

Image Denoising and the Generative Accumulation of Photons

Alexander Krull^{1,2}, Hector Basevi^{2,3}, Benjamin Salmon¹, Andre Zeug⁴, Franziska Müller⁴,
 Samuel Tonks¹, Leela Muppala¹, and Aleš Leonardis¹

¹School of Computer Science, University of Birmingham, UK

²Centre of Membrane Proteins and Receptors, Universities of Nottingham and Birmingham, UK

³Metabolism and Systems Research, University of Birmingham, UK

⁴Medizinische Hochschule Hannover, Germany

Abstract

We present a fresh perspective on shot noise corrupted images and noise removal. By viewing image formation as the sequential accumulation of photons on a detector grid, we show that a network trained to predict where the next photon could arrive is in fact solving the minimum mean square error (MMSE) denoising task. This new perspective allows us to make three contributions: *i.* We present a new strategy for self-supervised denoising, *ii.* We present a new method for sampling from the posterior of possible solutions by iteratively sampling and adding small numbers of photons to the image. *iii.* We derive a full generative model by starting this process from an empty canvas. We call this approach generative accumulation of photons (GAP). We evaluate our method quantitatively and qualitatively on 4 new fluorescence microscopy datasets, which will be made available to the community. We find that it outperforms its baselines or performs on-par.

1. Introduction

Scientific imaging techniques such as fluorescence microscopy have to limit the amount of light used to avoid damaging or destroying their sample [13]. As a result, the recorded images inevitably suffer from a certain degree of noise which has to be addressed in the downstream analysis. Images can be subject to a variety of different types [18] of noise which can be alleviated by various technical means (e.g. [7]). However, there is a type of noise which is physically inevitable for most imaging setups in low-light conditions. It is referred to as Poisson *shot noise*.

Shot noise is the result of the particle nature of light. Even high-end scientific detectors and cameras that can accurately count the precise number of photons hitting each pixel cannot record a noise-free image. For a given light

intensity the number of photons arriving at the detector is itself inherently random and follows a Poisson distribution. The effect is especially severe in microscopy applications, operating in low-light conditions.

The last decade has seen a number of deep learning-based computational methods designed to reduce noise after images have been recorded in order to allow for improved analysis of the data [18]. One of the first proposed methods, known as content-aware image restoration (CARE) [36], is based on training convolutional neural networks (CNNs) to learn a mapping from noisy images to clean images. Unfortunately, the method requires pairs of corresponding noisy and clean images during training, which can be hard to acquire in practice, rendering it inapplicable in many situations. However, other works have expanded on this line of research, enabling training with noisy image pairs [20] and even with unpaired noisy images, e.g. [17, 3, 5, 29].

While achieving impressive results, these supervised and self-supervised methods share a common shortcoming: Denoising is inherently an ill-posed problem and given an image corrupted by a substantial amount of noise, it is not generally possible to recover the true underlying clean image. In fact, there is a posterior distribution of possible solutions that all might have led to the original noisy observation. When we view denoising as a regression problem like [36] aiming to learn a mapping to the clean image, we are in fact learning a mapping to a compromise between possible solutions, which may itself look different (often being more blurry) from real clean images.

This problem has been explored by Prakash *et al.* [28, 27], who proposed the idea of *diversity denoising* based on a variational autoencoder [16] (VAE). Instead of producing a single solution for each noisy image, Prakash *et al.* are able to sample possible solutions from an approximate posterior distribution of clean images.

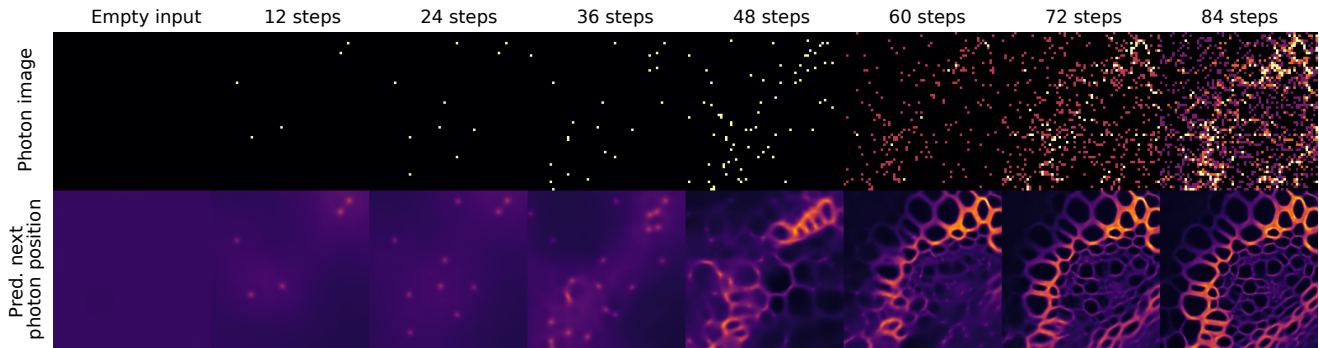


Figure 1. **Generative accumulation of photons (GAP):** Starting with an empty canvas, our method repeatedly predicts a map of probabilities of where the next photon might arrive and uses it to randomly place photons. We show the process for the *Neuro-PC* and *Conv-PC* datasets. Photon images have been down-sampled for better visibility.

Here, we take an entirely new perspective, focusing on shot noise and the denoising of shot noise corrupted images. Instead of viewing shot noise as a secondary corruption process applied to a clean image, we understand image formation as the sequential accumulation of photons and see any measured shot noise-affected image as a result of this process. We call this approach generative accumulation of photons (GAP).

We train a CNN to take a shot noise-affected image as input and predict a probability distribution over where the next photon might arrive. We show that, for normalized images, predicting the next photon position is identical to denoising the image.

Based on this insight, we derive a novel method of training a self-supervised model for image denoising. Additionally, by understanding image generation as the accumulation of photons, we describe a new method for diversity denoising: we iteratively predict the distribution of the next photon position and randomly sample photons accordingly. Finally, by starting with an empty photon-free canvas, we are able to derive a full generative image model. The process is illustrated in Figures 1 and 2.

We introduce four new shot noise-corrupted microscopy datasets for evaluation that will be made available to the community. We evaluate our method quantitatively and qualitatively and find that it yields competitive results.

2. Related Work

2.1. Supervised denoising

Supervised denoising methods usually train CNNs using pairs of noisy and clean images to learn a mapping between the two. A common choice for the loss function is the mean square error (MSE) between the prediction and clean ground truth. Considering that there is a distribution of possible solutions, minimising the MSE loss corresponds to finding the expected value. Unfortunately obtaining clean ground truth data can be challenging or impossible for many applications and so supervised methods are often not applicable in the context of scientific imaging.

Lehtinen *et al.* introduced *Noise2Noise* [20], a partial so-

lution to this problem. They showed that it is possible to replace the clean ground truth target with a second noisy version which might be more readily available. A network trained with this type of data will still find the same MMSE solution. While this presented a big step forward with respect to applicability, Noise2Noise still requires training pairs, which have to be collected for this purpose.

2.2. Self-supervised blind-spot denoising

Self-supervised blind-spot methods [17, 3, 12] suggest a training strategy that can do without paired training data, *i.e.* allowing training directly on the data that should be denoised, while still obtaining the same MMSE solution. The main idea is to block out individual pixels in order to use them as noisy targets (similar to Noise2Noise). These strategies rely on the assumption that imaging noise is conditionally pixel-independent given the underlying clean signal, making it not possible to predict the noise in a pixel from its surroundings. The downside of this approach is that, when making a prediction for a pixel, the network cannot make use of the pixel value itself, thus it is not making optimal use of the available information.

Our photon-based self-supervised denoising strategy is related to the blind-spot idea in that it removes part of the input image to use it as the target. However, instead of removing pixels, we are only removing individual photons, which means we are not facing the same problem of disregarded information.

2.3. VAE-based denoising

Another approach to image denoising has been suggested in [28]. The core idea is to use a variational autoencoder to describe the distribution of noisy images. By including a statistical model of the imaging noise as part of the decoder, the method allows us to: i. sample from an approximate posterior distribution of possible clean images, and ii. to sample clean images from scratch, functioning as a full generative model.

An extended method with a more powerful network architecture was presented in [27] under the name HDN. We

see this method as our main competitor as it can be trained from unpaired noisy data and, similarly to our method, can function as a generative model.

Unlike GAP which produces an MMSE denoising result in a single step, HDN produces MMSE results by repeated sampling and averaging from the posterior distribution.

2.4. Generative image models

Generative image models aim to describe a probability distribution over images, a highly challenging task, due to: i. the high dimensionality of the random variable (the number of pixels) and ii. due to the complex higher-order correlations between pixel values at different locations. As a result of ii. the distribution cannot easily be factorised into lower order terms and attempts to factorise using methods such as Markov random fields (MRFs) [21] have led to overly simplistic results that do not realistically describe the image distribution.

In recent years, a number of approaches to this problem have been highly successful. Latent variable models, such as generative adversarial networks (GANs) [8] and VAEs [16], or normalising flows [30], describe difficult distributions indirectly by starting with an easily modelled high dimensional latent variable (usually following a normal distribution) which is then deformed using convolutional neural networks (CNNs) to describe the distributions of interest.

A different approach to this is autoregressive modelling, as proposed by Van Oord *et al.* [35]. By viewing image generation as a sequential process in which the pixels of an image are thought to be generated one-at-a-time conditioned on all previous pixels. In this setup, the whole model can be formulated as a product of 1D conditional distributions over each pixel’s intensity value. Our method can be viewed as an autoregressive approach as we model image generation as a sequential process. However, we sample images by sequentially placing individual photons instead of drawing pixel values.

Finally, the current state-of-the-art approach to image modelling, *denoising diffusion models* [11], follows a similar approach by describing image generation as a sequence of steps. The process is inspired by physics and considers an image as a particle in a high dimensional space, diffusing away from its original position according to some noise distribution. To generate an image the denoising diffusion approach reverses the diffusion process by applying a sequence of denoising steps.

Denoising diffusion models iteratively reverse a diffusion process on clean images which typically involves Gaussian noise. This noise can be applied directly to the image [33, 11], or instead to a latent representation of the image corresponding to a pre-existing autoencoder [31]. In both cases, the diffusion noise distribution is unrelated to

the noise distribution of the training data. The diffusion model learns to sample from the *noisy* training data distribution and so its samples contain this noise. In contrast, GAP learns to sample from a noisy training data distribution *and* to denoise this distribution. In addition, every iteration of GAP results in a physically valid noisy image.

Recent works have explored generalisations of diffusion models to broader families of corruption processes. Bansal *et al.* [2] focused on deterministic image corruptions. Daras *et al.* [6] focused on image corruptions which are linear with respect to the clean image. GAP focuses on shot noise, which is neither deterministic nor linear.

3. Method




Figure 2. **Sampling algorithm:** Starting with a noisy or empty image, we repeatedly predicts a map of where the next photon might arrive and use it to sample a small number of photons.

3.1. Image formation and shot noise

When we record an image, we usually project light onto a digital sensor, such as a CMOS or a CCD chip.¹ These chips contain many detector elements measuring the amount of light arriving at different locations on the chip. In our simplified model we assume that each of these detector elements corresponds to one pixel of the final image. When measuring the amount of light in each pixel, we treat light as discrete particles called photons. In an ideal case with a perfect detector, each pixel value in the final image corresponds to the number of photons that fell onto the pixel.

The result of this process is a *shot noise* corrupted image $\mathbf{x} = (x_1, \dots, x_n)$, where the photon count x_i in each pixel i , is independently drawn from a Poisson distribution

$$p(x_i | s_i) = \frac{s_i^{x_i} \exp(-s_i)}{x_i!}, \quad (1)$$

where s_i refers to the expected number of photons hitting the pixel i during the exposure, *i.e.* to the light intensity

¹Some imaging technologies, especially those capable of counting photons, work by scanning the sample and recording one pixel at a time. Since this does not affect our model we will focus our explanation on camera-based systems for simplicity.

at the pixel – the quantity we were originally interested in measuring. We will refer to the vector $\mathbf{s} = (s_1, \dots, s_n)$ as the signal or as a clean image.

Since photons are hitting each pixel independently given a signal, we can describe the probability of observing a noisy image \mathbf{x} given a signal \mathbf{s} as

$$p(\mathbf{x}|\mathbf{s}) = \prod_{i=1}^n p(x_i|s_i). \quad (2)$$

We can now think of image formation as a two-step process. We can imagine an image being created by first drawing a clean image \mathbf{s} from a distribution $p(\mathbf{s})$ and then applying shot noise by drawing photon counts from Eq. 2 to create the shot noise-corrupted version.

3.2. The denoising task

Given noisy observation \mathbf{x} , denoising is defined as finding an estimate $\hat{\mathbf{s}}$ for the unknown clean image \mathbf{s} .

However, considering the process of image generation described above, finding the true signal may not be possible since many clean images can lead to the same noisy observation. We can use Bayes' theorem to write down a posterior distribution over possible clean images for a given noisy observation

$$p(\mathbf{s}|\mathbf{x}) \propto p(\mathbf{x}|\mathbf{s})p(\mathbf{s}). \quad (3)$$

Deep learning-based approaches (*e.g.* [36, 17]) often view denoising as a regression problem and use CNNs to try to directly learn a mapping from \mathbf{x} to \mathbf{s} . When such methods are trained with a mean squared error (MSE) loss function the optimal solution is the expectation

$$\hat{\mathbf{s}} = \int p(\mathbf{s}|\mathbf{x})\mathbf{s} ds. \quad (4)$$

We call this the minimum mean squared error (MMSE) solution. This is a sensible way to find an estimate, but we should be aware that it constitutes a compromise between all possible \mathbf{s} .

3.3. Image generation from photon sequences

Here, we take an alternative view on image generation. Instead of thinking of our pixel values as being drawn from a Poisson distribution, we will derive an equivalent description, viewing image generation as a sequential process. Remembering that our observation \mathbf{x} is created by photons hitting our detector, we can imagine that it was created by an ordered sequence of photons $\mathbf{i} = (i_1, \dots, i_T)$, where i_t is the index of the pixel where photon t hit the detector. The index t simply refers to the position of the photon in the sequence, with the first photon arriving at $t = 1$ and the last one at $t = T$.




Figure 3. **Photon splitting:** We propose a new way of generating training pairs that requires only noisy data. We split the noisy image by randomly assigning photons to the input or target image. The number of photons assigned to the input is drawn from a binomial distribution, with a parameter p controlling the noise level. Remaining photons are assigned to the target.

Assuming a known sequence of photons, we compute the resulting image by counting the number of photons hitting each pixel as

$$x_i = \sum_{t=1}^T \mathbb{1}(i_t = i), \quad (5)$$

where $\mathbb{1}(i_t = i)$ is the indicator function.

Considering a given signal \mathbf{s} and a known number T of photons, we can compute the probability of a sequence \mathbf{i} as

$$p(\mathbf{i}|\mathbf{s}, T) = \begin{cases} \prod_{t=1}^T p(i = i_t|\mathbf{s}) & T = |\mathbf{i}| \\ 0 & T \neq |\mathbf{i}| \end{cases}, \quad (6)$$

with the probability being 0 where the length $|\mathbf{i}|$ of the sequence does not match the number of photons T . Since the photons hit the detector independently, their position in the sequence does not matter and we can rewrite the probability as a product over pixels as

$$p(\mathbf{i}|\mathbf{s}, T) = \begin{cases} \prod_{i=1}^n p(i|\mathbf{s})^{x_i} & T = |\mathbf{i}| \\ 0 & T \neq |\mathbf{i}| \end{cases}. \quad (7)$$

where the probability $p(i|\mathbf{s})$ for a photon to hit a particular pixel i , given the signal, should be proportional to the light intensity at the pixel. Thus, we can compute it as the normalised signal at that pixel

$$p(i_t = i|\mathbf{s}) = \frac{s_i}{\sum_{j=1}^n s_j}. \quad (8)$$

However, if the clean signal is unknown the distribution will no longer factorise as easily as Eq. 7. Instead, we have to compute the probability of a sequence as

$$p(\mathbf{i}|T) = \begin{cases} \prod_{t=1}^T p(i = i_t|i_1, \dots, i_{t-1}, T) & T = |\mathbf{i}| \\ 0 & T \neq |\mathbf{i}| \end{cases}, \quad (9)$$

where the distribution $p(i = i_t|i_1, \dots, i_{t-1}, T)$ of the next possible photon location now depends on all previous photons. The order in which photons i_1, \dots, i_{t-1} arrived does

not provide any information regarding the next photon position. By additionally considering that the next photon position does not depend on total photon number T nor on the order of previous photons, we can write

$$p(i = i_t | i_1, \dots, i_{t-1}, T) = p(i = i_t | \mathbf{x}_{t-1}), \quad (10)$$

where \mathbf{x}_{t-1} is the observed image at step $t - 1$ according to Eq. 5.

Equation 10 refers to the distribution over the next possible photon locations given a photon image \mathbf{x}_{t-1} . Before taking a closer look at how it can be computed, we would like to point out its significance. Together, Eq. 9 and Eq. 10 provide not only a way to calculate the probability of a sequence but also an iterative way to sample a sequence of photons and therefore images \mathbf{x}_T . Furthermore, for large T , we can expect \mathbf{x}_T to approach the clean image \mathbf{s} , when scaled correctly, so that Eq. 9 and Eq. 10 hold the key to the generation of clean images as well.

3.4. Predicting the next photon location is MMSE denoising for normalised signals

Let us now take a closer look at the distribution of possible next photon locations $p(i = i_t | \mathbf{x}_{t-1})$. We can rewrite Eq. 10 by marginalising over the unknown signal and using Eq. 8 as

$$p(i = i_t | \mathbf{x}_{t-1}) = \int p(\mathbf{s} | \mathbf{x}_{t-1}) p(i_t = i | \mathbf{s}, \mathbf{x}_{t-1}) d\mathbf{s} \quad (11)$$

$$= \int p(\mathbf{s} | \mathbf{x}_{t-1}) \frac{s_i}{\sum_{j=1}^n s_j} d\mathbf{s}. \quad (12)$$

We can see that the result is a weighted average of the possible normalized signals. We should expect that the distribution will be high entropy for small t , *i.e.*, when we have not yet observed many photons, and that it should become more concentrated and low entropy for large t . For very large t , the distribution should approach a normalised version of the signal (Eq. 8), because \mathbf{x}_t will give us more and more information on the underlying signal.

Interestingly, Eq. 12 closely resembles Eq. 4. In fact, if we were to consider only normalized signals with $\sum_{j=1}^n s_j = 1$ the two equations are identical, meaning that the task of predicting the next photon location is identical to denoising the image in an MMSE sense.

We will use a CNN to approximate $f_\theta(\mathbf{x}_{t-1}) \approx p(i = i_t | \mathbf{x}_{t-1})$, where θ are the network parameters. In section 3.5, we will discuss how we can train the CNN to achieve this task.

3.5. Learning to predict the next photon location

Based on the insight from section 3.4, we know that any model trained for MMSE denoising can approximate the distribution over the next photon location $p(i = i_t | \mathbf{x}_{t-1})$.

Starting with normalised clean training images \mathbf{s}^k , the traditional way of creating training pairs is to simulate the corresponding noisy version \mathbf{x}^k . We can then train a denoiser network using a standard quadratic loss function, with \mathbf{x}^k as input and \mathbf{s}^k as target.

However, in many cases clean data is unavailable. Considering the task of predicting the next photon location suggests an alternative self-supervised approach by viewing the problem as a classification task learning the categorical distribution of possible photon positions. By using a softmax layer over pixels at the output of our network to ensure that outputs sum to one, we can use the standard cross-entropy loss. In principle, this would require only individual photon positions as target for each training image, just as classifiers are frequently trained using individual class labels for each training example. We could easily create such training pairs from unpaired noisy images \mathbf{x}^k by randomly removing a single photon and using it as target. The corresponding cross entropy loss is

$$L(\theta) = - \sum_{k=1}^m \sum_{i=1}^n \ln f_i(\mathbf{x}_{\text{inp}}^k; \theta) x_{\text{tar}, i}^k, \quad (13)$$

where m is the number of training images, $\mathbf{x}_{\text{inp}}^k$ is the training image with one photon randomly removed and $\mathbf{x}_{\text{tar}}^k$ is a one-hot representation of the removed photon position.

However, we require training data at multiple noise levels to enable our network to predict an accurate approximation of $p(i = i_t | \mathbf{x}_{t-1})$ at different times t . To achieve this, we use a control parameter p and split the image \mathbf{x}^k into two parts, $\mathbf{x}_{\text{inp}}^k$ and $\mathbf{x}_{\text{tar}}^k$. We can think of this process as simulating a shorter exposure time during image acquisition. Considering that \mathbf{x}^k was recorded with a certain exposure time τ , we can imagine what would be the result if we had instead recorded two images consecutively, with the first image being exposed for $p\tau$ and the second being exposed for $(1 - p)\tau$. Considering, that the underlying signal remained fixed during the entire time, each of the photons that make up \mathbf{x}^k would end up in the first image with probability p and in the second image with probability $(1 - p)$. To efficiently sample a split for parameter value $0 < p < 1$, we can determine each pixel value $x_{\text{inp}, i}^k$ by drawing from binomial distribution using p and x_i^k as the distributions parameters, for success probability and number of trials, respectively. We can then compute the number photons in the target image as $x_{\text{inp}, i}^k = x_i^k - x_{\text{inp}, i}^k$. By changing the value p we can control the number of photons that are on average assigned to the input or target image respectively. The process is illustrated in Figure 3. We use a randomly selected p for each training patch to cover all levels of noise. We show in the Supplementary material that the loss formulation in Eq. 13 can still be used to maximise the likelihood of the training data even when $\mathbf{x}_{\text{tar}}^k$ is not a one-hot encoding

of a single photon position but an image that contains an arbitrary number photons. In practice, we use a normalized variant that still maximizes likelihood of the data

$$L(\theta) = - \sum_{k=1}^m \frac{1}{n|\mathbf{x}_{\text{tar}}^k|} \sum_{i=1}^n \ln f_i(\mathbf{x}_{\text{imp}}^k; \theta) x_{\text{tar},i}^k, \quad (14)$$

where $|\mathbf{x}_{\text{tar}}^k|$ is the sum of photons in $\mathbf{x}_{\text{tar}}^k$.

3.6. Inference

MMSE denoising: To compute the MMSE denoising result $\hat{\mathbf{s}}$ for a noisy input image \mathbf{x} , we can simply apply our trained CNN. As shown in section 3.4, the resulting probability distribution corresponds to the MMSE estimate. However, since our network uses a final softmax layer, the output can only tell us about the normalized pixel intensities and not about the absolute ones. To obtain a scaled version, that is comparable to the results from N2V2 or HDN, we multiply our output with the number of photons in the input image.

Diversity denoising: To obtain a sample from the posterior of clean images, given a shot noise corrupted input, we use the iterative procedure illustrated in Figure 2. Starting with the original image, we repeatedly apply our network to obtain the distribution for the next photon position and add photons drawn from this distribution. Even though Eq. 9 contains a product over individual photons and we should in principle draw only a single photon at a time, we find that we can add multiple photons simultaneously while maintaining acceptable quality. In practice, we add 10% of the current photon count in each step, increase the total number of photons exponentially. A more detailed description of photon sampling can be found in the Supplementary materials.

Image generation: To obtain samples from our generative model, we follow the same process as diversity denoising, but start with a blank image.

4. Experiments

4.1. Network architecture and training

Here, we will only give a brief overview of the architecture and training procedure used for the experiments on microscopy data. A more detailed description can be found in the Supplementary material.

For all our experiments on microscopy datasets, we use a modified UNet [32] consisting of 6 levels, with a residual block at each level and skip connections. We use 28 feature channels in the first level and double the number of feature channels at each subsequent level. All our networks are trained using the *ADAM* [15] optimizer for 100 epochs. We use randomly cropped patches of 256×256 pixels, which are augmented 8-fold, using random flips and transpose operations. We use a batch size of 32.

4.2. Baselines

Supervised denoising uses the same network architecture as our method except for the softmax layer at the end. It is trained with the same hyperparameters but an MSE loss function. We use 8-fold data augmentation.

N2V2 uses the implementation from *et al.* [12], with default hyper parameters and the default 64×64 training patch size.

HDN uses the implementation from Prakash *et al.* [27], with default hyper parameters and the default 64×64 training patch size. HDN requires a model of the imaging noise, which is usually trained from data. Instead, because we know our data contains pure shot noise, we added an analytical Poisson noise model, accounting for shot noise.

HDN256 uses the implementation from Prakash *et al.* [27] but with increased network complexity to allow for a fairer comparison to our method. Specifically, we increase the dimensionality of the latent variables from 32 feature channels to 70, and the number of deterministic filters in the hidden units from 64 to 140. The method uses 256×256 pixel training patches. We use the same noise model as for HDN.

4.3. Photon counting datasets

While a number of denoising datasets are available in the microscopy domain (*e.g.* [38, 9]), none of them show purely shot noise corrupted data. To address this gap, we introduce four new quantitative datasets, including High-SNR ground truth data and one additional qualitative dataset that does not contain ground truth.

We use two photon-counting datasets that will be made available to the community. As a result, the recorded pixel intensities give a very accurate approximation of the photons hitting each pixel during the exposure.

The Conv-PC dataset We image 5 fields of view (FOV) repeatedly, 512 times at a resolution of 512×512 pixels. Each of the individual frames contains a substantial amount of shot noise. By summing the 512 images for each FOV, we obtain the high-SNR version. Four FOVs were used as training data for supervised denoising, the remaining one was used as test data. This amounts to a total $512 \times 4 = 2048$ noisy images.

The Neuro-PC dataset contains images of mouse neurons. The dataset is created from a z-stack of 2048×2048 pixels by using 2×2 binning in x- and y-direction and 4 times binning in the z-direction. We divided the images into non-overlapping 320×320 regions and rejected empty ones. To produce the corresponding low-SNR versions we reduced the photon count to simulate a 1000-fold shorter exposure by using a binomial distribution with $p = 0.001$. We use every fourth frame as test set and keep the rest as training set for the supervised baseline. All in all, this amounts to 133 images of size 320×320 , 33 of which are test images.




Figure 4. **Qualitative denoising results:** We show MMSE denoising results for all datasets and methods. All baselines except for the supervised one have been trained purely on low-SNR input data.

4.4. Single molecule localisation microscopy

Single molecule localisation microscopy (SMLM) [4] data is produced differently from photon counting data but is subject to the same type of shot noise corruption. It uses a large set of images of the same field of view to detect and localise individual fluorescent emitters in each image. The resulting emitter locations are then stored in a list and can be binned in x and y to produce a 2D histogram/image containing the number of emitters in each bin/pixel.

The NPC-SM dataset was derived from single molecule localisation data published by Löschberger *et al.* in [23]. It shows the arrangement of the *gp210* protein around the nuclear pore complex (NPC).

The MT-SM dataset was derived from single molecule localisation data published by Jimenez *et al.* in [14]. It shows the arranged cells labeled for microtubules. To create the dataset, we binned the detected emitter locations using a bin size of $28\text{nm} \times 28\text{nm}$ to produce the high-SNR data.

4.5. Denoising performance on microscopy images

To evaluate the denoising performance of our method we train one network for our method and one for each baseline (N2V2, HDN, and HDN256). Since these methods do not require clean data, we can train them on the full low-SNR data, including the section used for testing. The supervised baseline, which requires clean training data, is trained only on the designated training section of the data. Quantitative and qualitative results can be found in table 1 and Figure 4.

Our method is on-par or outperforms the baselines and even the supervised approach. The reason for this might be that, depending on the data split, supervised methods might suffer from a mismatch between training and test distributions, which might be especially the case for the *Conv-PC* dataset where test and training data consist of different FOVs showing slightly different patterns.

Table 1. Average PSNRs in dB (higher is better). We show results for microscopy datasets (top) and natural images [25] (bottom) at different peak photon levels λ_{peak} (higher values correspond to less noise). Best results are bold. Best un/self-supervised results are underlined.

| | Superv. | N2V | HDN | HDN256 | Ours |
|----------|---------|-------|-------|--------|--------------|
| NPC-SM | 39.09 | 39.00 | 39.01 | 38.73 | 39.17 |
| MT-SM | 36.59 | 36.17 | 34.32 | 36.43 | 36.64 |
| Conv-PC | 20.42 | 23.85 | 23.77 | 24.32 | 24.84 |
| Neuro-PC | 31.74 | 32.61 | 32.41 | 32.56 | 32.63 |

| λ_{peak} | Supervised | | Un/self-supervised | | |
|-------------------------|------------|--------------|--------------------|--------|--------------|
| | DnCNN | VBP | PURE-LET | IVBM3D | Ours |
| 1 | 21.56 | 22.2 | 20.49 | 21.33 | <u>22.12</u> |
| 2 | 22.67 | 23.25 | 21.35 | 22.18 | <u>23.15</u> |
| 4 | 23.76 | 24.38 | 22.27 | 23.23 | <u>24.21</u> |
| 8 | 24.88 | 25.6 | 23.29 | 24.36 | <u>25.37</u> |
| 20 | 26.63 | 27.39 | 24.88 | 26.06 | <u>27.01</u> |

4.6. Denoising performance on natural images

We additionally evaluate our method on the BSD68 [25] dataset and compare our method to the strongest baselines published in [22]. In Table 1, we show numbers for two supervised methods: VBP [22] and DNCNN [37], as well as two unsupervised methods IVBM3D [1], PURE-LET [24]. More baselines and training details can be found in the supplementary material. Our method outperforms all baselines that can cope without paired data (un/self-supervised).

4.7. Diversity denoising

In Figure 5, we qualitatively evaluate the performance of our method for diversity denoising, that is, its ability to sample diverse possible clean images from single noisy input. To show the full range of possible results, we trained our method on the high-SNR data of the *Conv-PC* dataset.

We generate six different shot noise corrupted versions of an image at different noise levels/photon numbers and use them as input for the sampling procedure described in

Figure 2, to generate three possible clean versions for each noisy input image. Noisy images with low photon counts can be explained by a broad range of possible clean images and yield highly diverse results. Increasing the photon count of the input image, the differences in the sampled clean images become more and more subtle.




Figure 5. **Diverse solutions for shot noise removal:** GAP models can be used to remove shot noise by taking the noisy image as starting point and sequentially adding additional photons, until a clean image is produced. Less noisy inputs lead to less diverse predictions as more information about the clean image becomes available. The last column depicts a zoomed in region indicated by the dashed box. Arrows highlight differences in the samples.




Figure 6. **Comparison of sample quality:** We show randomly selected generated samples for the *Conv-PC* and *MT-SM* datasets compared to randomly cropped real high- and low-SNR patches.

4.8. Image generation performance

Finally, we evaluate our method for the use as a generative image model. We are especially interested in the setting where only low-SNR data is available for training and want to investigate how the distribution of the generated images will compare to the clean high-SNR data. We train our model as well as the HDN and HDN256 baselines on the low-SNR data for each dataset. We generate 10000 images of 256×256 pixels using our method (Figure 2) and HDN256. We then compute the FID [10] score against 10000 random crops of the augmented high SNR-data. We compute the scores using the *clean FID* [26].

Table 2. We show the FID score [10] of 10k generated images compared to 10k random crops of high-SNR data. Note that all methods have been trained on low-SNR training data.

| | 64x64 pixels | | 256x256 pixels | |
|----------|---------------|---------------|----------------|--------------|
| | HDN | Ours | HDN256 | Ours |
| NPC-SM | 146.48 | 146.92 | 204.34 | 95.32 |
| MT-SM | 130.20 | 138.72 | 132.55 | 84.49 |
| Conv-PC | 200.70 | 134.20 | 163.28 | 86.82 |
| Neuro-PC | 111.23 | 115.06 | 77.90 | 64.16 |

For a fair comparison against the HDN baseline, trained on 64×64 pixel patches, we compute the FID using 64×64 pixel patches against 64×64 crops of the high-SNR data. To compare against our method, we use random 64×64 crops from the 256×256 pixel patches generated by our method. Quantitative results can be found in Table 2. Qualitative results for 256×256 patches are shown in Figure 6. We find that our method visually outperforms HDN and HDN256 and consistently achieves lower FID scores for 256×256 patches. For the smaller 64×64 patches FID results are less clear. We believe, that this is due to the fact, that larger structures are not captured at this patch size, and that our high-SNR data contains residual noise, which seems to be better represented by HDN.

5. Discussion and Conclusion

We have introduced a new perspective on shot noise-affected imaging, which we believe, might open the door to new applications in areas of microscopy where only shot noise-affected data is available. We also believe that our method can be extended to be used in a conditional setting for image-to-image translation, such as the prediction of fluorescence channels from bright-field images – a topic that has received much attention in the recent years [19, 34]. While our method is currently limited to data purely affected by shot noise, we hope that future work can extend the approach to be applicable in a more general setting. Finally, we applied GAP to two natural image datasets (see Supplementary material) with encouraging visual results, suggesting that GAP might be applicable as a generative model beyond microscopy.

Acknowledgements:

We would like to thank Jeremy Pike for pointing us the single molecule localisation data and for the helpful discussions we had. Computations were performed using the University of Birmingham’s BlueBEAR HPC service, which provides a High Performance Computing service to the University’s research community (<http://www.birmingham.ac.uk/bear>).

References

- [1] Lucio Azzari and Alessandro Foi. Variance stabilization for noisy+ estimate combination in iterative poisson denoising.

- IEEE signal processing letters*, 23(8):1086–1090, 2016.
- [2] Arpit Bansal, Eitan Borgnia, Hong-Min Chu, Jie S Li, Hamid Kazemi, Furong Huang, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Cold diffusion: Inverting arbitrary image transforms without noise. *arXiv preprint arXiv:2208.09392*, 2022.
- [3] Joshua Batson and Loic Royer. Noise2self: Blind denoising by self-supervision. In *ICML*, pages 524–533. PMLR, 2019.
- [4] Eric Betzig, George H Patterson, Rachid Sougrat, O Wolf Lindwasser, Scott Olenych, Juan S Bonifacino, Michael W Davidson, Jennifer Lippincott-Schwartz, and Harald F Hess. Imaging intracellular fluorescent proteins at nanometer resolution. *science*, 313(5793):1642–1645, 2006.
- [5] C. Broaddus, A. Krull, M. Weigert, U. Schmidt, and G. Myers. Removing structured noise with self-supervised blind-spot networks. In *ISBI*, pages 159–163, 2020.
- [6] Giannis Daras, Mauricio Delbracio, Hossein Talebi, Alexandros G Dimakis, and Peyman Milanfar. Soft diffusion: Score matching for general corruptions. *arXiv preprint arXiv:2209.05442*, 2022.
- [7] Donal J Denvir and Emer Conroy. Electron-multiplying ccd: the new iccd. In *Low-Light-Level and Real-Time Imaging Systems, Components, and Applications*, volume 4796, pages 164–174. SPIE, 2003.
- [8] Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *stat*, 1050:10, 2014.
- [9] Guy M Hagen, Justin Bendesky, Rosa Machado, Tram-Anh Nguyen, Tanmay Kumar, and Jonathan Ventura. Fluorescence microscopy datasets for training deep neural networks. *GigaScience*, 10(5):giab032, 2021.
- [10] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS*, 30, 2017.
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 33:6840–6851, 2020.
- [12] Eva Höck, Tim-Oliver Buchholz, Anselm Brachmann, Florian Jug, and Alexander Freytag. N2v2-fixing noise2void checkerboard artifacts with modified sampling strategies and a tweaked network architecture. In *ECCV workshops*, pages 503–518. Springer, 2022.
- [13] Jaroslav Icha, Michael Weber, Jennifer C Waters, and Caren Norden. Phototoxicity in live fluorescence microscopy, and how to avoid it. *BioEssays*, 39(8):1700003, 2017.
- [14] Angélique Jimenez, Karoline Friedl, and Christophe Leterrier. About samples, giving examples: optimized single molecule localization microscopy. *Methods*, 174:100–114, 2020.
- [15] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. 3rd international conference on learning representations, iclr 2015. In *ICLR*, 2015.
- [16] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *stat*, 1050:1, 2014.
- [17] Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. Noise2void-learning denoising from single noisy images. In *CVPR*, pages 2129–2137, 2019.
- [18] Romain F Laine, Guillaume Jacquemet, and Alexander Krull. Imaging in focus: an introduction to denoising bioimages in the era of deep learning. *The International Journal of Biochemistry & Cell Biology*, 140:106077, 2021.
- [19] Gyuhyun Lee, Jeong-Woo Oh, Mi-Sun Kang, Nam-Gu Her, Myoung-Hee Kim, and Won-Ki Jeong. Deepfcs: bright-field to fluorescence microscopy image conversion using deep learning for label-free high-content screening. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2018: 21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part II 11*, pages 335–343. Springer, 2018.
- [20] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2noise: Learning image restoration without clean data. In *International Conference on Machine Learning*, pages 2965–2974, 2018.
- [21] Stan Z Li. *Markov random field modeling in image analysis*. Springer Science & Business Media, 2009.
- [22] Hao Liang, Rui Liu, Zhongyuan Wang, Jiayi Ma, and Xin Tian. Variational bayesian deep network for blind poisson denoising. *Pattern Recognition*, 143:109810, 2023.
- [23] Anna Löschberger, Sebastian van de Linde, Marie-Christine Dabauvalle, Bernd Rieger, Mike Heilemann, Georg Krohne, and Markus Sauer. Super-resolution imaging visualizes the eightfold symmetry of gp210 proteins around the nuclear pore complex and resolves the central channel with nanometer resolution. *Journal of cell science*, 125(3):570–575, 2012.
- [24] Florian Luisier, Cédric Vonesch, Thierry Blu, and Michael Unser. Fast interscale wavelet denoising of poisson-corrupted images. *Signal processing*, 90(2):415–427, 2010.
- [25] 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 *CVPR*, volume 2, pages 416–423. IEEE, 2001.
- [26] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in gan evaluation. In *CVPR*, 2022.
- [27] Mangal Prakash, Mauricio Delbracio, Peyman Milanfar, and Florian Jug. Interpretable unsupervised diversity denoising and artefact removal. In *International Conference on Learning Representations*, 2022.
- [28] Mangal Prakash, Alexander Krull, and Florian Jug. Fully unsupervised diversity denoising with convolutional variational autoencoders. In *International Conference on Learning Representations*, 2020.
- [29] Yuhui Quan, Mingqin Chen, Tongyao Pang, and Hui Ji. Self2self with dropout: Learning self-supervised denoising from single image. In *CVPR*, pages 1890–1898, 2020.
- [30] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR, 2015.
- [31] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684–10695, 2022.

- [32] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [33] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 07–09 Jul 2015. PMLR.
- [34] Samuel Tonks, Chih Hsu, Steve Hood, Ryan Musso, Ceriden Hopely, Minh Doan, Erin Edwards, Alexander Krull, and Iain Styles. Evaluation of virtual staining for high-throughput screenings. In *20th IEEE International Symposium on Biomedical Imaging*. IEEE, 2023.
- [35] Aaron Van Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. In *International conference on machine learning*, pages 1747–1756. PMLR, 2016.
- [36] Martin Weigert, Uwe Schmidt, Tobias Boothe, Andreas Müller, Alexandr Dibrov, Akanksha Jain, Benjamin Wilhelm, Deborah Schmidt, Coleman Broaddus, Siân Culley, et al. Content-aware image restoration: pushing the limits of fluorescence microscopy. *Nature methods*, 15(12):1090–1097, 2018.
- [37] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017.
- [38] Yide Zhang, Yinhao Zhu, Evan Nichols, Qingfei Wang, Siyuan Zhang, Cody Smith, and Scott Howard. A poisson-gaussian denoising dataset with real fluorescence microscopy images. In *CVPR*, 2019.