxation of the Nonlinear Problem

Let $\{\sigma_t\}_{t=0}^T$ be a sequence of standard deviations, set to $\sigma_0 > \sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_T > \sigma_{T+1} = 0$. Let

$$\Re(e_t^{\text{part}}), \Im(e_t^{\text{part}}) \sim N\left[0, \frac{1}{4}(\sigma_t^2 - \sigma_{t+1}^2)\right] \quad (33)$$

be statistically independent Gaussian variables. As described in Section 2.2 and in Eq. (3), for $\mathbf{H} = \mathbf{I}$, $n^{\text{meas}} \sim N(0, \sigma_0^2 |o|^2)$. Using $|o|^2$ from Eq. (30), $n^{\text{meas}} \sim N[0, \sigma_0^2 \Re(o)^2 + \sigma_0^2 \Im(o)^2]$. In analogy to Eq. (16), we use Eq. (33) and design dummy variables $\Re(e_t^{\text{part}})$, $\Im(e_t^{\text{part}})$ such that,

$$n^{\text{meas}} = 2\Re(o) \sum_{\tau=0}^T \Re(e_{\tau}^{\text{part}}) + 2\Im(o) \sum_{\tau=0}^T \Im(e_{\tau}^{\text{part}}). \quad (34)$$

Notice that Eq. (34) now has a form similar to the coupling term e_t^{couple} in Eq. (31). To make computations tractable, we assume $o \approx \tilde{o}_t$ in Eq. (34). We do not use \tilde{o}_t as a de-noised version yet. An approximate⁴ measurement noise is

$$\tilde{n}_t^{\text{meas}} = 2\Re(\tilde{o}_t) \sum_{\tau=0}^T \Re(e_{\tau}^{\text{part}}) + 2\Im(\tilde{o}_t) \sum_{i=0}^T \Im(e_{\tau}^{\text{part}}). \quad (35)$$

Using Eqs. (31) and (35), define

$$\begin{aligned} e_t^{\text{remain}}(\tilde{o}_t) &= \tilde{n}_t^{\text{meas}} - e_t^{\text{couple}} \\ &= 2\Re(\tilde{o}_t) \sum_{\tau=0}^{t-1} \Re(e_{\tau}^{\text{part}}) + 2\Im(\tilde{o}_t) \sum_{\tau=0}^{t-1} \Im(e_{\tau}^{\text{part}}). \end{aligned} \quad (36)$$

The term z_t in Eq. (32) is designed to address the squared annealing error term $|e_t^{\text{annea}}|^2$. We detail z_t in the supplementary material, here we give the final result,

$$\begin{aligned} z_t^{\text{remain}} &= |e_t^{\text{annea}}|^2 - \sigma_t^2 z_t \\ &= \left[\sum_{\tau=0}^{t-1} \Re(e_{\tau}^{\text{part}}) \right]^2 + \left[\sum_{\tau=0}^{t-1} \Im(e_{\tau}^{\text{part}}) \right]^2. \end{aligned} \quad (37)$$

Inserting Eqs. (36) and (37) into the relaxed model in Eq. (32), yields

$$y = |\tilde{o}_t|^2 + e_t^{\text{remain}}(\tilde{o}_t) + z_t^{\text{remain}}. \quad (38)$$

Using Eq. (38), $\nabla_{\tilde{o}_t} \log p(y|\tilde{o}_t)$ can be derived analytically. Here we give the results: the derivation is detailed in the supplementary material. Let I_0, I_1 be the modified Bessel functions of the first kind [51]. Then,

$$\nabla_{\tilde{o}_t} \log p(y|\tilde{o}_t) = \frac{\tilde{o}_t}{2(\sigma_0^2 - \sigma_t^2)} \left[\mathcal{I}(\tilde{o}_t, y) \frac{\sqrt{y}}{|\tilde{o}_t|} - 1 \right], \quad (39)$$

⁴From Eq. (28), the approximation becomes accurate as $e_t^{\text{annea}} \rightarrow 0$, where $\tilde{o}_t \rightarrow o$ and $\tilde{n}_t^{\text{meas}} \rightarrow n^{\text{meas}}$.

CelebA 64x64	$\sigma_0 = 0.025$			$\sigma_0 = 0.05$			$\sigma_0 = 0.1$			$\sigma_0 = 0.2$		
	PSNR \uparrow	SSIM% \uparrow	FID \downarrow	PSNR \uparrow	SSIM% \uparrow	FID \downarrow	PSNR \uparrow	SSIM% \uparrow	FID \downarrow	PSNR \uparrow	SSIM% \uparrow	FID \downarrow
I + VST + BM3D [23]	39.1 \pm 1.0	99 \pm 0.1	2.80	35.8 \pm 1.1	98 \pm 0.2	8.90	30.9 \pm 1.2	96 \pm 1.0	25.4	27.2 \pm 1.1	93 \pm 2.0	63.2
PnP RED [50]	39.8 \pm 0.9	99 \pm 0.1	4.00	36.9 \pm 1.0	98 \pm 0.3	7.50	31.8 \pm 0.9	96 \pm 1.0	12.4	27.8 \pm 1.2	93 \pm 2.0	28.4
Bahjat et al. [48]	34.2 \pm 1.1	97 \pm 0.9	2.10	34.1 \pm 0.7	97 \pm 1.0	6.70	30.6 \pm 0.9	95 \pm 2.0	11.1	27.1 \pm 1.2	91 \pm 4.0	15.2
MMSE	42.5 \pm 1.0	99 \pm 0.1	<u>1.32</u>	38.4 \pm 1.1	99 \pm 0.2	<u>3.60</u>	34.5 \pm 1.1	98 \pm 0.5	<u>7.50</u>	30.9 \pm 1.2	96 \pm 1.0	11.9
Ours	41.5 \pm 1.1	99 \pm 0.0	1.25	37.1 \pm 1.0	99 \pm 0.4	2.70	32.8 \pm 1.1	98 \pm 0.5	6.10	28.6 \pm 1.3	94 \pm 1.0	<u>13.0</u>
FFHQ 256x256												
DPS [42]	36.8 \pm 4.7	99 \pm 0.1	<u>2.00</u>	33.8 \pm 3.8	98 \pm 0.4	<u>5.47</u>	31.5 \pm 2.7	98 \pm 2.0	<u>12.3</u>	28.6 \pm 2.0	93 \pm 2.0	23.8
Ours	41.9 \pm 1.0	99 \pm 0.0	0.72	37.9 \pm 1.3	99 \pm 0.1	2.50	33.5 \pm 1.1	96 \pm 1.0	10.0	29.0 \pm 1.0	93 \pm 1.0	<u>33.6</u>

Table 1: Results on 1000 images of CelebA and FFHQ datasets. For $\sigma_0 \leq 0.1$, our algorithm has better FID and similar SSIM with the MMSE denoiser and DPS [42]. The MMSE denoiser inherently achieves superior PSNR.

where $\mathcal{I}(\tilde{o}_t, y) = I_1\left(\frac{|\tilde{o}_t|\sqrt{y}}{\sigma_0^2 - \sigma_t^2}\right) / I_0\left(\frac{|\tilde{o}_t|\sqrt{y}}{\sigma_0^2 - \sigma_t^2}\right)$. The gradient $\nabla_{\tilde{o}_t} \log p(\tilde{o}_t)$ is given by

$$\nabla_{\tilde{o}_t} \log p(\tilde{o}_t) = \mathbb{E} \left[\frac{\mathbf{o} - \tilde{o}_t}{\sigma^2} \middle| \tilde{o}_t \right]. \quad (40)$$

See the supplemental material for the proof of Eq. (40).

Hence again, we pre-train⁵ a DNN $h_{\Theta}^{\text{complex}}$ having learnable parameters Θ to estimate $\nabla_{\tilde{o}_t} \log p(\tilde{o}_t)$

$$\hat{\Theta} = \arg \min_{\Theta} \mathbb{E} \left[\left\| \frac{\mathbf{o} - \mathbf{o}'}{\sigma^2} - h_{\Theta}^{\text{complex}}(\mathbf{o}' | \sigma) \right\|_2^2 \right]. \quad (41)$$

Overall, the complex-valued derivative is

$$\begin{aligned} \nabla_{\tilde{o}_t} \log p(\tilde{o}_t | \mathbf{y}) &= \nabla_{\tilde{o}_t} \log p(\mathbf{y} | \tilde{o}_t) + h_{\hat{\Theta}}^{\text{complex}}(\tilde{o}_t | \sigma_t) = \\ &= \frac{\tilde{o}_t}{2(\sigma_0^2 - \sigma_t^2)} \odot \left[\mathcal{I}(\tilde{o}_t, \mathbf{y}) \odot \frac{\sqrt{y}}{|\tilde{o}_t|} - 1 \right] + h_{\hat{\Theta}}^{\text{complex}}(\tilde{o}_t | \sigma_t). \end{aligned} \quad (42)$$

4. Experiments

To evaluate recoveries we use the FID, SSIM and PSNR⁶ criteria. In all experiments we used the NCSNv2 architecture [49] as the DNN. The bottleneck of the proposed algorithm is iteratively running this DNN $T \sim 1000$ times, where time complexity increases with image size and linearly with T . We apply our algorithm on both real biological data [18, 11] and on synthetic datasets, CelebA (64×64) and FFHQ (256×256). We stress that during testing on synthetic scenes, we synthesized data using **actual Poisson noise** (Eq. 1). Moreover, 8-bit quantization was applied on the simulated electron count.

4.1. Denoising Poissonian Image Intensities

We compare our algorithm with [23, 42, 48, 50] and an MMSE solution. Fig. 3 and Table 1 present Poisson denois-

⁵The network input is a two-channel tensor of $\Re(\mathbf{o}')$, $\Im(\mathbf{o}')$.

⁶For complex-valued objects, the amplitude is evaluated by standard PSNR metric. The recovered object's phase $\angle \mathbf{o}$ has an inherent ambiguity. We evaluate phase recovery by $\text{PSNR}^{\text{phase}}(\angle \hat{\mathbf{o}}) = 10 \log_{10} \left[(2\pi)^2 [\text{MSE}^{\text{phase}}(\angle \hat{\mathbf{o}})]^{-1} \right]$, where $\text{MSE}^{\text{phase}}(\angle \hat{\mathbf{o}}) = \arg \min_{\theta} \|\angle \mathbf{o} - (\angle \hat{\mathbf{o}} + \theta)\|_2^2 / q$.

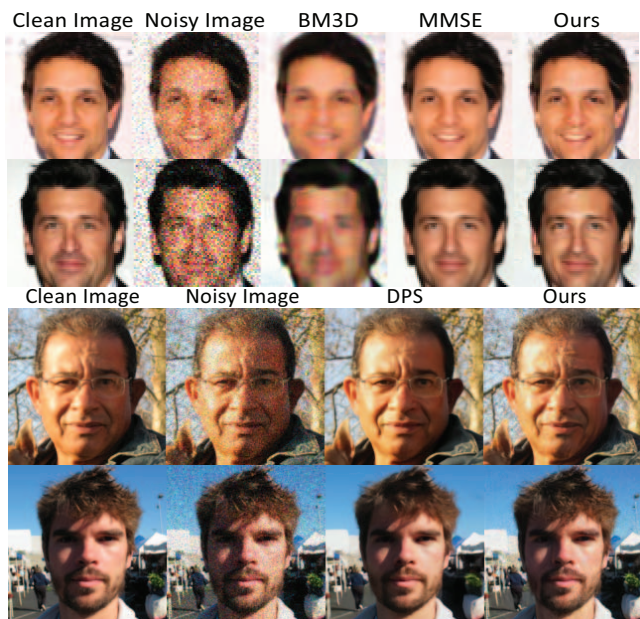


Figure 3: Poisson denoising results. From top to bottom: CelebA $\sigma_0 = 0.1$, $\sigma_0 = 0.2$. FFHQ $\sigma_0 = 0.2$.

ing results. Across all tests, our method outperforms [23], a non-learning-based state-of-the-art method for Poisson denoising. Furthermore, we outperform Ref. [48], which utilizes annealed Langevin Dynamics for *Gaussian* noise. This indicates the benefit of addressing the natural property of imaging noise. PnP RED [50] is a method to solve general inverse problems. It uses Gaussian denoiser as its prior⁷. Our method surpasses it, emphasizing the significance of employing a customized denoiser for Poisson denoising.

Typically, the SNR in standard optical systems satisfies $\sigma_0 \leq 0.1$. In that regime, we have better FID results and similar SSIM results than the MMSE denoiser and DPS [42]. MMSE denoisers are designed to minimize the MSE criterion, that is, maximizing the PSNR. Thus, unsurprisingly, an MMSE denoiser yields a superior PSNR than

⁷The data term derives from a Poissonian model.

$\sigma_0 =$	2/255	3/255	4/255
HIO [52]	22.2	19.9	17.9
BM3D-prGAMP [54]	24.5	23.2	21.0
BM3D-ADMM [53]	27.5	24.5	21.8
DnCNN-ADMM [53]	29.3	26.0	22.0
prDeep [53]	28.4	28.5	26.4
Ours	33.1	32.6	28.9

Table 2: Phase retrieval PSNR results for the “natural” images dataset in [53], using different noise levels σ_0 .

other algorithms⁸, however, with blur outputs (see Fig. 3).

4.1.1 Poisson Denoising of Real Data

Fig. 1 shows our method on real data [18]. It contains noisy colored images of fixed BPAE cells, obtained using two-photon microscopy. The imaged sample was split into 20 non-overlapping patches, 19 for training, and the remaining for testing. In [18], each patch was repeatedly captured 50 times, as 50 noise realizations. The clean ground-truth image is based on temporal averaging of the 50 images. The DNN trains only on 19 images with synthetic Poisson noise, while testing on raw data involves real noise. The authors in Ref. [18] modeled the imaging system noise by two additive components with zero means: A signal-dependent Poisson noise, and a negligible⁹ signal-independent Gaussian noise. Therefore, we apply our denoising algorithm as the measured image \mathbf{y} can be described by Eq. (4). The PSNR and SSIM of the noisy raw data are 26 and 0.67, respectively, while ours are 33 and 0.92.

4.2. Phase Retrieval

As described in Section 2.2, phase retrieval can be modeled using $\mathbf{H} = \mathbf{F}$, that is, $\mathbf{y} = \mathcal{N}(|\mathbf{F}\mathbf{x}|^2)$. Here \mathcal{N} contains photon noise and quantization, the latter using 8 bits per pixel. Note that even if \mathbf{x} is a real-valued image, the measurement \mathbf{y} has no phase. This task is ill-posed. The Hybrid Input-Output (HIO) [52] algorithm is frequently used in phase retrieval. Let ψ be a random Gaussian vector. HIO initializes hundreds of possible candidate solutions $\{\mathbf{x}_l\}$ by $\mathbf{F}\mathbf{x}_l = \sqrt{\mathbf{y}} \exp(j2\pi\psi_l)$. Then, HIO iteratively uses negative feedback in Fourier space, in order to progressively force the candidate solutions to conform to the problem constraints. The candidate that holds best these constraints is the evaluated HIO solution.

We compare our method (Algorithm 1) for phase retrieval, vs. prior art. Here, we initialize our algorithm

⁸An alternative method averages outputs of DDMs, to approximate an MMSE result, followed by PSNR comparison [40].

⁹Let $\mathbf{x} \in [0, 1]$ be the true normalized image. Let n^G, n^P be Gaussian noise and Gaussian approximation for the Poisson noise, with corresponding standard deviations $\sigma_G = 5 \times 10^{-4}$ and $\sigma_P = \sqrt{x}(2.5 \times 10^{-2})$. Then, the measured image \mathbf{y} is approximated by $\mathbf{y} = \mathbf{x} + n^G + n^P$.

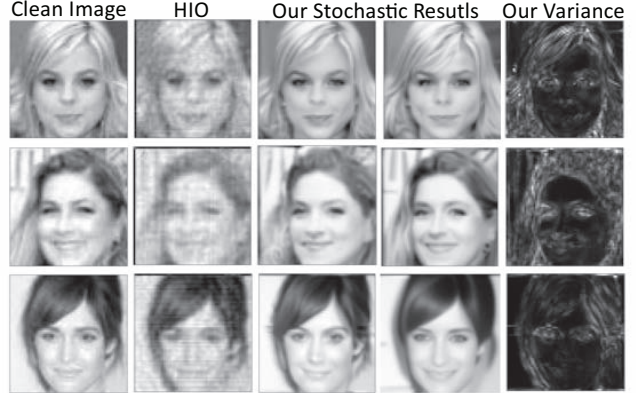


Figure 4: Phase retrieval results on CelebA images. A variance image is obtained by using 50 stochastic solutions.



Figure 5: Phase retrieval qualitative results for the dataset suggested in [53] ($\sigma_0 = 4/255$).

with an extremely noisy version of the HIO solution. We found that the noisy initialization gives us better robustness during iterations, but still stochastically converges to a variety of possible solutions (see Fig. 4). Additionally, we follow the experiment suggested in [53]. First, we trained $h_{\Theta}^{\text{Complex}}(\mathbf{o}'|\sigma)$ with overlapping patches drawn from 400 images in the Berkeley Segmentation Dataset [55]. Then, we evaluate our method on a test set of six “natural” images[53], three images are presented in Fig. 5. Table 2 compares the results. In the supplementary material, we present more results and further details.

4.3. Fourier Ptychography: Real Experiment

We further demonstrate our approach on the Fourier ptychography problem, described in Section 2.2. We use data of real-world publicly available complex-valued biological samples [56]. Using a Fourier ptychography imaging system, the amplitude and phase of a sample of human bone cells are recovered [11]. We randomly divided the recovered sample into non-spatially overlapping patches, each of

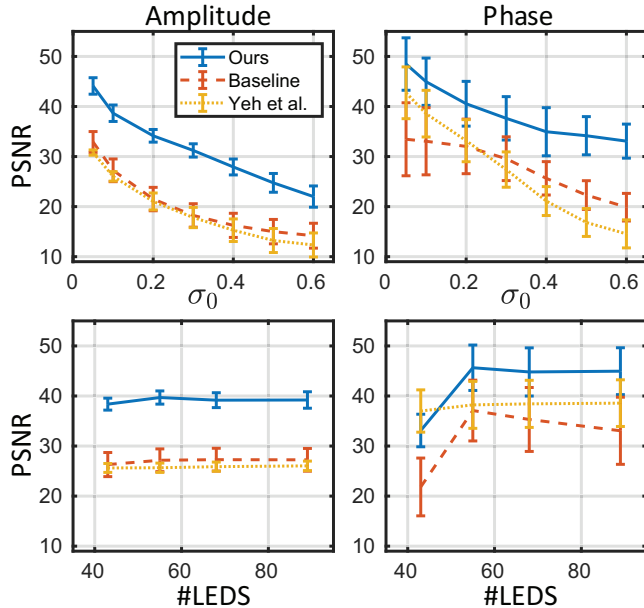


Figure 6: Fourier ptychography. Average PSNR for 10 complex-valued objects: [Top] Varying Poisson noise levels σ_0 and fixed $M = 89$. [Bottom] Varying number of LEDs M and a fixed noise level $\sigma_0 = 0.1$. [Baseline] Algorithm 1 excluding the prior term $h_{\hat{\Theta}}^{\text{complex}}$. The bars stand for PSNR standard deviations.

size 256×256 pixels. Then, the patches were split randomly into training and test sets. We train our denoiser over the training set, according to Eq. (41). For each patch \mathbf{o} in the test set, we simulate M noisy measured images, each corresponding to a single LED in the optical system, according to Eqs. (5) and (6). Figs. 6 and 7 show the results using Algorithm 1. Further details about this test and more results are in the supplementary material.

5. Discussion

We generalize annealed Langevin Dynamics to realistic optical imaging, which involves complex-valued objects and Poisson noise. Our method is generic and can be applied to a variety of imaging transformations \mathbf{H} . Furthermore, the pre-trained DNN in Eq. (41) is independent of the imaging system parameters σ_0 and \mathbf{H} of Eq. (7). Hence, the same DNN can be used for various optical systems.

Dealing with non-Gaussian noises, is essential in real-world imaging. Hence, noise statistics create new and interesting challenges, specifically in low-light conditions. There are more challenging forward models in imaging. These include imaging by coherent illumination (with speckles), partial-coherence, and models that are nonlinear with respect to the illumination (e.g. two-photon microscopy). Other computer vision problems are highly nonlinear in the unknowns, such as tomography scatter




Figure 7: Fourier ptychography. Our results of a complex-valued restoration, for $M = 89$ LEDs, $\sigma_0 = 0.1$ and 8 bit quantization. [Top] The amplitude and phase images' dynamic range $\in [0, 1]$ and $[-\pi, \pi]$ respectively. [Bottom] Scatter plots for the amplitude and phase values, the dashed red lines represent a perfect recovery.

fields [57, 58]. Therefore, motivated by this paper, we hope that more research expands annealed Langevin Dynamics to additional challenging models in real-world imaging.

6. Acknowledgments

We thank Laura Waller for referring us to data and code. Yoav Schechner is the Mark and Diane Seiden Chair in Science at the Technion. He is a Landau Fellow supported by the Taub Foundation. His work was conducted in the Ollendorff Minerva Center. Minvera is funded through the BMBF. This research was partially supported by the Israel Science Foundation (ISF) under Grants 335/18 and 2514/23, and by the Council For Higher Education - Planning & Budgeting Committee.

References

- [1] Dorian Chan, Srinivasa G Narasimhan, and Matthew O’Toole. Holocurtains: Programming light curtains via binary holography. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17886–17895, 2022. **1**
- [2] Kazunori Akiyama, Antxon Alberdi, Walter Alef, Keiichi Asada, Rebecca Azulay, Anne-Kathrin Baczko, David Ball, Mislav Baloković, John Barrett, Dan Bintley, et al. First M87 Event Horizon Telescope results. IV. imaging the central supermassive black hole. *The Astrophysical Journal Letters*, 875(1):L4, 2019. **1**
- [3] Rakib Hyder, Zikui Cai, and M Salman Asif. Solving phase retrieval with a learned reference. In *European Conference on Computer Vision (ECCV)*, pages 425–441. Springer, 2020. **1**
- [4] Duoduo Xue, Ziyang Zheng, Wenrui Dai, Chenglin Li, Junni Zou, and Hongkai Xiong. On the convergence of non-convex phase retrieval with denoising priors. *IEEE Transactions on Signal Processing*, 70:4424–4439, 2022. **1**
- [5] Feilong Zhang, Xianming Liu, Cheng Guo, Shiyi Lin, Junjun Jiang, and Xiangyang Ji. Physics-based iterative projection complex neural network for phase retrieval in lensless microscopy imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10523–10531, 2021. **1**
- [6] Angélique Drémeau, Antoine Liutkus, David Martina, Ori Katz, Christophe Schülke, Florent Krzakala, Sylvain Gigan, and Laurent Daudet. Reference-less measurement of the transmission matrix of a highly scattering material using a dmd and phase retrieval techniques. *Optics Express*, 23(9):11898–11911, 2015. **1**
- [7] Yicheng Wu, Manoj Kumar Sharma, and Ashok Veeraraghavan. Wish: wavefront imaging sensor with high resolution. *Light: Science & Applications*, 8(1):44, 2019. **1**
- [8] John Rodenburg and Andrew Maiden. Ptychography. *Springer Handbook of Microscopy*, pages 819–904, 2019. **1**
- [9] Eunju Cha, Chanseok Lee, Mooseok Jang, and Jong Chul Ye. Deepphasecut: Deep relaxation in phase for unsupervised fourier phase retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(12):9931–9943, 2021. **1**
- [10] Fahad Shamshad, Asif Hanif, Farwa Abbas, Muhammad Awais, and Ali Ahmed. Adaptive ptych: Leveraging image adaptive generative priors for subsampled fourier ptychography. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (WICCV)*, pages 0–0, 2019. **1**
- [11] Lei Tian, Xiao Li, Kannan Ramchandran, and Laura Waller. Multiplexed coded illumination for Fourier ptychography with an LED array microscope. *Biomedical Optics Express*, 5(7):2376–2389, 2014. **1, 6, 7**
- [12] Atreyee Saha, Salman S Khan, Sagar Sehrawat, Sanjana S Prabhu, Shanti Bhattacharya, and Kaushik Mitra. Lwgnnet-learned wirtinger gradients for fourier ptychographic phase retrieval. In *European Conference on Computer Vision (ECCV)*, pages 522–537. Springer, 2022. **1**
- [13] Guoan Zheng, Roarke Horstmeyer, and Changhui Yang. Wide-field, high-resolution fourier ptychographic microscopy. *Nature Photonics*, 7(9):739–745, 2013. **1**
- [14] Regina Eckert, Zachary F Phillips, and Laura Waller. Efficient illumination angle self-calibration in fourier ptychography. *Applied Optics*, 57(19):5434–5442, 2018. **1**
- [15] Benjamin Attal and Matthew O’Toole. Towards mixed-state coded diffraction imaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. **1**
- [16] Tao Huang, Songjiang Li, Xu Jia, Huchuan Lu, and Jianzhuang Liu. Neighbor2neighbor: Self-supervised denoising from single noisy images. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14781–14790, 2021. **1**
- [17] Prashanth G Kumar and Rajiv Ranjan Sahay. Low rank poisson denoising (LRPD): A low rank approach using split bregman algorithm for poisson noise removal from images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (WCVPR)*, pages 0–0, 2019. **1**
- [18] Yide Zhang, Yin hao Zhu, Evan Nichols, Qingfei Wang, Siyuan Zhang, Cody Smith, and Scott Howard. A Poisson-Gaussian denoising dataset with real fluorescence microscopy images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11710–11718, 2019. **1, 6, 7**
- [19] Jaeseok Byun, Sungmin Cha, and Taesup Moon. FBI-denoiser: Fast blind image denoiser for poisson-gaussian noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5768–5777, 2021. **1**
- [20] Ben Moseley, Valentin Bickel, Ignacio G López-Francos, and Loveneesh Rana. Extreme low-light environment-driven image denoising over permanently shadowed lunar regions with a physical noise model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6317–6327, 2021. **1**
- [21] Wesley Khademi, Sonia Rao, Clare Minnerath, Guy Hagen, and Jonathan Ventura. Self-supervised poisson-gaussian denoising. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2131–2139, 2021. **1**
- [22] Ben Mildenhall, Jonathan T Barron, Jiawen Chen, Dillon Sharlet, Ren Ng, and Robert Carroll. Burst denoising with kernel prediction networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2502–2510, 2018. **1**
- [23] Lucio Azzari and Alessandro Foi. Variance stabilization for noisy+ estimate combination in iterative poisson denoising. *IEEE signal processing letters*, 23(8):1086–1090, 2016. **1, 6**
- [24] Antoni Buades, Bartomeu Coll, and J-M Morel. A non-local algorithm for image denoising. In *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR’05)*, volume 2, pages 60–65. Ieee, 2005. **1**

- [25] Michael Elad and Michal Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image processing*, 15(12):3736–3745, 2006. 1
- [26] Shuhang Gu, Lei Zhang, Wangmeng Zuo, and Xiangchu Feng. Weighted nuclear norm minimization with application to image denoising. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2862–2869, 2014. 1
- [27] Ding Liu, Bihan Wen, Yuchen Fan, Chen Change Loy, and Thomas S Huang. Non-local recurrent network for image restoration. *Advances in neural information processing systems*, 31, 2018. 1
- [28] Yunjin Chen and Thomas Pock. Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1256–1272, 2016. 1
- [29] Jan Funke, Fabian Tschopp, William Grisaitis, Arlo Sheridan, Chandan Singh, Stephan Saalfeld, and Srinivas C Turaga. Large scale image segmentation with structured loss based deep learning for connectome reconstruction. *IEEE transactions on pattern analysis and machine intelligence*, 41(7):1669–1680, 2018. 1
- [30] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155, 2017. 1
- [31] Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. Noise2void-learning denoising from single noisy images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2129–2137, 2019. 1
- [32] Daniel Zoran and Yair Weiss. From learning models of natural image patches to whole image restoration. In *2011 international conference on computer vision*, pages 479–486. IEEE, 2011. 1
- [33] Yochai Blau and Tomer Michaeli. The perception-distortion tradeoff. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6228–6237, 2018. 1
- [34] Zhou Wang and Alan C Bovik. Mean squared error: Love it or leave it? a new look at signal fidelity measures. *IEEE Signal Processing Magazine*, 26(1):98–117, 2009. 1
- [35] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. *arXiv preprint arXiv:2111.08005*, 2021. 2
- [36] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. *arXiv preprint arXiv:2108.02938*, 2021. 2
- [37] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. *Advances in Neural Information Processing Systems*, 35:23593–23606, 2022. 2
- [38] Hyungjin Chung, Byeongsu Sim, Dohoon Ryu, and Jong Chul Ye. Improving diffusion models for inverse problems using manifold constraints. *Advances in Neural Information Processing Systems*, 35:25683–25696, 2022. 2
- [39] Zahra Kadkhodaie and Eero Simoncelli. Stochastic solutions for linear inverse problems using the prior implicit in a denoiser. *Advances in Neural Information Processing Systems (NIPS)*, 34:13242–13254, 2021. 2
- [40] Bahjat Kawar, Gregory Vaksman, and Michael Elad. SNIPS: Solving noisy inverse problems stochastically. *Advances in Neural Information Processing Systems (NIPS)*, 34:21757–21769, 2021. 2, 3, 4, 7
- [41] Hyungjin Chung, Byeongsu Sim, and Jong Chul Ye. Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12413–12422, 2022. 2
- [42] Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*, 2022. 2, 6
- [43] Nadav Torem, Roi Ronen, Yoav Y Schechner, and Michael Elad. Towards a most probable recovery in optical imaging. *arXiv preprint arXiv:2212.03235*, 2022. 2
- [44] Netanel Ratner and Yoav Y Schechner. Illumination multiplexing within fundamental limits. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8. IEEE, 2007. 2
- [45] Yoav Shechtman, Yonina C Eldar, Oren Cohen, Henry Nicholas Chapman, Jianwei Miao, and Mordechai Segev. Phase retrieval with application to optical imaging: A contemporary overview. *IEEE Signal Processing Magazine*, 32(3):87–109, 2015. 3
- [46] Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the International Conference on Machine Learning (ICML)*, pages 681–688, 2011. 3
- [47] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems (NIPS)*, 32, 2019. 3
- [48] Bahjat Kawar, Gregory Vaksman, and Michael Elad. Stochastic image denoising by sampling from the posterior distribution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1866–1875, 2021. 3, 4, 6
- [49] Yang Song and Stefano and Ermon. Improved techniques for training score-based generative models. *Advances in Neural Information Processing Systems (NIPS)*, 33:12438–12448, 2020. 3, 6
- [50] Yaniv Romano, Michael Elad, and Peyman Milanfar. The little engine that could: Regularization by denoising (RED). *SIAM Journal on Imaging Sciences*, 10(4):1804–1844, 2017. 6

- [51] Frank Bowman. *Introduction to Bessel Functions*. Courier Corporation, 2012. 5
- [52] James R Fienup. Phase retrieval algorithms: A comparison. *Applied Optics*, 21(15):2758–2769, 1982. 7
- [53] Christopher Metzler, Phillip Schniter, Ashok Veeraraghavan, and Richard Baraniuk. Prdeep: Robust phase retrieval with a flexible deep network. In *International Conference on Machine Learning*, pages 3501–3510. PMLR, 2018. 7
- [54] Christopher A Metzler, Arian Maleki, and Richard G Baraniuk. BM3D-PRGAMP: Compressive phase retrieval based on BM3D denoising. In *2016 IEEE International Conference on Image Processing (ICIP)*, pages 2504–2508. IEEE, 2016. 7
- [55] David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision (ICCV)*, volume 2, pages 416–423. IEEE, 2001. 7
- [56] Lei Tian, Xiao Li, Kannan Ramchandran, and Laura Waller. (Dataset) *LED array 2D Fourier Ptychography*. Available online. <http://gigapan.com/gigapans/170956>. 7
- [57] Roi Ronen, Vadim Holodovsky, and Yoav Y Schechner. Variable imaging projection cloud scattering tomography. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 8
- [58] Roi Ronen, Yoav Y Schechner, and Eshkol Eytan. 4D cloud scattering tomography. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5520–5529, 2021. 8