



# Noise2Ghost: self-supervised deep convolutional reconstruction for ghost imaging

MATHIEU MANNI,<sup>1,2</sup> DMITRY KARPOV,<sup>3</sup>  KEES JOOST BATENBURG,<sup>4</sup> SHARON SHWARTZ,<sup>2</sup> AND NICOLA VIGANÓ<sup>3,\*</sup> 

<sup>1</sup>*Algorithms and Scientific Data Analysis Group, ESRF — The European Synchrotron, Grenoble, 38000, France*

<sup>2</sup>*Physics Department and Institute of Nanotechnology and Advanced Materials, Bar Ilan University, Ramat Gan, 52900, Israel*

<sup>3</sup>*Université Grenoble Alpes, CEA, IRIG-MEM, Grenoble, 38000, France*

<sup>4</sup>*Leiden Institute of Advanced Computer Science, Leiden Universiteit, 2333, CA Leiden, The Netherlands*  
\**nicola.vigano@cea.fr*


**Abstract:** We present a self-supervised deep-learning-based ghost imaging (GI) reconstruction method, which provides a significant improvement in reconstruction quality for noisy acquisitions among unsupervised methods. We present the supporting mathematical framework and results from both theoretical and real data (from a recent synchrotron X-ray GI experiment). Self-supervision removes the need for clean reference data while offering strong noise reduction. This provides the necessary tools for addressing signal-to-noise ratio concerns for GI acquisitions in emerging and cutting-edge low-light GI scenarios. Notable examples include micro- and nano-scale x-ray emission imaging, e.g., x-ray fluorescence imaging of dose-sensitive samples. Their applications include in-vivo and in-operando case studies for biological samples and batteries.

© 2026 Optica Publishing Group under the terms of the [Optica Open Access Publishing Agreement](#)

## 1. Introduction

Ghost imaging (GI) is a paradigm-changing technique that is mostly known and anticipated for its promise to cut radiation doses [1]. However, these possible radiation dose advantages rely on a much more fundamental concept: The ability to choose the trade-off between spatial resolution, field-of-view (FoV), and acquisition time somewhat independently of the incident beam size. Traditionally, to spatially resolve diffused signals like x-ray fluorescence (XRF), one had to raster scan the object of interest with focused illumination (pencil-beam, PB, Fig. 1(a)). The PB focal spot size determines the spatial resolution. Classical GI achieves that by probing extended regions of a sample at once with structured illumination (Fig. 1(b)). This provides spatially resolved information, whose resolution is independent of the beam size and detector pixel size, but depends instead on the size of the beam structures. Hereafter, we will refer to classical GI simply as GI, which, unlike quantum GI [2], does not require pairs of quantum entangled photons.

A detector that does not observe the illumination beam collects the signals of interest. The acquired GI realizations, composed of the detected signals with each associated illumination beam structure, are then combined and computationally reconstructed into a two-dimensional projection image of the probed object. This means that the acquired data mathematically resides in a different space from the reconstructed image space, and the reconstruction problem is a so-called *inverse problem* [3]. The algorithm used to perform this inversion, known as the *reconstruction method*, significantly influences the quality and accuracy of the reconstructed image, particularly under noisy conditions. GI is independent of the contrast and probe type used, and it has been demonstrated with radiation across the whole electromagnetic spectrum [4,5],



**Fig. 1.** Schematic representation of diffused emission signal acquisitions (e.g., x-ray fluorescence imaging) using pencil raster beam scanning (a) and ghost imaging (b). The former uses a point beam to scan every pixel to form an image, while the latter illuminates the sample with a series of structured beams. The abbreviations “Img.” and “Det.” stand for “Imaging” and “Detector”, respectively.

neutrons [6], electrons [7], and atoms [8]. Thus, it can be used for the same applications where PB raster scanning is used.

The advantages of GI compared to PB acquisitions stem from the inherent small- and large-scale correlations found in natural images, which make them compressible under suitable representations [9], while non-structured noise is not. Thanks to structured illumination, which can capture unique large- and small-scale information for each measurement, GI enables image reconstruction with fewer data points than the number of reconstructed pixels. Thus, GI unlocks potential gains in acquisition speed or deposited dose, compared to raster scanning, by simply reducing the number of acquired data points [10].

GI offers a specific advantage over PB scanning for dose-sensitive samples. High-flux focused PBs create significant excess charges in small localized regions at each exposure from radiation-induced local ionization of the samples. This is one of the main contributors to damage in biological samples and the skewing of functioning parameters in batteries during acquisitions. On the contrary, GI creates this same charge on a much larger area (i.e., the whole illumination field-of-view). The latter can recombine and/or disperse faster than the former, leading to lower degradation over time. For instance, in cryogenic electron microscopy (also known as cryo-EM), reduced dose rates have been observed to produce less damage compared to higher dose rates at the same total dose [11].

Thanks to all these advantages, GI has the potential to spark many breakthroughs in applications in various fields ranging from biomedical imaging to remote sensing. However, GI has only seen limited applicability so far, especially in the above-mentioned nano-scale imaging applications. In these cases, the observed signals are affected by overwhelmingly high Poisson noise levels due to their very small photon fluxes. This noise type does not enjoy the intrinsic noise reduction of GI, known as Felgett’s advantage, thus reducing GI’s advantage over PB [10]. Both traditional methods, based on regularized variational reconstructions (e.g., Total Variation minimization - TV-min [12]), and modern machine learning (ML) techniques (both supervised [13] and unsupervised [14]) are not specifically designed for this scenario. Here, we propose a GI reconstruction method particularly suited for working with acquisitions affected by high noise levels, regardless of whether the noise type is Gaussian or Poisson. Our method leverages the latest developments in unsupervised ML techniques, so it does not require high-quality reference data. This work is a step towards practical applications of GI, particularly for photon-limited scenarios like nano-scale x-ray emission imaging [15].

## 2. Method

In this section, we present the proposed method, Noise2Ghost (N2G). First, we introduce related methods. Then, we describe the signal formation (forward) model and present our proposed reconstruction method. Finally, we show how this method addresses random acquisition noise. The code of the algorithms presented in this section is available at [Code 1](#), Ref. [16].

### 2.1. Existing methods

GI's ability to obtain images with more pixels than the acquired realizations results in an undersampled reconstruction problem. These reconstructions require some prior knowledge (regularization) to recover accurate images: Unregularized least-squares (LS) reconstructions present specific high-frequency large-scale artifacts that resemble high noise levels [17,18]. Methods employing convolutional neural networks (CNNs) are currently the best-performing tools for image processing applications. They can capture small- and large-scale correlations and learn the most suitable representations to capture the features of the treated signals (acting like regularizers) [19]. Thus, they are also better suited for GI reconstructions than traditional variational methods, e.g., TV-min, which use strong but also rather simple *a priori* assumptions on the types of features present in the reconstructed images.

Both ML methods and CNN architectures have seen an explosion in development and applications over the last decade. Quite a few methods have been specifically developed for GI reconstructions. Notable examples are supervised methods [13], where a CNN (e.g., a U-net [33]) compares LS GI reconstructions from undersampled acquisitions against their corresponding high-quality reconstructions. The network learns to identify and correct the corrupted features in the LS images. This approach heavily depends on large amounts of known examples that present strong similarity to the images of interest. It also requires specific training data for each acquisition parameter variation (e.g., noise levels, image size, levels of undersampling, etc).

Untrained generator-network methods do not require pre-training against high-quality reference images, thereby removing the need for large databases [34,35]. They can incorporate knowledge of the image formation model in the learning algorithm (e.g., physics-informed ML), and train the model to produce an image that fits the acquired data. Various implementations of this approach have been proposed [26–28], including the method called GIDC (shortened from GI using Deep neural network Constraint) [14], which is based on the deep image prior (DIP) [19]. Further iterations have shown improved results with larger and/or more complex models [29–31], and more complex priors [32]. As a drawback, these methods exhibit lower reconstruction quality for high noise levels in the acquired data, because they delegate noise reduction solely to regularization and early stopping. Similar to DIP-based methods, implicit neural representations (INRs) -based methods use untrained multi-layer perceptron (MLP) NNs to reconstruct an image without high-quality reference data [15,36]. They accept pixel coordinates instead of input images. During training, the MLP learns a representation of the imaged object, which produces the desired image for the input coordinates.

Self-supervised methods like Noise2Noise (N2N) use the same training procedure as supervised methods, but they use noisy data as target [20]. They only require two noisy realizations of the same measurement (e.g., two noisy images of the same scene) and exploit the fact that the only difference between the two should be the noise. For each of the two measurements, they train the CNN to predict the other. As previously mentioned, CNNs can capture small and large-scale correlations, leading them to learn the features of the measurements while discarding their uncorrelated parts, i.e., the noise. Similarly, single-image self-supervised denoising methods have also emerged, including Noise2Void (N2V) [21], and a few more recent developments [22–25]. Compared to N2N, these methods only exploit the spatial correlation of imaging data. N2N can be adapted to work with inverse problems like computed tomography [37], while single-image methods like N2V can be directly applied to least-squares reconstructions. Recently,

the resulting ML model from N2V has been used as prior for iterative denoising of GI LS images [38]. However, as we will see in Sec. 2.5, they do not handle the so-called *missing realization* problem, which refers to artifacts arising from not having enough realizations to represent the imaged objects, and that is usually an important source of noise in GI reconstructions. A brief comparison of all the mentioned methods is given in table 1.

**Table 1. Comparison of GI Reconstruction Methods**

Method	Description	Advantages	Limitations
Unregularized Least-Squares (LS) [17,18]	Direct reconstruction without regularization	Simple, no prior assumptions	High-frequency artifacts; high noise levels; no compression
Regularized Variational (e.g., TV) [4]	Uses strong <i>a priori</i> assumptions (e.g., sparsity, smoothness)	Robust; cheap; reference in literature	Oversimplified assumptions; limited feature capture
Supervised denoising (e.g., DeepGhost) [13]	Uses pairs of LS and high-quality reconstructions	Effective for known image types	Dependent on large, similar training data; sensitive to parameter variations
Self-Supervised denoising (e.g., N2V, N2N) [20–25]	No pre-training; denoising in image space	No clean data required	Unable to deal with missing realization artifacts
Unsupervised (e.g., GIDC, INRs) [14,15,26–32]	No pre-training; physics-informed learning	No clean data required; best results with high compression	Reduced quality with high noise; relies on regularization/early stopping

## 2.2. Forward model and variational methods

Here, we provide the mathematical description of the signal formation and the deep image prior (DIP) reconstruction, which will serve as a basis for our method's derivation. Let us represent the discretized expected image of the studied object as the vector  $x^* \in \mathbb{R}^N$ , where  $N$  is the number of its pixels. We probe  $x^*$  with a set of structured illumination patterns  $\mathcal{W} = \{w_1, w_2, \dots, w_m\}$ , which form the acquisition matrix  $W = [w_1, w_2, \dots, w_m] \in \mathbb{R}^{M \times N}$ , where  $M$  is the number of measurements. If we assume that the interaction between the property represented by  $x^*$  and the illumination patterns  $W$  is linear, then the recorded signal  $y \in \mathbb{R}^M$ , corrupted by an uncorrelated zero-mean noise  $\epsilon$  is given by:

$$y = Wx^* + \epsilon = b + \epsilon \quad (1)$$

where  $b \in \mathbb{R}^M$  is the non corrupted signal. The reconstruction aims to recover the vector  $\hat{x}$  that is the most probable estimate of the clean image  $x^*$ . Regularized variational methods seek  $\hat{x}$  by solving the following minimization problem:

$$\hat{x} = \operatorname{argmin}_x \frac{1}{2} \|Wx - y\|_2^2 + \lambda R(x) \quad (2)$$

where  $R : \mathbb{R}^N \rightarrow \mathbb{R}$  denotes a regularization term that imposes some *prior knowledge* on the reconstruction. This prior knowledge is often related to the structural properties of the to-be-reconstructed image, and it is supposed to enhance its signal-to-noise ratio. A common choice of regularization is the so-called total variation (TV) minimization, which promotes reconstructed images with sparse gradients. This, in turn, imposes short-range correlations in the reconstructions.

Unsupervised generative machine learning methods have taken inspiration from Eq. (2) [34], and developed an image recovery scheme that seeks to solve a very similar minimization problem,

but on the output of a machine-learning model (typically a CNN), as follows:

$$\hat{\theta} = \operatorname{argmin}_{\theta} \frac{1}{2} \|WN_{\theta}(r) - y\|_2^2 + \lambda R(N_{\theta}(r)) \quad (3a)$$

$$\hat{x} = N_{\hat{\theta}}(r) \quad (3b)$$

where  $r \in \mathbb{R}^N$  is any input image, which could be randomly chosen,  $N_{\theta} : \mathbb{R}^N \rightarrow \mathbb{R}^N$  is the model, and  $\theta \in \mathbb{R}^T$  is its parameters vector with  $T$  different parameters. In [14],  $r = W^{\dagger}y$  where  $W^{\dagger}$  is the pseudo-inverse of  $W$ , making  $r$  the least-squares reconstruction. The model  $N_{\theta}$  learns to produce the image  $\hat{x}$  that satisfies both the forward model Eq. (1) and the regularization  $R$ , weighted by the  $\lambda$  parameter. Neural networks have notions of non-locality and non-linearity in their response, which can behave as an additional prior on top of  $R$  [19] when the training process is stopped before convergence. The  $\lambda$  parameter is commonly determined through cross-validation, by creating a small leave-out set of measurements, and choosing the  $\lambda$  that makes the model fit them best [39].

While Eq. (3) has proven to help with the reconstruction of under-determined acquisitions [14], it is not tuned to deal with strong random noise. Following the same derivation steps as in [37], the expected prediction error of the data fitting term in Eq. (3a) is:

$$\mathbb{E}_y \|WN_{\theta}(r) - y\|_2^2 = \mathbb{E}_{b,\epsilon} \|WN_{\theta}(W^{\dagger}(b + \epsilon)) - (b + \epsilon)\|_2^2 \quad (4)$$

which is intuitively minimized when  $N_{\theta} = I$  (the identity matrix). Thus, noise reduction is solely delegated to the regularization term  $R$ . Therefore, this reconstruction technique might not be able to provide an advantage over traditional variational methods in this specific scenario.

### 2.3. Proposed method: Noise2Ghost (N2G)

Here, we propose our CNN-based noise-tolerant unsupervised GI reconstruction method. It is inspired by DIP and N2N while addressing their limitations, namely accounting for the random noise in the optimization technique and the missing realization artifacts, respectively. The former is addressed in Sec. 2.4, and the latter in Sec. 2.5. We start from the assumption that by selecting a subset of the original realization set (i.e., the collection of all the acquired signals paired with the corresponding illumination patterns), we should obtain a degraded reconstruction due to a decrease in captured information. However, the degraded reconstruction should still represent the same signal as the reconstruction from the full realization set. This means that if we partition the realizations into subsets with equal numbers of realizations, we should obtain equivalent but different and degraded sub-reconstructions.

We partition the realizations in  $k$  splits  $k \in [1, K]$ , producing  $K$  different sub-reconstructions (i.e., a  $k$ -tuple of reconstructions). Each sub-reconstruction in the said  $k$ -tuple has unique noise characteristics compared to the other reconstructions, and it provides different information on the FoV. From the partitioned data vectors  $y_k$  we compute the sub-reconstructions  $x_k$  as follows (Fig. 2(a) and algorithm 1):

$$x_k = W_k^{\dagger}y_k = W_k^{\dagger}(b_k + \epsilon_k) \quad (5)$$

where  $W_k^{\dagger}$  is pseudo-inverse of the partition forward model  $W_k$ , and  $b_k$  is the clean data. The  $x_k$  solutions can be reconstructed with a common least-squares solver.

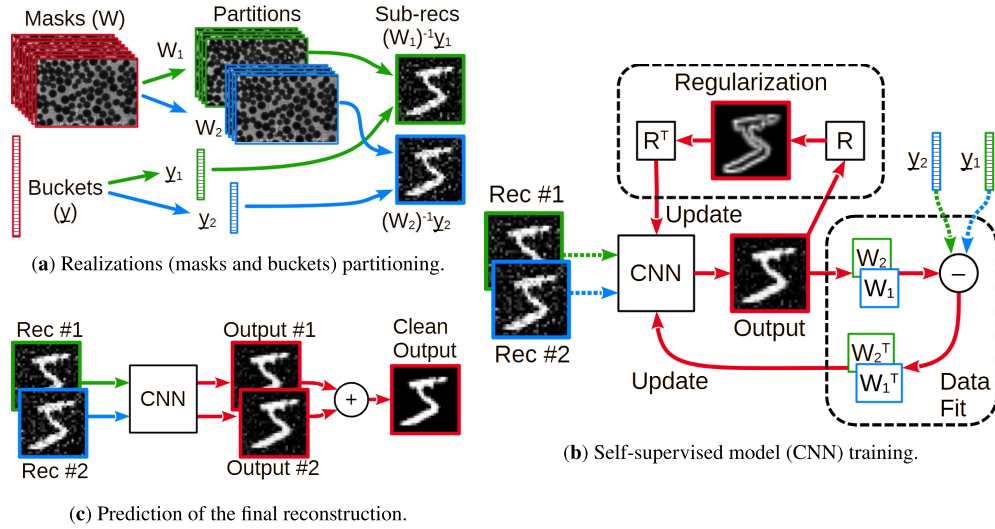
Each sub-reconstruction  $x_k$  represents the same signal but has a different noise. In N2G, we feed each of these sub-reconstructions to the model, and then we optimize the following minimization problem (Fig. 2(b) and algorithm 2) by comparing them against the realizations

**Algorithm 1. Split Realizations and Least-Squares Reconstructions**

**Require:** Realizations  $W$ , measurements  $y$ , number of splits  $K$

**Ensure:** Sub-reconstructions  $\{x_k\}_{k=1}^K$

- 1:  $s \leftarrow \text{Shuffle}(\text{Range}(1, M))$  ▷ Randomly shuffle bucket indices [1, M]
- 2: **for**  $k = 1$  to  $K$  **do**
- 3:    $s_k \leftarrow s[M(k/K) + 1, M((k+1)/K)]$  ▷ Slice  $s$
- 4:    $W_k, y_k \leftarrow W[s_k], y[s_k]$  ▷ Partition  $W$  and  $y$
- 5:    $x_k \leftarrow W_k^\dagger y_k$  ▷ Least-squares reconstruction
- 6: **end for**
- 7: **return**  $\{x_k\}_{k=1}^K, \{W_k, y_k\}_{k=1}^K$



**Fig. 2.** Schematic representation of the proposed method: (a) The partitioning of the realizations set, generating the partial reconstructions (sub-recs), (b) the training procedure, and (c) the prediction of the final reconstruction.

that were not used to reconstruct them:

$$\hat{\theta} = \arg \min_{\theta} \frac{1}{2} \sum_k \sum_{i \neq k} \|W_i N_{\theta}(x_k) - y_i\|_2^2 + \lambda R(N_{\theta}(x_k)) \quad (6a)$$

$$\hat{x} = \frac{1}{K} \sum_k N_{\hat{\theta}}(x_k). \quad (6b)$$

Since the different sub-reconstructions  $x_k$  should represent the same signal, we would like the model to produce the same result for each. Thus, the model should learn to differentiate between the actual signal and random noise. This is in opposition to Eq. (3a), where the same realizations were used for the input reconstruction and as targets in the learning loss. Finally, Fig. 2(c) presents how Eq. (6b) produces the reconstructed images

#### 2.4. Self-supervised random noise reduction

We now derive the noise-reduction properties of Eq. (6a). Reconstruction noise comes from two main sources: Measurement noise; and the lack of information from a low number of realizations, which creates the so-called *missing realization* artifacts. These noise sources can be represented

## Algorithm 2. N2G Training Loop

---

**Require:** Sub-reconstructions  $\{x_k\}_{k=1}^K$ , realizations  $\{W_k, y_k\}_{k=1}^K$ , regularization  $R$ , parameter  $\lambda$   
**Ensure:** Trained model parameters  $\hat{\theta}$

- 1: Initialize model  $N_\theta$  randomly
- 2: **for** each epoch **do**
- 3:     **for**  $k = 1$  to  $K$  **do**
- 4:         **for**  $i = 1$  to  $K, i \neq k$  **do**
- 5:              $\mathcal{L}_{k,i} = \|W_i N_\theta(x_k) - y_i\|_2^2$  ▷ Self-supervised losses
- 6:         **end for**
- 7:          $\mathcal{L}_{\text{TV},k} = R(N_\theta(x_k))$  ▷ Prior (TV) losses
- 8:          $\mathcal{L}_k \leftarrow (\frac{1}{2} \sum_{i \neq k} \mathcal{L}_{k,i}) + \lambda \mathcal{L}_{\text{TV},k}$  ▷ Sub-reconstruction losses
- 9:         **end for**
- 10:         $\theta \leftarrow \theta - \nabla_\theta(\sum_k \mathcal{L}_k)$  ▷ Update parameters
- 11:     **end for**
- 12: **return**  $\hat{\theta}$

---

in the sub-reconstructions  $x_k$  as follows:

$$x_k = x + v_k + W_k^\dagger \epsilon_k = x + v_k + t_k \quad (7a)$$

$$W_k v_k = 0 \quad (7b)$$

$$W_k x_k = y_k + \epsilon_k \quad (7c)$$

where  $W_k^\dagger \epsilon_k = t_k$  is the pseudo-inverse of the bucket measurement noise, while the vector  $v_k \subseteq \ker W_k$  describes spurious solutions that could be added to the reconstruction, and still fit the data. Our method focuses on the random measurement noise, and the desired noise reduction effect is given by the specific construct of the data fitting term from Eq. (6a). When projecting the predicted image onto the target data from the other splits, the projected noise from the  $k$  split will be inconsistent with the random noise from the other  $K - 1$  splits. This can be seen by substituting Eq. (7a) and Eq. (5) into the data fitting term of Eq. (6a):

$$\frac{1}{2} \sum_{i \neq k} \|W_i N_\theta(x + v_k + W_k^\dagger \epsilon_k) - (b + \epsilon_i)\|_2^2. \quad (8)$$

As  $\epsilon_k$  is independent of any other  $\epsilon_i$  and their domains do not overlap (Appendix A from [37]), the expected prediction error of Eq. (8) becomes:

$$\mathbb{E}_{b, \epsilon, i \neq k} \|W_i N_\theta(x + v_k + W_k^\dagger \epsilon_k) - b_i\|_2^2 + \mathbb{E}_{b, \epsilon, i \neq k} \|b_i - (b_i - \epsilon_i)\|_2^2 \quad (9a)$$

$$= \mathbb{E}_{b, i \neq k} \|W_i N_\theta(x_k) - b_i\|_2^2 + \mathbb{E}_{\epsilon, i \neq k} \|\epsilon_i\|_2^2 \quad (9b)$$

where the term  $\|W_i N_\theta(x_k) - b_i\|_2^2$  is the *supervised reconstruction error*, and the term  $\|\epsilon_i\|_2^2$  is the signal noise variance. In the supervised reconstruction term, the acquisition noise is not present. This means that it is equivalent to having noiseless realizations as the targets of our N2G reconstruction, thus providing the noise suppression characteristics of Noise2Ghost. The signal noise variance is just a constant in the objective function, which only depends on the acquisition noise and should not contribute to the objective function gradient. In reality, it may introduce some noise in the optimization gradients at very high noise levels. This could interfere with the reconstruction, guiding it towards local minima. As for N2N and derived methods, this often results in excessive blur and resolution loss in the reconstructions.

### 2.5. Relationship with Noise2Inverse

The Noise2Inverse method addresses the problem in Eq. (6) directly in the reconstruction domain [37], by solving the following problem:

$$\hat{\theta} = \operatorname{argmin}_{\theta} \frac{1}{2} \sum_k \|N_{\theta}(x_k) - \frac{1}{K-1} \sum_{i \neq k} (x_i)\|_2^2 \quad (10a)$$

$$\hat{x} = \frac{1}{K} \sum_k N_{\hat{\theta}}(x_k) \quad (10b)$$

where  $x_i = W_i^{\dagger} y_i$  is the pseudo-inverse (LS reconstruction) of the partial set of masks  $W_i$  and corresponding buckets  $y_i$  with  $i \in [1, K]$ . Theoretical grounds support this formulation, and it should retrieve the same result as Eq. (6), regarding its noise reduction aspect. However, it does not address the missing realization artifacts and suffers from strong technical limitations. Hereafter, we justify these statements by first noting that, in ghost imaging, we aim to reduce the number of realizations as much as possible, which results in faster acquisition times and potentially reduced dose deposition. This exacerbates the rise of related artifacts. Missing realization and random noise artifacts are strongly structured and have long-range correlations, i.e., they generally cover the whole FoV. Commonly used models like CNNs have a limited receptive field (i.e., the region around each pixel that the CNN probes and that provides the context for the said pixel) that may not extend to the entire FoV. Depending on the model, this could have a strong negative impact on the denoising performance.

More precisely, about the missing realization artifacts, in [37] it is stated that Eq. (10) is unable to cope with them. The LS sub-reconstructions cope with the missing realizations by acting as if their bucket values were set to 0. This selects one specific solution  $v_k$  for each  $x_k$ , out of the infinitely many solutions in  $\ker W_k$  that satisfy Eq. (1). This specific choice of  $v_k$  is probably incorrect, but is imposed by N2I as the learning target in Eq. (10). The model learns to discard the inconsistent features of different  $v_k$  functions, selecting a specific  $v \in \ker W$  as a common ground.

The formulation we propose in (Eq. (6)) circumvents this limitation by using the measured buckets as “target”. Thus, N2G does not embed any specific  $v_k$  into the learning target. This allows the model to select a  $v \in \ker W$  that satisfies both the measured buckets and the prior knowledge injected by the regularization term  $R$  and the regularization properties of CNNs [19].

A possible approach to overcome the just-discussed limitations of Eq. (10) would be to use variational methods (like in Eq. (2)) to compute the sub-reconstructions. These non-linear methods are known to remove or strongly reduce these long-range artifacts. However, as discussed in [37], their non-linear nature contrasts with the assumptions of Eq. (10). This may then result in a degraded reconstruction quality.

### 2.6. Data augmentation

Thanks to the assumption that any permutation (reordering) of the acquired realizations provides the same reconstruction, we can achieve simple and effective data augmentation: We can increase the number of training examples for the self-supervised method from the measured data. We permute our dataset in  $p$  arrangements with  $p \in [1, P]$  and partition each arrangement in  $k$  splits  $k \in [1, K]$ , producing  $KP$  different sub-reconstructions (i.e.,  $P$  k-tuples of reconstructions). In practice, the permutations  $P$  provide unique mixtures of missing realization artifacts and random noise, while the  $K$  splits are still the underlying mechanism for noise reduction. This modifies

Eq. (6) into:

$$\hat{\theta} = \operatorname{argmin}_{\theta} \frac{1}{2} \sum_{p,k} \|WN_{\theta}(x_{p,k}) - y\|_2^2 + \lambda R(N_{\theta}(x_{p,k})) \quad (11a)$$

$$\hat{x} = \frac{1}{PK} \sum_{p,k} N_{\hat{\theta}}(x_{p,k}). \quad (11b)$$

## 2.7. Models

Here, we briefly discuss the machine learning models, in particular neural networks (NN), that N2G could leverage. The model should accept multiple images of the same object as input. This is incompatible with INRs, which learn coordinate-based representations of the signal and thus take coordinates as inputs. CNNs and other MLP-based NNs are instead good candidates. Both NN types take images as input and can learn short- and long-range pixel correlations. Thus, N2G is not tied to a specific NN architecture and can take advantage of future developments in that area. MLP-based NNs examples include the Vision Transformer (ViT) [40] and the swin-transformer [41]. CNNs examples include U-Net [33], DnCNN [42], and MS-D net [43]. Regarding the model size, one should try to work with small NNs to minimize the computation time and resources and to limit overfitting. For instance, we suggest models with parameter numbers in the order of 100k for a  $100 \times 100$  pixels image. However, some models like MS-D net could achieve similar results with 10-20k parameters for similar image sizes. In our tests, we had the best results with U-Net and DnCNN with around 100-300k parameters. Further details will be provided with the results (Sec. 3.2).

## 2.8. Computational performance & requirements

N2G presents specific computational and memory demands that require to be considered. It requires a  $KP$ -fold increase in operations, compared to GIDC (e.g.,  $\times 24$  for the tests performed in Sec. 3). This does not correspond to a proportional increase in computational time, thanks to batching and GPU parallelism, as detailed in the [Supplement 1 \(SI\)](#), Sec. S6. The actual increase in reconstruction time is approximately 2.1 to 2.5 times. On the other hand, memory usage increases proportionally, requiring  $\times KP$  activation weights during backpropagation. This is critical for GPUs, where memory constraints are typically more stringent compared to CPU RAM. The increased memory requirements can limit the size of images and the extent of noise reduction that can be processed with N2G. To mitigate this, gradient accumulation can be used, where backpropagation weights are accumulated over smaller batches of model passes. While this approach can impact reconstruction performance, it effectively allows for the processing of larger reconstruction sizes. Looking ahead, this strategy could enable implementing multi-GPU reconstructions, further enhancing the capabilities of N2G. Furthermore, adjustments of “activation checkpointing” could further yield reductions in GPU memory consumption.

Note that, due to the intrinsic iterative nature of the methods discussed here, their use in real-time applications should be considered relative to the acquisition time. As presented in Sec. S6 of the SI, reconstructions can take from several seconds to several minutes. This is incompatible with high-throughput acquisitions, where full datasets are produced short time-frames of less than a second or millisecond. However, given the sequential nature of GI acquisitions, this scenario is unlikely. In particular, for x-ray GI, the current acquisition time is on the order of a few hours per image [39]. In this context, N2G’s increased computational time requirements over the existing methods are negligible, even when performing a full hyperparameter grid search.

## 3. Results

We now analyze N2G’s performance against existing unsupervised methods for reconstructing both synthetic and real data. The reference methods are LS, TV, the GIDC algorithm from [14]

(based on DIP), and INRs. For the synthetic test cases, we first derive a method to modulate the noise and compare different noise levels. Additionally, for the synthetic study cases, we also compare the results against N2V and N2I, two common self-supervised denoising methods.

### 3.1. Synthetic data generation & quantification

We first evaluate the performance of the proposed method on synthetic data, as it provides the ground truth and precisely allows us to control the noise added to the data. In Sec. S7 of the SI, we present and analyze additional synthetic test cases, with increasing image complexity.

We extracted a  $100 \times 100$  pixels image to be our ground truth (phantom, e.g., from Figs. 6(h) and 7(h)) from [44], a curated collection of chromosome images. We generated synthetic random masks with a normalized half-Gaussian distribution and forward projected the phantom to create the corresponding buckets. We added Poisson noise to the buckets by first multiplying them by the maximum expected emitted number of photons per pixel. We fed the resulting mean expected photons per realizations to the Poisson distribution function. We then normalized the signals by dividing the noisy buckets by the same number of expected photons, resulting in signals:

$$y_m = \mathcal{P} \left( C \sum_n^N w_{mn} x_n \right) / C. \quad (12)$$

This expression is derived in Sec. S3 of the SI for the specific case of XRF imaging. We use the constant  $C$  to modulate the noise and evaluate its effect on the reconstruction performance of the proposed method. The advantage of this noise model is that it gives emphasis to the number of photons delivered to the sample.


We evaluate the reconstruction algorithms both visually and with multiple objective metrics: the mean square error (MSE), the peak signal-to-noise ratio (PSNR), the structural similarity index (SSIM), and the resolution, which is estimated from the Fourier ring correlation (FRC). Their formulas are provided in Sec. S1 of the SI, while here we shortly describe their meaning. The MSE evaluates the average pixel-wise difference between the phantom and the reconstruction intensity values. The SSIM calculates the similarity between these two images by comparing the inter-dependencies of spatially neighboring pixels over various window sizes [29]. PSNR is based on MSE, but it focuses on the peak intensity power instead of its average. Finally, the FRC computes the similarity in Fourier space between the phantom and the reconstruction for all the sampled frequencies. We estimate the resolution of the reconstruction by finding the intersection of the resulting curve with the threshold function defined in [45]. However, the meaning of the FRC goes beyond the simple resolution estimation, and higher areas under the FRC curve signify higher reconstruction quality even in the case of similar resolutions.

### 3.2. Synthetic data reconstructions across various noise levels

We present three different study cases: In the first case (Fig. 3), we present the performance of the above-mentioned different algorithms (LS, TV, GIDC, INR, and N2G) across multiple noise levels (i.e., photons per pixel) using only 2000 realizations, corresponding to 20% of the total number of pixels. The simulated maximum emitted photons per pixel per realization (i.e., constant  $C$ ) intensities span the range of  $[10^0, 10^4]$ . The lower margin means that each pixel is illuminated on average by around one photon. The upper margin was empirically chosen well above the threshold where noise is not more the limiting factor of the reconstruction quality (i.e., it becomes the missing realizations). To add some statistical significance to our tests, we generated five different sets of realizations for each noise level. We then independently reconstruct these five sets and compute their individual evaluation metrics (e.g., mean error against the phantom, etc). The presented results report the average and standard deviation of the said computed metrics for each noise level.

For the TV, GIDC, INR, and N2G reconstructions, the regularization weight  $\lambda$  was chosen by cross-validation, by putting aside a set of realizations (10% of the total) and using them

to assess the best value over a wide range [39]. For TV, we use a  $\lambda \in [10^{-5}, 10^{-1}]$ , while for GIDC, INR, and N2G we used a  $\lambda \in [10^{-7}, 10^{-4}]$ . This difference in range is due to the different data normalization used by the different algorithm implementations. As regularization  $R$  for GIDC, INR and N2G, we also used the TV term. Although stronger and more advanced priors could have been used, we wanted a direct comparison against the well-known and used TV-min-based variational regularization. Moreover, the TV term has very simple and widely understood implementation and behavior. To improve the procedure's robustness, we ran it three times with different cross-validation sets for each run and averaged the cross-validation loss values between those different runs. For the INR reconstructions, we used a 3-layer feed-forward NN with 512 neurons per layer, sinusoidal activation functions [36], and 256 Fourier embeddings [46]. Regarding specifically GIDC and N2G, we used a U-net [33] with 20 features and 3 decomposition levels ( $\sim 192k$  parameters). We trained the U-net and INR models with 5000 and 7000 epochs of the Adam optimizer [47], respectively, with a learning rate of  $3 \cdot 10^{-4}$  and weight decay of  $10^{-2}$  for both architectures. We selected the model weights corresponding to the epoch in which the cross-validation loss function is minimized. For N2G, we used 4 partitions and 6 permutations ( $K = 4$ , and  $P = 6$ ). For a more detailed discussion of the choice of the  $K$  and  $P$  parameters, we refer to Sec. S5 of the SI.




**Fig. 3.** Synthetic data GI reconstructions comparison (chromosomes phantom), for  $5\times$  compression ratio, and maximum emitted photons per pixel per realization in the range  $[10^0, 10^4]$ . We compare peak signal-to-noise ratio (PSNR), structural similarity index (SSIM), and bandwidths (Nyquist frequency) against the ground truth. Higher values indicate higher reconstruction quality. We also show the corresponding image resolutions in pixels.

The images returned by all the reconstruction methods are scaled and shifted (i.e., adjusted reconstruction mean) to quantitatively fit the acquired bucket data, following the procedure from Sec S3 of the SI. In other words, we adjust the zero-frequency and frequency power of each reconstruction to remove the influence of data normalization.

On the top left side of Fig. 3, we present the peak signal-to-noise ratio (PSNR) of the reconstructions against the phantom. It is evident that N2G consistently achieves superior reconstructions, with PSNRs surpassing those of competing methods by 1–2 dB, highlighting its robustness across diverse imaging conditions. Nonetheless, its advantage seems to diminish under low-photon scenarios. While counterintuitive, this is easily explained by the fact that the lower frequencies are less affected by Poisson noise. PSNR takes all frequencies into account, resulting in a bias towards lower frequencies with low photon counts. The SSIM, by construction, is instead less affected by lower frequencies. For this reason, on the top right side of Fig. 3, we see that N2G has a growing lead in SSIM for lower photon counts. We also observe that for high photon numbers (equating to higher signal-to-noise ratio), the compression ratio becomes the limiting factor for reconstruction accuracy, and both PSNR and SSIM reach a plateau. From Fig. 3, we have instead a perspective on how the resolution is degraded with increasing noise levels. We observe that N2G better preserves the reconstruction resolution, by showing a drop-off at lower photon counts than other methods. To complement and better explain the results from Fig. 3, Fig. 4 presents the bias and variance of the reconstruction methods. As mentioned in

Sec. S1 for the SI, the reconstruction bias and variance are two components of MSE, and the reconstruction bias metric in particular describes the intrinsic behavior of the reconstruction method with respect to noisy features. While the reconstruction variance of N2G is most of the times equal or better than the competing methods, the reconstruction bias is significantly lower. This should be attributed to the fact that N2G puts a lower emphasis on the regularization for dealing with noise, thanks to its inherent noise reduction mechanics (as seen in Sec. 2.4).



**Fig. 4.** Reconstruction bias and variance comparison for the chromosomes phantom, with  $5\times$  compression ratio, and maximum emitted photons per pixel per realization in the range  $[10^0, 10^4]$ . Lower values indicate higher reconstruction quality.

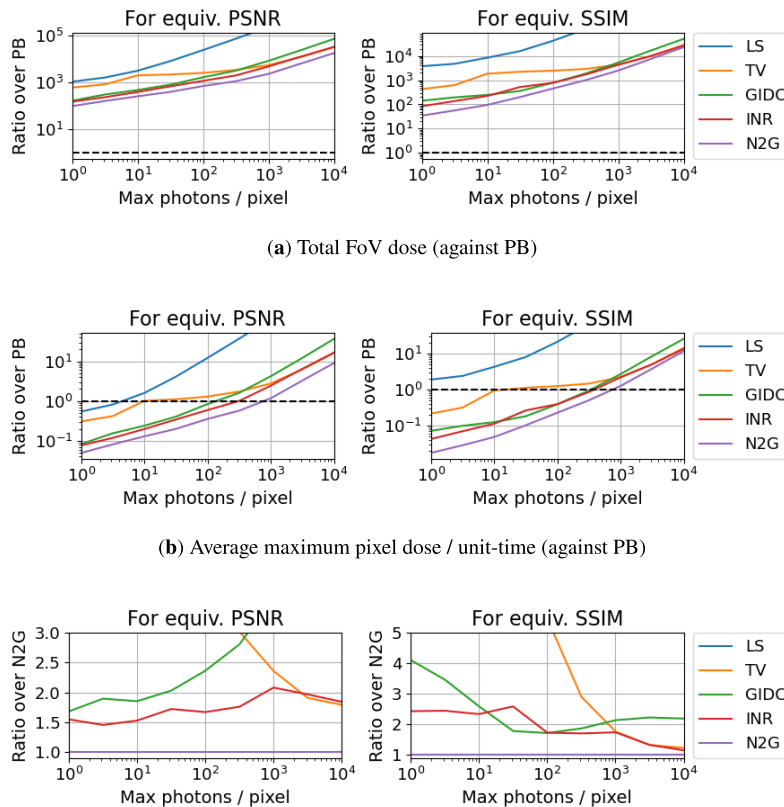
Finally, we explore dose considerations by comparing GI with PB and assess N2G's performance against LS, TV, GIDC, and INR. In particular, we evaluate the total sample and average maximum pixel dose necessary to achieve a GI image with PSNR or SSIM quality comparable to a PB scan. The procedure is presented in Sec. S4 of the SI. From figure Fig. 5(a), we observe that in terms of total dose, PB always outperforms GI in the chosen compression ratio, regardless of the method used, although N2G narrows the difference. Concerning the average maximum pixel dose per exposure (Fig. 5(b)), GI generally exceeds PB's performance when the main noise source is detection noise (except when LS is used). However, N2G still decreases dose requirements compared to the other GI reconstruction techniques and widens the region where GI is preferable over PB. Figure 5(c) clearly shows this trend by plotting the required dose ratio from LS, TV, GIDC, and INR to obtain the same performance as N2G.

### 3.3. In-depth analysis of noisy reconstructions

We now provide an in-depth analysis of two examples: (a) high compression ratio and moderate noise and (b) low compression ratio and very high noise (higher than signal).

We present the results for (a) in Fig. 6. We produced 1000 realizations, resulting in a  $10\times$  compression ratio. The applied Poisson noise of  $10^2$  maximum emitted photons per pixel per realization (i.e., constant  $C = 10^2$ ) resulted in fluctuations equal to 24.5% of the mean clean bucket fluctuations. Comparing reconstructions from Figs. 6(a) to 6(g) against Fig. 6(h), we note that except for the LS reconstruction, all other algorithms, including N2G, can correctly reconstruct the shape of the phantom. The only visible difference is that N2G's reconstruction presents much less high-frequency noise. This results in the N2G reconstruction having the highest resolution (in pixels), PSNR, and SSIM values, and the lowest MSE among all reconstructions. These results are summarized in Fig. 6(j).

We present the results of (b) in Fig. 7. Here, we produced 3333 realizations, resulting in a  $3\times$  compression ratio. The applied Poisson noise of 2 maximum emitted photons per pixel per realization (i.e., constant  $C = 2$ ) resulted in fluctuations equal to  $\sim 164\%$  of the mean clean bucket fluctuations. All the reconstructed images are noticeably affected by noise with a larger bandwidth than in the previous examples. N2G can suppress high-frequency noise, but cannot completely correct lower-frequency noise. Despite that, by looking at the FRC in Fig. 7(i), we notice that it preserves much more bandwidth than the other algorithms and presents a much better resolution. This again results in the N2G reconstruction having the highest PSNR and



**Fig. 5.** Dose requirements for the chromosomes phantom, with  $5\times$  compression ratio, and varying noise levels: (a) total required dose for each algorithm to obtain equivalent PSNR and SSIM as a PB acquisition; (b) same plot as (a) for the average maximum pixel dose per illumination; and (c) for N2G.

SSIM values, and the lowest MSE among all reconstructions. These results are summarized in Fig. 7(j).



As a final remark, we observe that in both the presented cases, N2G tends to provide even smoother backgrounds than the phantom itself. The overly smoothed features are usually high-frequency and low-signal features, which are much more overwhelmed by noise than the high-signal parts of the image. As we discussed in Sec. 2, the suppression of high-frequency noise can result in oversmoothing in these cases.

### 3.4. Demonstration with experimental data

Finally, we compare the reconstruction performance of N2G on real GI data acquired at the beamline ID19 of the ESRF — The European Synchrotron [48]. More details on the setup are provided in Sec. S10 of the SI. The considered dataset, available at Dataset 1, Ref. [48], consists of  $42 \times 87$  pixels images, with  $24\mu\text{m}$  pixel size, and 896 realizations, where the buckets are the emitted XRF signals from a sample composed of a glass capillary and three wires. Two of the three wires are made of Cu, and they will be the focus of this reconstruction. We used the signal from their  $K_\alpha$  emission line. At each realization, we exposed the sample to a 26 keV incoming X-ray beam for 0.1 seconds and acquired each realization 32 times. The accumulated 3.2-second exposures serve as high-SNR buckets. Reducing the number of accumulated exposures artificially reduces the collected number of photons (photon flux), thus simulating lower dose depositions.

The XRF signals are collected with a single-pixel hyper-spectral detector, whose spectral output is discretized into 4096 bins of 150 eV energy steps. Each XRF peak has a Lorentzian shape around the mean energy  $E_{e,l}$ . This means that we observe photon counts in adjacent bins around the expected mean signal energy, which decrease  $\sim \gamma_{e,l}/((E_{e,l} - E)^2 + \gamma_{e,l}^2)$  for bins at energy  $E$  further away from it, where  $\gamma_{e,l}$  is a constant that depends on the emission line. These neighboring bins are either summed, or the area under their curve is fitted to increase the signal's SNR. Therefore, we can simulate further dose reductions by summing fewer and lower-count bins. This is especially important if we consider that at the moment there exists only one dataset of nano-scale XRF-GI [15]. Our dose reduction simulation technique allows us to emulate cases of much lower fluxes such as nano-scale acquisitions.

We present the results of the XRF-GI reconstructions in Fig. 8. Figure 8(b) presents an X-ray transmission image (radiograph) of the sample, while Fig. 8(c) indicates the regions that give rise to the XRF signal of interest. As mentioned above, we consider the 0.1-second exposures as the “100% dose” exposures, while the 3.2-second accumulated images serve as ground truth (high-quality reference images). The other two images in Fig. 8(a) were produced using the five highest-counts and the single highest-counts XRF bins, which provide an average of 12.9% and 2.6% of the reference dose per realization, respectively. These two reduced-flux datasets could thus correspond to 12.9 ms and 2.6 ms exposures per realization, respectively. The corresponding average photons per bucket and per bucket & pixel can be found in Fig. 8(d).





(i) FRC

	LS	N2V	N2I	TV	GIDC	INR	N2G
MSE ↓	0.400	0.021	0.021	0.004	0.004	0.002	<b>0.001</b>
PSNR ↑	2.50	15.30	15.27	22.44	22.70	25.00	<b>26.89</b>
SSIM ↑	0.030	0.235	0.246	0.656	0.672	0.756	<b>0.827</b>
Resolution ↓	9.6	12.6	8.3	4.4	5.2	4.3	<b>4.0</b>

(j) Performance comparison.

**Fig. 6.** GI reconstructions of the chromosomes phantom, with 10× compression ratio and moderate Poisson noise: Mean noise fluctuations  $\sim 24.5\%$  of the mean clean bucket fluctuations. From (a) to (g) we show the reconstructions of (h) with LS, TV, GIDC, INR and N2G respectively. In (i) and in (j) we present the Fourier ring correlation against (h) and various performance metrics of each reconstruction, respectively.



(i) FRC

	LS	N2V	N2I	TV	GIDC	INR	N2G
MSE ↓	0.150	0.040	0.042	0.047	0.013	0.012	<b>0.006</b>
PSNR ↑	6.75	12.53	12.27	11.81	17.29	17.79	<b>20.58</b>
SSIM ↑	0.034	0.122	0.142	0.186	0.337	0.380	<b>0.693</b>
Resolution ↓	12.1	19.7	11.5	14.1	10.5	10.3	<b>4.3</b>

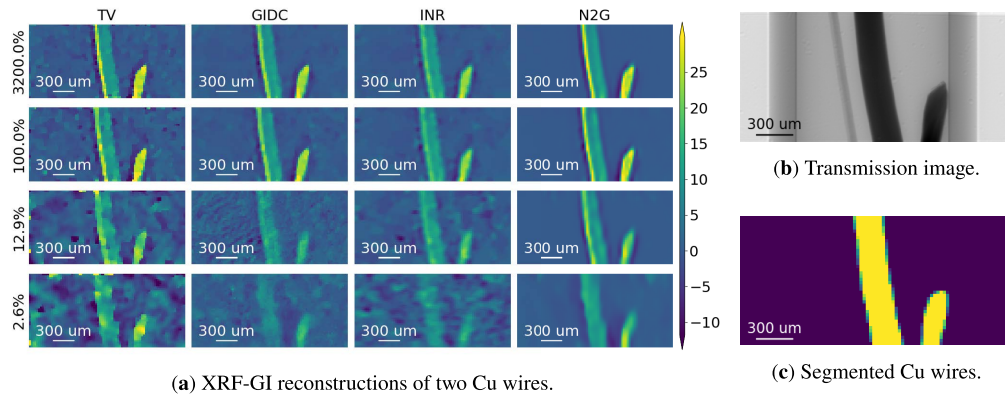
(j) Performance comparison.

**Fig. 7.** GI reconstructions of the chromosomes phantom, with  $3\times$  compression ratio and very high Poisson noise: Mean noise fluctuations  $\sim 164\%$  of the mean clean bucket fluctuations. From (a) to (g) we show the reconstructions of (h) with LS, TV, GIDC, INR and N2G respectively. In (i) and in (j) we present the Fourier ring correlation against (h) and various performance metrics of each reconstruction, respectively.

The table from Fig. 8(e) presents the estimated image quality degradation with increasing noise levels, according to PSNR and SSIM. Each result is compared against its own method-related “reference” reconstruction because we lack a PB scan of the object, which would serve as shared ground truth. Despite that, N2G provides the highest noise reduction according to PSNR and SSIM across all the reconstruction methods considered.

#### 4. Conclusions & outlook

In this article, we present a self-supervised deep-learning GI reconstruction method specifically targeted at dealing with random noise in the acquired data. If provided with enough high-quality reference data on the studied samples, existing supervised approaches might exhibit high noise suppression performance, with possibly higher PSNR and/or resolution than our approach. However, gathering high-quality reference data can be expensive and requires extensive and organized efforts, thus sometimes proving to be a difficult or even impossible task. This particularly applies to cutting-edge and niche applications in micro- and nano-imaging, dealing with rare (or even unique) samples with high radiation-dose sensitivity (e.g., biological specimens, battery cells) in in-vivo/in-operando conditions. In those cases, high-quality reference data are usually simply unavailable. Our approach solves this problem with self-supervision, resulting in higher reconstruction performance against the state-of-the-art unsupervised approaches across



(a) XRF-GI reconstructions of two Cu wires.

(c) Segmented Cu wires.

	PSNR			SSIM		
	100%	12.9%	2.6%	100%	12.9%	2.6%
Avg. photons	32000%	100%	12.9%	2.6%		
Bucket	2.07e+05	6.47e+03	8.38e+02	1.69e+02		
Bucket & pixel	3.48	0.109	0.0141	0.00283		
TV	30.24	24.34	21.19	0.680	0.319	0.192
GIDC	32.76	23.23	21.51	0.721	0.315	0.217
INR	33.14	24.73	19.43	0.771	0.389	0.182
N2G	<b>35.33</b>	<b>27.75</b>	<b>23.76</b>	<b>0.933</b>	<b>0.796</b>	<b>0.501</b>

(d) Photon counts of each acquisition.

(e) PSNR and SSIM of the GI reconstructions.

**Fig. 8.** Real data XRF-GI reconstructions: (a) reconstructions at different photon counts (and corresponding noise levels); (b) and (c) an x ray transmission image of the sample at 26 keV and the segmentation of the Cu wires, respectively; (d) the photoncounts of each acquisition; and (e) PSNR and SSIM of each reconstruction against their high-SNR versions in the top row of (a).

multiple metrics like MSE, PSNR, SSIM, and image resolution. From both the synthetic test cases and the real data reconstructions, we observe that our method preserves higher reconstruction quality than existing unsupervised methods as the noise increases, i.e., at faster acquisition speeds and/or lower deposited dose, across a suite of case studies. This translates into the following relevant aspects: For the same image quality, we could be looking at a reduction in the acquisition time and/or delivered dose by a factor of  $\sim 1.5 - 2$ , compared to the best existing unsupervised methods, when working with the lower deposited doses; and the range of deposited doses where GI offers an advantage over PB is also expanded.

The above-mentioned aspects also mean that our method could enable operating GI scans at  $\sim 1.5 - 2$  shorter dwell time, resulting in faster GI acquisitions. This aspect is crucial for real case scenarios in two alternative ways: Either increasing acquisition stability or the ability to acquire more realizations for the same dose. In the first case, when paired with the reduced number of GI realizations compared to PB acquisitions, we further reduce the total acquisition times of GI acquisitions, strongly mitigating the impact of sample drifts in our GI scans. Moreover, compared to PB scans, where the sample moves across the field of view, in a GI acquisition, the sample is fixed [39]. This further consolidates GI as a very attractive technique for reducing sample drifts and positioning errors in nano-scale imaging. Alternatively, the reduced exposure time for each GI realization would also enable more realizations to be acquired within a fixed acquisition time budget. This, in turn, is the only way to reduce the missing-realization noise, which is the ultimate limit in GI reconstruction quality, as seen earlier in the text.

N2G focuses on the reduction of noise from acquired measurements. It delegates dealing with the missing realization artifacts to the generalization power of the underlying model (NN) and the rather simplistic prior knowledge given by the TV term. Future work could tackle this problem with a combination of a few possible approaches, including using stronger or more complex priors than TV (e.g., multi-level undecimated wavelet minimization) and multi-channel/multi-modal information from correlated signals (e.g., the transmission signal in [15]), as demonstrated in computed tomography [49].

**Funding.** Israel Science Foundation (2208/24); Pazy Foundation (49579).

**Disclosures.** The authors declare no conflicts of interest.

**Data availability.** Data underlying the experimental results are available in [Dataset 1](#), Ref. [48]. The raw data is available in Ref. [50].

**Supplemental document.** See [Supplement 1](#) for supporting content.

## References

1. T. Tian, S. Oh, and C. Spielmann, "Multi-wavelength ghost imaging: a review," *Vicinagearth* **2**(1), 4 (2025).
2. T. B. Pittman, Y. H. Shih, D. V. Strekalov, *et al.*, "Optical imaging by means of two-photon quantum entanglement," *Phys. Rev. A* **52**(5), R3429–R3432 (1995).
3. A. Kirsch, *An Introduction to the Mathematical Theory of Inverse Problems*, vol. 120, of *Applied Mathematical Sciences* (Springer International Publishing, Cham, 2021).
4. O. Katz, Y. Bromberg, and Y. Silberberg, "Compressive ghost imaging," *Appl. Phys. Lett.* **95**(13), 4192–4200 (2009).
5. D. Pelliccia, A. Rack, M. Scheel, *et al.*, "Experimental X-Ray Ghost Imaging," *Phys. Rev. Lett.* **117**(11), 113902 (2016).
6. A. M. Kingston, G. R. Myers, D. Pelliccia, *et al.*, "Neutron ghost imaging," *Phys. Rev. A* **101**(5), 053844 (2020).
7. S. Li, F. Cropp, K. Kabra, *et al.*, "Electron Ghost Imaging," *Phys. Rev. Lett.* **121**(11), 114801 (2018).
8. R. I. Khakimov, B. M. Henson, D. K. Shin, *et al.*, "Ghost imaging with atoms," *Nature* **540**(7631), 100–103 (2016).
9. S. L. Brunton and J. N. Kutz, *Data-Driven Science and Engineering* (Cambridge University Press, 2019).
10. T. J. Lane and D. Ratner, "What are the advantages of ghost imaging? Multiplexing for x-ray and electron imaging," *Opt. Express* **28**(5), 5898 (2020).
11. M. Karuppasamy, F. Karimi Nejadasl, M. Vulovic, *et al.*, "Radiation damage in single-particle cryo-electron microscopy: effects of dose and dose rate," *J. Synchrotron Radiat.* **18**(3), 398–412 (2011).
12. L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D* **60**(1–4), 259–268 (1992).
13. S. Rizvi, J. Cao, K. Zhang, *et al.*, "DeepGhost: real-time computational ghost imaging via deep learning," *Sci. Rep.* **10**(1), 11400 (2020).
14. F. Wang, C. Wang, M. Chen, *et al.*, "Far-field super-resolution ghost imaging with a deep neural network constraint," *Light:Sci. Appl.* **11**(1), 1 (2022).
15. J. Li, S. Chen, D. Ratner, *et al.*, "Nanoscale chemical imaging with structured X-ray illumination," *Proc. Natl. Acad. Sci.* **120**(49), e2314542120 (2023).
16. <https://github.com/CEA-MetroCarac/noise2ghost>.
17. F. Ferri, D. Magatti, L. A. Lugiato, *et al.*, "Differential Ghost Imaging," *Phys. Rev. Lett.* **104**(25), 253603 (2010).
18. B. Sun, S. S. Welsh, M. P. Edgar, *et al.*, "Normalized ghost imaging," *Opt. Express* **20**(15), 16892 (2012).
19. V. Lempitsky, A. Vedaldi, and D. Ulyanov, "Deep Image Prior," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, (IEEE, 2018), pp. 9446–9454.
20. J. Lehtinen, J. Munkberg, J. Hasselgren, *et al.*, "Noise2Noise: Learning Image Restoration without Clean Data," in *Proceedings of the 35th International Conference on Machine Learning*, vol. 80, of *Proceedings of Machine Learning Research* J. Dy and A. Krause, eds. (PMLR, 2018), pp. 2965–2974.
21. A. Krull, T.-O. Buchholz, and F. Jug, "Noise2Void - learning denoising from single noisy images," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, (2019), pp. 2129–2137.
22. J. Lequyer, R. Philip, A. Sharma, *et al.*, "A fast blind zero-shot denoiser," *Nature Machine Intelligence* **4**(11), 953–963 (2022).
23. T. Huang, S. Li, X. Jia, *et al.*, "Neighbor2Neighbor: A Self-Supervised Framework for Deep Image Denoising," *IEEE Trans. on Image Process.* **31**, 4023–4038 (2022).
24. J. Li, L. Wei, and X. Zhao, "Self-supervised denoising for enhanced volumetric reconstruction and signal interpretation in two-photon microscopy," *Photonics Res.* **13**(8), 2418 (2025).
25. W. Wu, Y. Long, Z. Gao, *et al.*, "Multi-Level Noise Sampling From Single Image for Low-Dose Tomography Reconstruction," *IEEE J. Biomed. Health Inform.* **29**(2), 1256–1268 (2025).
26. S. Liu, X. Meng, Y. Yin, *et al.*, "Computational ghost imaging based on an untrained neural network," *Optics and Lasers in Engineering* **147**, 106744 (2021).
27. J. Li, B. Wu, T. Liu, *et al.*, "URNet: High-quality single-pixel imaging with untrained reconstruction network," *Optics and Lasers in Engineering* **166**, 107580 (2023).

28. K. Liu, L. Yin, T. Liu, *et al.*, “A super-resolution algorithm of Ghost Imaging using CNN with Grouped orthonormalization algorithm Constraint,” *Opt. Laser Technol.* **181**, 111847 (2025).
29. Z. Wang, A. Bovik, H. Sheikh, *et al.*, “Image Quality Assessment: From Error Visibility to Structural Similarity,” *IEEE Trans. on Image Process.* **13**(4), 600–612 (2004).
30. G. Lei, W. Lai, H. Jia, *et al.*, “Low-sampling and noise-robust single-pixel imaging based on the untrained attention U-Net,” *Opt. Express* **32**(17), 29678 (2024).
31. Y. Chen, H. An, Z. Sun, *et al.*, “Large model enhanced computational ghost imaging,” *Sci. China Technol. Sci.* **68**(11), 2120403 (2025).
32. X. Chang, Z. Wu, D. Li, *et al.*, “Self-supervised learning for single-pixel imaging via dual-domain constraints,” *Opt. Lett.* **48**(7), 1566 (2023).
33. O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional Networks for Biomedical Image Segmentation,” in *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2015*, (2015), pp. 234–241.
34. A. Habring and M. Holler, “Neural-network-based regularization methods for inverse problems in imaging,” *GAMM-Mitteilungen* **47**(4), e202470004 (2024).
35. A. Qayyum, I. Ilahi, F. Shamshad, *et al.*, “Untrained Neural Network Priors for Inverse Imaging Problems: A Survey,” *IEEE Trans. Pattern Anal. Mach. Intell.* **45**(5), 1–20 (2022).
36. V. Sitzmann, J. N. P. Martel, A. W. Bergman, *et al.*, “ in Implicit Neural Representations with Periodic Activation Functions,” *Proc. NeurIPS*, (2020).
37. A. A. Hendriksen, D. M. Pelt, and K. J. Batenburg, “Noise2Inverse: Self-Supervised Deep Convolutional Denoising for Tomography,” *IEEE Trans. Comput. Imaging* **6**, 1320–1335 (2020).
38. S. Oh, T. Tian, Z. Sun, *et al.*, “Efficient high-resolution microscopic ghost imaging via sequenced speckle illumination and deep learning from a single noisy image,” *Adv. Opt. Technol.* **14** (2025).
39. M. Manni, A. Ben Yehuda, Y. Klein, *et al.*, “Synchrotron-based X-ray Fluorescence Ghost Imaging,” *Opt. Lett.* **48**(23), 6271–6274 (2023).
40. A. Dosovitskiy, L. Beyer, A. Kolesnikov, *et al.*, “An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale,” (2021).
41. Z. Liu, Y. Lin, Y. Cao, *et al.*, “Swin Transformer: Hierarchical Vision Transformer using Shifted Windows,” in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, (IEEE, 2021), pp. 9992–10002.
42. K. Zhang, W. Zuo, Y. Chen, *et al.*, “Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising,” *IEEE Trans. on Image Process.* **26**(7), 3142–3155 (2017).
43. D. M. Pelt and J. A. Sethian, “A mixed-scale dense convolutional neural network for image analysis,” *Proc. Natl. Acad. Sci.* **115**(2), 254–259 (2018).
44. J.-J. Tseng, C.-H. Lu, J.-Z. Li, *et al.*, “An Open Dataset of Annotated Metaphase Cell Images for Chromosome Identification,” *Sci. Data* **10**(1), 104 (2023).
45. M. van Heel and M. Schatz, “Fourier shell correlation threshold criteria,” *J. Struct. Biol.* **151**(3), 250–262 (2005).
46. M. Tancik, P. Srinivasan, B. Mildenhall, *et al.*, “Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains,” in *Advances in Neural Information Processing Systems*, vol. 33, H. Larochelle, M. Ranzato, R. Hadsell, *et al.*, eds. (Curran Associates, Inc., 2020), pp. 7537–7547.
47. D. P. Kingma and J. L. Ba, “Adam: A method for stochastic optimization,” in *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, Y. Bengio and Y. Lecun, eds. (2015).
48. M. Manni and N. Vigano, “X-ray Fluorescence Ghost Imaging - CuSn mask - Three Wires (Fe & Cu),” Zenodo 2023 <https://zenodo.org/records/7828494>.
49. D. S. Rigie and P. J. La Rivière, “Joint reconstruction of multi-channel, spectral CT data via constrained total nuclear variation minimization,” *Phys. Med. Biol.* **60**(5), 1741–1762 (2015).
50. M. Manni, S. Schwartz, and N. Vigano, “X-ray Fluorescence Ghost Imaging - CuSn mask - Three Wires (Fe Cu),” European Synchrotron Radiation Facility 2026 <https://doi.esrf.fr/10.15151/ESRF-DC-2440457455>.