

Practical Phase Retrieval Using Double Deep Image Priors

Zhong Zhuang¹, David Yang², Felix Hofmann², David A. Barmherzig³, Ju Sun⁴

¹Electrical & Computer Engineering, University of Minnesota, Minneapolis, USA

²Department of Engineering Science, University of Oxford, Oxford, UK

³Center for Computational Mathematics, Flatiron Institute, New York City, USA

⁴Computer Science & Engineering, University of Minnesota, Minneapolis, USA

Abstract

Phase retrieval (PR) concerns the recovery of complex phases from complex magnitudes. We identify the connection between the difficulty level and the number and variety of symmetries in PR problems. We focus on the most difficult far-field PR (FFPR), and propose a novel method using double deep image priors. In realistic evaluation, our method outperforms all competing methods by large margins. As a single-instance method, our method requires no training data and minimal hyperparameter tuning, and hence enjoys good practicality. Our paper is also available at: <https://arxiv.org/abs/2211.00799>.

Introduction

In scientific imaging, observable physical quantities about the object of interest are often complex-valued, e.g., when diffraction happens (1). However, practical detectors can only record complex magnitudes, but not phases, resulting in phaseless observations. Phase retrieval (PR), broadly defined, is the nonlinear inverse problem (NIP) of estimating the object of interest from the phaseless observations. PR is central to coherent diffraction imaging ((B)CDI) (2; 3), image-based wavefront sensing (4), radar and sonar sensing (5); see the recent survey (6).

Which phase retrieval (PR)? Without loss of generality, consider a 2D object of interest $\mathbf{X} \in \mathbb{C}^{m \times n}$, and a physical observation model \mathcal{A} that leads to an ideal complex-valued observation $\mathcal{A}(\mathbf{X}) \in \mathbb{C}^{m' \times n'}$. However, the detector can only record $\mathbf{Y} = |\mathcal{A}(\mathbf{X})|^2$, where $|\cdot|^2$ denotes the elementwise squared magnitudes—corresponding to the photon flux the detector is able to capture. In **far-field (Fraunhofer) PR (FFPR) that stems from far-field propagation and is also the focus of this paper**, \mathcal{A} is the oversampled 2D Fourier transform \mathcal{F} with $m' \geq 2m - 1$, and $n' \geq 2n - 1$ to ensure recoverability. For notational simplicity, we assume m, n are odd numbers, and \mathbf{X} is centered at $(0, 0)$ so that index ranges are $[-(m-1)/2, (m-1)/2]$ and $[-(n-1)/2, (n-1)/2]$, respectively. Numerous other \mathcal{A} 's have been studied in the literature, notable ones including:

- **Generalized PR (GPR):** $\mathcal{A}(\mathbf{X}) = \{\langle \mathbf{A}_i, \mathbf{X} \rangle\}_{i=1}^k$ where \mathbf{A}_i 's are iid Gaussian or randomly-masked Fourier basis matrices (7; 8). These elegant mathematical models do not correspond to physically feasible imaging systems so far;
- **Near-Field (Fresnel) PR (NFPR):** $\mathcal{A}(\mathbf{X}) = \mathcal{F}(\mathbf{X} \odot [e^{i\pi\beta(\ell^2 + \kappa^2)}]_{\ell, \kappa})$ (9; 10), where the constant $\beta > 0$ depends on the sampling intervals, wavelength, and imaging distance (11),

comes from near-field propagation. Note that FFPR corresponds to $\beta \rightarrow 0$, and PR problems solved in image-based wavefront sensing for astronomical applications correspond to multi-plane near-field propagation with sequential optical aberrations (12);

- **Holographical PR (HPR):** $\mathcal{A}(\mathbf{X}) = \mathcal{F}([\mathbf{X}, \mathbf{R}])$, where \mathbf{R} is a known reference that is placed side-by-side with the object of interest \mathbf{X} (13); depending on the propagation distance, near-field versions are also possible (1, Chapter 11);
- **Ptychography (PTY):** \mathbf{X} is raster-scanned by a sharp illumination pattern \mathbf{W} that is focused over a local patch of \mathbf{X} each time. Now \mathbf{Y} is the set of magnitude measurements $\mathbf{Y}_i = |\mathcal{F}(\mathbf{W} \odot \mathbf{X}(\mathbf{p}_i))|^2$, where \mathbf{p}_i indexes the raster grid (14; 15).




Figure 1. Illustration of the three intrinsic symmetries in FFPR on simulated complex-valued crystal data (see Experiments Results for details). Any composition of 2D conjugate flipping, translation, and global phase, when applied to \mathbf{X} , leads to the same set of magnitudes \mathbf{Y} .

Symmetry matters Identifiability in PR is often up to intrinsic symmetries. For example, any global phase factor $e^{i\theta}$ added to \mathbf{X} leaves \mathbf{Y} unchanged for FFPR, NFPR, GPR, and PTY, i.e., **global phase symmetry**. While this is the only symmetry for NFPR, GPR, and PTY, FFPR has two other symmetries: translation and 2D conjugate flipping, as demonstrated in Fig. 1 (16). A crucial empirical observation is **the difficulty level of a PR problem is proportional to the number of its symmetries**. To see the point, consider a natural least-squares (LS) formulation of PR²

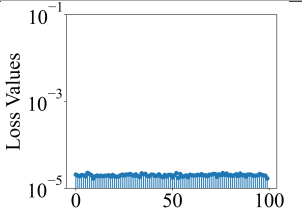
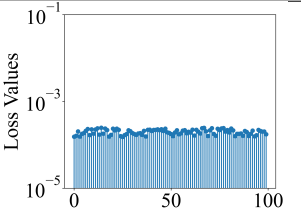
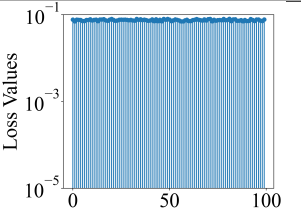
$$\min_{\mathbf{Z} \in \mathbb{C}^{m \times n}} \frac{1}{m'n'} \|\sqrt{\mathbf{Y}} - |\mathcal{A}(\mathbf{Z})|\|_F^2, \quad (0.1)$$

with the groundtruth complex-valued 2D crystal sample in Fig. 1 as the target \mathbf{X} . On GPR with Gaussian, NFPR, and FFPR, we run subgradient descent with 100 random starts, respectively, and record their final convergent losses. As is evident from Table 1, while we can consistently find numerically satisfactory solutions for GPR and NFPR, we always find bad local solutions

²We consider this inner-unsquared version as it has better noise stability than the inner-squared one in practice (17).

¹3D phase retrieval problems can be formalized similarly.

Table 1: Comparison of GPR, NFPR, and FFPR in terms of their symmetries and numerical solvability with the least-squares formulation in Eq. (0.1) combined with gradient descent.

| PR model | GPR | NFPR | FFPR |
|--|---|--|---|
| Symmetry | global phase | global phase | global phase, translation, 2D conjugate flipping |
| Final loss of solving LS in Eq. (0.1) using gradient descent from 100 random initializations |  |  |  |

for FFPR—which has three symmetries. Similarly, for FFPR, the

form:

$$\mathcal{F}(\mathbf{X}) = \mathbf{F}_{m'} \begin{bmatrix} \mathbf{X} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{F}_{n'}^\top, \quad (0.2)$$

where $\mathbf{F}_{m'}$ and $\mathbf{F}_{n'}$ are normalized discrete Fourier matrices. We always assume that $m' \geq 2m - 1$ and $n' \geq 2n - 1$, which is necessary to ensure recoverability.

Prior arts on FFPR Since we focus on practical FFPR, here we only discuss methods that have been tested on FFPR with at least partial success.

- **Classical iterative methods:** Due to the failure of the LS in Eq. (0.1), most (if not all) classical methods tackle the over-parameterized feasibility reformulation:

$$\text{find } \mathbf{Z} \in \mathbb{C}^{m' \times n'} \text{ s.t. } |\mathcal{F}(\mathbf{Z})|^2 = \mathbf{Y}, \mathcal{L}(\mathbf{Z}) = \mathbf{0}, \quad (0.3)$$

where \mathcal{L} restricts \mathbf{Z} to the zero-padding locations defined by the oversampling in Eq. (0.2). **More refined support information, if available, can be naturally incorporated into the support constraint $\mathcal{L}(\mathbf{Z}) = \mathbf{0}$.** Most of these classical methods are based on generalized alternating projection for solving Eq. (0.3), represented by error-reduction (ER), hybrid input-output (HIO) (24), reflection average alternating reflectors (RAAR) (25), difference map (DM) (26), and oversampling smoothness (OSS) (27). They are empirically observed to find good solutions for FFPR, provided that the support specification for \mathbf{Z} is tight and hyperparameters are properly tuned. Alternative formulations solved by second-order methods (28; 29) are less sensitive to hyperparameters. However, all these methods require tight support specification to avoid translation symmetry—failing so leads to spurious solutions that look like the superposition of translated copies; see Fig. 3. The tight-support difficulty is partially addressed by the popu-




Figure 2. HIO to solve FFPR with vs without precise support. We plot the final least-squares losses defined in Eq. (0.1) over 100 random starts. \mathbf{X} is the groundtruth in Fig. 1.

gold-standard hybrid input-output (HIO) algorithm can typically solve the problem when provided with tight support specification³—translation symmetry is killed, but it fails when the support is loose—translation symmetry remains; see Fig. 2. Moreover, our prior works (18; 19; 20; 21) also show the learning difficulties caused by these symmetries when developing data-driven methods for solving FFPR. See FFPR: Formulation and Prior Arts for details of HIO and related algorithms.

Our focus on practical FFPR methods We have stressed that symmetries largely determine the difficulty level of PR. However, in previous research, there are often simplifications to FFPR, including (1) **randomized model** \mathcal{S} that only keeps the global phase symmetry, (2) **evaluation on natural images** that removes the translation symmetry and simplifies the global phase symmetry into sign symmetry (8; 22; 23). These simplifications invariably lead to FFPR methods that do not work on practical data. **The goal of this paper is to develop practical methods for FFPR that involve all three symmetries.** In particular, we propose a novel FFPR method based on double deep image priors (see Our method: FFPR using double DIPs), and validate its superiority over state-of-the-art (SOTA) on realistic datasets (see Experiments Results).

FFPR: Formulation and Prior Arts

FFPR model The object of interest is $\mathbf{X} \in \mathbb{C}^{m \times n}$, and $\mathbf{Y} = |\mathcal{F}(\mathbf{X})|^2 \in \mathbb{R}_+^{m' \times n'}$, where \mathcal{F} is the oversampled 2D Fourier trans-

³Specifically, we mean that the tightest rectangular bounding box of the support is provided.




Figure 3. Two failure examples when solving FFPR using classical iterative methods without precise support specification and without shrinkwrap.

lar **shrinkwrap** trick (30) in practice, which refines the support by smoothing-and-thresholding over iterations.

- **Data-driven methods:** The **first family** represents the inverse mapping from \mathbf{Y} to \mathbf{X} by a deep neural network (DNN) g_{θ} , which is trained over either an extensive training set $\{(\mathbf{Y}_i, \mathbf{X}_i)\}_i$ or unpaired $\{\mathbf{Y}_i\}_i$ and $\{\mathbf{X}_i\}_i$ only following the cycle-consistency constraint: $|\mathcal{F}(g_{\theta}(\mathbf{Y}))|^2 \approx \mathbf{Y}$ (31; 32; 33; 34; 35; 36; 18; 19; 20; 21; 37; 38). But, as discussed in our prior work (18; 19; 20; 21), symmetries in the problem cause substantial learning difficulties, as any \mathbf{Y} maps to a set of equivalent \mathbf{X} 's. The **second family** (39; 40; 41) is tied to specific iterative methods for solving FFPR and replaces certain components of these methods with trainable DNNs. A typical limitation of this line is the reliance on good initialization that is obtained from classical iterative methods. Therefore, this family can be viewed as a final refinement of the results obtained from classical methods and does not address the essential difficulty of solving FFPR. **Both families** suffer in generalization when the training data are not sufficiently representative.

Our method overcomes the limitations of both classical and data-driven methods and features (1) **no training set**: it works with a single problem instance each time, with zero extra training data; (2) **no shrinkwrap**: we can specify the size of \mathbf{X} directly as $\lfloor m'/2 \rfloor \times \lfloor n'/2 \rfloor$, i.e. the information-theoretic recovery limit, without worrying about translation symmetry; (3) **minimal tuning**: mostly we only need to tune 2 learning rates as hyperparameters, vs. the 5 or 6 hyperparameters used in HIO + ER + shrinkwrap (HES) commonly used in practice (42).

Our method: FFPR using double DIPs

Deep image prior (DIP) for visual inverse problems DIP and variants (43) parameterize visual objects as outputs of DNNs—typically structured convolutional networks to favor spatially smooth structures, i.e., $\mathbf{x} = G_{\theta}(\mathbf{z})$, where \mathbf{z} is normally a random but fixed seed, and G_{θ} is a trainable DNN parameterized by θ . For a visual inverse problem of the form $\mathbf{y} \approx f(\mathbf{x})$ where \mathbf{y} is the observation and f is the observation model, the classical regularized data-fitting formulation

$$\min_{\mathbf{x}} \ell(\mathbf{y}, f(\mathbf{x})) + \lambda \Omega(\mathbf{x}) \quad (0.4)$$

can now be empowered by DIP and turned into

$$\min_{\theta} \ell(\mathbf{y}, f \circ G_{\theta}(\mathbf{z})) + \lambda \Omega \circ G_{\theta}(\mathbf{z}). \quad (0.5)$$

This simple idea has recently claimed numerous successes in computer vision and computational imaging; see, e.g., the recent survey (44), and our recent work addressing practicality issues around DIP (45; 46; 47). **A salient feature of DIP is the strong structured prior it imposes through DNNs, with zero extra data!** Although the theoretical understanding of DIP is still far from complete, current theories attribute its success to two aspects: (1) **structured priors** imposed by convolutional and up-sampling operations, and (2) **global optimization** due to significant overparameterization and first-order methods (48; 49).

Applying DIP to FFPR As shown in Table 1, solving the LS formulation in Eq. (0.1) using gradient descent always gets

trapped in bad local minimizers. It is then tempting to try DIP, as (1) the objects we try to recover in scientific imaging are visual objects and probably can be blessed by the structured priors enforced by DIP, and (2) more importantly, the issue we encounter in solving the LS is exactly about global optimization, which could be eliminated by overparameterization in DIP⁴. In fact, we have tried the single-DIP version in our preliminary work (50):

$$\min_{\theta} \|\sqrt{\mathbf{Y}} - |\mathcal{F} \circ G_{\theta}(\mathbf{z})|\|_F^2, \quad G_{\theta}(\mathbf{z}) \in \mathbb{C}^{m \times n}. \quad (0.6)$$

Systematic evaluation in Figs. 4 and 5 shows that it is already competitive compared to the gold-standard HES, although it struggles to reconstruct complicated complex phases.

Double DIPs boost the performance For FFPR applications such as CDI, \mathbf{X} as a complex-valued object can often be naturally split into two parts with disparate complexity levels. For example, in Bragg CDI on crystals, the magnitude part on the support is known to have uniform values, but the phase part can have complex spatial patterns due to strains (51; 52; 53); in CDI on live cells, the nonnegative real part contains useful information, and the imaginary part acts like small-magnitude noise (54). In these cases, due to the apparent asymmetry in complexity, it makes sense to parameterize \mathbf{X} as two separate DIPs (55; 47) instead of one:

$$\mathbf{X} = G_{\theta_1}^1(\mathbf{z}_1) e^{iG_{\theta_2}^2(\mathbf{z}_2)}, \text{ or } \mathbf{X} = G_{\theta_1}^1(\mathbf{z}_1) + iG_{\theta_2}^2(\mathbf{z}_2). \quad (0.7)$$

This can be justified as balancing the learning paces: with a single DIP, “simple” part is learned much faster than the “complex” part; with double DIPs, we can balance the learning paces by making the learning rate for the “simple” part relatively small compared to that for the “complex” part. We observe a substantial performance boost in Figs. 4 and 5 due to the double-DIP parametrization.

Experiments Results

Evaluation on 2D simulated Bragg CDI crystal data We first compare our Double-DIP method with multiple SOTA methods for FFPR, including Naive (37), CGAN (56), Passive (33), prDeep (39), HIO+ER, HIO+ER+Shrinkwrap (HES), and (single-)DIP on simulated 2D data for Bragg CDI on crystals. The final form of our learning objective for this task is:

$$\min_{\theta_1, \theta_2} \|\sqrt{\mathbf{Y}} - |\mathcal{F} \circ G_{\theta_1}^1(\mathbf{z}_1) e^{iG_{\theta_2}^2(\mathbf{z}_2)}|\|_F^2, \\ G_{\theta_1}^1(\mathbf{z}_1) \in \mathbb{R}_+^{m \times n}, G_{\theta_2}^2(\mathbf{z}_2) \in [0, 2\pi]^{m \times n}. \quad (0.8)$$

To ensure that the evaluation data reflect real-world complexity, we simulate 2D complex-valued crystal data in Bragg CDI applications (3). The dataset is generated by first creating 2D convex and nonconvex shapes based on random scattering points in a 110×110 grid on a 128×128 background. The complex magnitudes are uniformly 1, and the complex phases are determined by projecting simulated 2D displacement fields (due to crystal defects) onto the corresponding momentum transfer vectors. To maximize the diversity, the dataset contains diverse shapes and different numbers and densities of crystal defects that

⁴Interestingly, the feasibility reformulation in Eq. (0.3) solved by classical methods can also be understood as performing overparameterization.

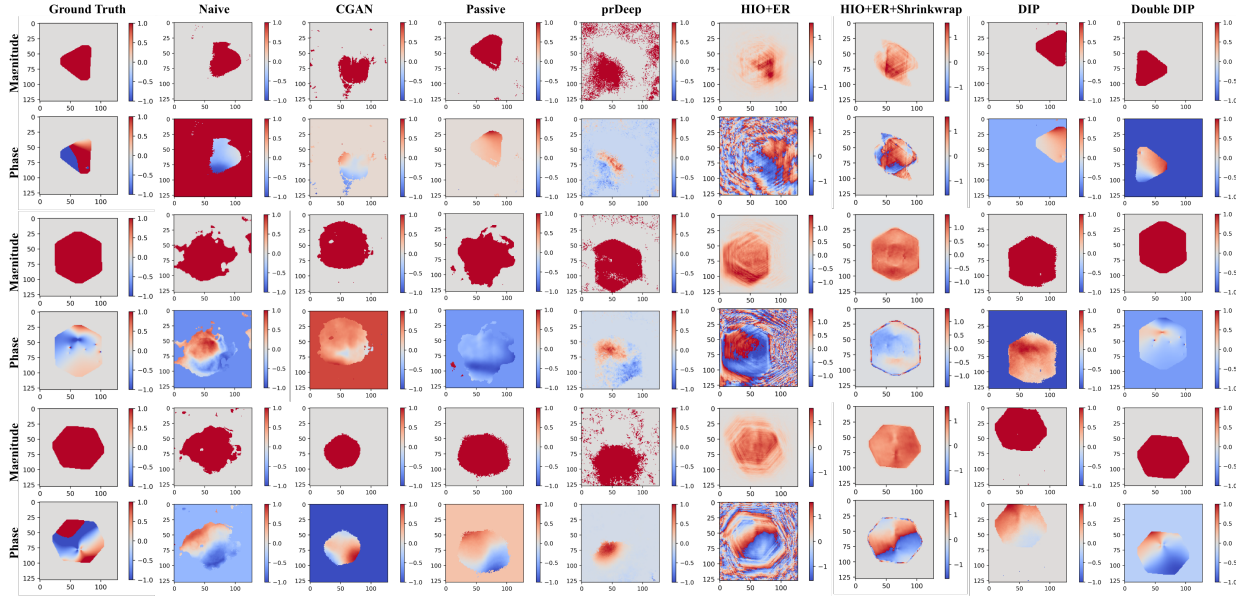


Figure 4. Visual comparison of reconstruction results by different methods on 2D simulated crystal data




Figure 5. Quantitative comparison of reconstruction results by different methods on 2D simulated crystal data by symmetry-adjusted MSE




Figure 6. Visual comparison of reconstruction results by HIO+ER with Shrinkwrap and our method on a 3D simulated crystal instance

dictate the complexity of the phases. **Although our double-DIP method is a single-instance method that requires no training data, the dataset is large enough to support data-driven methods, such as Passive and prDeep.** For methods that require a training set, we provide 8000 samples. All methods are tested on 50 randomly selected samples not from the 8000 training samples.

From both visual (Fig. 4) and quantitative (Fig. 5) results, it is evident that: (1) all data-driven methods, including Naive, CGAN, Passive, prDeep, perform poorly. We believe that this is due to either the learning difficulty caused by the three symmetries (18; 19; 20; 21) or the bad initialization given by HIO (i.e., for prDeep); (2) HES, DIP, and our double-DIP are the top three methods. HES deals with translation symmetry by iteratively refining the support, whereas DIP and ours do not need tight support estimation at all, substantially reducing parameter tuning; (3) Our method wins HES and DIP by a large margin. Although the latter two perform reasonably well in magnitude estimation, their phase estimations are typically off for complicated instances.

Evaluation on 3D simulated Bragg CDI crystal data We will exclude data-driven methods, due to their clear performance deficiency on 2D data and the considerable cost to obtain sufficiently representative training sets for 3D. We only compare HES, which is the gold-standard used in Bragg CDI practice, with our double-DIP method. Since both methods can work with single instances and need no training data, here we quickly compare their performance qualitatively on a single 3D simulated crystal instance (the simulation process is similar to the 2D case), as shown in Fig. 6. It is obvious that even with Shrinkwrap iteratively refining support, HES still struggles to get the support right. By contrast, our double-DIP method obtains a sharp support recovery and good phase estimation.

References

- [1] J. W. Goodman, *Introduction to Fourier Optics*, 4th ed. Freeman & Company, W. H., 2017.
- [2] J. Miao, P. Charalambous, J. Kirz, and D. Sayre, "Extending the methodology of x-ray crystallography to allow imaging of micrometre-sized non-crystalline specimens," *Nature*, vol. 400, no.

- 6742, pp. 342–344, 1999.
- [3] I. K. Robinson, I. A. Vartanyants, G. J. Williams, M. A. Pfeifer, and J. A. Pitney, “Reconstruction of the shapes of gold nanocrystals using coherent x-ray diffraction,” *Physical Review Letters*, vol. 87, no. 19, p. 195505, oct 2001.
 - [4] D. R. Luke, J. V. Burke, and R. G. Lyon, “Optical wavefront reconstruction: Theory and numerical methods,” *SIAM Review*, vol. 44, no. 2, pp. 169–224, jan 2002.
 - [5] P. Jaming, “Phase retrieval techniques for radar ambiguity problems,” *The Journal of Fourier Analysis and Applications*, vol. 5, no. 4, pp. 309–329, jul 1999.
 - [6] Y. Shechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev, “Phase retrieval with application to optical imaging: A contemporary overview,” *IEEE Signal Processing Magazine*, vol. 32, no. 3, pp. 87–109, may 2015.
 - [7] E. J. Candes, Y. C. Eldar, T. Strohmer, and V. Voroninski, “Phase retrieval via matrix completion,” *SIAM review*, vol. 57, no. 2, pp. 225–251, 2015.
 - [8] A. Fannjiang and T. Strohmer, “The numerics of phase retrieval,” *Acta Numerica*, vol. 29, pp. 125–228, may 2020.
 - [9] F. Wang, Y. Bian, H. Wang, M. Lyu, G. Pedrini, W. Osten, G. Barbastathis, and G. Situ, “Phase imaging with an untrained neural network,” *Light: Science & Applications*, vol. 9, no. 1, may 2020.
 - [10] Y. Zhang, M. A. Noack, P. Vagovic, K. Fezzaa, F. Garcia-Moreno, T. Ritschel, and P. Villanueva-Perez, “PhaseGAN: a deep-learning phase-retrieval approach for unpaired datasets,” *Optics Express*, vol. 29, no. 13, p. 19593, jun 2021.
 - [11] T. Shimobaba and T. Ito, *Computer Holography Acceleration Algorithms and Hardware Implementations*. Taylor & Francis Group, 2019.
 - [12] J. R. Fienup, “Phase-retrieval algorithms for a complicated optical system,” *Applied Optics*, vol. 32, no. 10, p. 1737, apr 1993.
 - [13] D. A. Barmherzig and J. Sun, “Towards practical holographic coherent diffraction imaging via maximum likelihood estimation,” *Optics Express*, vol. 30, no. 5, p. 6886, feb 2022.
 - [14] P. Thibault, M. Dierolf, A. Menzel, O. Bunk, C. David, and F. Pfeifer, “High-resolution scanning x-ray diffraction microscopy,” *Science*, vol. 321, no. 5887, pp. 379–382, jul 2008.
 - [15] S. Marchesini, A. Schirotzek, C. Yang, H. tieng Wu, and F. Maia, “Augmented projections for ptychographic imaging,” *Inverse Problems*, vol. 29, no. 11, p. 115009, oct 2013.
 - [16] T. Bendory, R. Beinert, and Y. C. Eldar, “Fourier phase retrieval: Uniqueness and algorithms,” in *Compressed Sensing and its Applications*. Springer International Publishing, 2017, pp. 55–91.
 - [17] L.-H. Yeh, J. Dong, J. Zhong, L. Tian, M. Chen, G. Tang, M. Soltanolkotabi, and L. Waller, “Experimental robustness of fourier ptychography phase retrieval algorithms,” *Optics express*, vol. 23, no. 26, pp. 33 214–33 240, 2015.
 - [18] K. Tayal, C.-H. Lai, V. Kumar, and J. Sun, “Inverse problems, deep learning, and symmetry breaking,” *arXiv:2003.09077*, Mar. 2020.
 - [19] R. Manekar, K. Tayal, Z. Zhuang, C.-H. Lai, V. Kumar, and J. Sun, “Breaking symmetries in data-driven phase retrieval,” in *OSA Imaging and Applied Optics Congress 2021 (3D, COSI, DH, ISA, pcAOP)*. Optica Publishing Group, 2021.
 - [20] K. Tayal, C.-H. Lai, R. Manekar, Z. Zhuang, V. Kumar, and J. Sun, “Unlocking inverse problems using deep learning: Breaking symmetries in phase retrieval,” in *NeurIPS 2020 Workshop on Deep Learning and Inverse Problems*, 2020. [Online]. Available: <https://openreview.net/forum?id=oyhGIyTV1S>
 - [21] R. Manekar, Z. Zhuang, K. Tayal, V. Kumar, and J. Sun, “Deep learning initialized phase retrieval,” in *NeurIPS 2020 Workshop on Deep Learning and Inverse Problems*, 2020. [Online]. Available: <https://openreview.net/forum?id=gV4I5IJHP>
 - [22] D. R. Luke, S. Sabach, and M. Teboulle, “Optimization on spheres: Models and proximal algorithms with computational performance comparisons,” *SIAM Journal on Mathematics of Data Science*, vol. 1, no. 3, pp. 408–445, jan 2019.
 - [23] S. Marchesini, “Invited article: A unified evaluation of iterative projection algorithms for phase retrieval,” *Review of Scientific Instruments*, vol. 78, no. 1, p. 011301, jan 2007.
 - [24] J. R. Fienup, “Phase retrieval algorithms: a comparison,” *Applied optics*, vol. 21, no. 15, pp. 2758–2769, 1982.
 - [25] D. R. Luke, “Relaxed averaged alternating reflections for diffraction imaging,” *Inverse problems*, vol. 21, no. 1, p. 37, 2004.
 - [26] V. Elser, “Phase retrieval by iterated projections,” *JOSA A*, vol. 20, no. 1, pp. 40–55, 2003.
 - [27] J. A. Rodriguez, R. Xu, C.-C. Chen, Y. Zou, and J. Miao, “Oversampling smoothness: an effective algorithm for phase retrieval of noisy diffraction intensities,” *Journal of applied crystallography*, vol. 46, no. 2, pp. 312–318, 2013.
 - [28] S. Marchesini, “Phase retrieval and saddle-point optimization,” *JOSA A*, vol. 24, no. 10, pp. 3289–3296, 2007.
 - [29] Z. Zhuang, G. Wang, Y. Travadi, and J. Sun, “Phase retrieval via second-order nonsmooth optimization,” in *ICML Workshop on Beyond First-Order Methods for Machine Learning*, 2020.
 - [30] S. Marchesini, H. He, H. N. Chapman, S. P. Hau-Riege, A. Noy, M. R. Howells, U. Weierstall, and J. C. H. Spence, “X-ray image reconstruction from a diffraction pattern alone,” *Physical Review B*, vol. 68, no. 14, p. 140101, oct 2003.
 - [31] A. Sinha, J. Lee, S. Li, and G. Barbastathis, “Lensless computational imaging through deep learning,” *Optica*, vol. 4, no. 9, pp. 1117–1125, 2017.
 - [32] A. Scheinker and R. Pokharel, “Adaptive 3d convolutional neural network-based reconstruction method for 3d coherent diffraction imaging,” *Journal of Applied Physics*, vol. 128, no. 18, p. 184901, 2020.
 - [33] H. Chan, Y. S. Nashed, S. Kandel, S. O. Hruszkewycz, S. K. Sankaranarayanan, R. J. Harder, and M. J. Cherukara, “Rapid 3d nanoscale coherent imaging via physics-aware deep learning,” *Applied Physics Reviews*, vol. 8, no. 2, p. 021407, 2021.
 - [34] R. Harder, “Deep neural networks in real-time coherent diffraction imaging,” *IUCrJ*, vol. 8, no. Pt 1, p. 1, 2021.
 - [35] L. Wu, S. Yoo, A. F. Suzana, T. A. Assefa, J. Diao, R. J. Harder, W. Cha, and I. K. Robinson, “Three-dimensional coherent x-ray diffraction imaging via deep convolutional neural networks,” *npj Computational Materials*, vol. 7, no. 1, pp. 1–8, 2021.
 - [36] Y. Yao, H. Chan, S. Sankaranarayanan, P. Balaprakash, R. J. Harder, and M. J. Cherukara, “Autophasenn: unsupervised physics-aware deep learning of 3d nanoscale bragg coherent diffraction imaging,” *npj Computational Materials*, vol. 8, no. 1, pp. 1–8, 2022.
 - [37] M. J. Cherukara, Y. S. Nashed, and R. J. Harder, “Real-time coherent diffraction inversion using deep generative networks,” *Scientific reports*, vol. 8, no. 1, pp. 1–8, 2018.
 - [38] Y. Zhang, M. A. Noack, P. Vagovic, K. Fezzaa, F. Garcia-Moreno, T. Ritschel, and P. Villanueva-Perez, “Phasegan: a deep-learning phase-retrieval approach for unpaired datasets,” *Optics express*, vol. 29, no. 13, pp. 19 593–19 604, 2021.
 - [39] C. Metzler, P. Schniter, A. Veeraraghavan *et al.*, “prdeep: Robust

- phase retrieval with a flexible deep network,” in *International Conference on Machine Learning*. PMLR, 2018, pp. 3501–3510.
- [40] Ç. Işıl, F. S. Oktem, and A. Koç, “Deep iterative reconstruction for phase retrieval,” *Applied optics*, vol. 58, no. 20, pp. 5422–5431, 2019.
- [41] Y. Wang, X. Sun, and J. Fleischer, “When deep denoising meets iterative phase retrieval,” in *International Conference on Machine Learning*. PMLR, 2020, pp. 10 007–10 017.
- [42] M. C. Newton, S. J. Leake, R. Harder, and I. K. Robinson, “Three-dimensional imaging of strain in a single ZnO nanorod,” *Nature Materials*, vol. 9, no. 2, pp. 120–124, dec 2009.
- [43] D. Ulyanov, A. Vedaldi, and V. Lempitsky, “Deep image prior,” *International Journal of Computer Vision*, vol. 128, no. 7, pp. 1867–1888, 2020.
- [44] A. Qayyum, I. Ilahi, F. Shamshad, F. Boussaid, M. Bennamoun, and J. Qadir, “Untrained neural network priors for inverse imaging problems: A survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [45] T. Li, Z. Zhuang, H. Liang, L. Peng, H. Wang, and J. Sun, “Self-validation: Early stopping for single-instance deep generative priors,” *arXiv preprint arXiv:2110.12271*, 2021.
- [46] H. Wang, T. Li, Z. Zhuang, T. Chen, H. Liang, and J. Sun, “Early stopping for deep image prior,” *arXiv preprint arXiv:2112.06074*, 2021.
- [47] Z. Zhuang, T. Li, H. Wang, and J. Sun, “Blind image deblurring with unknown kernel size and substantial noise,” *arXiv preprint arXiv:2208.09483*, 2022.
- [48] R. Heckel and M. Soltanolkotabi, “Denoising and regularization via exploiting the structural bias of convolutional generators,” *arXiv preprint arXiv:1910.14634*, 2019.
- [49] —, “Compressive sensing with un-trained neural networks: Gradient descent finds a smooth approximation,” in *International Conference on Machine Learning*. PMLR, 2020, pp. 4149–4158.
- [50] K. Tayal, R. Manekar, Z. Zhuang, D. Yang, V. Kumar, F. Hofmann, and J. Sun, “Phase retrieval using single-instance deep generative prior,” in *OSA Optical Sensors and Sensing Congress 2021 (AIS, FTS, HISE, SENSORS, ES)*. Optica Publishing Group, 2021.
- [51] J. Clark, X. Huang, R. Harder, and I. Robinson, “High-resolution three-dimensional partially coherent diffraction imaging,” *Nature communications*, vol. 3, no. 1, pp. 1–6, 2012.
- [52] F. Hofmann, E. Tarleton, R. J. Harder, N. W. Phillips, P.-W. Ma, J. N. Clark, I. K. Robinson, B. Abbey, W. Liu, and C. E. Beck, “3d lattice distortions and defect structures in ion-implanted nanocrystals,” *Scientific reports*, vol. 7, no. 1, pp. 1–10, 2017.
- [53] D. Yang, M. T. Lapington, G. He, K. Song, M. Zhang, C. Barker, R. J. Harder, W. Cha, W. Liu, N. W. Phillips *et al.*, “Refinements for bragg coherent x-ray diffraction imaging: Electron backscatter diffraction alignment and strain field computation,” *arXiv preprint arXiv:2203.17015*, 2022.
- [54] G. Van Der Schot, M. Svenda, F. R. Maia, M. Hantke, D. P. De-Ponte, M. M. Seibert, A. Aquila, J. Schulz, R. Kirian, M. Liang *et al.*, “Imaging single cells in a beam of live cyanobacteria with an x-ray laser,” *Nature communications*, vol. 6, no. 1, pp. 1–9, 2015.
- [55] Y. Gandelsman, A. Shocher, and M. Irani, ““double-DIP”: Unsupervised image decomposition via coupled deep-image-priors,” in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, jun 2019.
- [56] T. Uelwer, A. Oberstraß, and S. Harmeling, “Phase retrieval using conditional generative adversarial networks,” in *2020 25th International Conference on Pattern Recognition (ICPR)*. IEEE, 2021, pp. 731–738.