LoDIP: Low-dose phase retrieval with deep image prior

Raunak Manekar[1*] **Elisa Negrini**[2*] **Minh Pham**[2] **Daniel Jacobs**[3]
Jaideep Srivastava[1] **Stanley J. Osher**[2] **Jianwei Miao**[3]

[1]Computer Science and Engineering, University of Minnesota, Twin Cities
[2]Mathematics, University of California, Los Angeles
[3]Physics and Astronomy, and California NanoSystems Institute, University of California, Los Angeles
{manek009,srivasta}@umn.edu {danieljacobs,miao}@physics.ucla.edu
{enegrini,sjo}@math.ucla.edu minhrose@ucla.edu

Abstract

Phase retrieval under very low dose conditions is a challenging problem as all the phase retrieval algorithms become unstable with the presence of very high Poisson noise. To mitigate this problem, in-situ coherent diffractive imaging (CDI) has been previously proposed which places a high-dose static region next to the sample region while imaging. Iterative phase retrieval algorithms are then used to reconstruct both regions from the diffraction patterns with high signal to noise ratio. While numerical simulations have indicated that in-situ CDI can reduce radiation dose by one to two orders of magnitude over conventional CDI, it requires multiple measurements with a common high-dose static region. Here we demonstrate low-dose phase retrieval with deep image prior, termed LoDIP, for in-situ CDI. Using both numerical and experimental data, we demonstrate that LoDIP outperfroms popular iterative phase retrieval algorithms under low-dose conditions. Our results show that LoDIP is not sensitive to the choice of the static structure nor to the geometric arrangement between the two objects. Additionally, unlike previous successful work with in situ CDI, LoDIP does not depend on multiple measurements with a common static region. We expect that the combination of deep-learning phase retrieval with in situ CDI will create numerous opportunities for high-resolution quantitative phase imaging for dose-sensitive materials, such as biological samples, polymers, organic semiconductors, and energy materials.

1 Introduction and Background

Low-dose CDI Coherent diffractive imaging (CDI) is a lensless imaging technique (Miao et al. [1999]) that has found broad applications across different disciplines (Miao et al. [2015]). In CDI a highly coherent wave is incident on an object. The scattered wave by the object produces a diffraction pattern which is then collected by a detector. If the diffraction pattern is sufficiently oversampled, i.e. the independently measured points are more than the unknown variable (Miao et al. [1998]), then the phase can be in principle retrieved from the diffraction pattern via iterative phase-retrieval(PR) algorithms (Shechtman et al. [2015]). However, many samples of interest for CDI, such as biological material, polymers or organic semiconductors, require to be imaged with minimal radiation exposure (low-dose CDI) to prevent damage during data acquisition. As a consequence, the inherent low signal-to-noise ratio in low-dose measurements poses a challenge for classical phase retrieval algorithms, as they become severely unstable under heavy Poisson noise.

*These authors contributed equally to this work

Machine Learning and the Physical Sciences Workshop, NeurIPS 2023.

Methods for Low-dose CDI For low-dose CDI, the most successful method is in-situ CDI Lo et al. [2018]. The experimental setup for in-situ CDI is described in Section 2. While successful, this requires a time-series of diffraction patterns, with the need to be collected with a time-invariant static region constraint (see Section 2). In this work, we propose a new method LoDIP which adopts this experimental setup but uses an implicit neural prior in the iterative phase-retrieval(PR) algorithm. Experiments show that this combination gives similar performance with a single diffraction pattern. This also bypasses the need for ensuring a time-invariant static constraint, thus greatly simplifying the experimental process.

Existing PR algorithms Many algorithms have been developed to solve the PR problems with single diffraction pattern only Fienup [1978], Miao et al. [2000], Bauschke et al. [2002], Luke [2004], Rodriguez et al. [2013], Pham et al. [2019]. While powerful, these algorithms require tuning of several algorithmic parameters and expert strategies. And most algorithms face severe performance degradation in low-dose CDI due to very high Poisson noise.

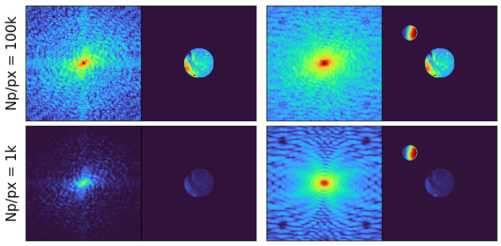

Figure 1: LoDIP experimental setup schematic: For imaging samples in low-light, a static region is added in the field of view and exposed to the full fluence of the source illumination. Whereas the fluence on the dose sensitive sample is reduced to a tolerable limit as measured in photons per pixel (Np/px). The first two columns are the diffraction pattern and the image sample without a static structure. And the last two columns are the same sample with a static region. At lower illumination (bottom row), the cell is very dark and it is hard to make out high resolution details.

Recently, data-driven methods (i.e. supervised learning) have shown great potential for solving inverse problems in computational imaging Kamilov et al. [2015], Wang et al. [2019], Wu et al. [2016], Goy et al. [2018], Li et al. [2019], Xue et al. [2019] including phase retrieval (for CDI) Chang et al. [2023]. This method, while successful, necessitates a substantial amount of labeled data. Moreover, depending on the network architecture and the amount of data, the training process can take several days.

Deep Learning for single-instance PR In this work, inspired by the work on Deep Image Prior (DIP) Ulyanov et al. [2018], we propose to use the deep image prior method (untrained neural prior) for single-instance phase retrieval. Similar to DIP, LoDIP does not require training om large labelled dataset. We note that the DIP framework alone has been used in some computational imaging applications before Anirudh et al. [2018], Liu et al. [2019], Jagatap and Hegde [2019]. However, it was never applied to the problem of low-dose imaging exploiting directly the physical set up of in-situ CDI. In our experiments we show that the use of a high dose static region greatly improves the quality of the sample reconstruction especially in the low dose regime and obtains results comparable or superior to state of the art methods.

2 Proposed Method

LoDIP: Experimental setup The method proposed in this work combines deep image prior [Ulyanov et al., 2018] and in-situ CDI [Lo et al., 2018] for low-dose phase retrieval. Our method exploits the power of an untrained neural prior, which does not require a large dataset for training, together with the experimental setup of in-situ CDI. In in-situ CDI, the sample of interest is placed within a finite support next to a heavily scattering, dose-tolerant static region such as a lacey Au pattern on an optical stage. A source of coherent X-rays provides plane-wave incident illumination at an energy suited to the sample thickness and K absorption edges of atomic species present within the sample. For biological specimens, a beam energy of 530 eV corresponds well to the window between the carbon and oxygen absorption edges and is the energy used in the simulations presented in this work [Kirz et al., 1995].

The fluence(radiation dose) on the dose-sensitive sample is reduced to a tolerable limit by the presence of an attenuator while the static region is exposed to the full fluence of the incident illumination. Far-field diffraction patterns are recorded from this setup containing Poisson noise relative to the

total fluence on the detector. An untrained neural prior combined with an accurate forward model for the imaging setup is then used to perform phase retrieval and recover the exit wave from the sample, exploiting the interference in Fourier space between the high-fluence static region and the sample to increase the signal to noise ratio (SNR) in measurements.

PR with untrained neural prior Given an image of an object $\boldsymbol{X} \in \mathbb{C}^{n\times n}$, the measurement process to capture a diffraction pattern $\boldsymbol{Y} \in \mathbb{R}^{m\times m}$ can be described as: $\boldsymbol{Y} = |\mathcal{F}(\boldsymbol{X})|^2$ To meet the oversampling criteria, the original object $\boldsymbol{X} \in \mathbb{C}^{n\times n}$ is zero-padded to size $m \times m$, where $m = 2 \times n$. And a support $\boldsymbol{S}_0$ represents the location of the object in this empty background. The overall goal of the phase retrieval problem is to recover the image $\boldsymbol{X}$ from the captured diffraction pattern $\boldsymbol{Y}$. In the presence of a non-overlapping static structure $\boldsymbol{U} \in \mathbb{C}^{n\times n}$ to the above setup, the optimization problem can be specified as:

$$\min_{\hat{X}\in\mathbb{C}^{n\times n}} \ell\left(\boldsymbol{Y}, \left|\mathcal{F}\left(\hat{X}+\boldsymbol{U}\right)\right|^2\right), s.t.\,(1-\boldsymbol{S}_0)\odot\hat{\boldsymbol{X}} = [\boldsymbol{0}]_{m\times m} \tag{2.1}$$

Here, the objective function represents the data consistency term; while the support constraint utilizes the known support, requiring the off-support values to be zero.

To impose an implicit neural prior on the variable $\hat{X}$, we reparametrize it with a neural network $g_{\boldsymbol{W}}(\boldsymbol{z})$. $\boldsymbol{W}$ represents the learnable parameters i.e. weights of the neural network and $\boldsymbol{z}$ is a fixed input, in our case the given diffraction pattern. Following the literature on this topic, we have used a U-Net Ronneberger et al. [2015] with skip connections and ReLU activation functions.

We observe that the results are not sensitive to the relative size or relative location of the sample and the static structure, nor to the specific choice of the static structure. The method also works with an inaccurate estimate of the support. This makes LoDIP a very general and robust method which has potential to be applied to a wide range of experimental data. Finally, unlike Fourier holography McNulty et al. [1992], Barmherzig et al. [2019], the proposed method works with different illuminations of the cell and the static structure.

To compare the methods, we use (1) physically accurate simulations of a gold lacey for the static structure and a biological cell for the sample region Fig. 2, as well as (2) diffraction patterns experimentally measured with a 534 nm HeNe laser Fig. 3.

3 Experimental Results and Discussion

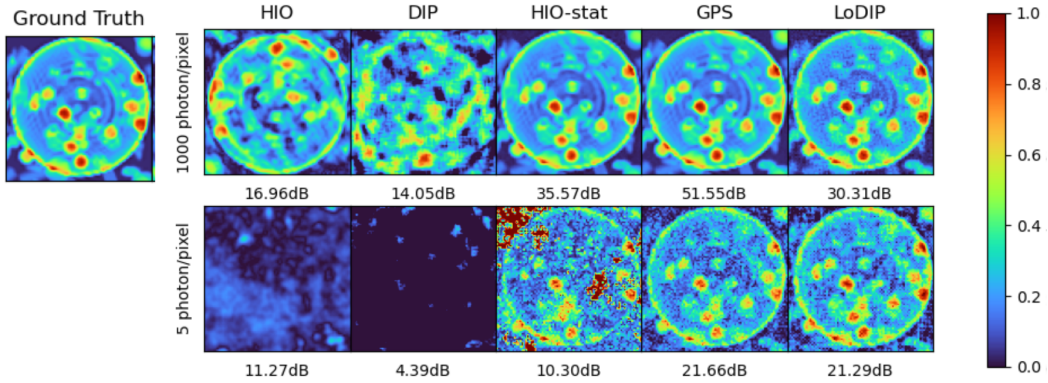

Figure 2: Experimental Results on biological cell sample. **(Top row)** Reconstruction at high photon count (1000 photon/pixel). **(Bottom row)** Reconstruction at low-photon counts (5 photon/pixel). For each method we report the peak signal-to-noise ratio (PSNR), larger the better.

Data: simulated and experimental The simulated data uses a simulated 20-nm thick gold lacey pattern as the static structure and a simulated cell consisting of a vesicle containing water and protein

aggregates. The illumination on the static structure has been fixed at 10^{10} photons per μm^{-2}. While the illumination on the sample has been varied from 10^3 to 10^7 photons per μm^{-2}. The lighting conditions have been graded in number of incident photons per pixel (Np/px in Fig. 2) and Poisson noise has been applied based on the total number of incident photons.

For the optical laser data, the static structure is a 100 µ pinhole exposed to the same incident illumination as the sample. This data was reused with permission from Lo et al. [2018]. Further information about the generation of simulated data and optical laser data collection can be found in their previous work on In-Situ CDI [Lo et al., 2018].

Reconstruction of biological sample In Fig. 2, we can see the relative performance of the proposed method vs. popular methods which can be used in this setup. HIO-ref is a modification of the most popular phase retrieval algorithm Hybrid input-output (HIO) Fienup [1982] to use the static region constraint. Generalized proximal smoothing (GPS) Pham et al. [2019] is the state of the art method for phase retrieval.

For the high illumination case of 1000 photons/pixel, it can be seen that both HIO-ref and LoDIP perform comparably both in terms of PSNR and visually. At low-light conditions, the performance of all methods degrades. But, LoDIP performs comparable to GPS and better than all other methods. GPS performs very well on simulated data. But no available version of GPS can work on experimental data in the presence of non-idealities such as a probe function. This prevents it from being used on the experiments in the next section.

Reconstruction from experimental data As a proof-of-concept experiment, we demonstrate the proposed method on experimentally captured diffraction patterns (shown in Fig. 3 left column). The object is complex-valued and the optical setup includes a probe. The in-situ CDI method Lo et al. [2018] uses 50 diffraction patterns with a fixed time-invariant reference object in all the images. HIO-ref and LoDIP use a single diffraction pattern. Since there is no ground truth available, we use the relative error (R-factor) calculated in the Fourier domain: The R-factor R_F is a measure of the similarity between the captured diffraction pattern $\boldsymbol{Y}$ to the Fourier magnitudes of the reconstruction $\hat{\boldsymbol{X}}$. The table in Fig. 3 shows the average R_F from 20 independent reconstructions of a single diffraction pattern. LoDIP gives performance comparable to in-situ CDI without requiring multiple diffraction patterns. Fig. 3 shows the means and variance image obtained from the top 5 reconstructions out of 20 independent runs. The LoDIP reconstructions have lower variance and thus are more consistent. The reconstruction from HIO-ref also has comparable R_F but shows strong visible artefacts.

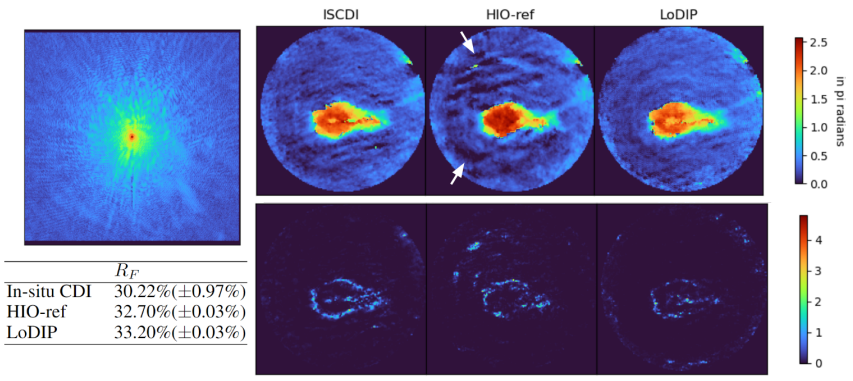

	R_F
In-situ CDI	30.22%(±0.97%)
HIO-ref	32.70%(±0.03%)
LoDIP	33.20%(±0.03%)

Figure 3: Reconstruction using experimental data. The proposed method LoDIP uses a single diffraction pattern to produce a reconstruction of comparable quality as in-situ CDI which uses 50 samples. The arrow indicates the artefacts in the HIO-ref reconstructions.**(Left)** (Top) Sample of experimentally captured diffraction pattern. (Bottom) Quantitative comparison. Mean and standard deviation of R-factor(R_F, **lower the better**). **(Right)** Means (Top) and variance (Bottom) calculated of the top five reconstructions from 20 independent runs.

Conclusion The LoDIP method leveraging an untrained neural prior offers a robust approach for phase retrieval at low photon counts. As demonstrated by our experiments, LoDIP outperforms other iterative methods such as HIO in image reconstruction under low-dose conditions, and, unike GPS, it can be successfully applied to experimental data. We anticipate applications of the LoDIP method to X-ray imaging of dose-sensitive samples of importance in a variety of fields, such as organic semiconductors relevant to modern perovskite solar cells and battery materials, as well as biological samples concerning cellular interactions and life cycles.

References

Rushil Anirudh, Jayaraman J Thiagarajan, Bhavya Kailkhura, and Timo Bremer. An unsupervised approach to solving inverse problems using generative adversarial networks. *arXiv preprint arXiv:1805.07281*, 2018.

David A Barmherzig, Ju Sun, Po-Nan Li, Thomas Joseph Lane, and Emmanuel J Candes. Holographic phase retrieval and reference design. *Inverse Problems*, 35(9):094001, 2019.

Heinz H. Bauschke, Patrick L. Combettes, and D. Russell Luke. Phase retrieval, error reduction algorithm, and fienup variants: a view from convex optimization. *Journal of the Optical Society of America A*, 19(7):1334, jul 2002. doi: 10.1364/josaa.19.001334.

Dillan J Chang, Colum M O'Leary, Cong Su, Daniel A Jacobs, Salman Kahn, Alex Zettl, Jim Ciston, Peter Ercius, and Jianwei Miao. Deep-learning electron diffractive imaging. *Physical review letters*, 130(1):016101, 2023.

J. R. Fienup. Phase retrieval algorithms: a comparison. *Applied Optics*, 21(15):2758, aug 1982. doi: 10.1364/ao.21.002758.

James R Fienup. Reconstruction of an object from the modulus of its fourier transform. *Optics letters*, 3(1):27–29, 1978.

Alexandre Goy, Kwabena Arthur, Shuai Li, and George Barbastathis. Low photon count phase retrieval using deep learning. *Physical review letters*, 121(24):243902, 2018.

Gauri Jagatap and Chinmay Hegde. Phase retrieval using untrained neural network priors. In *NeurIPS 2019 Workshop on Solving Inverse Problems with Deep Networks*, 2019.

Ulugbek S Kamilov, Ioannis N Papadopoulos, Morteza H Shoreh, Alexandre Goy, Cedric Vonesch, Michael Unser, and Demetri Psaltis. Learning approach to optical tomography. *Optica*, 2(6): 517–522, 2015.

Janos Kirz, Chris Jacobsen, and Malcolm Howells. Soft x-ray microscopes and their biological applications. *Quarterly reviews of biophysics*, 28(1):33–130, 1995.

Xiu Li, Huaiyu Qi, Shaowei Jiang, Pengming Song, Guoan Zheng, and Yongbing Zhang. Quantitative phase imaging via a cgan network with dual intensity images captured under centrosymmetric illumination. *Optics Letters*, 44(11):2879–2882, 2019.

Jiaming Liu, Yu Sun, Xiaojian Xu, and Ulugbek S Kamilov. Image restoration using total variation regularized deep image prior. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7715–7719. Ieee, 2019.

Yuan Hung Lo, Lingrong Zhao, Marcus Gallagher-Jones, Arjun Rana, Jared J. Lodico, Weikun Xiao, BC Regan, and Jianwei Miao. In situ coherent diffractive imaging. *Nature communications*, 9(1): 1826, 2018.

D Russell Luke. Relaxed averaged alternating reflections for diffraction imaging. *Inverse problems*, 21(1):37, 2004.

Ian McNulty, Janos Kirz, Chris Jacobsen, Erik H Anderson, Malcolm R Howells, and Dieter P Kern. High-resolution imaging by fourier transform x-ray holography. *Science*, 256(5059):1009–1012, 1992.

Jianwei Miao, David Sayre, and HN Chapman. Phase retrieval from the magnitude of the fourier transforms of nonperiodic objects. *JOSA A*, 15(6):1662–1669, 1998.

Jianwei Miao, Pambos Charalambous, Janos Kirz, and David Sayre. Extending the methodology of x-ray crystallography to allow imaging of micrometre-sized non-crystalline specimens. *Nature*, 400(6742):342–344, 1999.

Jianwei Miao, J Kirz, and D Sayre. The oversampling phasing method. *Acta Crystallographica Section D: Biological Crystallography*, 56(10):1312–1315, 2000.

Jianwei Miao, Tetsuya Ishikawa, Ian K Robinson, and Margaret M Murnane. Beyond crystallography: Diffractive imaging using coherent x-ray light sources. *Science*, 348(6234):530–535, 2015.

Minh Pham, Penghang Yin, Arjun Rana, Stanley Osher, and Jianwei Miao. Generalized proximal smoothing (gps) for phase retrieval. *Optics Express*, 27(3):2792–2808, 2019.

Jose A Rodriguez, Rui Xu, C-C Chen, Yunfei Zou, and Jianwei Miao. Oversampling smoothness: an effective algorithm for phase retrieval of noisy diffraction intensities. *Journal of applied crystallography*, 46(2):312–318, 2013.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.

Yoav Shechtman, Yonina C. Eldar, Oren Cohen, Henry Nicholas Chapman, Jianwei Miao, and Mordechai Segev. Phase retrieval with application to optical imaging: A contemporary overview. *IEEE Signal Processing Magazine*, 32(3):87–109, may 2015. doi: 10.1109/msp.2014.2352673.

Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9446–9454, 2018.

Fei Wang, Hao Wang, Haichao Wang, Guowei Li, and Guohai Situ. Learning from simulation: An end-to-end deep-learning approach for computational ghost imaging. *Optics express*, 27(18): 25560–25572, 2019.

Gang Wu, Thomas Nowotny, Yongliang Zhang, Hong-Qi Yu, and David Day-Uei Li. Artificial neural network approaches for fluorescence lifetime imaging techniques. *Optics letters*, 41(11): 2561–2564, 2016.

Yujia Xue, Shiyi Cheng, Yunzhe Li, and Lei Tian. Reliable deep-learning-based phase imaging with uncertainty quantification. *Optica*, 6(5):618–629, 2019.