

Noise-robust latent vector reconstruction in ptychography using deep generative models

Seifert, Jacob; Shao, Yifeng; Mosk, Allard P.

DOI

[10.1364/OE.513556](https://doi.org/10.1364/OE.513556)

Publication date

2024

Document Version

Final published version

Published in

Optics Express

Citation (APA)

Seifert, J., Shao, Y., & Mosk, A. P. (2024). Noise-robust latent vector reconstruction in ptychography using deep generative models. *Optics Express*, 32(1), 1020-1033. <https://doi.org/10.1364/OE.513556>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.



Noise-robust latent vector reconstruction in ptychography using deep generative models

JACOB SEIFERT,^{1,*}  YIFENG SHAO,²  AND ALLARD P. MOSK¹ 

¹*Nanophotonics, Debye Institute for Nanomaterials Science and Center for Extreme Matter and Emergent Phenomena, Utrecht University, P.O. Box 80000, 3508 TA Utrecht, The Netherlands*

²*Optics Research Group, Imaging Physics Department, Delft University of Technology, The Netherlands*
**j.seifert@uu.nl*

Abstract: Computational imaging is increasingly vital for a broad spectrum of applications, ranging from biological to material sciences. This includes applications where the object is known and sufficiently sparse, allowing it to be described with a reduced number of parameters. When no explicit parameterization is available, a deep generative model can be trained to represent an object in a low-dimensional latent space. In this paper, we harness this dimensionality reduction capability of autoencoders to search for the object solution within the latent space rather than the object space. We demonstrate what we believe to be a novel approach to ptychographic image reconstruction by integrating a deep generative model obtained from a pre-trained autoencoder within an automatic differentiation ptychography (ADP) framework. This approach enables the retrieval of objects from highly ill-posed diffraction patterns, offering an effective method for noise-robust latent vector reconstruction in ptychography. Moreover, the mapping into a low-dimensional latent space allows us to visualize the optimization landscape, which provides insight into the convexity and convergence behavior of the inverse problem. With this work, we aim to facilitate new applications for sparse computational imaging such as when low radiation doses or rapid reconstructions are essential.

© 2023 Optica Publishing Group under the terms of the [Optica Open Access Publishing Agreement](#)

1. Introduction

Obtaining clear and accurate images under noisy conditions is paramount in many scientific and industrial applications. Whether in medical diagnostics, materials analysis, or semiconductor inspection, noise-robust imaging techniques can mean the difference between precise understanding and potential misinterpretation. Typical noise removal techniques, like median filtering [1], anisotropic diffusion [2,3] and BM3D [4], have been shown to be effective and applicable across various domains, but are usually limited to image postprocessing and restoration. In the field of computational imaging, where an image is algorithmically recovered from noisy intensity measurements, an interesting option becomes available: noise-robustness can be intrinsically included in the process of solving the inverse problem. For example, techniques such as accurate modeling of the underlying noise statistics [5,6], sparse modeling [7–9], deep denoiser priors [10–12], and regularization by denoising [13–17] have been explored in this context.

Ptychography, a computational imaging method, has seen exponential growth in the number of related publications in recent years [18]. The origin of ptychography dates back to Hoppe's 1969 work, where they introduced a method for phase retrieval from electron diffraction interference [19]. This foundational concept was refined and first named "ptychography" in a subsequent paper the following year [20]. The process involves illuminating a thin object with a localized and coherent beam. The illumination field is diffracted by the object and propagates in free space to form a diffraction pattern on a camera sensor [21]. By laterally translating the object or illumination field to overlapping regions, the object's phase and amplitude can be retrieved using iterative algorithms [22]. Variations of this method include Fourier ptychography (FP), which uses a microscope objective to collect diffracted light and computationally synthesizes an image

in the spatial domain [23,24]. Interestingly, reciprocity relations allow for the conversion of acquired data between both modalities, opening doors for the integration and mutual enhancement of various reconstruction algorithms [25]. In the past decade, numerous algorithmic extensions have been developed to leverage redundancy in ptychographic measurements, enhancing image quality by recovering parameters such as the illumination field [26–28], scanning position errors [29,30], object-camera distance [31], and multiple incoherent modes [32,33].

The inherent data redundancy in ptychography also provides a unique opportunity for data-driven reconstruction techniques such as Automatic Differentiation Ptychography (ADP), which aims to optimize a loss function using gradient-descent and differentiable modeling [34–37]. This loss function is derived from the intensity prediction of a physics-based forward model, actual data, and regularization terms. Utilizing automatic differentiation (AD), this approach enables the simultaneous and joint reconstruction of multiple relevant parameters [38]. ADP enhances the reconstruction process by offering portability, flexibility, and adaptability to changes in the forward model. This includes fusing multiple camera sensors [39], adjusting the loss function for mixed noise statistics [40], or tailoring a maximum-information illumination scheme [41]. An intriguing extension of data-driven image retrieval involves integrating deep neural networks (DNNs) with the reconstruction algorithm. One approach is combining physics knowledge and machine learning to achieve optimal experimental designs [42–44]. Another approach is end-to-end mapping, where DNNs learn a direct mapping between the object and diffraction image domains [45–52]. This method reduces the computationally slow and costly reconstruction process, even enabling real-time inference in some cases [53–55]. However, it requires a large amount of training data. While the training data can be generated numerically in simulation, Sinha *et al.* [45] acquire them *in situ* by projecting thousands of phase objects onto an SLM, whereas in Cherukara *et al.* [50], the training data is obtained from iterative phase retrieval of experimental data. Although these methods allow the model to learn physical system inaccuracies, they are time-consuming and not portable to different optical systems. Given that wave propagation is well-described by the Helmholtz equation, requiring a DNN to learn this may be unnecessary. This insight leads to physics-informed and deep learning (DL) assisted computational imaging. For instance, Goy *et al.* [10] found that physics-informed phase retrieval outperformed end-to-end methods for low photon counts. Similarly, Metzler *et al.* [15] utilize a convolutional neural network (CNN) known as DnCNN [56] as a denoising regularizer, improving the reconstruction quality by exploiting the natural absence of additive Gaussian noise in images, while Chang *et al.* [57] advance this approach by incorporating a complex-domain neural network that leverages the latent correlations between amplitude and phase for improved coherent imaging reconstructions. A recent development in this field is the realization that the structure of deep generative networks can capture image statistics even before learning [58], improving ptychography reconstruction quality under the concept of deep image priors without requiring a preceding training procedure [12,59,60].

In this paper, we introduce a method that combines a fully physics-based ptychography reconstruction framework with a pre-trained deep generative model. When prior knowledge indicates that a sample is sparse in an unknown basis, we show that the learned low-dimensional representation in latent space enables accurate image reconstruction even under extremely challenging noise conditions. The integration of the deep generative model serves two distinct functions. First, in a pre-training step, an under-complete autoencoder [61] learns an implicit parameterization of images belonging to a specific class (e.g., MNIST [62]). Then, this learned model is utilized to successfully reconstruct images from ill-posed diffraction data. We empirically demonstrate noise-robust latent vector reconstruction using experimental data from a photolithographically manufactured sample. The compact space spanned by the latent vectors allows us to visualize and study the optimization landscape in ptychography through principal component analysis, a novel approach to the best of our knowledge. Lastly, we quantify the

reconstruction quality as a function of the total number of photons in the illumination field through numerical simulations. We find that latent vector reconstruction begins to produce faithful reconstructions with an average of 0.001 photons per camera pixel, even in the presence of readout noise. However, as the total number of photons approaches the degrees of freedom in conventional reconstruction, the performance of our approach diminishes in comparison due to its inability to produce comparably high-definition images. We provide the raw ptychography data and an open-source implementation of ptychographic latent vector reconstruction underlying this paper in [63].

2. Methods

2.1. Optical setup

We employ a ptychography setup in transmission geometry as illustrated in Fig. 1(A). A continuous-wave laser (Cobolt Jive 100) with a wavelength of $\lambda = 561$ nm is coupled into a polarization-maintaining single-mode fiber. Then, a fiber-coupled collimator (60FC-L-0-M75-26, Schäfter+Kirchhoff) expands the beam to 25 mm in diameter, illuminating a 500- μm pinhole. This pinhole is imaged onto the object using a 2-lens system with a magnification of $M = 3$, resulting in a circular illumination field with a uniform phase at the object plane. The object, a binary hand-drawn digit, is laterally scanned through the beam using stepper motor actuators (ZFS25B, Thorlabs). Diffraction patterns are recorded 6.5 cm downstream of the object using a CMOS camera sensor (acA2440-35um, Basler) with a pixel size of 3.45 μm and 1024×1024 total pixels. For calibration of the illumination field, object-camera distance, and actuator position inaccuracies, we employ a Fermat spiral scanning pattern [64] with 96 positions and an overlap of 80%. At 16 uniformly distributed positions during this calibration scan, we capture an additional evaluation series of diffraction patterns at a lower signal-to-noise ratio (SNR), varying the camera exposure time between 300 ms to 0.03 ms, and with an approximate illumination overlap of 67%. By acquiring the calibration and evaluation data within a single scan trajectory, we ensure that the scanning position correction remains valid for the evaluation reconstructions. In practice, calibrating the illumination field can be effectively achieved using any object that ensures high-SNR diffraction data and reliable reconstruction quality. The binary photomask sample, the reconstructed illumination field, and the scanning pattern are shown in Supplement 1.

2.2. Reconstruction procedure

In this work, we employ an ADP framework integrated with a pre-trained deep generative model, as illustrated in Fig. 1(B). A comprehensive description of the physics-informed ADP framework can be found in [38]. We use TensorFlow [65] to model the optical system, leveraging its differentiable programming capabilities to seamlessly incorporate the deep generative model into our forward model. Specifically, the decoder of a pre-trained autoencoder serves as this deep generative model. This enables us to represent the object as a compact latent vector rather than a conventional pixel-based image.

For each scanning position, the object is illuminated by a coherent light field. The exit field in the object plane is computed using the projection approximation [66] and propagated to the detection plane via a band-limited angular spectrum method [67]. This yields a set of predicted diffraction patterns I_k for any given object patch and illumination field. The optimization goal is to minimize a loss function $L(\theta)$ designed to maximize the likelihood that the parameter set θ accurately represents the observed diffraction patterns X_k . The loss function is defined as follows: [40]

$$L(\theta) = \sum_{k=1}^N \left(\ln[I_k(\theta) + \sigma_k^2] + \frac{[X_k - I_k(\theta)]^2}{I_k(\theta) + \sigma_k^2} \right), \quad (1)$$




Fig. 1. (A) Schematic of the optical setup used for ptychography. A 500- μm pinhole is illuminated with coherent light at $\lambda = 561$ nm and relayed onto the object using a 2-lens system. The object is moved laterally through the beam using a computer-controlled XY stage, and a CMOS camera sensor records the diffraction intensities 6.5 cm downstream from the object. (B) Diagram of the Automatic Differentiation Ptychography (ADP) framework, which models the physical system beginning from the object illumination. In the conventional mode, the object is represented by complex-valued pixels. With a pre-trained autoencoder for a specific class of objects, the decoder can be integrated into the ADP framework as a deep generative model, allowing the object to be represented as a latent vector and significantly reducing the number of free parameters.

where N is the total number of camera pixels, and σ_k^2 is the camera sensor readout noise determined from 300 dark measurements. In the high-SNR calibration reconstruction, the parameter θ encompasses the illumination field, object, object-camera distance, and the scanning positions. Conversely, when we assess our method's noise-robustness using latent vector reconstruction, we rely on the calibrated data for all parameters except the object. In this scenario, θ consists solely of the latent vector representing the object. While the optical setup and pre-calibrated illumination field are sampled at a 1024×1024 resolution, the decoder maps only to a 32×32 image output. Hence, we employ the Mitchell-Netravali cubic filter to resize the decoder accordingly. This filter is chosen empirically for its high-quality output with smooth gradients and minimal aliasing artifacts [68].

All reconstructions are executed on a commercial Nvidia RTX A6000 GPU using the Adam optimizer [69] with a randomized order of diffraction patterns. The process typically completes within 100 epochs, taking approximately 20 minutes for our datasets. The learning rate α serves as a hyperparameter that controls the step size for the gradient descent within the loss landscape. We find that learning rates in the range of $\alpha = 0.1$ to 1.0 with an exponentially decaying schedule $\alpha_n = \alpha \times 0.97^n$ for the n -th epoch yield optimal convergence. While regularization terms are included for conventional reconstructions as described in [40], reconstructing the latent vectors does not necessitate any additional regularization for optimal convergence.

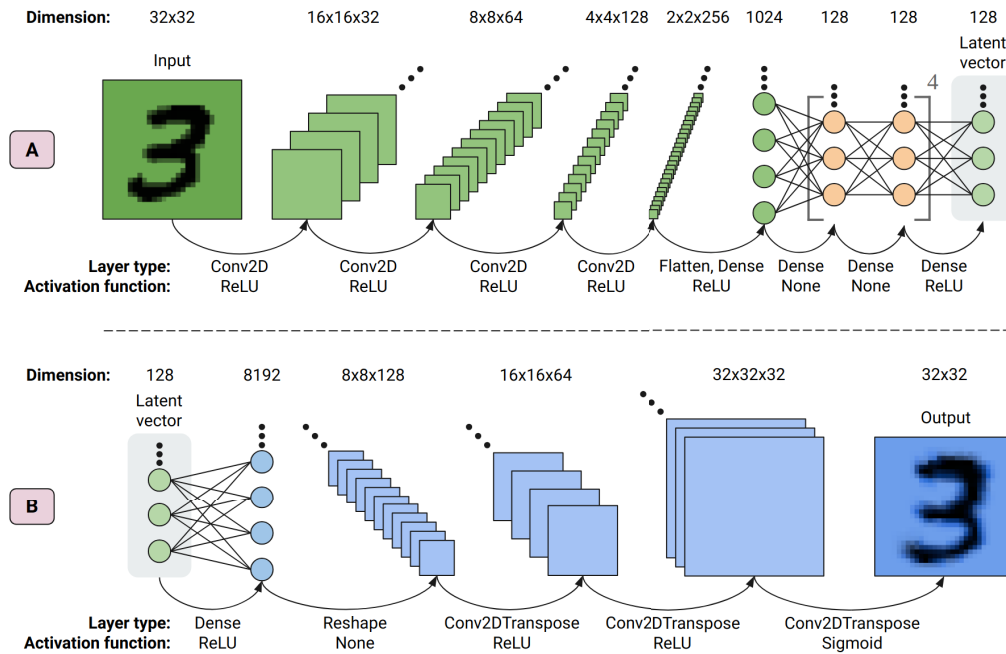


Fig. 2. Illustration of the autoencoder architecture, detailing output dimensions, activation functions, and layer types. (A) Schematic of the encoder model, designed to map input data into a lower-dimensional latent space. Convolutional layers employ 4×4 kernels and a stride of 2, halving the dimensions at each Conv2D layer, and are followed by rectified linear unit (ReLU) activation functions. Linear layers upstream to the latent vector facilitate rank-minimization. (B) Schematic of the decoder model, tasked with reconstructing the original input from the latent representation. Post-training on MNIST, the decoder serves as a deep generative model. The final sigmoid activation function constrains the output to the range $[0, 1]$, making it well-suited for amplitude transmission functions in ptychography.

2.3. Deep generative model

In computational imaging, incorporating machine learning models, particularly deep generative models, offers a compelling route for enhancing image reconstruction capabilities. Deep generative models, such as autoencoders, are neural networks trained to learn a compressed, yet informative, representation of unlabeled data. This learned representation, often referred to as the latent space, captures the essential features of the data while discarding noise and redundancies. This section delves into the architecture, training, and characterization of the autoencoder model used in this study, elucidating how it integrates with the ptychographic reconstruction framework and with a routine to image objects of a class that is known *a priori*.

The autoencoder architecture, adapted and implemented in TensorFlow from [70], is depicted in Fig. 2. An autoencoder network typically aims to map input data into a lower-dimensional latent space and reconstruct it back to the original form [61]. Mathematically, an encoder function f maps an input x to a latent vector $\mathbf{h} = f(x)$. A decoder function g then maps \mathbf{h} back to the reconstructed input $\hat{x} = g(\mathbf{h})$. The autoencoder is trained using the Adam optimizer and MNIST, a dataset containing 60,000 training and 10,000 validation images of handwritten digits, to minimize the binary cross-entropy loss function $L(x, \hat{x} = g(f(x)))$. The training is a one-time process and requires only 50 epochs which take about 20 minutes on an Nvidia RTX A6000 GPU. With respect to the image size of our object, the design of our under-complete autoencoder requires a latent space of significantly smaller dimensionality. Otherwise, the autoencoder would

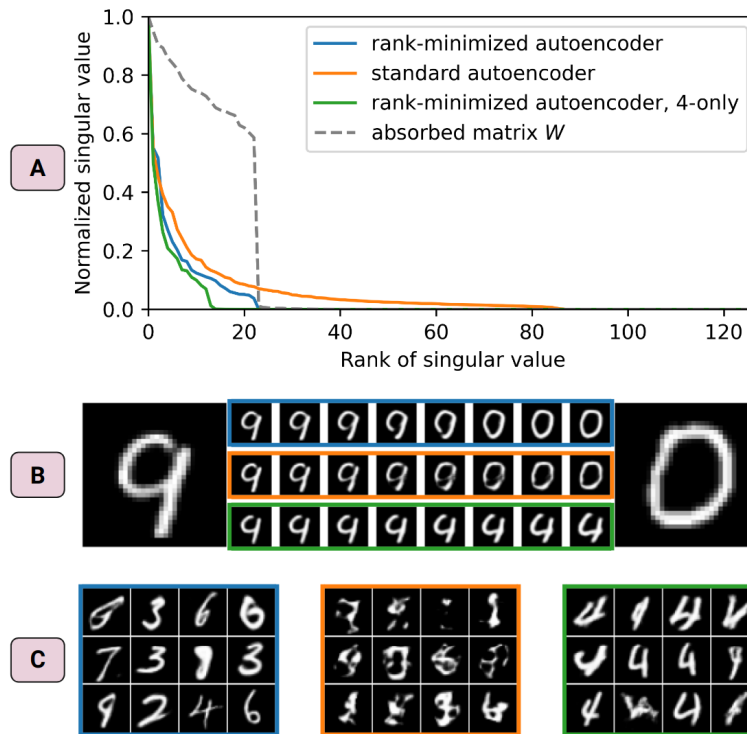


Fig. 3. (A) Singular values of the covariance matrix for latent vectors obtained from MNIST validation examples, revealing the effective rank of the feature representation. Different autoencoder designs and training sets are indicated by color (blue/green: implicit rank-minimized autoencoder; orange: standard autoencoder; blue/orange: full MNIST dataset; green: MNIST dataset filtered to include only '4's). (B) Linear interpolation in the latent space between two objects, illustrating the well-structured feature representation. (C) Images generated from multivariate Gaussian noise input to the decoder, highlighting the decoder's capability to produce meaningful handwritten digits. The color coding for panels (B) and (C) follows the legend provided in panel (A).

trivially learn the identity function, failing to capture the most salient features of the training data. This raises the question of selecting the optimal latent space dimension. Rather than relying on a heuristic trial-and-error approach, we employ the implicit rank-minimizing autoencoder (IRMAE) model [70]. The IRMAE includes eight additional linear layers W_1, W_2, \dots, W_8 at the end of the encoder network, which are randomly initialized. Deep linear networks have been shown to induce implicit regularization, leading to low-rank solutions [71]. Hence, we can choose the latent dimension of \mathbf{h} to be reasonably large (128 in our case) and let the training process automatically find the lowest rank.

In Fig. 3, we evaluate different autoencoder architectures across various scenarios. Fig. 3(A) illustrates the singular value decomposition on the covariance matrix of MNIST validation examples to assess the effective rank needed for feature representation. Our findings indicate that the rank-minimized autoencoder has an effective rank of 22. This is corroborated by the same rank of the matrix absorbed into the encoder, calculated as $W = \prod_{i=1}^8 W_i$. In contrast, we find that the singular values for a standard autoencoder, where W is an identity matrix, can only be neglected above a rank of 86. Hence, the rank-minimized autoencoder spans a more compact latent space. Additionally, we examine the impact of simplifying the training set by using only

the 982 handwritten '4's from the MNIST dataset, which results in a lower rank of 13. This adjustment in training samples can be viewed as imposing a stronger prior belief about the object in the context of latent vector reconstruction in computational imaging.

In Fig. 3(B), we further explore the properties of the latent space. Specifically, we illustrate latent vector arithmetic through linear interpolation between two latent vectors representing a '9' and a '0'. Using the rank-minimized autoencoder, we observe that the latent space interpolation captures smooth and meaningful transitions between different images, indicating a well-structured feature representation. However, this is less evident in the standard autoencoder, which results in a more ambiguous interpolation between images. In the scenario where the model is trained only on a subset of the MNIST dataset, images that are not part of the training set are not as accurately represented. This is by design, as the deep generative model can be intentionally optimized to reconstruct objects within its trained class, in this case, the digit '4'. This targeted approach offers an advantage for applications where more concrete prior knowledge about the object class is available.

Lastly, we sample images generated from multivariate Gaussian noise input to the decoder in Fig. 3(C). Remarkably, images generated using rank-minimized decoders predominantly resemble handwritten digits rather than random patterns that are observed with the standard decoder. This is particularly advantageous for our ptychography application, as it implies that the reconstruction process is intrinsically guided towards generating images that align with prior knowledge – here, the class of handwritten digits. We seek out this property to enhance the robustness of our method, especially when dealing with ill-posed data, by reducing the likelihood of spurious reconstructions. An extended evaluation of the autoencoder's performances across various inputs is available in [Supplement 1](#).

3. Results

3.1. Experimental amplitude transmission reconstructions

We present the main experimental result of this paper in Fig. 4. We illuminate an amplitude-only sample shaped like the digit '4' and adjust the camera's exposure time over four orders of magnitude. As a result, we acquire sets of diffraction patterns ranging from high SNR to extremely low SNR. These datasets are then used for ptychographic reconstructions. In conventional reconstruction, using a 468×468 pixel basis, the object's amplitude transmission function is pristine at the highest SNR but deteriorates to noise at a $30 \mu\text{s}$ exposure. In contrast, switching to latent space reconstruction with a pre-trained deep generative model significantly improves low-SNR performance. The reduced rank of the latent space, previously shown to be 22, cuts the number of free parameters by approximately $10,000\times$ compared to the conventional reconstruction method. This allows for successful object determination with remarkably fewer photons. Moreover, the model trained specifically on the digit '4' shows more stable convergence and slightly better reconstructions overall.

In the high-SNR scenario, the conventional reconstruction outperforms our latent vector approach in terms of image sharpness. When a sufficient number of detected photons is available, reducing the number of parameters that represent the object offers no advantage. Indeed, this becomes a drawback when the lower-resolution output of our deep generative model is resized to match the higher resolution used in our ADP framework, resulting in limited image sharpness. This can be interpreted as a trade-off for the immense reduction of free parameters.

During the optimization process, we observe that the latent vector reconstruction based on training with the full MNIST dataset can occasionally get stuck in local minima, even when using diffraction data with high SNR. Due to the randomization of the order of the diffraction patterns in the stochastic gradient descent, the first update step of the latent vector can walk in a different direction within the loss function optimization landscape. The latent space is non-injective, meaning that multiple different latent vectors can map onto the same output. Therefore, the

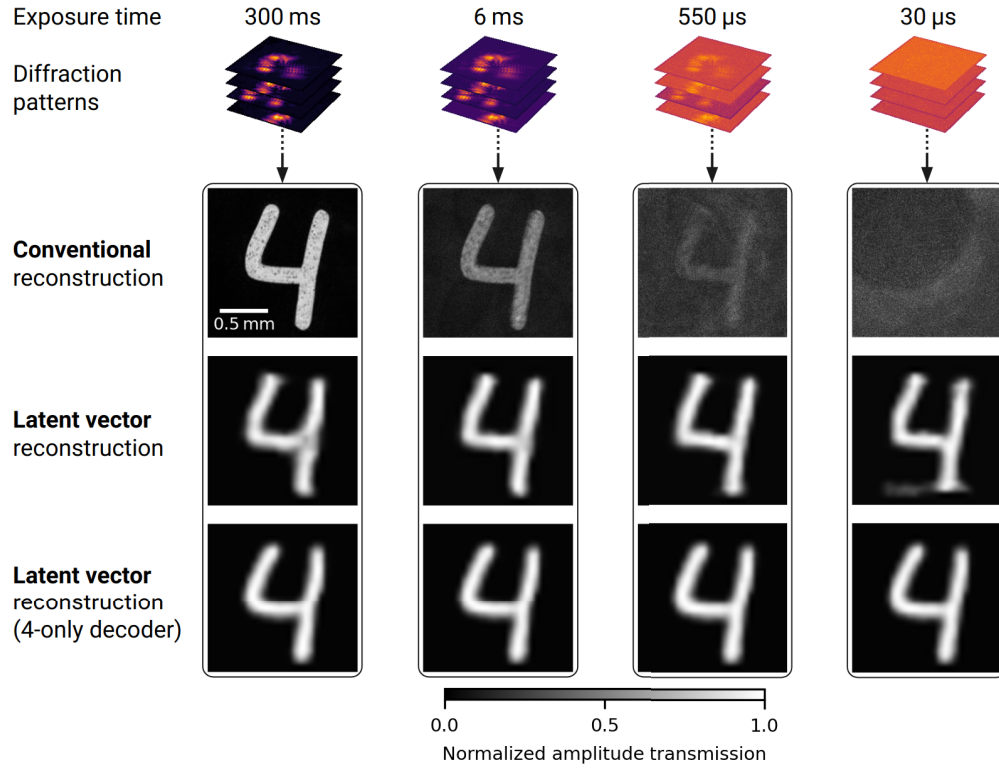


Fig. 4. Comparison of ptychographic amplitude image reconstruction results under varying signal-to-noise ratios (SNR). The top row displays stacks of diffraction patterns used for reconstruction, with exposure times decreasing from left to right, leading to a corresponding decrease in SNR. The second row presents results from conventional reconstruction. Subsequent rows feature latent vector reconstructions using a pre-trained deep generative model, first trained on the full MNIST dataset and secondly on a filtered MNIST dataset containing only images resembling the digit '4'.

optimization procedure is sensitive to the initial state and first gradient step in particular. In practice, we find that this sensitivity can be mitigated by initializing the latent vector as the vector average $\bar{\mathbf{h}}$ obtained using the pre-trained encoder function $f(x)$ and all training examples x_{train} expressed as

$$\bar{\mathbf{h}} = \frac{1}{N_x} \sum_{i=1}^{N_x} f(x_{\text{train},i}), \quad (2)$$

where N_x is the total number of training samples, and $f(x_{\text{train},i})$ is the latent vector corresponding to the i^{th} training sample. This initialization approach is practical as the training samples are readily available from the pre-training phase. Moreover, for the case of the deep generative model trained on the filtered MNIST dataset only including '4's, this vector average initialization is not required. In this case, the convergence is consistently stable and fast, and the latent vector can be initialized as random or uniform.

3.2. Numerical reconstructions and quantitative comparisons

To quantitatively assess the noise robustness of our method and its ability to generalize to another object, we simulate ptychographic reconstructions using a known amplitude-only hand-drawn '3'

as ground truth. We generate diffraction patterns using the probe and scanning pattern shown in Supplement 1, with an object-camera distance of 8 cm, while all other simulation parameters are consistent with the experimental setup. We vary the total photon count in the illumination field from 10 to 10^6 photons, assuming a uniform camera readout noise of $\sigma_k = 0.3$ photons, $\forall k$.

To evaluate reconstruction quality in Fig. 5, we use the Peak Signal-to-Noise Ratio (PSNR) between the reconstructed image \hat{x} and the ground-truth \hat{x}_{gt} . As the maximum pixel value for our decoder output is equal to one, we can write

$$\text{PSNR} = -10 \log_{10} \left(\frac{1}{N} \sum_{i=1}^N (\hat{x}_i - \hat{x}_{\text{gt},i})^2 \right). \quad (3)$$

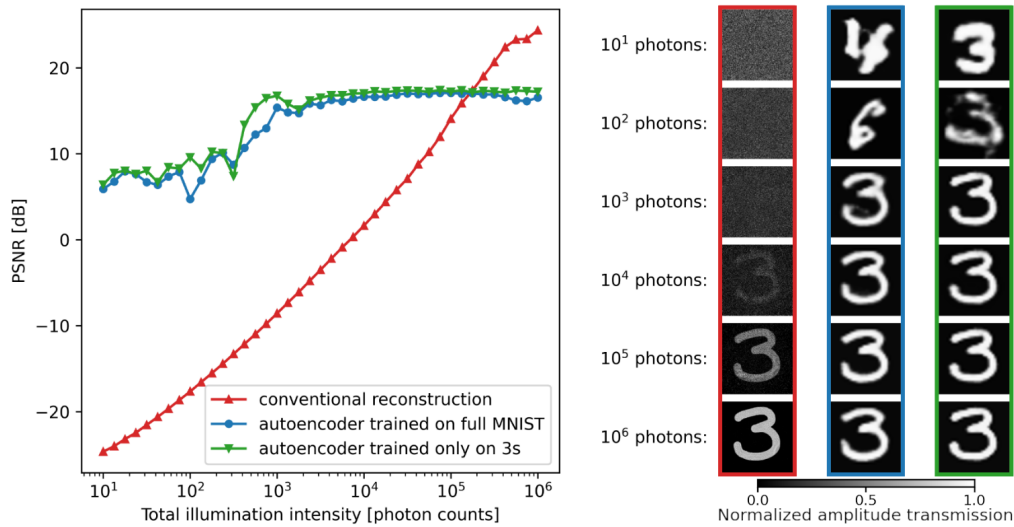


Fig. 5. Comparison of reconstruction quality for conventional and latent vector ptychographic methods across varying signal-to-noise ratios (SNRs), obtained from numerical simulations. The Peak Signal-to-Noise Ratio (PSNR) serves as the quality metric and is plotted against the total photon count in the illumination field. Selected object amplitude reconstructions are displayed for various photon counts to highlight the trade-offs between the methods. A latent vector reconstruction using a deep generative model trained on filtered data is also included for comparison.

In the double-logarithmic plot, a linear relationship between illumination intensity and conventional reconstruction quality suggests a power-law behavior. The curve plateaus at 10^6 photons, indicating the inverse problem becomes well-posed. This is further supported by the amplitude reconstruction visually nearing its optimum in sharpness and contrast at the highest simulated SNR.

For the latent vector reconstruction, successful convergence occurs at around 10^3 photons for both the full and filtered MNIST-trained models. This corresponds to an average photon count of 0.001 per camera pixel. Below this threshold, the PSNR is inflated due to the deep generative model's inability to generate noisy outputs matching the readout noise, leading to spurious correlations with the ground truth. The model trained exclusively on '3's demonstrates stable convergence at marginally lower SNRs, specifically around a few hundred photons in the illumination field.

In summary, our method excels at low SNRs where the conventional method struggles to produce meaningful reconstructions, although it is inherently limited at high SNRs due to the

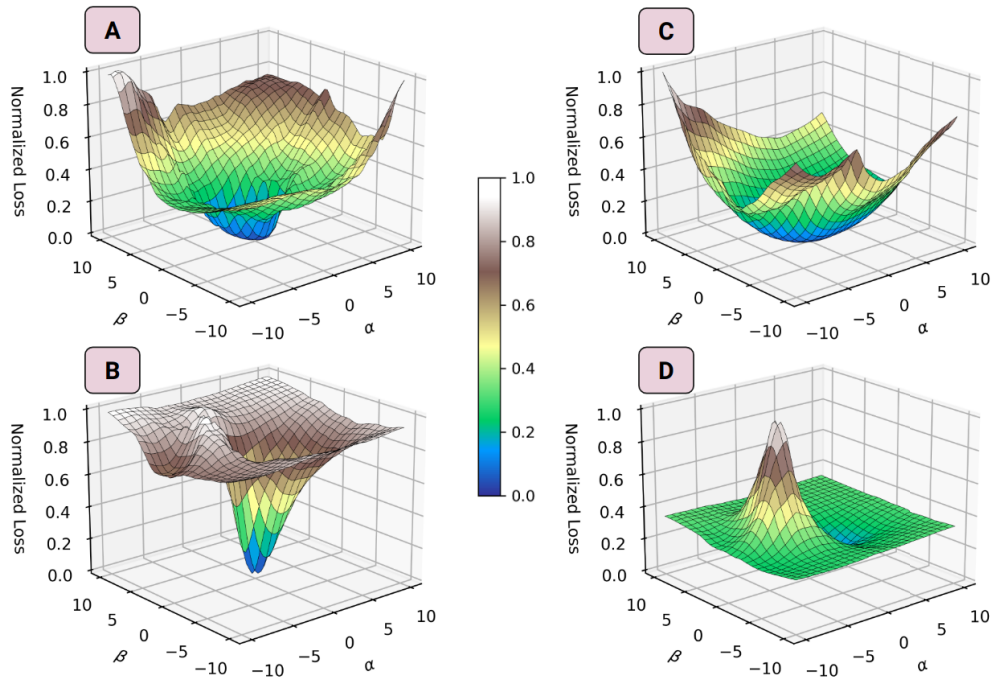


Fig. 6. Visualization of the optimization loss landscapes for different scenarios. α and β are coefficients for the two leading principal components of the latent space used for ptychographic reconstruction. (A) The landscape for high signal-to-noise ratio (SNR) and training on the full MNIST dataset. (B) The landscape when reconstructing from low-SNR (high-noise) diffraction data. (C) The landscape after training the deep generative model on a filtered MNIST dataset containing only the digit '4'. (D) The landscape when optimizing using a Poisson-only loss function at high SNR, for comparison with the mixed Poisson-Gaussian loss from all other panels.

reduced degrees of freedom in the deep generative model. Hence, when sufficient information is available in the diffraction data, the conventional reconstruction method should be preferred.

3.3. Loss landscapes

Given the compact nature of the latent space, we have the unique opportunity to approximate and visualize the loss landscape that is traversed during optimization (see Fig. 6). Utilizing the two leading orthogonal principal components of the latent space, we construct a three-dimensional representation of the loss landscape. We perform a principal component analysis on the covariance matrix of all latent vectors from the validation MNIST dataset to identify the two most informative directions associated with the two leading singular values, denoted as \mathbf{v}_1 and \mathbf{v}_2 . Given an optimal latent vector \mathbf{h}_{opt} obtained from experimental diffraction data, we explore the loss landscape by varying this optimal point along the directions of \mathbf{v}_1 and \mathbf{v}_2 :

$$\mathbf{h}(\alpha, \beta) = \mathbf{h}_{\text{opt}} + \alpha \mathbf{v}_1 + \beta \mathbf{v}_2. \quad (4)$$

Here, α and β range from -10 to 10 based on the extent of these variations in image space as illustrated in [Supplement 1](#). We then compute the loss value for each $\mathbf{h}(\alpha, \beta)$ to visualize the landscape.

In the high-SNR case with training on the full MNIST dataset (Fig. 6(A)), the loss landscape exhibits a distinct but non-convex and asymmetric minimum. This topography accounts for the sensitivity to the initial latent vector state. The non-convexity is exacerbated in low-SNR scenarios (Fig. 6(B)), confirming that convergence is more challenging when the signal is weak. Interestingly, the loss landscape exhibits a smoother and more convex topography when the model is trained exclusively on images resembling the digit '4' (see Fig. 6(C)). This finding aligns with our previous observation that such specialized training renders the optimization process less prone to getting trapped in local minima and more forgiving of suboptimal latent vector initialization. Finally, we explore the impact of using an alternative loss function based solely on Poisson statistics in Fig. 6(D). This loss function, defined as $L_P(\theta) = \sum_{k=1}^N \left(\sqrt{X_k} - \sqrt{I_k(\theta)} \right)^2$, reveals a loss landscape with large plateaus and a higher degree of non-convexity for the high-SNR, full MNIST-trained scenario. These characteristics align with our observation that the mixed Poisson-Gaussian loss function generally offers better convergence behavior in comparison.

4. Discussion and conclusion

We present a novel approach to ptychographic image reconstruction by integrating a deep generative model into a physics-informed automatic differentiation ptychography (ADP) framework. By incorporating prior knowledge about the object class of a specimen, our method significantly reduces the number of free parameters in the optimization problem, enabling robust reconstructions in low signal-to-noise ratio (SNR) scenarios. Since the pre-training of the generative model and the ADP reconstruction are separated, this approach is modular and portable to related AD-based imaging methods. For instance, the deep generative model can seamlessly be exchanged or reused in an adapted physics-informed imaging modality without retraining. As the field of generative artificial intelligence is currently undergoing an immense research interest, our work presents a straightforward way to incorporate future latent generative models into computational imaging.

One key observation of this paper is the inherent trade-off between noise robustness and the maximum achievable fidelity of the reconstructed image. This limitation is primarily due to the output resolution constraints of the pre-trained decoder. This opens up intriguing avenues for future research, including the exploration of alternative deep generative models for ptychographic reconstruction. While generative adversarial networks (GANs) and latent diffusion models have shown promise in related imaging contexts such as compressed superresolution imaging through multimode fibers [72] or high-resolution image reconstruction from human brain activity [73], they introduce their own set of challenges. These include increased computational complexity, less interpretable latent spaces [74], and the added model complexity required for capturing high-resolution or complex greyscale features, which could compromise our method's ability to reconstruct from severely ill-posed data. Indeed, architectures like the recently shown GigaGAN [75] show excellent generative abilities and a controllable latent space, but require up to 4700 days of training on a high-end A100 GPU, while the autoencoder design in this work trains within minutes on any commercial laptop.

Another promising direction for future research in latent vector reconstruction is the extension of the model to output complex-valued images. This would lift the current limitation of representing only amplitude objects, thereby broadening the applicability of our approach. Cherukara *et al.* have already made strides in this direction, demonstrating ptychographic imaging with a Y-shaped latent model that performs an end-to-end mapping from diffraction patterns to both phase and amplitude images [50]. This is particularly interesting for noise-robust latent vector reconstruction, as the shared latent representation could be leveraged to account for the high correlation typically observed between phase and amplitude in complex objects, akin to the implementation recently shown in [57]. Consequently, even though the output dimensionality

would effectively double to accommodate both phase and amplitude, the rank of the latent space may not necessarily need to double, thanks to this inherent correlation.

Our utilization of a compact latent space for object representation provides a unique opportunity to approximate and visualize the optimization loss landscape during ptychographic object retrieval. This offers valuable insights into the convergence behavior and the sensitivity to the initial state of the reconstruction process. However, the mapping from the latent space to the object space is non-injective, and our landscape visualization, therefore, is localized and truncated due to the omission of principal components associated with smaller but still relevant eigenvalues.

Our method's robustness in low-SNR conditions offers valuable applications in both biological and industrial settings where prior knowledge for pre-training is often available. It is particularly well-suited for medical imaging of delicate specimens, where minimizing radiation dose is a priority. The reduced computational complexity, thanks to fewer free parameters and the calibrated illumination field, also makes it ideal for real-time and specialized imaging scenarios like industrial quality control, where quick yet quality-assured reconstructions are desired. Indeed, we occasionally observe high-SNR latent vector reconstructions to converge within a single epoch, as compared to the multiple epochs typically required for conventional reconstructions. Furthermore, our method's noise resilience makes it a promising tool for imaging in photon-starved regimes, such as extreme ultraviolet (EUV) or X-ray applications, where imaging with minimal photon counts is often required.

In conclusion, our work represents a step forward in the field of computational imaging by demonstrating the power of integrating machine learning techniques with physics-based models for robust image reconstruction. As computational imaging continues to evolve, the integration of deep learning models with traditional imaging techniques promises to unlock new capabilities and applications across a wide range of scientific and industrial domains.

Funding. Nederlandse Organisatie voor Wetenschappelijk Onderzoek (Perspective P16-08).

Acknowledgments. We thank Dorian Bouchet for helpful discussions and Cees de Kok, Dante Killian, Jan Bonne Aans, Aron Opheij, Paul Jurrius and Arjan Driessen for technical support.

Disclosures. The authors declare no conflicts of interest.

Data Availability. Experimental raw data, synthetic data, and code underlying these results are available in [63].

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

References

1. C.-Y. Chang, C. Li, J.-W. Chang, *et al.*, "An unsupervised neural network approach for automatic semiconductor wafer defect inspection," *Expert Syst. Appl.* **36**(1), 950–958 (2009).
2. P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Trans. Pattern Anal. Machine Intell.* **12**(7), 629–639 (1990).
3. Y. Wang, W. Ren, and H. Wang, "Anisotropic second and fourth order diffusion models based on convolutional virtual electric field for image denoising," *Comput. & Math. with Appl.* **66**(10), 1729–1742 (2013).
4. K. Dabov, A. Foi, V. Katkovnik, *et al.*, "Image denoising by sparse 3-d transform-domain collaborative filtering," *IEEE Trans. on Image Process.* **16**(8), 2080–2095 (2007).
5. P. Thibault and M. Guizar-Sicairos, "Maximum-likelihood refinement for coherent diffractive imaging," *New J. Phys.* **14**(6), 063004 (2012).
6. X. Wei, H. P. Urbach, and W. M. J. Coene, "Cramér-rao lower bound and maximum-likelihood estimation in ptychography with poisson noise," *Phys. Rev. A* **102**(4), 043516 (2020).
7. V. Katkovnik and J. Astola, "Sparse ptychographical coherent diffractive imaging from noisy measurements," *J. Opt. Soc. Am. A* **30**(3), 367–379 (2013).
8. V. Katkovnik and K. Egiazarian, "Sparse phase imaging based on complex domain nonlocal bm3d techniques," *Digital Signal Process.* **63**, 72–85 (2017).
9. M. Schloz, T. C. Pekin, Z. Chen, *et al.*, "Overcoming information reduced data and experimentally uncertain parameters in ptychography with regularized optimization," *Opt. Express* **28**(19), 28306–28323 (2020).
10. A. Goy, K. Arthur, S. Li, *et al.*, "Low photon count phase retrieval using deep learning," *Phys. Rev. Lett.* **121**(24), 243902 (2018).
11. S. Aslan, Z. Liu, V. Nikitin, *et al.*, "Joint ptycho-tomography with deep generative priors," *Mach. Learn.: Sci. Technol.* **2**(4), 045017 (2021).

12. Q. Chen, D. Huang, and R. Chen, "Fourier ptychographic microscopy with untrained deep neural network priors," *Opt. Express* **30**(22), 39597 (2022).
13. S. V. Venkatakrisnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *2013 IEEE global conference on signal and information processing*, (IEEE, 2013), pp. 945–948.
14. Y. Romano, M. Elad, and P. Milanfar, "The little engine that could: Regularization by denoising (RED)," *SIAM J. Imaging Sci.* **10**(4), 1804–1844 (2017).
15. C. Metzler, P. Schniter, A. Veeraraghavan, *et al.*, "prDeep: Robust phase retrieval with a flexible deep network," in *Proceedings of the 35th International Conference on Machine Learning*, vol. 80 of *Proceedings of Machine Learning Research* (PMLR, 2018), pp. 3501–3510.
16. E. T. Reehorst and P. Schniter, "Regularization by denoising: Clarifications and new interpretations," *IEEE Trans. Comput. Imaging* **5**(1), 52–67 (2019).
17. Z. Wu, Y. Sun, J. Liu, *et al.*, "Online regularization by denoising with applications to phase retrieval," in *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, (IEEE, 2019), pp. 3887–3895.
18. T. Wang, S. Jiang, P. Song, *et al.*, "Optical ptychography for biomedical imaging: recent progress and future directions [invited]," *Biomed. Opt. Express* **14**(2), 489–532 (2023).
19. W. Hoppe, "Beugung im inhomogenen Primärstrahlwellenfeld. I. Prinzip einer Phasenmessung von Elektronenbeugungsinterferenzen," *Acta Cryst. A* **25**(4), 495–501 (1969).
20. R. Hegerl and W. Hoppe, "Dynamische Theorie der Kristallstrukturanalyse durch Elektronenbeugung im inhomogenen Primärstrahlwellenfeld," *Ber. Bunsenges. Phys. Chem.* **74**, 1148–1154 (1970).
21. J. Rodenburg and A. Maiden, "Ptychography," in *Springer Handbook of Microscopy*, P. W. Hawkes and J. C. H. Spence, eds. (Springer International Publishing, 2019).
22. J. M. Rodenburg and H. M. L. Faulkner, "A phase retrieval algorithm for shifting illumination," *Appl. Phys. Lett.* **85**(20), 4795–4797 (2004).
23. G. Zheng, R. Horstmeyer, and C. Yang, "Wide-field, high-resolution fourier ptychographic microscopy," *Nat. Photonics* **7**(9), 739–745 (2013).
24. P. C. Konda, L. Loetgering, K. C. Zhou, *et al.*, "Fourier ptychography: current applications and future promises," *Opt. Express* **28**(7), 9603–9630 (2020).
25. L. Loetgering, M. Du, D. Boonzajer Flaes, *et al.*, "PtyLab.m/py/jl: a cross-platform, open-source inverse modeling toolbox for conventional and fourier ptychography," *Opt. Express* **31**(9), 13763–13797 (2023).
26. A. M. Maiden and J. M. Rodenburg, "An improved ptychographical phase retrieval algorithm for diffractive imaging," *Ultramicroscopy* **109**(10), 1256–1262 (2009).
27. M. Du, L. Loetgering, K. S. E. Eikema, *et al.*, "Measuring laser beam quality, wavefronts, and lens aberrations using ptychography," *Opt. Express* **28**(4), 5022–5034 (2020).
28. M. Du, X. Liu, A. Pelekanidis, *et al.*, "High-resolution wavefront sensing and aberration analysis of multi-spectral extreme ultraviolet beams," *Optica* **10**(2), 255 (2023).
29. A. M. Maiden, M. J. Humphry, M. C. Sarahan, *et al.*, "An annealing algorithm to correct positioning errors in ptychography," *Ultramicroscopy* **120**, 64–72 (2012).
30. P. Dwivedi, A. P. Konijnenberg, S. F. Pereira, *et al.*, "Lateral position correction in ptychography using the gradient of intensity patterns," *Ultramicroscopy* **192**, 29–36 (2018).
31. L. Loetgering, M. Du, K. S. E. Eikema, *et al.*, "zPIE: an autofocusing algorithm for ptychography," *Opt. Lett.* **45**(7), 2030–2033 (2020).
32. P. Thibault and A. Menzel, "Reconstructing state mixtures from diffraction measurements," *Nature* **494**(7435), 68 (2013).
33. P. Li, T. Edo, D. Batey, *et al.*, "Breaking ambiguities in mixed state ptychography," *Opt. Express* **24**(8), 9038–9052 (2016).
34. Y. S. G. Nashed, T. Peterka, J. Deng, *et al.*, "Distributed automatic differentiation for ptychography," *Procedia Comput. Sci.* **108**, 404–414 (2017).
35. S. Ghosh, Y. S. G. Nashed, O. Cossairt, *et al.*, "ADP: Automatic differentiation ptychography," in *2018 IEEE International Conference on Computational Photography (ICCP)*, (2018), pp. 1–10.
36. S. Kandel, S. Maddali, M. Allain, *et al.*, "Using automatic differentiation as a general framework for ptychographic reconstruction," *Opt. Express* **27**(13), 18653–18672 (2019).
37. M. Du, Y. S. G. Nashed, S. Kandel, *et al.*, "Three dimensions, two microscopes, one code: Automatic differentiation for x-ray nanotomography beyond the depth of focus limit," *Sci. Adv.* **6**(13), 3700 (2020).
38. J. Seifert, D. Bouchet, L. Loetgering, *et al.*, "Efficient and flexible approach to ptychography using an optimization framework based on automatic differentiation," *OSA Continuum* **4**(1), 121–128 (2021).
39. K. Maathuis, J. Seifert, and A. P. Mosk, "Sensor fusion in ptychography," *Opt. Continuum* **1**(9), 1909 (2022).
40. J. Seifert, Y. Shao, R. van Dam, *et al.*, "Maximum-likelihood estimation in ptychography in the presence of Poisson-Gaussian noise statistics," *Opt. Lett.* **48**(22), 6027–6030 (2023).
41. D. Bouchet, J. Seifert, and A. P. Mosk, "Optimizing illumination for precise multi-parameter estimations in coherent diffractive imaging," *Opt. Lett.* **46**(2), 254–257 (2021).
42. A. Chakrabarti, "Learning sensor multiplexing design through back-propagation," *Adv. Neural Inf. Process. Syst.* **29** (2016).

43. M. R. Kellman, E. Bostan, N. A. Repina, *et al.*, “Physics-based learned design: optimized coded-illumination for quantitative phase imaging,” *IEEE Trans. Comput. Imaging* **5**(3), 344–353 (2019).
44. C. A. Metzler, H. Ikoma, Y. Peng, *et al.*, “Deep optics for single-shot high-dynamic-range imaging,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, (2020), pp. 1375–1385.
45. A. Sinha, J. Lee, S. Li, *et al.*, “Lensless computational imaging through deep learning,” *Optica* **4**(9), 1117–1125 (2017).
46. A. Kappeler, S. Ghosh, J. Holloway, *et al.*, “Ptychnet: CNN based Fourier ptychography,” in *2017 IEEE International Conference on Image Processing (ICIP)*, (2017), pp. 1712–1716.
47. K. H. Jin, M. T. McCann, E. Froustey, *et al.*, “Deep convolutional neural network for inverse problems in imaging,” *IEEE Trans. on Image Process.* **26**(9), 4509–4522 (2017).
48. Y. Sun, Z. Xia, and U. S. Kamilov, “Efficient and accurate inversion of multiple scattering with deep learning,” *Opt. Express* **26**(11), 14678–14688 (2018).
49. L. Boominathan, M. Maniparambil, H. Gupta, *et al.*, “Phase retrieval for fourier ptychography under varying amount of measurements,” (2018).
50. M. J. Cherukara, T. Zhou, Y. Nashed, *et al.*, “Real-time sparse-sampled ptychographic imaging through deep neural networks,” *arXiv*, arXiv:2004.08247 (2020).
51. R. Li, G. Pedrini, Z. Huang, *et al.*, “Physics-enhanced neural network for phase retrieval from two diffraction patterns,” *Opt. Express* **30**(18), 32680–32692 (2022).
52. Q. Ye, L.-W. Wang, and D. P. K. Lun, “Sisprnet: end-to-end learning for single-shot phase retrieval,” *Opt. Express* **30**(18), 31937–31958 (2022).
53. M. J. Cherukara, Y. S. G. Nashed, and R. J. Harder, “Real-time coherent diffraction inversion using deep generative networks,” *Sci. Rep.* **8**(1), 16520 (2018).
54. C. A. Metzler, F. Heide, P. Rangarajan, *et al.*, “Deep-inverse correlography: towards real-time high-resolution non-line-of-sight imaging,” *Optica* **7**(1), 63–71 (2020).
55. A. V. Babu, T. Bicer, S. Kandel, *et al.*, “AI-assisted automated workflow for real-time X-ray ptychography data analysis via federated resources,” *Electron. Imaging* **35**(11), 232-1–232-6 (2023).
56. 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).
57. X. Chang, R. Zhao, S. Jiang, *et al.*, “Complex-domain-enhancing neural network for large-scale coherent imaging,” *Adv. Photonics Nexus* **2**(04), 046006 (2023).
58. U. Dmitry, A. Vedaldi, and L. Victor, “Deep image prior,” *Int. J. Comput. Vis.* **128**(7), 1867–1888 (2020).
59. M. Du, X. Huang, and C. Jacobsen, “Using a modified double deep image prior for crosstalk mitigation in multislice ptychography,” *J. Synchrotron Rad.* **28**(4), 1137–1145 (2021).
60. S. Barutcu, D. Gürsoy, and A. K. Katsaggelos, “Compressive ptychography using deep image and generative priors,” (2022).
61. I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, 2016).
62. Y. LeCun, “The MNIST database of handwritten digits,” <http://yann.lecun.com/exdb/mnist/> (1998).
63. J. Seifert, Y. Shao, and A. P. Mosk, “Supplemental code and raw data for ‘noise-robust latent vector reconstruction in ptychography using deep generative models’,” Data publication platform of Utrecht University (2023). <https://doi.org/10.24416/UU01-AV4ZJT>.
64. X. Huang, H. Yan, R. Harder, *et al.*, “Optimization of overlap uniformness for ptychography,” *Opt. Express* **22**(10), 12634–12644 (2014).
65. M. Abadi, A. Agarwal, P. Barham, *et al.*, “Tensorflow: Large-scale machine learning on heterogeneous distributed systems,” *arXiv*, arXiv:1603.04467 (2016).
66. D. Paganin, *Coherent X-ray optics*, 6 (Oxford University Press, 2006).
67. K. Matsushima and T. Shimobaba, “Band-limited angular spectrum method for numerical simulation of free-space propagation in far and near fields,” *Opt. Express* **17**(22), 19662–19673 (2009).
68. D. P. Mitchell and A. N. Netravali, “Reconstruction filters in computer-graphics,” *ACM Siggraph Comput. Graph.* **22**(4), 221–228 (1988).
69. D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv* arXiv:1412.6980 (2014).
70. L. Jing, J. Zbontar, and Y. LeCun, “Implicit rank-minimizing autoencoder,” (2020).
71. S. Arora, N. Cohen, W. Hu, *et al.*, “Implicit regularization in deep matrix factorization,” in *Advances in Neural Information Processing Systems*, vol. 32 H. Wallach, H. Larochelle, A. Beygelzimer, *et al.*, eds. (Curran Associates, Inc., 2019).
72. W. Li, K. Abrashitova, G. Osnabrugge, *et al.*, “Generative adversarial network for superresolution imaging through a fiber,” *Phys. Rev. Appl.* **18**(3), 034075 (2022).
73. Y. Takagi and S. Nishimoto, “High-resolution image reconstruction with latent diffusion models from human brain activity,” *bioRxiv*, 517004 (2022).
74. A. Asperti and V. Tonelli, “Comparing the latent space of generative models,” *Neural Comput. Appl.* **35**(4), 3155–3172 (2023).
75. M. Kang, J.-Y. Zhu, R. Zhang, *et al.*, “Scaling up GANs for Text-to-Image Synthesis,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, (2023), pp. 10124–10134.