
Y-net: A Physics-constrained and Semi-supervised Learning Approach to the Phase Problem in Computational Electron Imaging

Nouamane Laanait

Computational Sciences and Engineering Division
Oak Ridge National Laboratory
Oak Ridge, TN, USA 37832 laanaitn@ornl.gov

Junqi Yin

Oak Ridge Leadership Computing Facility
Oak Ridge National Laboratory
Oak Ridge, TN, USA 37832 yinj@ornl.gov

Albina Borisevich

Center for Nanophase Materials Sciences
Oak Ridge National Laboratory
Oak Ridge, TN, USA 37832 albinab@ornl.gov

Abstract

The phase problem in diffraction physics is one of the oldest inverse problems in all of science. The central difficulty that any approach to solving this inverse problem must overcome is that half of the information, namely the phase of the diffracted beam, is always missing. In the context of electron microscopy, the phase problem is generally non-linear and solutions provided by phase-retrieval techniques are known to be poor approximations to the physics of electrons interacting with matter. Here, we show that a semi-supervised learning approach can effectively solve the phase problem in electron microscopy/scattering. In particular, we introduce a new Deep Neural Network (DNN), Y-net, which simultaneously learns a reconstruction algorithm via supervised training in addition to learning a physics-based regularization via unsupervised training. We demonstrate that this constrained, semi-supervised approach is an order of magnitude more data-efficient and accurate than the same model trained in a purely supervised fashion. In addition, the architecture of the Y-net model provides for a straightforward evaluation of the consistency of the model's prediction during inference and is generally applicable to the phase problem in other settings.

1 Introduction

Advances in materials have shaped the course of human civilization from the bronze age to the silicon-powered information age. Future advances in materials science depend critically on our ability to determine, with atomic resolution (10^{-10} m), a material's *local* electron density. This goal is within reach, for the first time, via a computational imaging technique commonly known as 4D-STEM (scanning transmission electron microscopy).

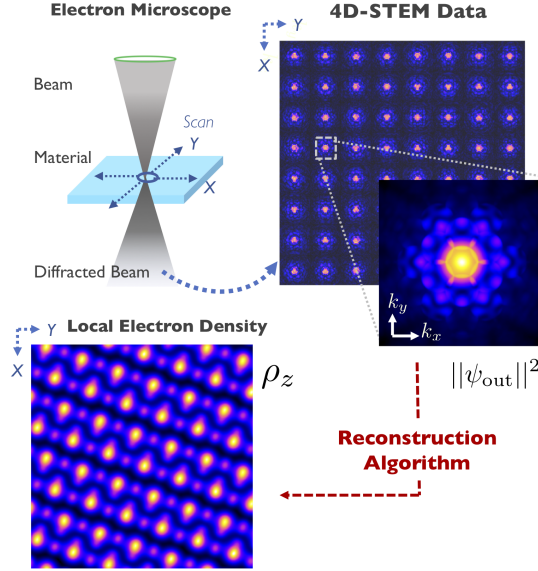


Figure 1: 4D-STEM is a computational imaging technique, where a picometer-size beam is scanned across a material and a diffraction pattern is collected at each spatial location. The resultant dataset is 4-D dimensional, where each “pixel” is indexed by (x, y, k_x, k_y) , where $k = \alpha/\lambda$, λ is the wavelength and α is the diffraction angle of the electron beam. Successful inversion of the 4D-STEM data should, in principle, provide local electron density maps of materials with higher spatial resolution and sensitivities than in another existing technique.

In 4D-STEM, a picometer-sized (10^{-12} m) electron beam is 2-D raster scanned across the material to collect a diffraction pattern with picometer spatial resolutions from each (x, y) position (see Figure 1). The resultant dataset is 4-D dimensional, where each “pixel” is indexed by (x, y, k_x, k_y) , where (k_x, k_y) are the wave-vectors associated with a diffracting beam. A 4D-STEM dataset encodes information about the material’s electron density from the vantage point of a single atom. Decoding electron diffraction patterns into the local electronic density is a longstanding inverse problem for two principal reasons Zuo & Spence (2013). First, the quantum interaction of electrons with matter is strong, which produces numerous interference processes and the resultant inverse problem is non-linear. Second, diffraction patterns provide incomplete data, since they only provide intensities as opposed to the complex-valued diffracted electron wavefunction.

Here, we introduce a semi-supervised and physics-constrained Deep Neural Network (DNN) to solve the phase problem in 4D-STEM, thereby reconstructing both the local electron density and the incident electron beam wavefunction. Estimating both of these quantities allows one to a posteriori quantify the reconstruction error. We also discuss how our approach is naturally extensible via differentiable programming.

2 Background

The interactions between the quantum wavefunction of an electron beam, ψ , and a material described by some scattering potential, V , is given by the fast-electron Schrödinger equation:

$$\frac{\partial \psi(\mathbf{r})}{\partial z} = \frac{i\lambda}{4\pi} \nabla_2^2 \psi(\mathbf{r}) + i\sigma V(\mathbf{r})\psi(\mathbf{r}) \quad (1)$$

where $\nabla_2^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$ is the 2-D Laplacian operator, and σ is a constant characterizing the strength of the interaction between the electron beam and the scattering potential of a material, V . Once V is determined, the local electronic density $\rho(\mathbf{r})$ is directly found (via an explicit solution of Poisson’s equation). Additionally, if either V or ρ and the initial value (i.e. incident wavefunction) $\psi_{in}(\mathbf{r})|_{\mathbf{r}=0}$ are known, Equation 1 can be solved using nearly-exact numerical methods (described in the next section).

Direct inversion of the complex-valued partial differential equation (PDE) in Equation 1, however, purely from experimental data is difficult. The central difficulty lies entirely in the fact that experimentally, one can only measure image intensities (i.e. diffraction patterns) of the exiting beam electrons, given by the squared modulus of the beam’s complex-value quantum wavefunction. Consequently, *half of the information to directly invert the PDE is always missing*, a challenge known as the phase problem Born & Wolf (2013). Various iterative phase-retrieval techniques have been developed to solve the phase problem Sayre (1952); Fienup (1978); Hoppe (1969); Miao et al. (1999, 2015); Rodenburg (2008). Unfortunately, most existing phase-retrieval techniques are not generally applicable to 4D-STEM due to the strong interaction between electrons and matter Zuo & Spence (2013).

3 Contribution

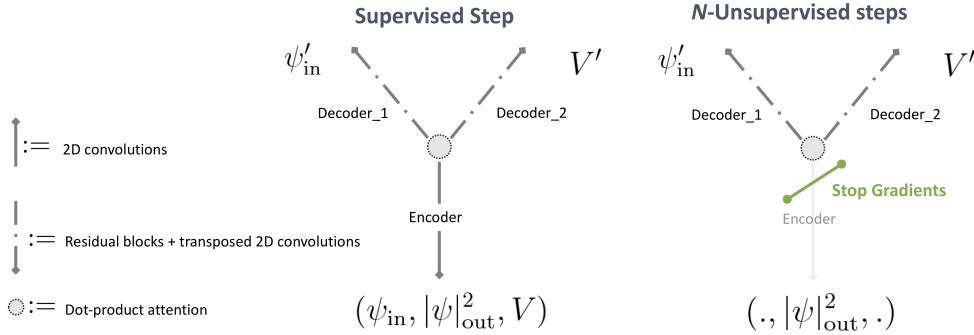


Figure 2: Illustration of the Y-net model architecture and its training policy. A training step of Y-net consists of alternating supervised learning with examples and labels sampled from the forward model, followed by multiple unsupervised learning steps, where the two decoder branches seek to minimize a learned regularization which relates the decoder outputs to the input data.

We begin by noting that the mathematical form of Equation 1 is prevalent in all of optics and diffraction, known as the scalar (inhomogeneous) Helmholtz equation, and describes propagation and interaction of a scalar wave-field (e.g. electrons, x-rays, optical lasers) with an object described by a scattering potential. Consequently, a complete solution to this central equation requires determining both the initial complex scalar wave-field, ψ_{in} as well as V .

Motivated by this universal property of the phase problem in optics, we devised a DNN architecture that learns to output both ψ_{in} and V . We achieved this by adopting a standard autoencoder architecture, except we split the decoder into two separate decoding paths or branches, as shown in Figure 3. The central idea of our approach consists of observing that the outputs of the decoder branches and the inputs to the model must all satisfy strict mathematical relations (dictated by fundamental laws of physics), and these physics-based constraints can be used to train Y-net in an unsupervised manner, in addition to, supervised learning on known ground truth data sampled from an explicit solution of the forward model of Equation 1. The overall loss function of Y-net is given by,

$$\mathcal{L} = \ell_2(\psi_{\text{in}} - \psi'_{\text{in}}) + \ell_2(V - V') + \varepsilon \ell_2(\psi_{\text{out}} - f(\psi'_{\text{in}}, V')), \quad (2)$$

the first 2 terms on the right hand side are inversion errors of the predicted incident wavefunction and the scattering potential from their true values (i.e. labels). The last term is the input reconstruction error of the autoencoder, where f is any function that outputs ψ'_{out} given ψ'_{in} and V' .

Our approach can be interpreted as combining previous Deep Learning approaches to solving inverse problems by learning both the inverse operator which maps observations to the solution (via supervised learning) as well as by learning a data-aware regularization that follows from the physics of the problem.

Here, we choose a prevalent approximation used in many phase-retrieval techniques known as the thin-object approximation which relates the inputs of the DNN to its outputs and is defined as


$$\psi_{\text{out}}(k_x, k_y) = \mathcal{F}\{\psi_{\text{in}}(x, y) \times e^{iV(x,y)}\} \quad (3)$$

where \mathcal{F} is a 2-dimensional Fourier transform. Within this approximation, $f(\psi'_{\text{in}}, V') = \mathcal{F}\{\psi'_{\text{in}} \times e^{iV'}\}$ in Equation 2. Given that our expression for f is an approximation, it can be effectively used as a regularization term on our solution (ψ'_{in}, V') to the phase problem. Note, however, it is a learnable regularization in the sense that it generates gradients we can use to train the two decoder branches.

Training of our DNN model, which we refer to as Y-net, proceeds via a supervised learning step, where the true (ψ_{in}, V) are provided as labels, followed by a number of unsupervised steps where gradient-descent is performed only on the input reconstruction loss (last term in Equation 2) and the gradients are not applied to the encoder branch. These two stages of training are then repeated across the entire training run. In practice, we found that 5-10 steps of unsupervised training steps following a supervised training step were sufficient.

In Figure 3 we show that the reconstruction quality of the Y-net model trained purely with supervised learning (ignoring the last term in Equation 2) is substantially inferior to Y-net trained in a semi-supervised fashion as described above.

We have quantified both the data efficiency and the reconstruction error for both type of training across a test data set spanning thousands of materials. We have also carried out various ablation studies to quantify the essential architectural details of Y-net and tested different learning policies. Finally, we compared Y-net’s reconstruction quality to model-based iterative reconstructions.



4 Conclusion

In summary, we found that by combining supervised and unsupervised learning to train a new DNN architecture, we can learn both a solution to an inverse problem as well as learn a physics-based regularization, leading to vastly improved reconstruction quality and data efficiency. We are currently extending the presented framework by substituting the learnable regularization with the full forward model in Equation 1 using techniques from differentiable programming.

Acknowledgments

This research was funded by a Lab Directed Research and Development project at Oak Ridge National Laboratory, a U.S. Department of Energy facility managed by UT-Battelle, LLC. An award of computer time was provided by the INCITE program. This research also used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

References

- Born, M. and Wolf, E. *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light*. Elsevier, 2013.
- Fienup, J. R. Reconstruction of an object from the modulus of its fourier transform. *Opt. Lett.*, 3(1):27–29, Jul 1978.
- Hoppe, W. Beugung im inhomogenen Primärstrahlwellenfeld. I. Prinzip einer Phasenmessung von Elektronenbeugungsinterferenzen. *Acta Crystallographica Section A*, 25(4):495–501, Jul 1969.
- Miao, J., Charalambous, P., Kirz, J., and Sayre, D. Extending the methodology of x-ray crystallography to allow imaging of micrometre-sized non-crystalline specimens. *Nature*, 400(6742):342, 1999.
- Miao, J., Ishikawa, T., robinson, I. K., and Murnane, M. M. Beyond crystallography: Diffractive imaging using coherent x-ray light sources. *Science*, 348(6234):530–535, 2015.
- Rodenburg, J. M. Ptychography and Related Diffractive Imaging Methods. *Advances in Imaging and Electron Physics*, 150:87–184, January 2008.
- Sayre, D. Some implications of a theorem due to Shannon. *Acta Crystallographica*, 5(6):843–843, November 1952.
- Zuo, J. and Spence, J. *Electron microdiffraction*. Springer Science & Business Media, 2013.