

OPTICS

Three dimensions, two microscopes, one code: Automatic differentiation for x-ray nanotomography beyond the depth of focus limit

Ming Du¹, Youssef S. G. Nashed², Saugat Kandel³, Doğa Gürsoy^{4,5}, Chris Jacobsen^{4,6,7*}

Conventional tomographic reconstruction algorithms assume that one has obtained pure projection images, involving no within-specimen diffraction effects nor multiple scattering. Advances in x-ray nanotomography are leading toward the violation of these assumptions, by combining the high penetration power of x-rays, which enables thick specimens to be imaged, with improved spatial resolution that decreases the depth of focus of the imaging system. We describe a reconstruction method where multiple scattering and diffraction effects in thick samples are modeled by multislice propagation and the 3D object function is retrieved through iterative optimization. We show that the same proposed method works for both full-field microscopy and for coherent scanning techniques like ptychography. Our implementation uses the optimization toolbox and the automatic differentiation capability of the open-source deep learning package TensorFlow, demonstrating a straightforward way to solve optimization problems in computational imaging with flexibility and portability.

INTRODUCTION

Depending on the photon energy used, x-rays are able to penetrate into samples with a thickness ranging from micrometers to centimeters. At the same time, x-ray microscopes are beginning to be able to deliver images with the sub-10-nm spatial resolution (1, 2). However, combining these characteristics is complicated by the fact that any imaging method with spatial resolution δ_r has a depth of focus (DOF) limit (3, 4) of

$$\text{DOF} = \frac{2}{0.61^2} \frac{\delta_r^2}{\lambda} \approx 5.4 \delta_r \frac{\delta_r}{\lambda} \quad (1)$$

This is straightforward to understand in lens-based imaging systems. However, even when lensless imaging methods involving wavefront recovery are used, the DOF limit of Eq. 1 gives the axial distance over which features can be considered to all lie within a common transverse plane before subsequent wavefield propagation effects are taken into account. That is, Eq. 1 represents the limit of validity of the pure projection approximation, within which a depth-extended object can be treated as producing a simple pure projection image when viewed from one illumination direction. For objects thicker than the DOF limit, one must instead account for wave propagation effects within the specimen. This will be especially important for fully exploiting the marked increases in coherent x-ray flux that the next generation of synchrotron light sources will provide (5).

One approach to simulate wave propagation in a complex object is the finite difference method (6). While it fits equally as the forward model in our reconstruction scheme, we use here another

wave propagation method named multislice propagation for our work. Multislice wave propagation (7) is a historic, simple, yet powerful method allowing one to account for wave diffraction in an inhomogeneous medium. The multislice simulation method subdivides the problem into a series of slab-wise refractive modulation and propagation operations and accounts for the change of the probe wave throughout the object instead of assuming a constant probe. Hence, it can provide accurate numerical results for propagation through a complicated object and remains valid over a much larger object thickness compared with diffraction tomography models assuming single scattering (8). The incorporation of multislice propagation provides a reliable strategy for the reconstruction of beyond-DOF objects.

We describe here an approach for imaging objects that extend beyond the DOF limit and within which multiple scattering might take place. We formulate the three-dimensional (3D) object reconstruction problem as a minimization problem that incorporates a data fidelity term (the L2 norm of predicted and measured data) and a regularization term (the L1 norm of the object and its gradient), where multislice wave propagation is used to accurately model the exit wave leaving the object. Because the new model better captures the wave-object interactions for any object size, the same model can be applied to reconstruct either near-field imaging with propagation phase contrast or ptychography (Fig. 1), without the need for any modification. We use the Adaptive Moment Estimation (Adam) optimizer that is implemented in TensorFlow, which is Google Brain's open-source software library. The automatic differentiation (AD) capability in TensorFlow allows us to solve for the two imaging methods with modest code branching. With this approach, we are able to use one computer code for two different types of microscopes to reconstruct 3D objects beyond the DOF limit.

It is worthwhile noticing that while there exist several multislice-based reconstruction methods, which have proven success in several imaging scenarios (9–12), our method differs from them in a few aspects. Our implementation provides both a ptychography mode and a full-field mode, while the above methods are concerned with ptychography alone. In addition, instead of requiring planes

Copyright © 2020
The Authors, some
rights reserved;
exclusive licensee
American Association
for the Advancement
of Science. No claim to
original U.S. Government
Works. Distributed
under a Creative
Commons Attribution
NonCommercial
License 4.0 (CC BY-NC).

¹Department of Materials Science, Northwestern University, Evanston, IL 60208, USA.

²Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, IL 60439, USA. ³Applied Physics Program, Northwestern University, Evanston, IL 60208, USA.

⁴Advanced Photon Source, Argonne National Laboratory, Lemont, IL 60439, USA.

⁵Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, IL 60208 USA.

⁶Department of Physics and Astronomy, Northwestern University, Evanston, IL 60208, USA.

⁷Chemistry of Life Processes Institute, Northwestern University, Evanston, IL 60208, USA.

*Corresponding author. Email: cjacobsen@anl.gov

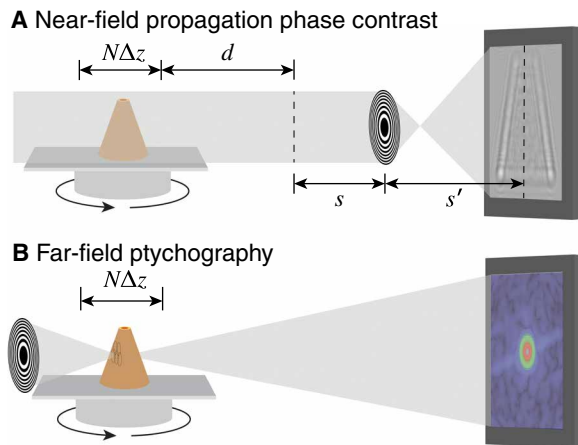


Fig. 1. Schematic representation of the two different microscope types used in our demonstration. In the full-field mode (A), phase contrast is incorporated by allowing the wavefield leaving the object to undergo Fresnel propagation over a distance d before that plane is imaged by a lens (with $1/f = 1/s + 1/s'$) onto a detector. In ptychography (B), a small coherent beam spot (probe) is scanned through the object, and the far-field diffraction intensities are recorded for each probe position. The schematic here shows Fresnel zone plates as lenses for quasi-monochromatic x-ray beams.

that are axially separated by 1 DOF or more, in our method, the spacing between slices can be equal to the lateral pixel width, which allows for an isotropic voxel size. Last, the method for updating the object function is different. In (9), slices are updated sequentially using an update function that resembles the modulus replacement operation in ePIE (extended ptychographic iterative engine), a reconstruction engine for 2D ptychography. In (10), the first method described is similarly based on modulus replacement, while the second method involves the minimization of a loss function that has a similar form of ours. However, in our approach to full-field microscopy, the loss equation is constructed also with a sparsity constraint, and non-negative and finite support constraints are applied to the object function throughout the minimization process. This also distinguishes our work from prior studies (11). Last, our use of AD through the widely used software package TensorFlow renders the implementation highly accessible on various computing platforms. Building upon prior work in 2D-phase retrieval using AD (13, 14), our results reinforce the vast potential of AD for a large variety of computational imaging tasks.

Imaging beyond the DOF limit

Present-day x-ray nanotomography is usually performed within the DOF limit of Eq. 1, such as with 1- μm resolution at 25 keV (giving $\lambda = 0.050$ nm and DOF = 110 mm) or 20-nm resolution at 6.2 keV (giving $\lambda = 0.20$ nm and DOF = 11 μm). In these cases, one can obtain measurements that represent a pure projection through the specimen at each rotation angle using standard-phase retrieval methods such as those based on the inversion of the transport-of-intensity equation (15); one can then use standard tomographic reconstruction algorithms such as filtered back-projection (FBP). For objects that are thicker and/or interact more strongly, the complete solution of the wave function of electromagnetic wave within an inhomogeneous scattering potential field results in a recursive equation. With the first iteration, one arrives at the first Born approximation, which physically accounts for single scattering within the sample. On this basis, one can approximate the imaging of thicker speci-

mens by acknowledging the fact that the far-field diffraction pattern of an object provides information on the surface of the Ewald sphere corresponding to the beam energy and viewing direction (16). This remapping of Fourier space information from a plane (pure projection), to the surface of the Ewald sphere, is used in filtered back-propagation algorithms in diffraction tomography (8). It has been applied in tomographic diffractive microscopy with visible light (17) and has been demonstrated in x-ray coherent diffraction imaging (18).

One approach that has been developed for imaging beyond the DOF limit is multislice ptychography (9). In standard ptychography (19), one scans a finite-sized coherent beam with overlap across a planar sample, records the set of far-field diffraction patterns, and separates or factorizes the probe from the optical modulation at each scan position. Multislice ptychography is based on utilization of the multislice method (7) to propagate a beam through a thick object, where the refractive effect of the first thin slab of the object is applied to the incident wavefield, the wavefield is free space propagated to the next slab position, and the process is repeated until one obtains the exit wave leaving the object (which can then be free space propagated to a far-field detector, for example). If the object is, in fact, composed of a series of discrete planes separated axially by 1 DOF or more, then one can factorize the probe from both transverse positions and axial planes. One can also account for violation of the Born approximation, in that the object-modulated exit wave from the upstream plane is propagated to the next axial plane in a recursive manner through all planes. This approach has been used with success in ptychography using visible light (20), x-rays (10, 21, 22), and electrons (23). It has also been used for tomographic imaging of more continuous specimens by assuming that the object could be represented by discrete axial planes separated by the DOF (24). However, this assumption is only approximately true, since one can often see image contrast variations with defocus settings of less than 1 DOF (or the separation of the slices), especially in phase contrast, which is the dominant contrast mechanism in transmission x-ray microscopy (25). In that case, variation of features along the beam axis between each two adjacent slices will not be captured.

Therefore, it can be advantageous to use reconstruction methods that use a forward model of multislice propagation in a continuous object and retrieve directly the refractive indices of the object instead of the phase of the exiting waves. This allows multiple scattering effects to be included and avoids a separate phase unwrapping. We have previously proposed an analytical model on the estimation of the degree to which multiple scattering alters the recorded data [see page 301 of (26)]. Using this model, it can be shown that one must begin to account for multiple scattering effects at a specimen thickness of a few micrometers in soft x-ray imaging at 0.5 keV and a few tens of micrometers in hard x-ray imaging at 15 keV. The need for the inclusion of multiple scattering effects is well known in optical diffraction microscopy (17), and we have shown that this approach can be used for x-ray microscopy as well (11). What we describe below is a new approach that is different from the above: It uses the method of AD to carry out the reconstruction, which offers greater flexibility on imaging method and for incorporating various constraints on the object as numerical optimization regularizers.

Formulation of the image reconstruction problem

Our approach is to treat image reconstruction of objects beyond the DOF limit as a numerical optimization problem. That is, we wish to

find the optimal parameter set \mathbf{n}_0 of the forward model f by minimizing an objective function L , leading to a solution of

$$\mathbf{n}_0 = \arg \min_{\mathbf{n}} L[f(\mathbf{n}), \mathbf{y}], \text{ subject to } \mathbf{n} \in \Phi \quad (2)$$

where the observable \mathbf{y} is the set of experimental measurements (near-field images or far-field diffraction patterns) and Φ is the manifold of constraints that \mathbf{n} is subject to. The parameter set \mathbf{n} will be defined in Eq. 10 below to be proportional to the x-ray refractive index distribution within the object's voxel grid positions \mathbf{r} . This refractive index is written as

$$n(\mathbf{r}) = 1 - \delta(\mathbf{r}) - i\beta(\mathbf{r}) \quad (3)$$

where the values of δ and β for various materials are readily obtained from tabulations (27). Except at photon energies right near x-ray absorption edges where anomalous dispersion effects can appear, δ and β have small positive values (typically $\delta \approx 10^{-4}$ and $\beta \approx 10^{-5}$ to 10^{-6}) so a positivity constraint can be applied to their solution. One can also apply a sparsity constraint for objects that are relatively discrete in space (28), and in most cases, tomography experiments are designed so that the object fits within the field of view so one can also apply a finite support constraint on the solution of $n(\mathbf{r})$.

Because the size of the optimization problem is large, efficient mechanisms must be used to find gradients for each iteration of the first-order solver used in optimizing Eq. 2. If one is always considering one type of imaging experiment, then one can calculate derivatives of the cost function, and this approach has been used with success for simulations of x-ray ptychographic reconstruction of objects beyond the DOF limit (11). However, if one wishes to be able to treat multiple imaging methods (so as to compare or benchmark their properties and performances, for example) and include a variety of regularizers, then other approaches that place the burden of finding minimization strategies on a computer rather than a scientist can have advantages. One approach is to represent multislice propagation with a computational architecture resembling a convolutional neural network and use mathematical formulations that are common in machine learning to solve for the object that matches the observations, as has been demonstrated for diffraction microscopy using visible light (29). AD (30) provides another approach that was suggested for use in phase retrieval problems (31) and then successfully implemented for x-ray ptychography (13). More recently, the adaptability of AD to a variety of coherent diffraction imaging methods has been demonstrated (14), and a variety of software toolkits are now available to implement this method (spurred on by their use for constructing the trainer module in supervised machine learning programs). We use this AD approach to reconstruct beyond-DOF imaging in two successful imaging methods and, thus, gain insight on their relative advantages and complications.

As noted above, in x-ray ptychography, one scans a finite-sized coherent beam through a series k of overlapping probe positions across the specimen and collects the far-field diffraction pattern from each. Because the extent of the far-field diffraction pattern is determined by the scattering properties of the object rather than the spatial resolution of the probe, one can obtain reconstructed x-ray images with a spatial resolution far finer than the size of the probe (32). In contrast, point projection x-ray microscopy (33), where an object is placed downstream of a point source of radiation, provides geometric magnification of the object with a penumbral blur limit

given by the source size and diffraction blurring that can be compensated for by near-field wave back-propagation (34, 35). Because near-field diffraction blurring is localized to a region given by the finest reconstructed feature size times the propagation distance divided by the wavelength, one can make use of illumination with a coherence width equal to this region rather than the entire illumination field as required for ptychography and, thus, make more complete use of partially coherent sources. Since the advancing of x-ray imaging instruments has granted considerable potential to both techniques in their imaging applications for thick samples with high resolution, it is of interest to understand the beyond-DOF imaging properties of both of these approaches.

We term our approach Adorym (Automatic Differentiation-based Object Retrieval with dYnamical Modeling). In Adorym, we wish to compare the present guess of the detected magnitude of $|f(\mathbf{n}, \theta, k, \Delta z, d)|$ against the square root $\sqrt{y_{\theta,k}}$ of the measured intensity at each angle θ and each scan position k in the experiment, with slice spacing Δz , and a free propagation distance d to the intensity measurement plane ($d \rightarrow \infty$ for ptychography). Then, a loss function gauging their disparity, which is expressed as

$$L = \frac{1}{N_{\theta} N_p N_k} \sum_{\theta,k} \| |f(\mathbf{n}, \theta, k, \Delta z, d)| - \sqrt{y_{\theta,k}} \|_2^2 + \alpha_{\delta} \|\mathbf{n}_{\delta}\|_1 + \alpha_{\beta} \|\mathbf{n}_{\beta}\|_1 + \gamma \text{TV}(\mathbf{n}_{\delta}) \quad (4)$$

is minimized. That is, the solution of the object function should be given by

$$\mathbf{n}_0 = \arg \min_{\mathbf{n}} (L) \text{ subject to } n_w = 0 \text{ for } n_w \notin \Theta \text{ and } n_w \geq 0 \text{ for } n_w \in \Theta \quad (5)$$

Further detail on the formulation of the forward model $f(\mathbf{n}, \theta, k, \Delta z, d)$ is provided in Materials and Methods. In Eqs. 4 and 5, $y_{\theta,k}$ is the measured intensity at orientation angle θ , N_{θ} is the number of projection angles, N_p is the number of pixels in each $y_{\theta,k}$, N_k is the number of probe positions for each projection angle ($N_k = 1$ for full field), and α_{δ} and α_{β} are the scalar normalizing coefficients added to the L1-norm regularization terms for the δ part and β part of \mathbf{n} , respectively. The separated regularization for the two parts of the object is necessary since $\delta(\mathbf{r})$ and $\beta(\mathbf{r})$ of the same material typically differ by a few orders of magnitude. Together, these two L1-norm terms enhance the sparsity of the object function, which is useful when the object is spatially discrete or contains a lot of empty space (such as a dispersion of cells or a hollow structure). Moreover, the anisotropic total variation $\text{TV}(\mathbf{n})$, weighted by coefficient γ , enhances the sparsity of the spatial gradient of the object function, which suppresses noise and unwanted heterogeneities (36). This regularizer is expressed as

$$\text{TV}(\mathbf{n}) = \sum_{l,m,n} [|x_{l+1,m,n} - x_{l,m,n}| + |x_{l,m+1,n} - x_{l,m,n}| + |x_{l,m,n+1} - x_{l,m,n}|] \quad (6)$$

where l , m , and n are indices along the three axes of \mathbf{n} . The TV regularizer is only applied to \mathbf{n}_{δ} because it usually carries higher contrast and better structural information than \mathbf{n}_{β} when hard x-rays are used. Last, Θ in Eq. 5 is a set corresponding to the 3D finite support constraint, which is explained in more detail in Materials and Methods.

RESULTS

Our Adorym approach was tested in simulations of thick objects that would normally involve multiple scattering along the beam path.

Two virtual samples were investigated: one hollow silicon cone with TiO₂ nanoparticle decoration that is described below (the details of the sample design are listed in Materials and Methods) and a protein sample that is presented in the Supplementary Materials. The reconstructions of the silicon cone sample shown here were performed on the computing cluster Cooley at the Argonne Leadership Computing Facility. Each node of this cluster is equipped with two 2.4-GHz Intel Haswell E5-2620 v3 central processing units (CPUs) (12 cores in total) and 384-gigabyte random access memory. Because of limits in graphics processing unit (GPU) memory, we did not use GPU acceleration on this machine. TensorFlow version 1.4.0 was used as the computational engine for our routine. For the Adam optimizer built with TensorFlow, we used a step size of $\lambda = 10^{-7}$ and first and second moment exponential decay rates of $\beta_1 = 0.9$ and $\beta_2 = 0.999$, respectively.

In x-ray full-field imaging, a variety of methods are available to obtain a phase-contrast image from detected intensities, but propagation-based phase contrast is the simplest to achieve experimentally. We therefore assumed that the x-ray wavefield exiting the object propagates downstream by a sample-detector distance d of 1 μm before the resulting wavefield is imaged without loss by a 1-nm resolution lens onto a detector, as shown in Fig. 1. This is beyond the state of the art with present-day x-ray optics, but as noted above, we chose parameters so as to substantially exceed the DOF limit within a small array size. Within this optical configuration, we simulated the recording of 500 images over a single-axis rotation range of 360° in one case and 180° in another case. The purpose of distinguishing the 360° and 180° rotations is that when diffraction is present in the sample, the images obtained at θ , as well as $\theta + 180$, can be different, unlike in conventional tomography. To acquire high-quality results, we conducted a series of experiments to determine that the optimal values for the regularizer weights in Eq. 4 were $\alpha_\delta = 1.5 \times 10^{-8}$ for the stronger phase-shifting part of the x-ray refractive index, and $\alpha_\beta = 1.5 \times 10^{-9}$ for the weaker absorptive part. The total variation minimization regularizer term of Eq. 6 was made small by setting $\gamma = 1 \times 10^{-11}$. The x-ray refractive index grid was initialized to a Gaussian distribution with a mean of $\delta = 8.7 \times 10^{-7}$ and $\beta = 5.1 \times 10^{-8}$ with SDs of about a 10th of the mean. These values are lower than the expected values but gave better reconstruction starts than values of zero; the reconstructions were not sensitive to the exact nonzero initialization values. In TensorFlow, we used a minibatch size of 10 and set the iterator to stop automatically once the incremental decrease in the total loss function of Eq. 4 fell below 3%. Parallelized with four threads, using three levels of multiscaling, and running on CPUs, the full-field reconstruction finished with 10, 10, and 6 epochs for the three passes with 4 \times , 2 \times , and 1 \times (original resolution) downsampling. The entire computation took approximately 5.0 hours of wall clock time and 120 core hours.

For ptychography, we assumed that an x-ray optic was used to focus a beam on the entrance of the object volume with a Gaussian profile beam profile with $\sigma_x = \sigma_y = 6$ nm and a maximum probe phase of 0.5 rad. A total of $23 \times 23 = 529$ probe positions were used to illuminate the specimen from each viewing angle, as shown in Fig. 1. The sparsity and smoothing constraints in ptychography are relaxed by probe overlap, so that a different set of regularizer values for Eq. 4 were used, with $\alpha_\delta = 1 \times 10^{-9}$, $\alpha_\beta = 1 \times 10^{-10}$, and $\gamma = 1 \times 10^{-9}$. The x-ray refractive index grid was initialized in the same way as the full-field case. The minibatch size was set to 1 rather than 10 so as to allow all data from one projection angle to fit in the com-

puter memory. For this larger dataset with far-field diffraction intensity recordings, the reconstruction was parallelized with 20 threads and required four epochs to yield a high-quality result over a wall clock time of 46 hours or 16,500 core hours.

Figure 2 shows the true object in row (A) and the reconstruction results for these various approaches in rows (B) through (E). In the 360° full-field reconstruction shown in row (B) and the 360° ptychographic reconstruction shown in row (D), the object boundaries are sharp, and features within the object are nicely reproduced. This is decidedly not the case for the error reduction (ER) + FBP or conventional tomographic reconstruction shown in row (E), where the small spheres on the outside of the object are poorly reproduced in the surface view of the fourth column, and the projection images of columns 2 and 3 do not accurately reproduce the true object. In visual appearance, one can argue that the ptychography reconstruction shown in row (D) is slightly sharper and has less “ghost” structure present in what are supposed to be empty voids inside the cone compared with the full-field reconstruction shown in row (B). This may be due to the fact that the large number of spatially separate illumination patterns used in ptychography help limit the regions that contribute scattering signal to each of the 529 individual data recordings acquired per rotation angle. On the other hand, the ptychographic reconstruction shows some slight fringe artifacts at the bottom of the cone, which might arise from the fact that the data are recorded in the far field rather than the near field. More quantitative comparisons will be presented below, using the Fourier shell correlation (FSC) method (37), which measures the consistency between images as a function of spatial frequency (resolution in the Fourier transform).

In conventional tomography within the pure projection approximation, projections obtained 180° apart are identical after projection reversal, so that data collected over a 180° range are sufficient for an accurate reconstruction. This is not the case when diffractive effects come into play, as has long been known in diffraction tomography (8) and as was observed in our previous study of ptychographic reconstructions of an object with DOF effects included (11). Modulations on the wavefield can be produced both by Fresnel diffraction from upstream features and refractive modulation from features at downstream planes; one cannot unambiguously distinguish between these two effects using a single viewing angle. To illustrate this, we have carried out a simulation of full-field imaging where the same 500 rotation angles used in the 360° case were instead distributed over a 180° angular range, giving the results shown in row (C). This leads to the presence of a number of voids in the reconstructed refractive index distribution, presumably because of the ambiguity noted above; the voids remained in the same position even if the 180° illumination angles were shifted to a different range, and they are not related to the shrink wrapping of the finite support. By removing the positivity constraint on the refractive index distribution and examining the intermediate object function as it was updated after each minibatch, we noticed that the values in the void regions became negative and kept decreasing. We therefore speculate that the voids might arise as the optimizer attempts to compensate for the rise of the loss function under information deficiency. The resolution of the 180° and 360° full-field reconstructions was evaluated using the FSC between two independent reconstructions with the same parameter settings, and the result shown in Fig. 3 shows the loss of resolution that results from using the same number of projection angles distributed over 180° only.

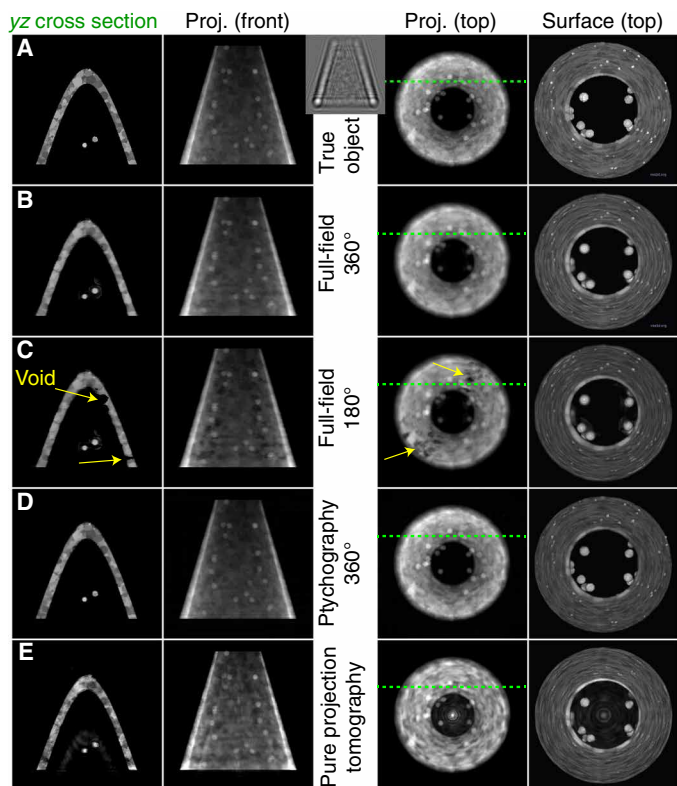


Fig. 2. Cone object used for computational experiments in beyond DOF x-ray imaging. The inset between columns 2 and 3 shows an example near-field diffraction image used in full-field reconstruction. The object (top row) and various reconstructed images (subsequent rows) are shown, displaying only the phase-shifting part δ of the x-ray refractive index since it provides higher contrast than the absorptive part β . In keeping with the convention of most synchrotron tomography experiments, the object is assumed to be rotated around the vertical or y axis. The first column shows a single-plane section of the object in the yz plane, with the location of the section indicated by the green dashed lines in the third column. The second column shows a projection or summation through the 3D grid in that same direction, while the third column shows a projection or summation through the 3D grid viewed from above. The fourth column shows the surface of the object, as rendered using Vaa3D (61). Row (A) shows the images of the ground truth. Both the full-field (B) and the ptychography (D) reconstructions of data acquired over a 360° rotation angle range reproduce the object with high fidelity. In the full-field reconstruction from 180° rotation data (C), voids appear where there is supposed to be material within the cone (in this case, the illumination was incident from top, to right, to bottom in the perspective of the third and fourth columns, although the same type of behavior was observed with different 180° illumination ranges). In the pure projection tomographic reconstruction shown in the bottom row (E), one obtains an image with lower resolution and greater differences from the true object: The fine TiO₂ spheres on the outside of the cone object are blurred out, and the reconstruction shows spurious material inside the cone. These images are all of a 256 × 256 nm² field of view. In the Supplementary Materials, we show the absolute value difference images between the yz cross section of column 1 for the true object and each experiment and reconstruction scheme.

To understand the robustness of our reconstruction method in the presence of noise due to limited exposure, we carried out full-field reconstructions where the recorded diffraction intensities were modified to incorporate noise. [Other studies have considered the noise robustness of simple coherent diffraction imaging against zone plate imaging (38) or against near-field imaging (39) with somewhat differing conclusions; we leave a comparison of full-field imaging

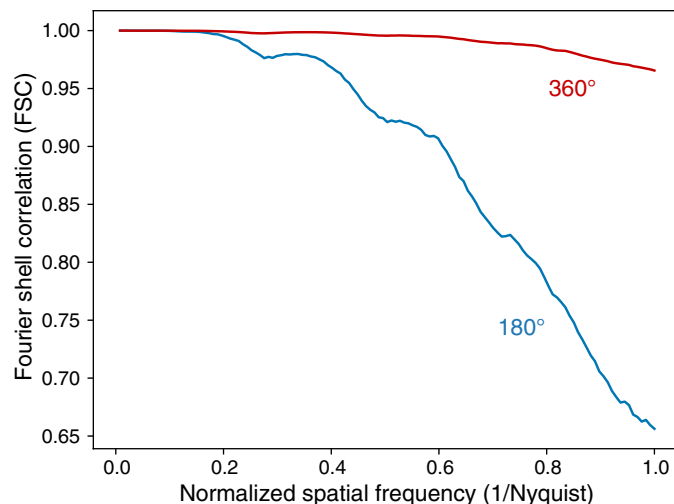


Fig. 3. FSC measurement of the resolution of full-field reconstructions using 500 rotation angles distributed over 360° (Fig. 2B) versus only 180° (Fig. 2C). The FSC measurement shows a considerable loss in spatial resolution over a wide range of spatial frequencies but especially at the highest spatial frequencies corresponding to the finest features. The FSC was calculated by comparing the phase-shifting part of the x-ray refractive index δ between two separate iterative reconstructions of the same full-field recording dataset.

and ptychography for future work.] This was performed by setting a quantity n_{ph} to be the total number of incident photons that intersect the object support with the object at each of the 500 projection angles. The object support is characterized by an area overdetermination ratio (AOR) of

$$\sigma_{AOR} = \frac{N^2}{N_{support}} \quad (7)$$

where N^2 is the number of pixels in the 2D object projection array and $N_{support}$ represents the total number of pixels within the finite support (40). The Gaussian-blurred “shrink wrap” procedure described above led to $\sigma_{AOR} \approx 57\%$ for the simulated cone object. With these factors considered, the number of photons $N_{pix,\theta}$ incident on each of the $N^2 = 256^2$ pixels is given by

$$N_{pix,\theta} = \frac{n_{ph}}{\sigma_{AOR} \times N^2} \quad (8)$$

for each of the N_θ viewing directions such that $n_{ph} = 1 \times 10^9$ corresponds to $N_{pix,\theta} = 2.7 \times 10^4$. The detected intensity images at each of the N_θ angles, generated by a normalized plane incident wave with unity magnitude, were scaled by $N_{pix,\theta}$ before Poisson noise was applied to them. We then obtained reconstructions from the Poisson-degraded datasets (Fig. 4) and compared them using the FSC method (37).

Given the normalized image intensity one would expect from a feature-present versus a feature-absent voxel, one can estimate the exposure required to see that object with a specified signal-to-noise ratio (SNR). Using the phase-contrast imaging expression of Eq. 39 of (26) for $t = 1$ -nm-thick Si at 5 keV, we obtain an exposure estimate of $N_{pix,\theta} = 5.0 \times 10^7$ for SNR = 5 imaging. Dose fractionation (41) tells us that this dose can be distributed over all N_θ viewing angles as 3D object statistics are built up from tomographic projections, so one would expect that to achieve full resolution at SNR = 5, one would require $N_{pix,\theta}/N_\theta = 1.0 \times 10^5$, which, from Eq. 8, translates to




Fig. 4. Evaluation on the noise robustness of the proposed algorithm. Near-field propagation images as would be recorded by the camera in Fig. 1A (top row) and reconstructed images (middle and bottom rows) in the presence of simulated Poisson noise with a total number of photons n_{ph} used in data recording. The middle row projection images correspond to the second column image of Fig. 2B, while the bottom row slice images correspond to the first column image of Fig. 2B. As can be seen, reducing the number of photons n_{ph} within the object (and, thus, the incident number of photons per pixel per viewing angle $N_{\text{pix},\theta}$ as given by Eq. 8) leads to a decrease in image fidelity and signal-to-noise ratio (SNR) at the single-pixel level. One can also evaluate this as a loss of spatial resolution with decreasing exposure, as shown in Fig. 5.

$n_{\text{ph}} = 3.7 \times 10^9$ per viewing angle. Because this exposure scales as SNR^2 , reducing n_{ph} from 3.7×10^9 to 1.0×10^9 corresponds to a decrease in SNR from 5 to $5/\sqrt{3.7/1.0} = 2.6$. Alternatively, because the radiation dose that must be necessarily imparted to a specimen for imaging at a specified SNR scales as the inverse fourth power of spatial resolution (26), a decrease in n_{ph} from 3.7×10^9 to $\{1 \times 10^7, 1 \times 10^8, 1 \times 10^9\}$ would be expected to correspond to a reduction in spatial resolution from 1 to $\{4.4, 2.5, 1.4\}$ nm. If one translates this into a fraction of the Nyquist sampling limit, the corresponding fractions are $\{0.23, 0.41, 0.72\}$. These fractions of the Nyquist sampling limit are shown via dashed lines in Fig. 5, and they are consistent with an FSC in the range of 0.5 to 0.6 as a measure of the spatial resolution limit.

DISCUSSION

We have shown here an approach (which we call Adorym, as noted above) whereby one can use AD and a multislice propagation forward model to reconstruct 3D objects in two different microscope types, with only minor branch points in one computer code base. Our approach has the following characteristics that are shared with another related non-AD approach (11), as well as with other multislice learning (29) and optimization-based (28) approaches:

1) In standard diffraction tomography approaches, one assumes that the 3D object can be decomposed into volume gratings so that data from one viewing angle are projected not onto a flat plane in Fourier space (as would be the case for a pure projection image) but onto the surface of the Ewald sphere. These assumptions are valid for the case where there is no multiple scattering in the specimen, but it can be shown [for example, see Figs. 2 and 3 in (26)] that

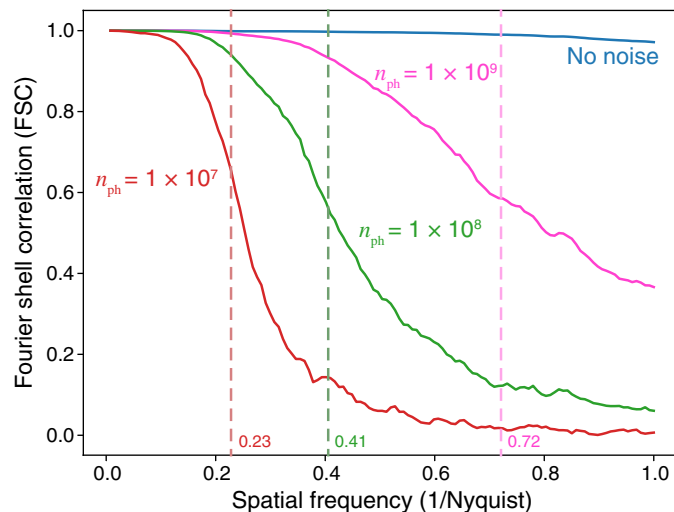


Fig. 5. The FSC for the full-field results with varying photon exposure n_{ph} as shown in Fig. 4. For each exposure, two different datasets were generated with different Poisson-distributed random noise, after which the same reconstruction algorithm with the same parameters was applied to each noisy dataset before performing the FSC. As discussed in the main text, one would expect the SNR to decline at normalized spatial frequencies of $\{0.23, 0.41, 0.72\}$ times the normalized spatial frequency for photon exposures of $n_{\text{ph}} = \{1 \times 10^7, 1 \times 10^8, 1 \times 10^9\}$. The dashed lines at these normalized spatial frequencies are all at roughly consistent decreases in the FSC to about 0.5 to 0.6 for the respective exposures, corresponding to a spatial resolution estimate of $\{4.4, 2.5, 1.4\}$ nm.

multiple scattering can play a role in x-ray imaging of thick specimens. That is, the illumination of downstream planes can be affected by the presence of strong features in upstream planes. Because our approach involves a full multislice forward calculation along each viewing angle, it can incorporate these effects correctly.

2) In previous multislice ptychography approaches (9), it has been assumed that the object can be decomposed into a set N_a of planes along the illumination direction with those planes separated by a DOF distance or more. This separation is required for allowing the combination of propagated probe and discrete object plane to be decomposed into sufficiently different results along the propagation direction. By allowing for an isotropic forward model where the plane separation distance can be the same as the transverse resolution distance, our approach is better able to represent the subtle contrast variations that occur in imaging over distances that are a reasonable fraction of the DOF. It should also be noted that the use of multislice methods to reconstruct N_a planes along the illumination direction means that one can reduce the number of illumination angles used (42) from what one might have expected based on the Crowther criterion (43).

3) Since refractive indices are being reconstructed directly in the proposed method, one can use additional constraints on the solution to address the phase unwrapping problem in reconstructions, so that a separate phase unwrapping process may become unnecessary. If the material is known, then this may be implemented by constraining the numerical range of \mathbf{n}_δ and \mathbf{n}_β . Alternatively, one can use an alternating direction method of multipliers (ADMM)-type solver for the joint optimization problem (44) that allows one to incorporate off-the-shelf phase unwrapping algorithms.

One would expect these characteristics to allow for the reconstruction of beyond-DOF 3D objects with greater fidelity than one

would have with multislice ptychographic tomography approaches; exploration of this hypothesis could be the topic of future work.

Our use of AD in a numerical optimization approach has several features:

1) For ptychography, it lets one phase the far-field Fourier magnitudes as was first suggested (31) and then demonstrated (13) in prior work. For near-field imaging, it avoids the approximations of uniform material type implicit in one commonly used approach (45).

2) It allows one to easily switch between different imaging modes (in this example, both near-field imaging and ptychography) within the same code framework, and it lets one explore different types of loss functions and regularizers without needing to rebuild the optimizer. This can be very useful for benchmarking different imaging and reconstruction techniques.

3) Several software packages that provide AD capabilities (such as TensorFlow and Autograd) are already built for parallelized operation on large compute clusters. As an example, an automated differentiation-based ptychography reconstruction code was demonstrated in (13). We carried out a direct test of our TensorFlow-based AD approach for ptychography against our previous result (11) using manual differentiation of the cost function and implementation in C++. The approach used here took 8.25 core hours per iteration per angle compared with 6.48 core hours per iteration per angle. One pays a modest penalty in computer time in this example but arguably uses less researcher time because AD does not require one to recalculate derivatives as the cost function is modified.

4) It also allows one to trivially compare and switch between synchronous and asynchronous schemes of optimization. In the synchronous scheme, object functions are broadcasted and synchronized among all threads for each several iterations. In the asynchronous scheme, each thread does the optimization on its own, and the object functions contained by them are only combined at the end. In the example discussed in (13), the synchronous approach took slightly longer to complete but gave more accurate results. The default scheme that we used to generate the above demonstrative data is a variant of the synchronous approach. Here, each thread keeps its own object function, but the gradient obtained by a thread is broadcasted and averaged along with the results of all other threads before being used to update the object function. The above stated characteristics have led to increasing attention to AD in the optics community. Other works have explored the use of AD for several other coherent diffraction imaging modalities (14).

In some cases, because of the special geometry of some samples, the range of rotation angles achievable in an experiment can be limited. We have shown that there is an improvement in reconstruction quality if one uses 360° rather than 180° rotation. We also plan on investigating other limits on rotation, such as a range from -70° to +70° and then +110° to -110°, which would be required for samples mounted on a flat substrate, as well as the alternative of laminography-type tilt and rotational sampling [see, for example, (46)] in a future study. In the cases where artifacts emerge due to missing rotation angles, regularizers can be easily added to suppress these unrealistic density concentrations.

Another point needing attention is that while the full-field mode and the ptychography mode are different in terms of acquisition method and processing wall time, the results they give are sometimes not equivalent as well. This is most obvious for diffusive features without clear boundaries, as in the case demonstrated in fig. S1. In this test case, a protein molecule (originally acquired using electron micro-

scope tomography) was numerically reconstructed by our algorithm using both full-field mode and ptychography mode. The result shows that the full-field reconstruction “throws away” the diffusive halo around the molecule (47), which, on the other hand, is correctly restored by ptychography reconstruction. Furthermore, we found that the reconstruction qualities of full-field and ptychography modes respond differently to changes in reconstruction hyperparameters (e.g., α_δ , α_β , γ , and step size). For the cone phantom, the ptychography mode is more robust to parameter changes and can provide good results even when the sparsity constraints and the TV regularizer are totally removed. The full-field mode is, however, more sensitive to parameter settings. With the sparsity weight α_β set always 10 times smaller than α_δ , the value of α_δ is good within half an order of magnitude around the setting used in the paper (1.5×10^{-8}). A value that is too low led to artifacts emerging around the contained particles (in the region not covered by the initial finite support mask), and a value that is too high caused the shrinkage of the reconstructed object. Even without noise in the raw data, the TV regularizer still helps suppress unreal heterogeneity inside the object's features due to the systematic error of the algorithm, but a too-high setting renders the result to be blocky. Hence, redetermination of the parameters might be necessary for different types of samples especially when one uses the full-field mode. While one can use a progressive strategy (optimize one parameter with others fixed and then move to the next one) to search the hyperparameter space, there are also dedicated packages for parameter searching such as DeepHyper (48).

While the proposed method has been implemented for both full-field microscopy and for ptychography, a special note should be given to the former. In the absence of the oversampling constraint in ptychography, we explored the use of finite support constraint and sparsity constraint in 3D space, which would provide more insights to the iterative retrieval of a bounded 3D object by solving an underdetermined system. The algorithm, when combined with non-scanning high-resolution imaging techniques, can potentially become the launchpad for a high-throughput imaging pipeline for measuring thick samples with sub-100-nm resolution. One of these possible paths is to apply the algorithm to point-projection microscopy (33, 49). While far-field diffraction patterns suffer a loss of speckle contrast as one goes from fully coherent to decreasingly partially coherent illumination (with the best results obtained when the coherence width of the beam equals the size of the object array), with near-field wave propagation, one only needs to have the spatial coherence match the distance $\lambda z/(\Delta x)$ over which one has the ability to record near-field fringes. Thus, point-projection near-field imaging is able to make use of a greater fraction of partially coherent sources, such as today's synchrotron light sources. At the same time, if one does have full spatial coherence, then the separation of subregions of the object into separate experimental recordings (diffraction patterns from limited-size illumination spots) gives reconstructions with better fidelity even with a more relaxed imposition of regularizers. Moreover, application of the proposed method to detection methods beyond x-rays—for example, broadband radiation used for atmospheric transmission (50)—might also be exploited to better model the dynamic diffraction of waves propagating in complicated media over long distances.

CONCLUSION

We have developed and demonstrated a novel 3D reconstruction algorithm for objects beyond the DOF limit. The algorithm uses

multislice propagation as the forward model and retrieves the refractive index map of the object by minimizing a loss function containing the squared difference between the amplitudes of the forward-propagated wavefront and the measured signals at all viewing angles. We implemented the algorithm for both full-field and ptychography imaging and compared them in terms of computation wall time and reconstruction fidelity. Investigation on the full-field version allowed us to explore the constraint requirements for reconstructing of bounded 3D objects with nonscanning imaging techniques, where sparsity and finite-support constraints are used to ensure a successful reconstruction. Another novelty of our method lies in the use of AD, which not only replaces the laborious manual differentiation involved in numerical optimization but also makes the computational model extremely adaptable and flexible. Numerical studies of the algorithm using simulated objects indicate that the proposed method is capable of recovering a thick object with high spatial resolution and good accuracy. Further validation of the approach with experimental data will be our subsequent step, which would examine the capability of the method in handling noise and help us identify practical challenges such as probe alignment (51). The ultimate goal is to combine the algorithm with high-resolution imaging techniques, which, for full-field imaging, could be point-projection microscopy. The combination of the two could potentially lead to the development of a realistic approach for imaging thick samples at high resolution.

MATERIALS AND METHODS

The forward model

We use the multislice method (7) to calculate the wave exiting the object, since it incorporates multiple scattering while also accounting for effects such as waveguide phenomena (52) (with the sole limitation of ignoring backscattering, which is negligible for all cases except Bragg diffraction from perfect crystals or synthetic multilayers). As illustrated in Fig. 6, in multislice propagation, the object is divided into J slices along the beam axis.

The wavefield $\psi_{j,k}(x, y, z_j)$ from probe position k (for full field, $k = 0$ for the first and only probe position) entering the j th slice with thickness Δz is modulated by the slice to yield a wavefield $\psi'_{j,k}(x, y, z_j)$ of

$$\begin{aligned} \psi'_{j,k}(x, y, z_j) &= \psi_{j,k}(x, y, z_j) \exp \left[-\frac{2\pi\Delta z}{\lambda} i [1 - \delta(x, y, z_j) - i\beta(x, y, z_j)] \right] \\ &= \psi_{j,k}(x, y, z_j) \exp \left(-\frac{2\pi\Delta z}{\lambda} i \right) \exp(n_j) \end{aligned} \quad (9)$$

with

$$n_j = \frac{2\pi(\Delta z)}{\lambda} [i\delta(x, y, z_j) - \beta(x, y, z_j)] \quad (10)$$

The actual implementation of Eq. 9 usually drops the constant phase factor $\exp(-2\pi\Delta z/\lambda)$ as we are only interested in the intensity that the detector collects. We will denote the refractive index distribution of the entire object by vector \mathbf{n} in the following text.

The wavefront is then free space propagated to the next slice according to the Fresnel diffraction integral given by

$$\psi_{j+1,k}(x, y, z_{j+1}) = \psi'_{j,k}(x, y, z_j) * h_{\Delta z}(x, y) \quad (11)$$




Fig. 6. Illustration of multislice propagation.

where $*$ denotes the convolution operator and $h_{\Delta z}(x, y)$ is the Fresnel propagator given by

$$h_{\Delta z}(x, y) = \exp \left[-i \frac{\pi}{\lambda \Delta z} (x^2 + y^2) \right] \quad (12)$$

This process is repeated for all J slices until one obtains the exit wave leaving the object.

Here, the slice thickness Δz can be equal to the transverse pixel size Δx , or for computational speed, one can combine multiple slices from the 3D volume together, provided one satisfies the condition of

$$\Delta z \leq \frac{2\tilde{n}Q(\Delta x)^2}{\pi \lambda} \quad (13)$$

where \tilde{n} is the mean refractive index, Δx is the pixel size, and Q is the Klein-Cook parameter for which values of $Q \lesssim 1$ represent the case of plane grating rather than volume grating diffraction (53).

Multislice propagation is only one of several steps that must be combined in the forward model of tomography beyond the DOF limit. Tomography acquisition requires the illumination of the object from multiple rotation angles θ . Our approach is to rotate the object onto a constant wave propagation direction rather than to rotate the illumination; this is performed with a rotation operator \mathbf{R}_θ . After carrying out the sequence of J multislice propagation steps through the rotated object, we then need to apply the operator \mathbf{P}_d to take the exit wave $\boldsymbol{\psi}_{J,k}$ from the object to the plane of the detector, either using free space propagation as described by $\boldsymbol{\psi}_{J,k} * h(x, y, d)$ for near-field propagation of distance d , or a simple Fourier transform for far-field propagation (in the Fraunhofer approximation). This leads us to a combined forward operation of

$$f(\mathbf{n}, \theta, k, \Delta z, d) = \mathbf{P}_d \mathbf{M}_{n,\theta,\Delta z} \boldsymbol{\psi}_{0,k} \quad (14)$$

In Eq. 14, $\mathbf{M}_{n,\theta,\Delta z}$ is the multislice propagation operator (Eqs. 11 and 12), which is a function of the 3D object and the incident probe. $\mathbf{M}_{n,\theta,\Delta z}$ describes the exit wave that leaves the depth-extended specimen and is equivalent to the recursive equations in Eqs. 9 and 11. It can be compactly written as

$$\mathbf{M}_{n,\theta,\Delta z} = \prod_j \mathbf{P}_{\Delta z} \mathbf{A}_{n,\theta,j} \quad (15)$$

with

$$\mathbf{A}_{n,\theta,j} = \exp[\text{diag}(\mathbf{S}_j \mathbf{R}_\theta \mathbf{n})] \quad (16)$$

where \mathbf{S}_j is a matrix that samples the j th slice of column vector $\mathbf{R}_\theta \mathbf{n}$. If the total number of voxels of the object function is N_v and the number of pixels in the detector is N_p , then \mathbf{n} is a $N_v \times 1$ column vector, \mathbf{R}_θ is a $N_v \times N_v$ square matrix, and \mathbf{S}_j is in the shape of $N_p \times N_v$, so that it yields an $N_p \times 1$ column vector, which is of the same size as the wavefront $\psi_{0,k}$. Multiplying diagonal matrix $\mathbf{A}_{n,\theta,j}$ with the wavefront vector $\psi_{0,k}$ is exactly the wavefront modulation given in Eq. 9. This forward model is used in the loss function of Eq. 4; the latter also includes sparsity and TV regularizers, and the minimization of loss function (Eq. 5) also involves the use of a finite support constraint, which will be explained in the next subsection.

Constrained loss function minimization

It has been pointed out for conventional 2D coherent diffraction imaging that since only magnitude information is available in the detected far-field diffraction pattern, a reconstructable object should be spatially isolated, with prior knowledge about the geometry of the object incorporated into the reconstruction through zeroing out pixels out of the object boundaries. This is known as a “finite support” constraint (54). Ptychography does not require a finite support constraint applied in object space because the bounded probe itself is already a form of finite support constraint, and furthermore, the overlap between adjacent probe positions supplies sufficient information to solve for all object unknowns.

In our work where a 3D object is retrieved, the same criteria are followed. As shown in Results, the ptychography reconstruction of the object using our algorithm does not need prior knowledge about the spatial extent of the sample. However, a finite support constraint was found to be necessary for the full-field case. The initial finite support mask is determined following the procedures below. First, single-distance near-field phase retrieval (45) is applied to all projection images to obtain a first guess of the weak phase contrast projection through the object. We then use the tomographic set of these projections to obtain a rough guess of the object support using the standard FBP tomographic reconstruction algorithm. The reconstructed volume is then Gaussian filtered to remove noise and local discontinuities. A Boolean mask is subsequently obtained by thresholding the filtered object, which yields a support mask denoted by set Θ . During the iterative reconstruction process, the finite support is contracted to exclude low-value pixels for every epoch, a technique known as shrink wrap in conventional CDI (coherent diffraction imaging) processing (55). The successful retrieval of \mathbf{n} requires the simultaneous pixel-wise update of it, guided by $\nabla_n L$, which is the gradient of loss function L in the parameter space of \mathbf{n} .

We use TensorFlow (56), a deep learning package first initiated by Google but now available as an open-source toolkit, for carrying out our AD reconstruction. It provides a user-friendly Python application programming interface and the ability to write a reconstruction code of relative simplicity and with easy implementation on a variety of computing platforms. The AD algorithm uses the so-called “back-propagation” method to derive the partial derivatives in a semianalytical fashion (30). Here, the loss function L is first evaluated in the forward direction using Eq. 4, during which the intermediate variables produced by every algebraic operation are computed and stored. After that, the algorithm calculates the deriv-

ative of L with regard to the intermediate variables immediately before L using the values saved in memory. This is repeated back through the entire computation model, and the gradient of L with regard to \mathbf{n} , $\nabla_n L$, is then found on the basis of the chain rule of differentiation. Compared with symbolic differentiation that attempts to acquire the closed-form expression of $\nabla_n L$ before doing any numerical calculation, AD is free from the problem of expression swell when the forward model is complicated. On the other hand, AD is also more accurate than the finite difference method, which approximates $\nabla_n L = (\partial f / \partial x_1, \dots, \partial f / \partial x_n)$ with $\partial f / \partial x_i \approx [f(\mathbf{n} + h\mathbf{e}_i) - f(\mathbf{n})] / h$ for a small h (30). We use the well-established optimization algorithm Adam to update \mathbf{n} (57). A brief description of the algorithm is provided in the Supplementary Materials.

Computational performance enhancements

Similar to conventional tomography, the dataset acquired for reconstruction would generally involve a large number of rotation angles N_θ [although in multislice reconstruction methods one can reduce N_θ from below the number one would have expected from the Crowther criterion (42)]. The large value of N_θ can lead to the use of considerable computation power in the iterative update of \mathbf{n} . To reduce this, we note that the first term in Eq. 4 is essentially an expectation value of error per pixel, and it can be adequately approximated by calculating the error over a subset of N_θ . This technique, known as “minibatching” [or “ordered subsets” in the tomography literature (58)], can speed up the convergence of the algorithm by several times. For each minibatch, the subset of N_θ to be processed is chosen randomly without replacement, so that the entire collection of N_θ will be fully gone through after a certain number of minibatches are completed. We hereafter refer to this process as an “epoch,” using terminology drawn from the machine learning community.

In an actual experiment, the presence of noise typically induces uncertainty in the “sub-loss function” obtained from each minibatch. In this case, a true global minimum is ill-defined, which causes \mathbf{n} to dangle at the late stage of the optimization and, thus, prevents a stable convergence. For this reason, users of TensorFlow have the option to aggregate the gradients calculated from several minibatches and apply them to the optimizer after the completion of all of these minibatches. The larger sample amount for gradient calculation reduces the statistical fluctuation induced by noise and guarantees a more stable solution.

A multiscale technique is also used in this work to further improve reconstruction speed and accuracy. While the algorithm requires that \mathbf{n} has the same number of lateral pixels as the measured data, instead of directly reconstructing the object with the same pixel size as the acquired projections, both \mathbf{n} and $\mathbf{y}_{\theta,k}$ are downsampled by a factor of 2^m , where m is an integer so that the lateral dimension of \mathbf{n} is not larger than 64 pixels. The voxel size of the object is, thus, accordingly enlarged by 2^m times. A first pass reconstruction of \mathbf{n} is therefore computed rapidly. The result is upsampled by a factor of 2 and then used as the initial guess for the next pass, where the scale of the object and projections are doubled. This process is repeated until the object is reconstructed with the acquired voxel size. By initializing \mathbf{n} with the result from a coarse pass for a higher-resolution pass, the optimization begins at a location closer to the global minimum in the parameter space of L , so that fewer iterations are required to converge. This may also reduce the chance for the optimizer to get trapped in a remote local minimum.

Parallelization

In view of the huge number of unknowns (3.4×10^7 for a 256^3 object because of the presence of both the δ and β parts of the refractive index) to be solved in our algorithm, parallelized computation is necessary to guarantee a reasonable computational wall time (within a few hours for a 256^3 object). We use a TensorFlow add-on called Horovod (59) to implement distributed parallelization using the message passing interface (MPI) standard for inter-rank (or inter-process) communication, rather than the TCP/IP protocol used by native TensorFlow (MPI is faster on tightly bound high-performance computer clusters). Each thread (or worker) initializes and keeps its own object function and processes a minibatch simultaneously with other threads. When a rank finishes its minibatch, it waits for other threads to finish theirs, after which the gradients obtained by all threads are averaged. The averaged gradient is then used to update the object functions in all threads. Since the volume of samples used for gradient calculation is effectively enlarged by a factor that equals the number of threads n_{threads} , this means that the actual learning rate (57) for the Adam optimizer is multiplied by n_{threads} .

The code used in this work has been made publicly available on GitHub in the repository named “Adorym.”

Reference reconstructions

To compare the outcomes of the proposed algorithm with methods that are conventionally used for phase retrieval, the full-field data demonstrated in this work are also processed and reconstructed by first performing an iterative 2D phase retrieval method termed ER [widely used in coherent diffraction imaging (54)] for every projection image. The workflow of ER can be summarized as follows:

1) Propagate the initial guess of the exit wavefront ψ_0 to the detector plane as

$$\psi'_0 = \mathbf{P}_d \psi_0$$

2) Replace the magnitude of the wavefront with the modulus of the measured intensity I_θ as

$$\psi''_0 = \frac{\psi'_0}{|\psi'_0|} \sqrt{I_\theta}$$

3) Backpropagate the wavefront to the exiting plane as

$$\psi'''_0 = \mathbf{P}_{-d} \psi''_0$$

4) Mask out the pixels of the wavefront that do not belong to the finite support Θ_2 by doing

$$\psi_1 = \begin{cases} \psi'''_0(\mathbf{r}), & \mathbf{r} \in \Theta_2 \\ 0, & \mathbf{r} \notin \Theta_2 \end{cases}$$

5) The above processes are then repeated until the mean square error between the calculated intensity and the measured intensity converges.

The FBP tomographic reconstruction method is then applied to phase-retrieved images to obtain a 3D reconstruction result. This ER + FBP approach will be subsequently referred to as “pure projection tomography.”

Phantom object design

We carried out computational experiments on a cone object created using the optical simulation package XDesign (60). To exceed the DOF limit within a moderate array size, we chose to create a (256^3)

voxel grid with 1-nm voxel size and to use 5-keV x-ray beam energy. As a result, the x-ray wavelength was 0.248 nm, and the DOF given by Eq. 1 was 21.8 nm. (Present-day x-ray microscopes achieve a spatial resolution of more typically 15 to 30 nm as noted in the introduction, but 1 nm represents a goal for the future). Therefore, the reconstruction grid was almost 12 times larger than the DOF, so the object significantly exceeds the DOF limit. Within this grid, a hollow cone of silicon was computationally created so as to resemble a thin-walled capillary heated and then pulled. The tube has a top diameter of 80 nm and a bottom diameter of 200 nm, so that neither end fits within the DOF limit. To examine the algorithm’s capability to restore fine details, we also placed 50 TiO₂ nanospheres with radii ranging from 2 to 4 nm on the outer wall of the tube, as well as 10 larger spheres (5 to 13 nm in radius) inside the tube. The refractive indices of both materials were generated using tabulated data (27). In addition, we added spherical bubbles or “grains,” whose refractive indices fluctuate within 30% of the original material, into the cone’s body as a means to test the ability of the algorithm to retrieve internal structure. These grains also served as marks in assessing the influence of photon noise on ptychography results as will be discussed below. The phase-shifting part δ of the x-ray refractive index for this computationally created object is shown in the top row of Fig. 2.

SUPPLEMENTARY MATERIALS

Supplementary material for this article is available at <http://advances.sciencemag.org/cgi/content/full/6/13/eaay3700/DC1>

Section S1. Linear algebraic notation of the forward model

Section S2. Studies on a protein sample

Section S3. FSC plots for all reconstruction results of the cone phantom

Section S4. Absolute error on the vertical cross section shown in Fig. 2

Fig. S1. The nucleosome assembly protein (NAP) object sectioned from the two xz planes (columns 1 and 2) marked by yellow dashed lines in the xy cross section (column 3).

Fig. S2. Line profiles across the phase-shifting part of the refractive index δ for the NAP protein sample.

Fig. S3. FSC plots for all reconstruction results of the cone object, including 360° full field, 180° full field, 360° ptychography, and ER + FBP.

Fig. S4. Absolute error on the vertical cross section shown in Fig. 2.

References (62, 63)

REFERENCES AND NOTES

- W. Chao, P. Fischer, T. Tyliczszak, S. Rekawa, Real space soft x-ray imaging at 10 nm spatial resolution. *Opt. Express* **20**, 9777–9783 (2012).
- S. Bajt, M. Prasciolu, H. Fleckenstein, M. Domaracký, H. N. Chapman, A. J. Morgan, O. Yefanov, M. Messerschmidt, Y. Du, K. T. Murray, V. Mariani, M. Kuhn, S. Aplin, K. Pande, P. Villanueva-Perez, K. Stachnik, J. P. Chen, A. Andrejczuk, A. Meents, A. Burkhardt, D. Pennicard, X. Huang, H. Yan, E. Nazaretski, Y. S. Chu, C. E. Hamm, X-ray focusing with efficient high-NA multilayer Laue lenses. *Light Sci. Appl.* **7**, 17162 (2018).
- M. Born, E. Wolf, *Principles of Optics* (Cambridge Univ. Press, Cambridge, ed. 7, 1999).
- Y. Wang, C. Jacobsen, J. Maser, A. Osanna, Soft x-ray microscopy with a cryo scanning transmission x-ray microscope: II. Tomography. *J. Microsc.* **197**, 80–93 (2000).
- M. Eriksson, J. F. van der Veen, C. Quitmann, Diffraction-limited storage rings – A window to the science of tomorrow. *J. Synchrotron Radiat.* **21**, (Pt 5) 837–842 (2014).
- L. Melchior, T. Salditt, Finite difference methods for stationary and time-dependent x-ray propagation. *Opt. Express* **25**, 32090–32109 (2017).
- J. M. Cowley, A. F. Moodie, The scattering of electrons by atoms and crystals. I. A new theoretical approach. *Acta Crystallogr.* **10**, 609–619 (1957).
- A. J. Devaney, A filtered backpropagation algorithm for diffraction tomography. *Ultrason. Imaging* **4**, 336–350 (1982).
- A. M. Maiden, M. J. Humphry, J. M. Rodenburg, Ptychographic transmission microscopy in three dimensions using a multi-slice approach. *J. Opt. Soc. Am. A* **29**, 1606 (2012).
- E. H. R. Tsai, I. Usov, A. Diaz, A. Menzel, M. Guizar-Sicairos, X-ray ptychography with extended depth of field. *Opt. Express* **24**, 29089–29108 (2016).
- M. A. Gilles, Y. S. G. Nashed, M. Du, C. Jacobsen, S. M. Wild, 3D x-ray imaging of continuous objects beyond the depth of focus limit. *Optica* **5**, 1078–1086 (2018).

12. W. Van den Broek, C. T. Koch, Method for retrieval of the three-dimensional object potential by inversion of dynamical electron scattering. *Phys. Rev. Lett.* **109**, 245502 (2012).
13. Y. S. G. Nashed, T. Peterka, J. Deng, C. Jacobsen, Distributed automatic differentiation for ptychography. *Procedia Comput. Sci.* **108**, 404–414 (2017).
14. S. Kandel, S. Maddali, M. Allain, S. O. Hruszkewycz, C. Jacobsen, Y. S. G. Nashed, Using automatic differentiation as a general framework for ptychographic reconstruction. *Opt. Express* **27**, 18653–18672 (2019).
15. A. Burvall, U. Lundström, P. A. C. Takman, D. H. Larsson, H. M. Hertz, Phase retrieval in x-ray phase-contrast imaging suitable for tomography. *Opt. Express* **19**, 10359–10376 (2011).
16. E. Wolf, Three-dimensional structure determination of semi-transparent objects from holographic data. *Optics Commun.* **1**, 153–156 (1969).
17. O. Haeberlé, K. Belkebir, H. Giovaninni, A. Sentenac, Tomographic diffractive microscopy: Basics, techniques and perspectives. *J. Mod. Opt.* **57**, 686–699 (2010).
18. H. N. Chapman, A. Barty, S. Marchesini, A. Noy, S. P. Hau-Riege, C. Cui, M. R. Howells, R. Rosen, H. He, J. C. H. Spence, U. Weierstall, T. Beetz, C. Jacobsen, D. Shapiro, High-resolution *ab initio* three-dimensional x-ray diffraction microscopy. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **23**, 1179 (2006).
19. H. M. L. Faulkner, J. M. Rodenburg, Movable aperture lensless transmission microscopy: A novel phase retrieval algorithm. *Phys. Rev. Lett.* **93**, 023903 (2004).
20. T. M. Godden, R. Suman, M. J. Humphry, J. M. Rodenburg, A. M. Maiden, Ptychographic microscope for three-dimensional imaging. *Opt. Express* **22**, 12513–12523 (2014).
21. A. Suzuki, S. Furutaku, K. Shimomura, K. Yamauchi, Y. Kohmura, T. Ishikawa, Y. Takahashi, High-resolution multislice x-ray ptychography of extended thick objects. *Phys. Rev. Lett.* **112**, 053903 (2014).
22. H. Öztürk, H. Yan, Y. He, M. Ge, Z. Dong, M. Lin, E. Nazaretski, I. K. Robinson, Y. S. Chu, X. Huang, Multi-slice ptychography with large numerical aperture multilayer Laue lenses. *Optica* **5**, 601 (2018).
23. S. Gao, P. Wang, F. Zhang, G. T. Martinez, P. D. Nellist, X. Pan, A. I. Kirkland, Electron ptychographic microscopy for three-dimensional imaging. *Nat. Commun.* **8**, 163 (2017).
24. P. Li, A. Maiden, Multi-slice ptychographic tomography. *Sci. Rep.* **8**, 2049 (2018).
25. G. Schmahl, D. Rudolph, Proposal for a phase contrast x-ray microscope, in *X-ray Microscopy: Instrumentation and Biological Applications*, P. C. Cheng, G. J. Jan, Eds. (Springer-Verlag, 1987), pp. 231–238.
26. M. Du, C. Jacobsen, Relative merits and limiting factors for x-ray and electron microscopy of thick, hydrated organic materials. *Ultramicroscopy* **184**, 293–309 (2018).
27. B. L. Henke, E. M. Gullikson, J. C. Davis, X-ray interactions: Photoabsorption, scattering, transmission, and reflection at $E=50\text{--}30,000$ eV, $Z=1\text{--}92$. *At. Data Nucl. Data Tables* **54**, 181–342 (1993).
28. U. S. Kamilov, I. N. Papadopoulos, M. H. Shoreh, A. Goy, C. Vonesch, M. Unser, D. Psaltis, Optical tomographic image reconstruction based on beam propagation and sparse regularization. *IEEE Trans. Comput. Imag.* **2**, 59–70 (2016).
29. U. S. Kamilov, I. N. Papadopoulos, M. H. Shoreh, A. Goy, C. Vonesch, M. Unser, D. Psaltis, Learning approach to optical tomography. *Optica* **2**, 517 (2015).
30. L. B. Rall, *Automatic Differentiation: Techniques and Applications* (Springer Berlin Heidelberg, 1981), vol. 120.
31. A. S. Jurling, J. R. Fienup, Applications of algorithmic differentiation to phase retrieval algorithms. *J. Opt. Soc. Am. A* **31**, 1348 (2014).
32. J. Rodenburg, A. C. Hurst, A. G. Cullis, B. R. Dobson, F. Pfeiffer, O. Bunk, C. David, K. Jefimovs, I. Johnson, Hard-x-ray lensless imaging of extended objects. *Phys. Rev. Lett.* **98**, 034801 (2007).
33. V. E. Cosslett, W. C. Nixon, X-ray shadow microscope. *Nature* **168**, 24–25 (1951).
34. P. Cloetens, W. Ludwig, J. Baruchel, D. van Dyck, J. van Landuyt, J. P. Guigay, M. Schlenker, Holotomography: Quantitative phase tomography with micrometer resolution using hard synchrotron radiation x-rays. *Appl. Phys. Lett.* **75**, 2912–2914 (1999).
35. R. Mokso, P. Cloetens, E. Maire, W. Ludwig, J. Y. Buffiere, Nanoscale zoom tomography with hard x rays using Kirkpatrick-Baez optics. *Appl. Phys. Lett.* **90**, 144104 (2007).
36. M. Grasmair, F. Lenzen, Anisotropic total variation filtering. *Appl. Math. Opt.* **62**, 323–339 (2010).
37. M. van Heel, M. Schatz, Fourier shell correlation threshold criteria. *J. Struct. Biol.* **151**, 250–262 (2005).
38. X. Huang, H. Miao, J. Steinbrener, J. Nelson, D. Shapiro, A. Stewart, J. Turner, C. Jacobsen, Signal-to-noise and radiation exposure considerations in conventional and diffraction x-ray microscopy. *Opt. Express* **17**, 13541–13553 (2009).
39. J. Hagemann, T. Salditt, The fluence–resolution relationship in holographic and coherent diffractive imaging. *J. Appl. Crystallogr.* **50**, 531–538 (2017).
40. J. Miao, D. Sayre, H. N. Chapman, Phase retrieval from the magnitude of the Fourier transforms of non-periodic objects. *J. Opt. Soc. Am. A* **15**, 1662 (1998).
41. R. Hegerl, W. Hoppe, Influence of electron noise on three-dimensional image reconstruction. *Z. Naturforsch. A* **31**, 1717–1721 (1976).
42. C. Jacobsen, Relaxation of the Crowther criterion in multislice tomography. *Opt. Lett.* **43**, 4811–4814 (2018).
43. R. A. Crowther, D. J. DeRosier, A. Klug, The reconstruction of a three-dimensional structure from projections and its application to electron microscopy. *Proc. R. Soc. Lond. A* **317**, 319 (1970).
44. S. Aslan, V. Nikitin, D. J. Ching, T. Bicer, S. Leyffer, D. Gürsoy, Joint ptycho-tomography reconstruction through alternating direction method of multipliers. *Opt. Express* **27**, 9128–9143 (2019).
45. D. Paganin, S. C. Mayo, T. E. Gureyev, P. R. Miller, S. W. Wilkins, Simultaneous phase and amplitude extraction from a single defocused image of a homogeneous object. *J. Microsc.* **206**, 33–40 (2002).
46. F. Xu, L. Helfen, T. Baumbach, H. Suhonen, Comparison of image quality in computed laminography and tomography. *Opt. Express* **20**, 794–806 (2012).
47. M. E. Kandel, M. Fanous, C. Best-Popescu, G. Popescu, Real-time halo correction in phase contrast imaging. *Biomed. Opt. Express* **9**, 623–635 (2018).
48. P. Balaprakash, M. Salim, T. Uram, V. Vishwanath, S. Wild, DeepHyper: Asynchronous hyperparameter search for deep neural networks, in *2018 IEEE 25th International Conference on High Performance Computing (HiPC)* (IEEE, 2018), pp. 42–51.
49. K. M. Sowa, B. R. Jany, P. Korecki, Multipoint-projection x-ray microscopy. *Optica* **5**, 577 (2018).
50. J. A. Rubio, A. M. Belmonte, A. Comerón, Numerical simulation of long-path spherical wave propagation in three-dimensional random media. *Opt. Eng.* **38**, 1462 (1999).
51. M. Guizar-Sicairos, J. R. Fienup, Phase retrieval with transverse translation diversity: A nonlinear optimization approach. *Opt. Express* **16**, 7264–7278 (2008).
52. K. Li, M. Wojcik, C. Jacobsen, Multislice does it all—Calculating the performance of nanofocusing x-ray optics. *Opt. Express* **25**, 1831–1846 (2017).
53. W. R. Klein, B. D. Cook, Unified approach to ultrasonic light diffraction. *IEEE Trans. Sonics Ultrason.* **14**, 123–134 (1967).
54. J. R. Fienup, Reconstruction of an object from the modulus of its Fourier transform. *Opt. Lett.* **3**, 27–29 (1978).
55. S. Marchesini, H. He, H. N. Chapman, S. P. Hau-Riege, A. Noy, M. R. Howells, U. Weierstall, J. C. H. Spence, X-ray image reconstruction from a diffraction pattern alone. *Phys. Rev. B* **68**, 140101 (2003).
56. M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mane, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viegas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, X. Zheng, Tensorflow: Large-scale machine learning on heterogeneous distributed systems. arXiv:1603.04467 [cs.LG] (14 March 2016).
57. D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, in *3rd International Conference on Learning Representations (ICLR)*, 2015, pp. 1–15.
58. H. M. Hudson, R. S. Larkin, Accelerated image reconstruction using ordered subsets of projection data. *IEEE Trans. Med. Imaging* **13**, 601–609 (1994).
59. A. Sergeev, M. D. Balso, Horovod: Fast and easy distributed deep learning in TensorFlow. arXiv:1802.05799 [cs.LG] (15 February 2018).
60. D. J. Ching, D. Gürsoy, XDesign: An open-source software package for designing x-ray imaging phantoms and experiments. *J. Synchrotron Radiat.* **24**, 537–544 (2017).
61. H. Peng, Z. Ruan, F. Long, J. H. Simpson, E. W. Myers, V3D enables real-time 3D visualization and quantitative analysis of large-scale biological image data sets. *Nat. Biotechnol.* **28**, 348–353 (2010).
62. M. P. Scheffer, L. Gonzalez-Gonzalez, A. Seybert, M. Ratera, M. Kunz, J. M. Valpuesta, I. Fita, E. Querol, J. Piñol, J. Martin-Benito, A. S. Frangakis, Structural characterization of the NAP; the major adhesion complex of the human pathogen *Mycoplasma genitalium*. *Mol. Microbiol.* **105**, 869–879 (2017).
63. R. A. London, M. D. Rosen, J. E. Trebes, Wavelength choice for soft x-ray laser holography of biological samples. *Appl. Optics* **28**, 3397–3404 (1989).

Acknowledgments: This research used resources of the Advanced Photon Source and the Argonne Leadership Computing Facility, which are U.S. Department of Energy (DOE) Office of Science User Facilities operated for the DOE Office of Science by Argonne National Laboratory under contract no. DE-AC02-06CH11357. We would like to thank T. Childers and C. J. Adams for support in setting up and testing Adorym on ALCF supercomputers. We also thank M. Gilles and S. Wild for the helpful discussions. **Funding:** We acknowledge the National Institute of Mental Health, NIH, for support under grant R01 MH115265. **Author contributions:** M.D. did the coding and testing. Y.S.G.N. contributed to the mathematical derivation and formulation. S.K. proposed and contributed to the improvement of code vectorization. C.J. proposed the initial idea, provided computing

resources, and participated in the data analysis. M.D. and C.J. wrote the manuscript with inputs from all authors. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** Code used to generate test samples and the main code of Adorym are available at <https://github.com/mdw771/adorym>, with the raw data for the cone test object available at <https://tinyurl.com/s5cdrn2>. All other data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials. Additional data related to this paper may be requested from the authors.

Submitted 12 June 2019
Accepted 3 January 2020
Published 27 March 2020
10.1126/sciadv.aay3700

Citation: M. Du, Y. S. G. Nashed, S. Kandel, D. Gürsoy, C. Jacobsen, Three dimensions, two microscopes, one code: Automatic differentiation for x-ray nanotomography beyond the depth of focus limit. *Sci. Adv.* **6**, eaay3700 (2020).