
Uncertainty quantification for ptychography using normalizing flows

Agnimitra Dasgupta

Sonny Astani Department of Civil and Environmental Engineering
University of Southern California
Los Angeles, CA 90089
adasgupt@usc.edu

Zichao Wendy Di

Mathematics and Computer Science Division
Argonne National Laboratory
wendydi@anl.gov

Abstract

Ptychography, as an essential tool for high-resolution and nondestructive material characterization, presents a challenging large-scale nonlinear and non-convex inverse problem; however, its intrinsic photon statistics create clear opportunities for statistical-based deep learning approaches to tackle these challenges, which has been underexplored. In this work, we explore normalizing flows to obtain a surrogate for the high-dimensional posterior, which also enables the characterization of the uncertainty associated with the reconstruction: an extremely desirable capability when judging the reconstruction quality in the absence of ground truth, spotting spurious artifacts and guiding future experiments using the returned uncertainty patterns. We demonstrate the performance of the proposed method on a synthetic sample with added noise and in various physical experimental settings.

1 Introduction

Ptychography is a special type of coherent diffraction imaging technique, wherein an object is scanned at a series of overlapping regions using a coherent beam of incident waves and reconstructed from the series of diffraction patterns [26, 10]. As a lensless imaging technique, ptychography has become an essential tool for high-resolution and nondestructive material characterization, with state-of-the-art techniques reaching below 10 nm in spatial resolution [22], across a wide variety of disciplines including materials science and biology; see [14] and references therein. However, the aforementioned advantages come at the expense of having to tackle a nonlinear and non-convex (i.e., ill-posed as presenting nonunique local minima) inverse problem. Traditionally, ptychographic reconstruction has been solved iteratively using deterministic optimization-based approaches. Due to the ill-posed nature, the solution quality returned by traditional approaches can contain artifacts and is usually highly sensitive to initialization, model accuracy, and the underlying optimization solver used, especially in the presence of noise [5]. Therefore, in the absence of ground truth, robust characterization (i.e., uncertainty quantification) of the different solutions, and any artifacts therein, is highly desirable.

End-to-end approaches using deep learning have become popular in ptychographic reconstruction [6] and other phase imaging methods [29, 16, 25, 19] with the focus to obtain a reconstruction efficiently. There has been limited work demonstrating uncertainty quantification on phase imaging using Bayesian convolution neural networks (e.g., [36, 34]); albeit in a supervised learning setting

such that any uncertainty descriptions are conditioned on the training set. Wei et al. [33] used the Cramér-Rao bound to derive analytical lower bound for uncertainty in ptychographic reconstruction. Their approach, however, remains limited to Poisson noise models and will be challenging to extend to more complicated models.

One potential candidate for a systematic and robust uncertainty quantification is using Bayesian approaches to obtain the posterior distribution and associated statistics. Typically, for nonlinear inverse problems, the posterior is approximated using samples generated with the help of Markov-chain Monte-Carlo (MCMC) methods [4]. However, MCMC can be inefficient for high dimensional problems such as ptychography; slower mixing of the Markov chains leading to drastically increased computational costs [7]. Normalizing flows (NFs) are a class of generative models that can transform a probability density, usually simple and easy to sample, into another more complicated probability density through a sequence of invertible mappings [24, 13, 21]. One benefit of flow-based generative models, compared to other types of generative models, is its capability of computing likelihood exactly and efficiently [12]. Therefore, NFs have been popular choices to approximate complicated posterior densities. For example, in computer science, NFs have been successfully applied in various problems including computer vision, speech synthesis, etc. [12, 3, 38, 23, 15], and recently been explored for various inverse problems in natural science ranging from cosmology to medical imaging [37, 18, 20, 28, 35, 30, 1, 2].

Inspired by the unsupervised deep probabilistic imaging framework proposed by Sun and Bouman [30], in this work, we explore the applicability of NFs for the task of reconstruction and uncertainty quantification on ptychography. More specifically, we construct a surrogate for the posterior density of the reconstructed object using flow-based generative models. Samples generated using such *trained* surrogates can be used to obtain reconstructions of the object and estimate standard deviations of the reconstructions. In this work, we use the standard deviation of the reconstructed phase and magnitude as the metric for uncertainty quantification, and demonstrate that the reconstruction quality obtained by the proposed NFs is comparable to an advanced ptychographic reconstruction methods "rPIE" [17] under various scan settings that reflect varied degree of ill-posedness. Furthermore, compared to traditional methods, our approach presents a unique capability of quantifying the solution quality in the absence of ground truth.

2 Methods

Forward model: Let $\mathbf{z} \in \mathbb{C}^{n^2}$ be a complex variable representing the object of interest, and $\mathbf{d}_j \in \mathbb{R}^{m^2}$ be the observed data (or intensities) measured from the j^{th} scanning position, where n^2 and m^2 are the dimensions of the vectorized object and data, respectively. A ptychography experiment is modeled by

$$\mathbf{d}_j = |\mathcal{F}(\mathbf{P}_j \mathbf{z})|^2 + \epsilon_j, \quad j = 1, \dots, N, \quad (1)$$

where N is the total number of scanning positions, $\mathcal{F}: \mathbb{C}^{m^2} \mapsto \mathbb{C}^{m^2}$ is the two-dimensional discrete Fourier operator, $\mathbf{P}_j \in \mathbb{C}^{m^2 \times n^2}$ is the given illumination probe at the j^{th} step, and $\epsilon_j \in \mathbb{R}^{m^2}$ is the noise corresponding to the j^{th} measurement. Then we have the ptychographic forward model as $f_j(\mathbf{z}) = |\mathcal{F}(\mathbf{P}_j \mathbf{z})|^2$ for the j^{th} scan, which is nonlinear and non-convex, and $f(\mathbf{z}) = \{f_j(\mathbf{z})\}_{j=1}^N$.

Approximating the posterior distribution using normalizing flows: A surrogate $q(\mathbf{z})$ for the posterior density $p(\mathbf{z}|\mathbf{d})$ can be constructed by letting $\mathbf{z} = \mathcal{G}_\theta(\mathbf{w})$ where $\mathbf{w} \sim \mathcal{N}(0, \mathcal{I}_{2n^2})$ ¹. \mathcal{G} is a generative model (in this case an NF), and θ is its parameters. For $q(\mathbf{z})$ to be a valid probability distribution, the mapping \mathcal{G} must be bijective, which gives us

$$\log q(\mathbf{z}) = \log p(\mathbf{w}) - \log \left| \det \frac{\partial \mathcal{G}_\theta(\mathbf{w})}{\partial \mathbf{w}} \right|,$$

due to the change of variables rule. Therefore, we use invertible neural networks (INNs) to form \mathcal{G}_θ to ensure its bijectiveness. The parameters θ can be learned by minimizing the Kullback-Liebler

¹In this work, due to the presence of complex number, the actual output from the NF is $2n^2$ real numbers which is divided into two chunks that are used separately as the real and imaginary part of each pixel to form \mathbf{z} .

distance between $q(z)$ and $p(z|\mathbf{d})$ which leads to the following minimization problem [30]

$$\theta^* = \arg \min_{\theta} \sum_{k=1}^B \mathcal{L}(\mathbf{d}|f(\mathcal{G}_{\theta}(\mathbf{w}_k))) + \lambda \mathcal{R}(\mathcal{G}_{\theta}(\mathbf{w}_k)) - \log \left| \det \frac{\partial \mathcal{G}_{\theta}(\mathbf{w})}{\partial \mathbf{w}} \right|_{\mathbf{w}=\mathbf{w}_k}, \quad (2)$$

where \mathbf{w}_k is the k^{th} sample, B is the batch size for training, and \mathcal{L} and \mathcal{R} are the data misfit and regularization terms, respectively, such that $p(z|\mathbf{d}) \propto \exp(-\mathcal{L}(\mathbf{d}|z) - \lambda \mathcal{R}(z))$ where $z = \mathcal{G}_{\theta}(\mathbf{w})$. λ is the regularization weight which requires tuning to be relatively optimal. Additionally, NFs do not require training data separately [30]. Instead the cost function (2) can be optimized to learn the weights and concurrently monitored to tune any hyperparameters.

There are various ways to model the data misfit in ptychography given Eqn. (1). Given σ_j as the assumed variance of ϵ_j , we focus on the amplitude-based error metric which is

$$\mathcal{L}(\mathbf{d}_j|f_j(z)) = \frac{1}{2\sigma_j^2} \left\| |\mathcal{F}(\mathbf{P}_j z)| - \sqrt{\mathbf{d}_j} \right\|_2^2, \quad \text{s.t. } \mathcal{L}(\mathbf{d}|f(z)) = \sum_{j=1}^N \mathcal{L}(\mathbf{d}_j|f_j(z)), \quad (3)$$

due to its stability, as compared to other options such as the intensity Gaussian metric [11].

Constraints: The only constraint we enforce is a $[0, 1]$ box constraint on the magnitude (i.e., $|z|$) of an object, as a common assumption corresponding to the absorption capability of material [31]. In our implementation, we constrain the magnitude, which is formed by querying the surrogate, by normalizing it via $z \leftarrow z'/\|z'\|_{\infty}$ where $z' = \mathcal{G}_{\theta}(\mathbf{w})$, before evaluating the cost function (2).

Coupling layers: We construct NFs using INNs that are based on the coupling architecture proposed by Dinh et al. [8] and later refined in [9]. The coupling layer (see Fig. 1) works as follows: an input \mathbf{u} is divided into two parts \mathbf{u}_1 and \mathbf{u}_2 which are transformed into corresponding parts \mathbf{v}_1 and \mathbf{v}_2 of the output \mathbf{v} using affine transformations as follows

$$\mathbf{v}_1 = \mathbf{u}_2 \odot \exp[s_1(\mathbf{u}_1)] + t_1(\mathbf{u}_1) \rightarrow \mathbf{v}_2 = \mathbf{u}_1 \odot \exp[s_2(\mathbf{v}_1)] + t_2(\mathbf{v}_1)$$

where s_1, s_2, t_1, t_2 are scale and shift operators modeled using standard neural networks (detailed in Sec. 3). A deep INN can then be constructed by composing K such coupling layers. The overall representation power of the deep INN can be enhanced by permuting the inputs to each coupling layer [2, 12]. The use of coupling layer constrains the Jacobian to be upper triangular; meaning $\det \partial \mathcal{G}(\mathbf{w})/\partial \mathbf{w}$ in Eqn. 2 can be computed efficiently [8].




Fig. 1: Schematic of a coupling layer

3 Numerical Results

We demonstrate the proposed approach on a synthetic sample (see Fig. 2) having the cameraman and baboon images as the magnitude ($|z|$) and phase ($\arg(z)$) components of z , respectively. The object is scanned 8 times along each axis (i.e., $N = 64$) using a probe of size $m = 36$. We investigate three different scanning settings corresponding to three different strides of the probe. Associated with each setting is a field of view (FOV) corresponding to the different region of the image that being illuminated (i.e., different n , see color-coded boundaries in Fig. 2), however, with the same number of data, i.e., \mathbf{d} is of dimension $64 \times 36 \times 36$ for all three settings. Furthermore, we add 1% Gaussian noise to the measurements to reflect practical scenarios. The sensor noise can be modeled as Gaussian when the photon counts are large due to the fact that the limiting distribution of the Poisson distribution is Gaussian.

For all the results shown herein, we use deep INNs with 16 blocks. Each block consists of two affine coupling layers, as shown in Fig. 1, and each coupling layer is preceded by an *activation normalization*






Fig. 2: Ground truth object: magnitude (left) and phase (right).

Table 1: Different scan settings and reconstruction results

Scan Setting	FOV n	Overlap Ratio	Recon. Mag. PSNR			Recon. Phase SSIM		
			NF Mode 1	NF Mode 2	rPIE	NF Mode 1	NF Mode 2	rPIE
<i>S1</i>	50	0.94	26.69	26.64	26.82	0.48	0.54	0.25
<i>S2</i>	78	0.83	22.77	23.66	23.08	0.20	0.39	0.18
<i>S3</i>	92	0.78	21.20	22.69	22.32	0.50	0.44	0.16

unit [12]. The dimensions are permuted and divided into two equal parts at the beginning of each block. The scale and shift transformations are modeled using dense neural networks with two $n^2/4$ -dimensional hidden layers. Leaky RELU activation is used and batch normalization is also applied. The INNs are trained for 2×10^4 epochs with a learning rate of 5×10^{-6} and a batch size of 64. Given \mathcal{L} as Eqn. (3), \mathcal{R} corresponds to a total variation prior [27] with $\lambda = 0.01$. Following [11], the reconstructed magnitudes are compared using PSNR whereas the reconstructed phases are compared using SSIM after the predicted phase is linearly mapped to the range ground truth region $[0, 0.4]$. Table 1 details the scanning settings with the corresponding reconstruction results using NFs and their comparison against rPIE. All experiments were performed on a computer with a 48GB NVIDIA Quadro RTX 8000 GPU and 64GB of RAM.


We first focus on setting *S3*, and perform an analysis of the samples drawn from the generative model using the t-distributed stochastic neighbor embedding (t-SNE) [32], which reveals bi-modality in the posterior (see Fig. 3(left)). These modes reflect the classical *bias-variance trade-off* of estimators. Figure 3(right) shows that mode 1 has higher bias than mode 2 as evidenced by the slightly right-shifted distribution of \mathcal{L} for mode 1. However, as seen later in Fig. 4, mode 1 has lower variance particularly for the phase reconstruction.

**Fig. 3:** Left: bi-modality revealed by t-SNE; right: distribution of \mathcal{L} within each mode with mode 1 having higher bias.**Fig. 4:** Posterior mean and standard deviation (SD) of the reconstructed magnitude and phase for both modes.

Moreover, the artifacts in the phase reconstruction (seen as the tripod leg), as a consequence of the cross talks between magnitude and phase, show higher standard deviations. This is an extremely useful result in order to judge the reconstruction quality in practice since the ground truth is often unavailable in real applications.

We further quantify the reconstruction qualities caused by different scanning settings. As the settings move from *S1* to *S3*, reconstruction task potentially becomes harder due to the less overlap. Consistently, the standard deviation increases over common areas of the three FOVs as the overlap ratio reduces, as shown in Fig. 5. The

In Fig. 4, we show the posterior mean and its associated uncertainty of the two modes (labeled as mode 1 and 2, respectively). Both modes yield equally good reconstructions of the object magnitude; however, the phase reconstructions vary across the modes (see Table 1). In particular, as expected, the edges (particularly corners) of the FOV are reconstructed with higher standard deviations due to the limited overlap of the scans near the edges.

**Fig. 5:** Logarithm (base 10) of standard deviation of phase reconstruction in mode 2 for *S1*, *S2* and *S3* over common regions.

lower SSIM of the reconstructed phases in setting S2 as compared to S3 in Table 1 may be an artefact that indicates the need for better hyperparameter tuning or longer training of the INNs.

4 Conclusions

We demonstrate that normalizing flows can be used as a framework for the statistical analysis of inverse problems arising in ptychography on a synthetic sample. First, the uncertainty quantification returned by the proposed NFs shows that the reconstruction quality deteriorates as the overlap ratio reduces. This is expected since the overlap ratio is a measure of oversampling, less the extent of oversampling leads to greater the ill-posedness of the inverse problem. Second, with phase reconstruction as the main goal of phase retrieval in general, the proposed NF outperforms traditional method by providing better phase reconstruction. Overall, the main contribution of the proposed NF is its capability of quantifying uncertainties associated with the reconstruction, which is essential in judging the solution quality in the absence of ground truth. Future direction includes a careful prior design to provide more guidance on the mode selection. The application of the proposed methodology on a larger object with more data will be explored in a future work.

Acknowledgments and Disclosure of Funding

This material was based upon work supported by the U.S. Department of Energy under contract DE-AC02-06CH11357. The first author also acknowledges the support of the Provost Fellowship from the University of Southern California.

References

- [1] M. Albergo, G. Kanwar, and P. Shanahan. Flow-based generative models for Markov chain Monte Carlo in lattice field theory. *Physical Review D*, 100(3):034515, 2019.
- [2] L. Ardizzone, J. Kruse, S. Wirkert, D. Rahner, E. W. Pellegrini, R. S. Klessen, L. Maier-Hein, C. Rother, and U. Köthe. Analyzing inverse problems with invertible neural networks. *arXiv preprint arXiv:1808.04730*, 2018.
- [3] L. Ardizzone, C. Lüth, J. Kruse, C. Rother, and U. Köthe. Guided image generation with conditional invertible neural networks. *arXiv preprint arXiv:1907.02392*, 2019.
- [4] J. M. Bardsley. Mcmc-based image reconstruction with uncertainty quantification. *SIAM Journal on Scientific Computing*, 34(3):A1316–A1332, 2012.
- [5] H. Chang, P. Enfedaque, J. Zhang, J. Reinhardt, B. Enders, Y.-S. Yu, D. Shapiro, C. G. Schroer, T. Zeng, and S. Marchesini. Advanced denoising for x-ray ptychography. *Optics express*, 27(8):10395–10418, 2019.
- [6] M. J. Cherukara, T. Zhou, Y. Nashed, P. Enfedaque, A. Hexemer, R. J. Harder, and M. V. Holt. AI-enabled high-resolution scanning coherent diffraction imaging. *Applied Physics Letters*, 117(4):044103, 2020.
- [7] T. Cui, J. Martin, Y. M. Marzouk, A. Solonen, and A. Spantini. Likelihood-informed dimension reduction for nonlinear inverse problems. *Inverse Problems*, 30(11):114015, 2014.
- [8] L. Dinh, D. Krueger, and Y. Bengio. NICE: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.
- [9] L. Dinh, J. Sohl-Dickstein, and S. Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- [10] B. Enders and P. Thibault. A computational framework for ptychographic reconstructions. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 472(2196):20160640, 2016.
- [11] S. W. Fung and Z. Di. Multigrid optimization for large-scale ptychographic phase retrieval. *SIAM Journal on Imaging Sciences*, 13(1):214–233, 2020.
- [12] D. P. Kingma and P. Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *arXiv preprint arXiv:1807.03039*, 2018.
- [13] I. Kobyzev, S. Prince, and M. Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [14] P. C. Konda, L. Loetgering, K. C. Zhou, S. Xu, A. R. Harvey, and R. Horstmeyer. Fourier ptychography: current applications and future promises. *Optics express*, 28(7):9603–9630, 2020.
- [15] M. Kumar, M. Babaeizadeh, D. Erhan, C. Finn, S. Levine, L. Dinh, and D. Kingma. Videoflow: A conditional flow-based model for stochastic video generation. *arXiv preprint arXiv:1903.01434*, 2019.
- [16] S. Li, M. Deng, J. Lee, A. Sinha, and G. Barbastathis. Imaging through glass diffusers using densely connected convolutional networks. *Optica*, 5(7):803–813, 2018.
- [17] A. Maiden, D. Johnson, and P. Li. Further improvements to the ptychographical iterative engine. *Optica*, 4(7):736–745, 2017.
- [18] S. Mo, N. Zabarar, X. Shi, and J. Wu. Deep autoregressive neural networks for high-dimensional inverse problems in groundwater contaminant source identification. *Water Resources Research*, 55(5):3856–3881, 2019.
- [19] G. Ongie, A. Jalal, C. A. Metzler, R. G. Baraniuk, A. G. Dimakis, and R. Willett. Deep learning techniques for inverse problems in imaging. *IEEE Journal on Selected Areas in Information Theory*, 1(1):39–56, 2020.
- [20] G. A. Padmanabha and N. Zabarar. Solving inverse problems using conditional invertible neural networks. *Journal of Computational Physics*, 433:110194, 2021.
- [21] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *Journal of Machine Learning Research*, 22(57):1–64, 2021.
- [22] F. Pfeiffer. X-ray ptychography. *Nature Photonics*, 12(1):9–17, 2018.

- [23] R. Prenger, R. Valle, and B. Catanzaro. Waveglow: A flow-based generative network for speech synthesis. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3617–3621. IEEE, 2019.
- [24] D. Rezende and S. Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR, 2015.
- [25] Y. Rivenson, Y. Zhang, H. Günaydin, D. Teng, and A. Ozcan. Phase recovery and holographic image reconstruction using deep learning in neural networks. *Light: Science & Applications*, 7(2):17141–17141, 2018.
- [26] J. M. Rodenburg. Ptychography and related diffractive imaging methods. *Advances in imaging and electron physics*, 150:87–184, 2008.
- [27] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena*, 60(1-4):259–268, 1992.
- [28] A. Siahkoobi and F. J. Herrmann. Learning by example: fast reliability-aware seismic imaging with normalizing flows. *arXiv preprint arXiv:2104.06255*, 2021.
- [29] A. Sinha, J. Lee, S. Li, and G. Barbastathis. Lensless computational imaging through deep learning. *Optica*, 4(9):1117–1125, 2017.
- [30] H. Sun and K. L. Bouman. Deep probabilistic imaging: Uncertainty quantification and multi-modal solution characterization for computational imaging. *arXiv preprint arXiv:2010.14462*, 9, 2020.
- [31] D. F. Swinehart. The beer-lambert law. *Journal of chemical education*, 39(7):333, 1962.
- [32] L. Van der Maaten and G. Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(11), 2008.
- [33] X. Wei, H. P. Urbach, and W. M. Coene. Cramér-rao lower bound and maximum-likelihood estimation in ptychography with poisson noise. *Physical Review A*, 102(4):043516, 2020.
- [34] Z. Wei and X. Chen. Uncertainty quantification in inverse scattering problems with Bayesian convolutional neural networks. *IEEE Transactions on Antennas and Propagation*, 2020.
- [35] M. J. Williams, J. Veitch, and C. Messenger. Nested sampling with normalizing flows for gravitational-wave inference. *Physical Review D*, 103(10):103006, 2021.
- [36] Y. Xue, S. Cheng, Y. Li, and L. Tian. Reliable deep-learning-based phase imaging with uncertainty quantification. *Optica*, 6(5):618–629, 2019.
- [37] Y. Zhu, N. Zabaras, P.-S. Koutsourelakis, and P. Perdikaris. Physics-constrained deep learning for high-dimensional surrogate modeling and uncertainty quantification without labeled data. *Journal of Computational Physics*, 394:56–81, 2019.
- [38] Z. Ziegler and A. Rush. Latent normalizing flows for discrete sequences. In *International Conference on Machine Learning*, pages 7673–7682. PMLR, 2019.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes] We propose to use NFs for the reconstruction and uncertainty quantification in the ptychography inverse problem. In Section 3, we show for a synthetic object that NFs can (1) provide good quality reconstructions compared to state-of-the-art iterative approaches, (2) can also help quantify uncertainties in the reconstruction, under different experimental settings. In fact, the phase retrieval capabilities of NFs were found to be better than rPIE.
 - (b) Did you describe the limitations of your work? [Yes] In Section 3, we show that the posterior density approximated using NFs can be bi-modal. Our proposed approach is unable to provide any guidance on which mode to be used for the final reconstruction. In this context, we emphasize the importance of prior selection.
 - (c) Did you discuss any potential negative societal impacts of your work? [No] We do not foresee any negative impacts of our work.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes] Our paper conforms to the ethics review guidelines set forth by the conference. We have tried our best to envisage any potential ethical dilemmas.
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [N/A] We do not present any new theoretical results.
 - (b) Did you include complete proofs of all theoretical results? [N/A] No theorems or theoretical results were stated in this work.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] The code and any pertinent data will be shared upon request by contacting the corresponding author.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] We have mentioned all the pertinent choices made in relation to hyperparameters and other parameters necessary to reproduce the results.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [No] Training INNs is computationally expensive. Hence the effects of random seed could not be studied. If requested, we will perform further ablation studies and report on the effects of random seeds.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] In Section 3, we clearly mention that we carried out all experiments on a computer equipped with a 48GB NVIDIA Quadro RTX 8000 GPU on top of a 64GB of RAM.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes] Our implementation of INN using coupling blocks is based on the publicly available code provided by Sun and Bouman [30], which we cite in the introduction, and can be found at <https://github.com/HeSunPU/DPI>. We have implemented ourselves the remaining necessary components including forward model and constraints.
 - (b) Did you mention the license of the assets? [N/A] There is no license attached to the aforementioned code base.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...

- (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A] We have not used any crowd sourced data or conducted any research involving human subjects.
- (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]