

SCIENTIFIC REPORTS



OPEN

Real-time coherent diffraction inversion using deep generative networks

Mathew J. Cherukara ^{1,3}, Youssef S. G. Nashed ² & Ross J. Harder¹

Phase retrieval, or the process of recovering phase information in reciprocal space to reconstruct images from measured intensity alone, is the underlying basis to a variety of imaging applications including coherent diffraction imaging (CDI). Typical phase retrieval algorithms are iterative in nature, and hence, are time-consuming and computationally expensive, making real-time imaging a challenge. Furthermore, iterative phase retrieval algorithms struggle to converge to the correct solution especially in the presence of strong phase structures. In this work, we demonstrate the training and testing of CDI NN, a pair of deep deconvolutional networks trained to predict structure and phase in real space of a 2D object from its corresponding far-field diffraction intensities alone. Once trained, CDI NN can invert a diffraction pattern to an image within a few milliseconds of compute time on a standard desktop machine, opening the door to real-time imaging.

Central to many imaging techniques including coherent X-ray diffraction imaging^{1,2}, electron microscopy³, astronomy⁴ and super-resolution optical imaging⁵, is the process of phase retrieval, or the recovery of phase information from the measured intensities alone. In particular, in X-ray coherent diffraction imaging (CDI), an object is illuminated with a coherent X-ray beam and the resulting far-field diffraction pattern is measured. This far-field diffraction pattern is the modulus of the Fourier transform of the object, and phase retrieval algorithms are used to reconstruct the measured object by recovering the lost phase information. As such, the imaging methods are extremely sensitive to any material properties that contribute a phase to the scattered beam⁶. In particular, when measured in the vicinity of a Bragg peak, the measured coherent far-field X-ray diffraction pattern encodes the strain within the object in the local asymmetry of the coherent diffraction pattern around the Bragg peak⁷. The strain induces distortion in the lattice, which manifests itself in the scattered beam as an additional phase. Upon successful inversion of the diffraction pattern to an image, the local distortion of the crystal lattice is then displayed as a phase in the complex image of the sample⁸. X-ray CDI and Bragg CDI (BCDI) in particular have been widely used to provide a unique 4D view of dynamic processes including phonon transport^{9,10}, transient melting¹¹, dissolution and recrystallization¹², phase transformations¹³, grain growth¹⁴ and device characterization^{15,16}. Notwithstanding its widespread use, reciprocal space phase retrieval algorithms suffer from several shortcomings. Firstly, the iterative phase retrieval algorithms that are commonly used, such as error-reduction (ER) and hybrid input-output (HIO)¹⁷ or difference map (DM)¹⁸ are time consuming, requiring thousands of iterations and multiple random initializations to converge to a solution with high confidence¹⁹. Furthermore, such algorithms often fail to converge in presence of strong phase structures for instance those associated with multiple defects in materials²⁰. Additionally, algorithmic convergence is often sensitive to certain iterative phase retrieval parameters such as initial guesses, shrinkwrap threshold²¹, choice of algorithms and their combinations²². Finally, a necessary mathematical condition for reciprocal space phase retrieval is that the measured intensities are oversampled by at least a factor of two²³. In practice, this requirement translates to necessitating a minimum of two pixels per coherent feature on the detector. Consequently, this limits the extent of reciprocal space that is accessible for a given detector size and x-ray wavelength.

Neural networks have been described as universal approximators, with the ability to represent a wide variety of functions²⁴. As such, they have been used for an enormous variety of applications ranging from natural language processing and computer vision to self-driving cars²⁵. More recent work has predominantly involved the

¹Advanced Photon Source, Argonne National Laboratory, Lemont, IL, 60439, USA. ²Mathematics and Computer Science, Argonne National Laboratory, Lemont, IL, 60439, USA. ³Center for Nanoscale Materials, Argonne National Laboratory, Lemont, IL, 60439, USA. Correspondence and requests for materials should be addressed to M.J.C. (email: mcherukara@aps.anl.gov)

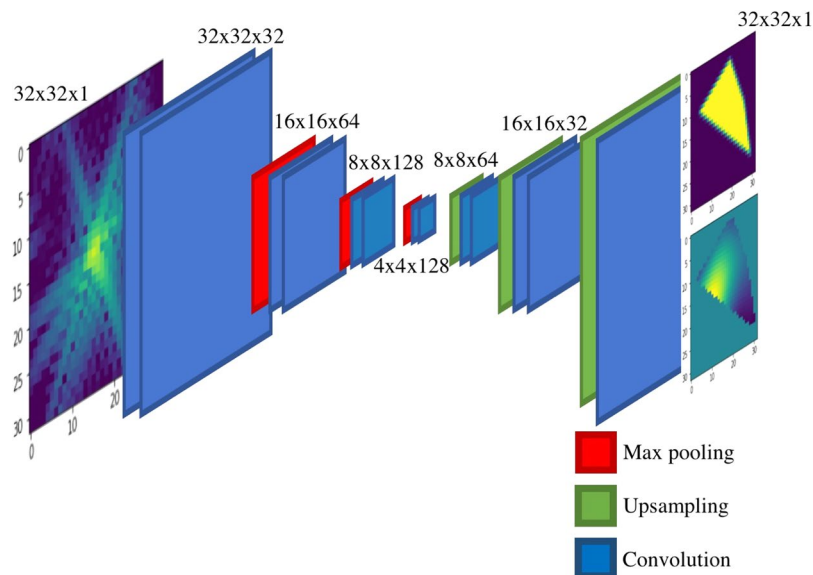


Figure 1. Structure of the deep generative network CDI NN. CDI NN is implemented using an architecture composed entirely of convolutional, max pooling and upsampling layers. All activations are rectified linear units (ReLU) except for the final convolutional layer which uses sigmoidal activations.

use of deep neural networks, so termed because of the manner in which they are structured to learn increasingly more complex features or hierarchal representations with successive layer of neurons²⁶. In particular, deep deconvolutional networks have found a variety of applications in various imaging techniques, ranging from automated image segmentation of electron microscopy images²⁷, image reconstruction from magnetic resonance imaging (MRI)²⁸, to the enhancement of images from mobile phone microscopes²⁹.

Specific to the problem of phase recovery, deep neural networks have been used in holographic image reconstruction³⁰, phase retrieval following spatial light modulation (SLM)³¹, optical tomography³², and as a denoiser in iterative phase retrieval³³. We note that none of these works represent an end-to-end solution to the far-field reciprocal space phase retrieval problem.

In this work, we train two deep deconvolutional networks to learn the mapping between 2D coherent diffraction patterns (which are the magnitudes of an object's Fourier transform) and the corresponding real-space structure and phase. While we have trained these neural networks (NNs) with the intention of applying them to CDI measurements, in particular to BCDI measurements, the approach outlined in this study is easily transferable to any imaging modality that requires reciprocal space phase retrieval. Once trained, these deconvolutional networks, which we term CDI NN can predict the structure and phase of test data within a few milliseconds on a standard desktop machine. This is thousands of times faster than what is achievable with iterative phase retrieval algorithms currently in use. Such real-time image reconstruction has the potential to revolutionize the various advanced imaging modalities that rely on phase retrieval and is essential to performing *in-situ* and *operando* characterization studies of rapidly evolving samples for experimental feedback.

Results

Coherent Diffraction Imaging (CDI). BCDI measurements are typically performed on compact objects such as isolated nanoparticles or single grains within a polycrystalline material. To simulate the compact structure associated with an isolated particle or a single grain, we use convex polygons of random size and shape. Points within these polygons are complex values with magnitude of 1, while points lying outside have a magnitude of 0. We also give the edges of the polygons a Gaussian transition from 1 to zero with a width of one pixel. Points within the polygon have a spatially-varying complex phase that simulates the distortion of a crystalline lattice due to strain within a material. Without loss of generality, this phase can represent any structural inhomogeneity that modifies the phase of a scattered beam. Finally, to obtain the diffraction signal corresponding to the object, we take the magnitudes of the two-dimensional (2D) Fourier transform (FT) of the complex valued compact object.

CDI NN's structure and training. Our generative network, CDI NN is a feed-forward network consisting of two parts, as shown in Fig. 1. The first part is a convolutional autoencoder that is responsible for finding a representation, or encoding, of an input image in feature space. This encoding in feature space represents the underlying manifold of the input data. This encoding is then subsequently propagated through a deconvolutional decoder to generate an output image. The deconvolutional decoder is composed of convolutional layers that learn the mapping from the encoded diffraction pattern to the corresponding real-space object and phase. The overall network is trained in a supervised fashion, where the output image is known a priori^{34,35}. We train two networks with identical architecture, one that takes diffraction amplitudes as input and produces object shape as output (structure CDI NN or sCDI NN), and a second that takes diffraction amplitudes as input while outputting real space phase information (phase CDI NN or pCDI NN). The convolutional and max pooling operations (max

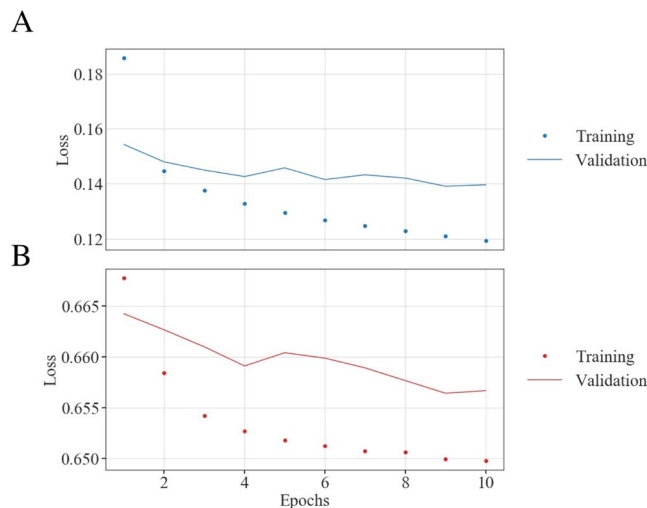


Figure 2. Training and validation. Training and validation loss as function of training epoch for (A) sCDI NN and (B) pCDI NN. Training was stopped after 10 epochs beyond which the validation loss diverged, suggesting overfitting beyond 10 epochs.

pooling is a binning/downsampling operation using the maximum value over a prespecified pixel neighborhood serve to transform the image data (in this case the diffracted amplitudes) into feature space, while the deconvolutional and upsampling operations serve to transform back from feature space into pixel space.

To train the two networks that together compose CDI NN (sCDI NN and pCDI NN), we exposed the networks to 180,000 training examples consisting of diffraction magnitudes and the corresponding real space structure and phase. Each instance of the training data was generated as described in methods. We set aside 20,000 instances from the generated training data for model validation at the end of each training pass (also see methods). While we have used 32×32 pixel images in the work described here, we note that modern deep learning frameworks use tensor algebra that is highly data-parallel, and are optimized to use many-core accelerators like GPUs and FPGAs, enabling much larger image sizes. Figure 2A shows the training and validation loss as a function of epochs for sCDI NN, while Fig. 2B shows the training and validation loss as a function of epoch for pCDI NN. Each epoch refers to one complete pass of the training data, while the loss (or error metric) for both training and validation is computed using cross-entropy. For both networks, we see that the weights converge within 10 epochs as evinced by the behavior of the validation loss. We found that training for more epochs causes the validation loss to diverge, suggesting that the network was beginning to overfit to training data beyond 10 training epochs.

CDI NN's performance on test data. To test the performance of the trained CDI NN networks, we evaluate their performance on a new set of 1000 test cases that was not shown to the networks at any point during training. In testing, we import the trained neural networks' topology and optimized weights and evaluate its ability to reconstruct real space structure and phase from input diffraction patterns. Figure 3 shows random samples of the performance of the network in testing. The first row (Fig. 3A) shows the input diffracted amplitudes. Figure 3B (second row) shows the corresponding ground truth objects, while Fig. 3C (third row) shows the structures predicted by sCDI NN. We observe an excellent match between the predicted and true object structures. Figure 3D shows the true phase structure, while Fig. 3E shows the phases predicted by pCDI NN. Again, we observe a good agreement between the prediction of CDI NN and the actual phase structure. We note that the images shown in Fig. 3E are bounded by the structure predicted by sCDI NN in Fig. 3C, i.e. phases outside of the predicted object shape are set to 0. We use a threshold of 0.1 to define the boundary of the object. We also draw the reader's attention to the third example, where we note that the actual (B) and predicted (C) objects are twin images of each other, and that they can be obtained from each other through a centrosymmetric inversion and complex conjugate operation. Both images are equivalent solutions to the input diffraction pattern. We observe several such instances where shape predicted by CDI NN is the twin image, especially when the phase structure is weak or constant and the corresponding diffraction image is nearly centrosymmetric. Finally, Fig. 3E,G show the structure and phase obtained after running iterative phase retrieval (see Methods). We note that while phase retrieval performs better than CDI NN, several of the recovered images possess phase wraps. On the other hand, CDI NN is trained on data with no phase wraps and so, is immune to phase wraps that can be challenging to mitigate in complex images from phase retrieval. CDI NN is also $\sim 500X$ faster, taking only a few milliseconds on a standard desktop machine, in contrast to phase retrieval which took ~ 1.5 seconds on the same machine for 620 iterations.

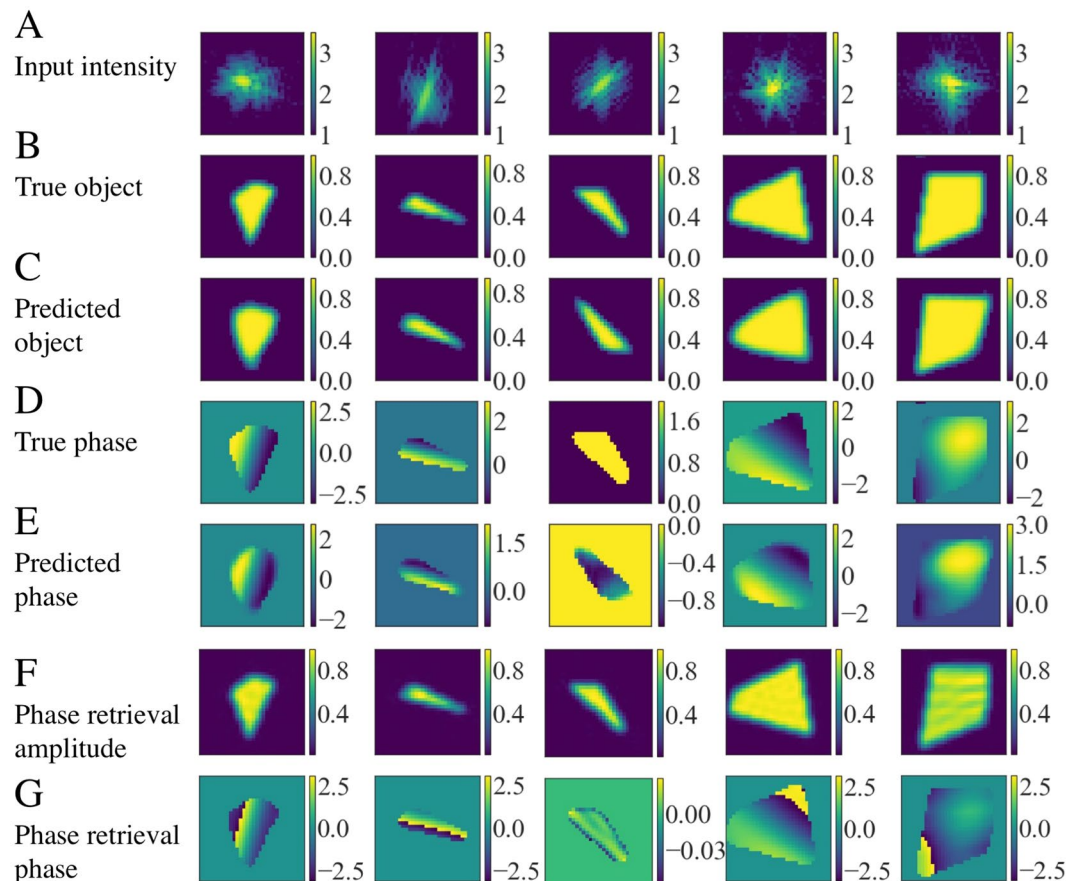


Figure 3. Examples of the performance of CDI NN in testing. (A) Input diffraction intensities to CDI NN. (B) Actual shape of the corresponding object and (C) The shape predicted by sCDI NN. (D) The actual phase structure of the object and (E), the phase structure predicted by pCDI NN. CDI NN successfully recovers structure and shape for a variety of different shapes and phase structures. (F,G) Show the structure and phase obtained after running iterative phase retrieval. Phase wrap artifacts are seen in a few of the recovered phases. Intensity units are dimensionless while phase is in radians.

Discussion

Strengths and weaknesses of BCDI NN. To gain some insight into the strengths and weaknesses of CDI NN in predicting an object's structure and phase from its diffraction pattern, we quantify the error in each prediction by comparing the prediction to the ground truth. We make the comparison in reciprocal space, where we take the retrieved image (structure and phase), compute the reciprocal space amplitudes and compare those with the amplitudes given to CDI NN.

In Fig. 4A we plot a histogram of the χ^2 error for each of the test cases. To compute this error for each of the test cases, we take the predictions from sCDI NN and pCDI NN and compute the FT to obtain the predicted diffraction intensity. The error χ^2 is then given by:

$$\chi^2 = \frac{\sum_{i=1}^N (\sqrt{I_p^i} - \sqrt{I_t^i})^2}{\sum_{i=1}^N I_t^i} \quad (1)$$

where I_t^i are the true diffraction intensities and I_p^i are the predicted diffraction intensities at each pixel. Figure 4B shows a zoom of the histogram at the lowest χ^2 error, i.e., where the predictions are the best. The panels on the right show the 5 best predictions as computed by the error metric. Figure 4C shows the input diffraction intensity, Fig. 4E the true object structure, Fig. 4F the predicted shape, Fig. 4G the actual phase structure and Fig. 4H the predicted phase structure. Figure 4D shows the predicted diffraction intensity which is obtained by taking an FT of the predicted shape and predicted phase. We observe that the best predictions as defined by the error in diffraction intensities, is found when the objects are large, with relatively weak phase structure. This is perhaps unsurprising since in these situations, the diffraction pattern is quite symmetric with most of the intensity centered around the central peak.

Conversely, Fig. 5 shows the 5 worst predictions as inferred by the χ^2 error metric. Figure 5B shows a zoomed in view of the histogram showing the error of the worst cases, with the maximum computed error being $\chi^2 \sim 0.5$.

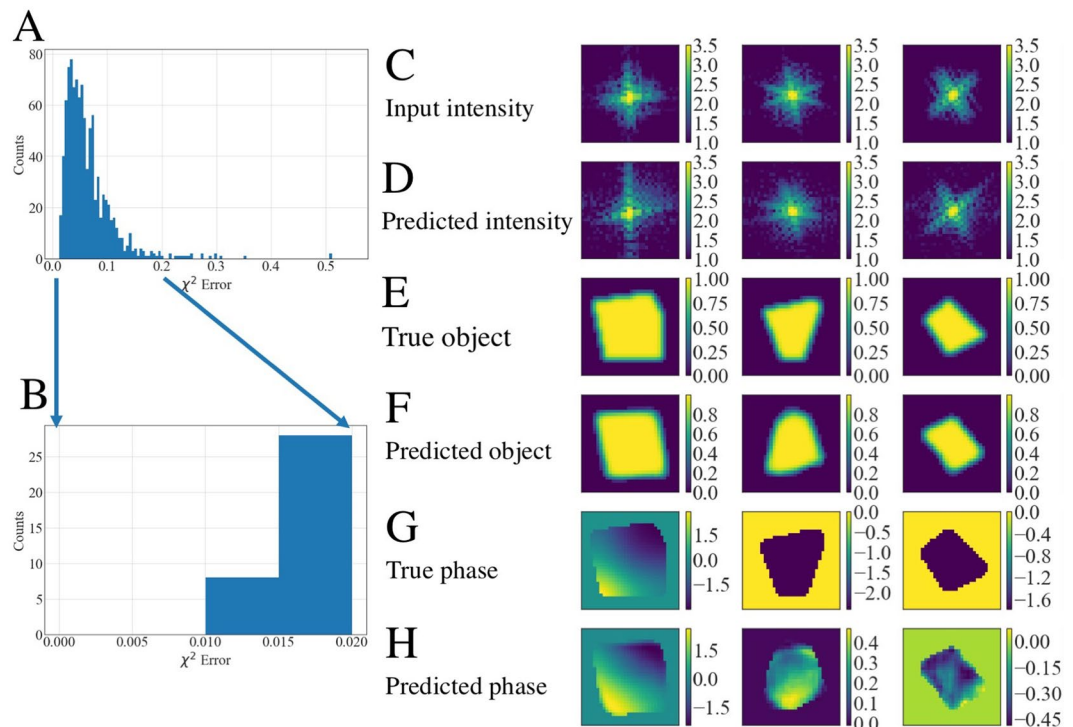


Figure 4. Best predictions of CDI NN. (A) Histogram of the error computed for the test samples. (B) Close-up view of the predictions with the lowest χ^2 error. (C) Input diffraction patterns, (D) predicted intensities computed by taking an FT of the complex object obtained from the predicted shape and phase. (E,F) Corresponding actual shape and predicted shape. (G,H) Corresponding actual phase and predicted phase. CDI NN fares well when faced with large objects with weak phase. Intensity units are dimensionless while phase is in radians.

The least accurate predictions of CDI NN correspond to smaller objects and strong phase structures (Fig. 5E–H). In such instances, the forward scattered peak center is poorly defined, and the diffracted intensity is spread across a large range of q , as seen in Fig. 5C,D. However, we see that even for the most difficult instances (where the signal to noise ratio is low), the predicted shape and phase structures are reasonable. We note however that all of the data presented in this paper is simulated, and we expect further complications to arise when dealing with experimental data, such as the effect of noise, pixel variation, beam stability, partial coherence etc.

Network activation maps. To investigate the nature of the features that the convolutional layers learn, we study the layer activations for different input diffraction patterns. Figure 6 shows average activation maps of the 2nd convolutional layer and 4th convolutional layer (see Fig. 1) for 3 different input diffraction patterns (Fig. 6A). These activation maps represent the average activations of the 32 convolutional filters that make up the 2nd convolutional layer (Fig. 6B,D), and an average of the 64 convolutional filters that make up the 4th convolutional layer (Fig. 6C,E) for the two networks (sCDI NN and pCDI NN). We note that the images shown in Fig. 6C,E have been interpolated from 16×16 to 32×32 to enable a straightforward comparison. For both networks, we observe that at the 2nd convolutional layer, the network focusses on regions close to the brightest pixels at the center of the forward scattered intensity. Interestingly, we observe differing behaviors of the two networks at the 4th convolutional layer. At the 4th convolutional layer, the structure network (sCDI NN) focusses solely on regions at higher spatial frequencies, i.e. choosing to focus on finer scale features. In slight contrast, the phase network (pCDI NN), continues to focus strongly on the center of the diffraction pattern (low spatial frequencies), while also paying more attention to pixels at higher spatial frequencies. In both networks, successive layers start paying attention to data at higher spatial frequencies, and this suggests that CDI NN progressively learns higher order features in the image in encoding the structure of the image.

Conclusion

In conclusion, to the best of our knowledge, this work is the first demonstration of an end-to-end deep learning solution to the phase retrieval problem in the far-field. We believe the results described in this manuscript have widespread ramifications for both BCDI experiments of the future for which this study was designed as well as other imaging modalities reliant on successful phase retrieval.

CDI NN is thousands of times faster than traditional phase retrieval and requires only modest resources to run. We note that while the training of CDI NN was performed on a dual GPU machine (~1 hour training time); once trained, CDI NN can easily be deployed at a standard desktop at the experiment's location. Indeed, the test cases in this manuscript were run on the CPUs of a 2013 Mac Pro desktop with Quad-Core Intel Xeon E5, where

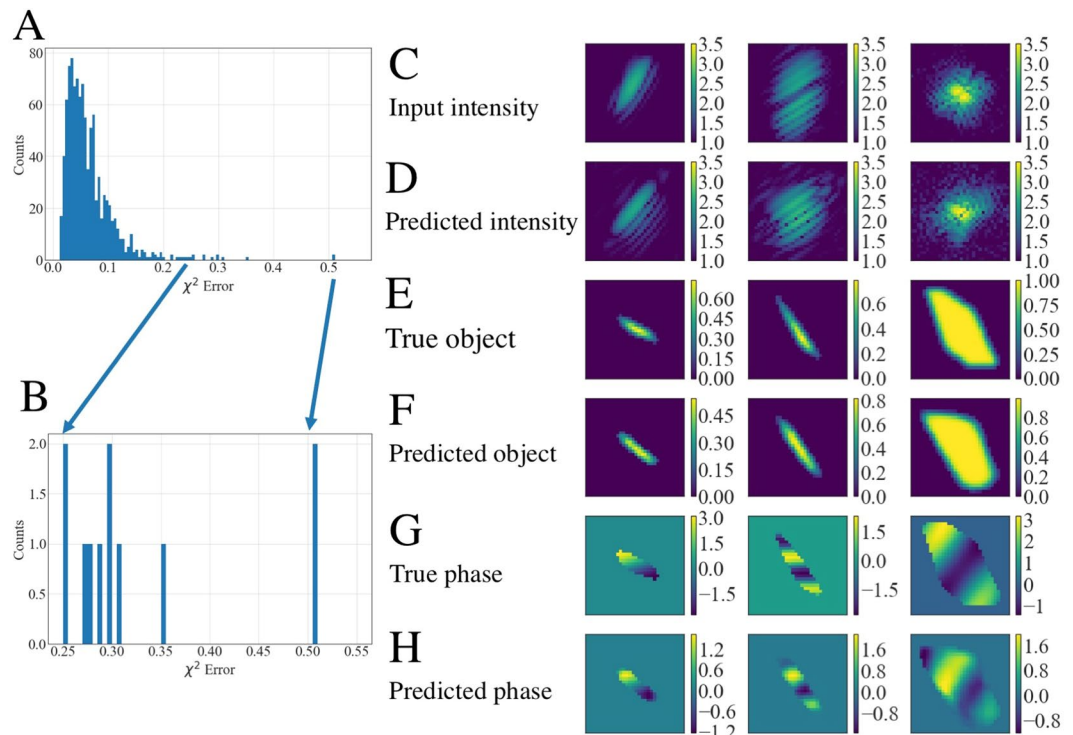


Figure 5. Worst predictions of CDI NN. (A) Histogram of the error computed for the test samples. (B) Close-up view of the predictions with the highest χ^2 error. (C) Input diffraction patterns, (D) predicted intensities computed by taking an FT of the complex object obtained from the predicted shape and phase. (E,F) Corresponding actual shape and predicted shape. (G,H) Corresponding actual phase and predicted phase. CDI NN performs worst when faced with diffraction patterns corresponding to small objects with strong phase. Intensity units are dimensionless while phase is in radians.

the prediction time was ~ 2.7 milliseconds. This represents a speed of ~ 500 times when compared to iterative phase retrieval (620 iterations of ER+HIO) run on the same machine. We expect that such real-time feedback will be crucial to coherent imaging experiments especially in the light of ongoing upgrades to major light sources such as Advanced Photon Source Upgrade project (APS-U), European Synchrotron Research Facility Extremely Brilliant Source ESRF-EBS and PETRA-III. Additionally, CDI NN was shown to be successful at recovering structure and phase even in the presence of strong phase structures that heavily distorts the coherent diffraction pattern about a Bragg peak, and this shows strong promise for the successful reconstruction of objects that have a high density of defects (if CDI NN is extended to 3D objects). We also note that while CDI NN's predictions are least accurate for small objects with strong phase structures (these cases have the lowest signal to noise ratio), the predictions still show reasonable agreement with the ground truth structure and phase.

Finally, whereas oversampling is a necessary condition for phase retrieval algorithms to work, CDI NN does not require oversampled data. CDI NN does not perform phase retrieval at all but instead learns the mapping between Fourier space intensity and phase and real-space structure and phase. Experimentally, the relaxation of this oversampling requirement will translate to several further advantages. For a given detector configuration (pixel size and distance), higher energies (that allow deeper penetration into material) and access to a higher volume in reciprocal space (that provides increased resolution) will become possible.

Methods

Simulated training data. Each instance of the training data was generated as follows; first a random convex object was created from a convex hull of a random scattering of points within a 32×32 grid. Array values within the object were set to 1, while values outside are set to 0. A Gaussian blur 1 pixel in width was applied to the object to smooth the edges. A second array (also 32×32) was created with a random, spatially varying phase field. This phase field was generated by mixing sine gratings and Gaussians with random weights, periods, positions and numbers. The peak to peak phase intensity is allowed to span the entire $-\pi$ to π range. For convenience, phases outside the compact object were set to 0. The corresponding diffraction pattern was then generated by taking the FT of a complex valued array created from the object's amplitudes and phases. Only the amplitude information from the computed diffraction patterns was retained for both training and testing of CDI NN. The generated training set contains a wide variety of structures and phase states, including some poorly defined complex objects.

CDI NN training. Training was performed in parallel on two NVIDIA K40 GPUs using the Keras package running the Tensorflow backend^{36,37}. We trained the networks for 10 epochs each using a batch size of 256. The training for each network took less than half an hour when trained in parallel across the two GPUs. At each step,

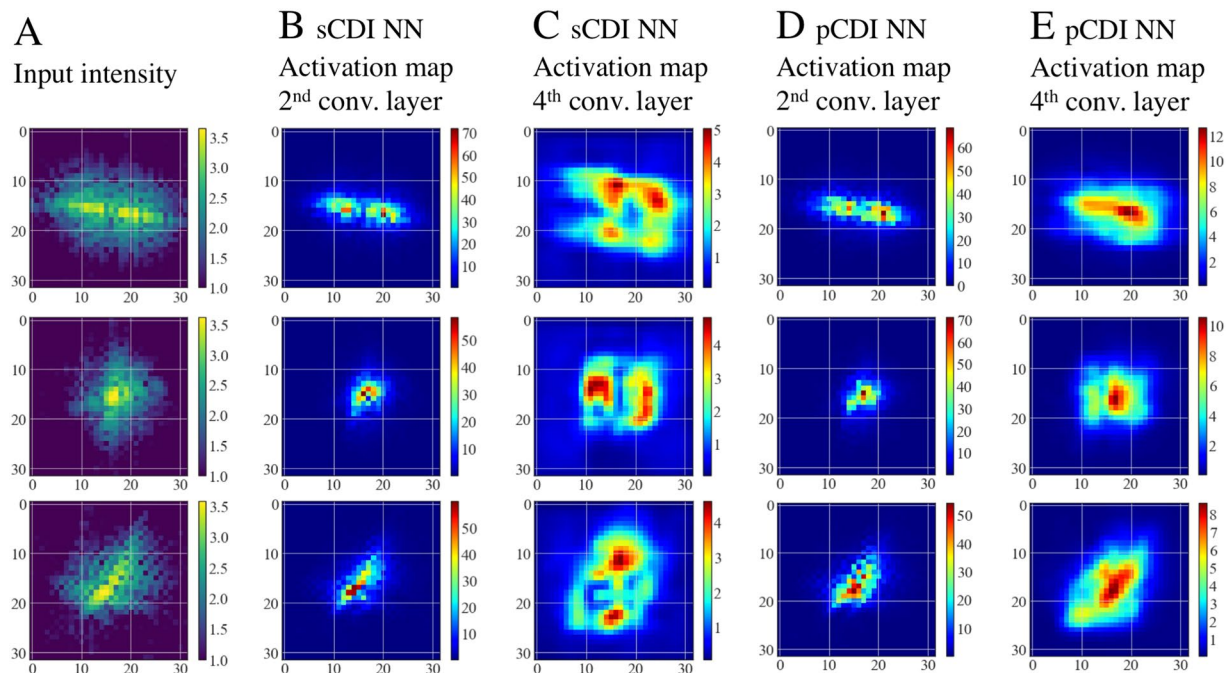


Figure 6. Activation maps of select convolutional layers. **(A)** Input diffraction patterns. **(B)** Average activation of the 2nd convolutional layer of sCDI NN. **(C)** Average activation of the 4th convolutional layer of sCDI NN. **(D)** Average activation of the 2nd convolutional layer of pCDI NN. **(E)** Average activation of the 4th convolutional layer of pCDI NN. Successive layers of CDI NN learn features at higher Q (higher resolution). Intensity units are arbitrary.

we used adaptive moment estimation (ADAM)³⁸ to update the weights while minimizing the per-pixel loss as defined by the crossentropy. We computed the performance of the network at the end of each training epoch using the validation set.

Phase retrieval. To perform phase retrieval, the complex object array was zero padded (to twice the size), Fourier transformed and the resulting amplitudes input to iterative phase retrieval that switched between error reduction (ER) and hybrid input-output (HIO)¹⁷. 620 iterations were performed using a shrink-wrapped support in real space²¹. The final 20 iterations were averaged over to obtain the final result.

Data Availability

The trained network, test data and accompanying Jupyter notebooks of Python code are available upon reasonable request from the corresponding author.

References

- Robinson, I. & Harder, R. Coherent X-ray diffraction imaging of strain at the nanoscale. *Nat. Mater.* **8**, 291–298 (2009).
- Miao, J., Ishikawa, T., Robinson, I. K. & Murnane, M. M. Beyond crystallography: Diffractive imaging using coherent x-ray light sources. *Science* (80-). **348**, 530–535 (2015).
- Zuo, J. M., Vartanyants, I., Gao, M., Zhang, R. & Nagahara, L. A. Atomic Resolution Imaging of a Carbon Nanotube from. *Science* (80-). **300**, 1419–1422 (2003).
- Dean, B. H., Aronstein, D. L., Smith, J. S., Shiri, R. & Acton, D. S. Phase retrieval algorithm for JWST Flight and Testbed Telescope. In *Proceedings of SPIE* (eds Mather, J. C., MacEwen, H. A. & de Graauw, M. W. M.) **626511**, 626511 (2006).
- Szameit, A. *et al.* Sparsity-based single-shot subwavelength coherent diffractive imaging. *Nat. Mater.* **11**, 455–459 (2012).
- Chapman, H. N. *et al.* High-resolution ab initio three-dimensional x-ray diffraction microscopy. *JOSA A* **23**(5), 1179–1200 (2006).
- Newton, M. C., Leake, S. J., Harder, R. & Robinson, I. K. Three-dimensional imaging of strain in a single ZnO nanorod. *Nat. Mater.* **9**, 120–124 (2010).
- Pfeifer, M. A., Williams, G. J., Vartanyants, I. A., Harder, R. & Robinson, I. K. Three-dimensional mapping of a deformation field inside a nanocrystal. *Nature* **442**, 63–66 (2006).
- Cherukara, M. J. *et al.* Ultrafast Three-Dimensional X-ray Imaging of Deformation Modes in ZnO Nanocrystals. *Nano Lett.* **17**, 1102–1108 (2017).
- Cherukara, M. J. *et al.* Ultrafast Three-Dimensional Integrated Imaging of Strain in Core/Shell Semiconductor/Metal Nanostructures. *Nano Lett.* **17**, 7696–7701 (2017).
- Clark, J. N. *et al.* Imaging transient melting of a nanocrystal using an X-ray laser. *Proc. Natl. Acad. Sci. USA* **112**, 7444–7448 (2015).
- Clark, J. N. *et al.* Three-dimensional imaging of dislocation propagation during crystal growth and dissolution. *Nat. Mater.* **14**, 780–784 (2015).
- Ulvestad, A. *et al.* Avalanching strain dynamics during the hydriding phase transformation in individual palladium nanoparticles. *Nat. Commun.* **6**, 10092 (2015).
- Yau, A., Cha, W., Kanan, M. W., Stephenson, G. B. & Ulvestad, A. Bragg Coherent Diffractive Imaging of Single-Grain Defect Dynamics in Polycrystalline Films. *Science* (80-). **742**, 739–742 (2017).

15. Ulvestad, A. *et al.* *In situ* strain evolution during a disconnection event in a battery nanoparticle. *Phys. Chem. Chem. Phys.* **17**, 10551–10555 (2015).
16. Cherukara, M. J. *et al.* Three-Dimensional Integrated X-ray Diffraction Imaging of a Native Strain in Multi-Layered WSe₂. *Nano Lett.* **18**, 1993–2000 (2018).
17. Fienup, J. R. Phase retrieval algorithms: a comparison. *Appl. Opt.* **21**, 2758–2769 (1982).
18. Elser, V., Rankenburg, I. & Thibault, P. Searching with iterated maps. *Proc. Natl. Acad. Sci. USA* **104**, 418–23 (2007).
19. Chen, C. C., Miao, J., Wang, C. W. & Lee, T. K. Application of optimization technique to noncrystalline x-ray diffraction microscopy: Guided hybrid input-output method. *Phys. Rev. B - Condens. Matter Mater. Phys.* **76**, 064113 (2007).
20. Ihli, J. *et al.* Strain-relief by single dislocation loops in calcite crystals grown on self-assembled monolayers. *Nat. Commun.* **7**, 11878 (2016).
21. Marchesini, S. *et al.* X-ray image reconstruction from a diffraction pattern alone. *Phys. Rev. B* **68**, 140101(R) (2003).
22. Ulvestad, A. *et al.* Identifying Defects with Guided Algorithms in Bragg Coherent Diffractive Imaging. *Sci. Rep.* **7**, 1–9 (2017).
23. Miao, J., Sayre, D. & Chapman, H. N. Phase retrieval from the magnitude of the Fourier transforms of nonperiodic objects. *J. Opt. Soc. Am. A* **15**, 1662 (1998).
24. Hornik, K. Approximation capabilities of multilayer feedforward networks. *Neural Networks* **4**, 251–257 (1991).
25. Lecun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
26. Chollet, F. *Deep Learning with Python*. (Manning Publications Co., 2017).
27. Fakhry, A., Zeng, T. & Ji, S. Residual Deconvolutional Networks for Brain Electron Microscopy Image Segmentation. *IEEE Trans. Med. Imaging* **0062**, 1–1 (2016).
28. Zhu, B., Liu, J. Z., Cauley, S. F., Rosen, B. R. & Rosen, M. S. Image reconstruction by domain-transform manifold learning. *Nature* **555**, 487–492 (2018).
29. Rivenson, Y. *et al.* Deep Learning Enhanced Mobile-Phone Microscopy. *ACS Photonics*, <https://doi.org/10.1021/acsp Photonics.8b00146> (2018).
30. Rivenson, Y., Zhang, Y., Günaydin, H., Teng, D. & Ozcan, A. Phase recovery and holographic image reconstruction using deep learning in neural networks. *Light Sci. Appl.* **7**, 17141 (2018).
31. Sinha, A., Lee, J., Li, S. & Barbastathis, G. Lensless computational imaging through deep learning. **4** (2017).
32. Kamilov, U. S. *et al.* Learning approach to optical tomography. *Optica*. **2**(6), 517–522 (2015).
33. Metzler, C. A., Schniter, P., Veeraraghavan, A. & Baraniuk, R. G. prDeep: Robust Phase Retrieval with Flexible Deep *Neural Networks*. doi:arXiv:1803.00212v1 (2018).
34. Nair, V. & Hinton, G. E. Rectified Linear Units Improve Restricted Boltzmann Machines. *Proc. 27th Int. Conf. Mach. Learn.* 807–814, doi:10.11.1.165.6419 (2010).
35. Shelhamer, E., Long, J. & Darrell, T. Fully Convolutional Networks for Semantic Segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**, 640–651 (2017).
36. Chollet, F. & others. Keras (2015).
37. M Abadi *et al.* Large-Scale Machine Learning on Heterogeneous Systems. (2015).
38. Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization. 1–15, <https://doi.org/10.1145/1830483.1830503> (2014).

Acknowledgements

This work was supported by Argonne LDRD 2018-019-N0 (A.I C.D.I: Atomistically Informed Coherent Diffraction Imaging). An award of computer time was provided by the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. This work was performed, in part, at the Center for Nanoscale Materials, a U.S. Department of Energy Office of Science User Facility, and supported by the U.S. Department of Energy, Office of Science, under Contract No. DE-AC02-06CH11357. This work also used computational resources at the Advanced Photon Source. Use of the Advanced Photon Source and Argonne Leadership Computing Facility, both Office of Science user facilities, was supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No. DE-AC02-06CH11357.

Author Contributions

M.J.C., Y.N. and R.H. designed the research. M.J.C. built, trained and tested the deep learning networks. All authors contributed to the analysis, discussion and writing of the manuscript.

Additional Information

Competing Interests: The authors declare no competing interests.

Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2018