

High-resolution *ab initio* three-dimensional x-ray diffraction microscopy

Henry N. Chapman, Anton Barty, Stefano Marchesini, Aleksandr Noy, and Stefan P. Hau-Riege

University of California, Lawrence Livermore National Laboratory, 7000 East Avenue, Livermore, California 94550

Congwu Cui, Malcolm R. Howells, and Rachel Rosen

Advanced Light Source, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, California 94720

Haifeng He, John C. H. Spence, and Uwe Weierstall

Department of Physics and Astronomy, Arizona State University, Tempe, Arizona 85287-1504

Tobias Beetz, Chris Jacobsen, and David Shapiro

Department of Physics and Astronomy, Stony Brook University, Stony Brook, New York 11794-3800

Received August 26, 2005; revised October 14, 2005; accepted October 18, 2005; posted October 25, 2005 (Doc. ID 64418)

Coherent x-ray diffraction microscopy is a method of imaging nonperiodic isolated objects at resolutions limited, in principle, by only the wavelength and largest scattering angles recorded. We demonstrate x-ray diffraction imaging with high resolution in all three dimensions, as determined by a quantitative analysis of the reconstructed volume images. These images are retrieved from the three-dimensional diffraction data using no *a priori* knowledge about the shape or composition of the object, which has never before been demonstrated on a nonperiodic object. We also construct two-dimensional images of thick objects with greatly increased depth of focus (without loss of transverse spatial resolution). These methods can be used to image biological and materials science samples at high resolution with x-ray undulator radiation and establishes the techniques to be used in atomic-resolution ultrafast imaging at x-ray free-electron laser sources. © 2006 Optical Society of America

OCIS codes: 340.7460, 110.1650, 110.6880, 100.5070, 100.6890, 070.2590, 180.6900.

1. INTRODUCTION

In many fields of science the ability to visualize the three-dimensional (3D) organization of component parts is proving crucial to our understanding of the mechanisms involved in atomic and molecular processes. This is occurring in fields as diverse as whole-cell imaging in biology, the study of the minimum energy pathway for crack propagation in brittle solids, and the internal structure of the new labyrinthine mesoporous structures developed by inorganic chemists for a wide range of applications.

The field of coherent x-ray diffraction imaging (also known as diffraction microscopy) is expected to make a significant contribution to this effort. In this method, first put forward and developed by Sayre *et al.*,^{1,2} an image is reconstructed from measurements of the far-field scattered intensity of an isolated and nonperiodic object. The resolution of this form of microscopy is limited by only the wavelength and the largest scattering angle recorded. Hence this method is being pursued as a method for high-resolution x-ray microscopy without the technological limitations of manufacturing high-resolution optical elements.^{3–8} The penetrating nature of x rays allows imaging of objects much thicker than those that can be examined in a transmission electron microscope (e.g., 10 μm), at resolutions much better than those of visible microscopes. Preliminary studies of radiation damage

suggest that 3D resolutions of about 10 nm should be achievable on frozen hydrated biological material.⁹ The method is also being pursued in order to push x-ray imaging to its resolution limits, namely ultrafast near-atomic-resolution imaging of macromolecules at x-ray free-electron laser sources^{10,11} and of laser-aligned molecules,¹² which will enable structure determination without the need for crystallizing material.

High-resolution imaging of thick objects can be attained only in the context of 3D measurement and reconstruction. In most cases, other than surface studies or imaging of man-made objects, the analysis of the structure can be properly interpreted only in three dimensions. Unless the object itself is a slice of material that is thinner than the depth of focus of a two-dimensional (2D) image, artifact-free structural analysis can be carried out only with knowledge of the surrounding material, or by applying imaging modalities whereby depth information is not strongly transferred to the image (such as confocal imaging). At resolution lengths very much larger than the wavelength, thickness effects do not play a significant role, since, at the correspondingly low numerical aperture (NA), the depth of focus may be much larger than the size of the object. This is certainly the case as one satisfies the projection approximation of high-energy x-ray tomography, where the depth of focus approaches infinity. Tomog-

raphic imaging in this mode is limited by detector pixel size or, if a diverging beam is used, by Fresnel diffraction effects. However, as one moves to higher resolution, the depth of focus decreases dramatically, with the ratio of transverse resolution length to longitudinal depth of focus given by the NA. For the classes of high-resolution microscopy and structure determination applications in which we are interested, that imaging can properly be carried out only in three dimensions.

Coherent 3D x-ray diffraction imaging is especially suited to high-resolution x-ray microscopy. With a collimated beam incident on an object, the far-field diffraction pattern (recorded on a flat CCD) represents diffracted intensities, which, in reciprocal space, sample points on the momentum- and energy-conserving Ewald sphere. By rotating the sample about an axis normal to the beam, this sphere, which passes through the origin, sweeps through almost all of the reciprocal space volume of continuous diffuse scattering from our nonperiodic object. In this way we collect the 3D distribution of scattered intensity in reciprocal space, which is phased by using the 3D implementations of iterative methods, as discussed below. Once the phases of the diffraction intensities in the diffraction volume have been determined, the 3D Fourier transform of the object is known and the 3D image can be obtained simply by an inverse Fourier transform. As will be demonstrated in this paper, such data sets can be used for artifact-free analysis of structures. This is also the case for crystallography but is not generally the case for imaging with a lens. Partially coherent tomographic imaging techniques, such as tomography in the scanning transmission x-ray microscope¹³ or transmission x-ray microscope,^{14,15} lead to a complicated transfer of object spatial frequencies into the measured image, and there is no longer a simple one-to-one mapping of a measurement on a detector pixel, for example, to a spatial frequency of the object. For some classes of object, such as pure phase or amplitude objects, it may be possible to deconvolve the 3D transfer function, but this is not generally assured.¹⁶ As with coherent diffraction imaging and crystallography, coherent imaging with a lens also leads to a direct mapping of spatial frequencies in the object to spatial frequencies of the image. Again, a tomographic reconstruction from coherent 2D images can be easily achieved for pure amplitude or phase objects but would otherwise require knowing the phase and amplitude of the image in order to transform into 3D reciprocal space. Coherent diffraction imaging essentially attempts to emulate coherent lens-based imaging, using a computer algorithm in place of a lens. The advantage, for tomography of complex objects, is that the diffraction amplitudes are measured and the phases retrieved from the oversampling of those amplitudes, so that a direct 3D Fourier synthesis of the object can be achieved.

In this paper we perform an important demonstration of the feasibility of high-resolution diffraction microscopy required for biological and materials characterization, as well as single-molecule imaging. Significantly this is done without the use of detailed *a priori* information about the sample structure or low-resolution data obtained by other means. We also demonstrate that a full 3D reconstruction can be produced on a 1024^3 or larger data cube in a rea-

sonable amount of time using currently available computational hardware. Three significant recent developments have enabled us to perform full 3D image reconstructions with high resolution in all three dimensions. The commissioning of a new diffraction tomography apparatus¹⁷ by Stony Brook University at an undulator beamline¹⁸ of the Advanced Light Source (ALS) allows us to acquire diffraction patterns at over 100 orientations of an object, with short exposure times, over angular ranges of more than $\pm 70^\circ$. The Shrinkwrap phase-retrieval algorithm that we developed⁶ has proven to be extremely robust and effective in performing phase retrieval on diffraction data sets with missing data (e.g., due to a beam stop) or limited angles. The algorithm retrieves images from the measured diffraction intensities. It does not require additional information about the object, such as a low-resolution image, and can retrieve phases of general complex-valued objects. In direct methods of crystallography the term *ab initio* is used to mean "from the experimental data, without any supplemental prior information,"¹⁹ and we employ this term to describe our method. The third advance is the ability to perform 3D fast Fourier transforms (FFTs) on the large 1024^3 -element arrays of diffraction data that are assembled from our measurements. Although the familiar increase of computer processing power has brought gigaelement FFTs in reach of today's computers, it has been the development of computer clusters and specific software for distributed computation of FFTs that has made feasible the 3D implementation of the Shrinkwrap algorithm. In particular, we utilize the `dist_fft` software²⁰ on a 16-node cluster of dual-processor Apple G5 Xserves, giving us a performance of 8.6 s per 1024^3 -element single-precision complex FFT. We note that this computational advance should also benefit the field of diffraction tomography (Ref. 21, Sect. 13.2), in which both the phase and amplitude of the scattered field are measured, as is possible with scattered ultrasonic waves.

We present here experimental results of high-resolution 3D x-ray diffraction imaging of a well-characterized test object to demonstrate the practical application of these advances and quantitatively assess the technique. We show the first full 3D x-ray diffraction images that have been reconstructed without prior knowledge of the sample. We believe that these are the highest-resolution 3D x-ray images of noncrystalline objects ever achieved, with a demonstrable resolution volume of $10\text{ nm} \times 10\text{ nm} \times 40\text{ nm}$. We demonstrate that artifact-free 2D images can be created from the 3D diffraction data set of objects very much thicker than the depth of focus. In Section 2 we review diffraction imaging, the experimental requirements for 3D image reconstructions, and our computer implementation to perform the 3D phase retrieval and Fourier synthesis of the image. Our sample preparation and characterization techniques are discussed in Subsection 3.A, and our particular experimental setup and methods are described in Subsections 3.B and 3.C. Image reconstruction results are presented in Section 4. The 3D images are visualized as isosurface renderings, extremely large depth-of-focus projection images, maximum value projections, and tomographic slices through the object. We also compare artifact-free 2D pro-

jections of 3D data with reconstructions of individual 2D views and illustrate the artifacts present in single-view 2D images of thick objects. In Section 5 we quantitatively assess our 3D image resolution.

2. THREE-DIMENSIONAL COHERENT DIFFRACTION IMAGING

The incident x-ray wave field interacts with a 3D periodic or nonperiodic object through the scattering potential of the object, $o(\mathbf{x}) = r_e \rho(\mathbf{x})$, where $\rho(\mathbf{x})$ is the complex electron density and r_e is the classical electron radius. This object scattering function may be decomposed into a Fourier representation of 3D spatial frequencies \mathbf{u} , with complex amplitudes

$$O(\mathbf{u}) = \mathcal{F}\{o(\mathbf{x})\} \equiv \int o(\mathbf{x}) \exp(2\pi i \mathbf{u} \cdot \mathbf{x}) d\mathbf{x}, \quad (1)$$

in which each spatial frequency can be thought of as a volume grating. In the case of coherent diffraction imaging a plane wave with wave vector \mathbf{k}_{in} is incident on the object, and the intensity of the scattered field in the direction of the wave vector \mathbf{k}_{out} is measured on a 2D pixellated detector (e.g., a bare CCD) in the diffraction far field. This detector is typically centered on the forward direction but in principle could be oriented in any angle to the incident beam (see Fig. 1). For elastic scattering, only the volume gratings that satisfy Bragg's law will scatter, and the wave-vector transfer $\mathbf{q} = \mathbf{k}_{\text{out}} - \mathbf{k}_{\text{in}}$ will be equal to the grating spatial frequency; that is, $\mathbf{q} = \mathbf{u}$. Since the magnitudes $|\mathbf{k}_{\text{out}}|$ and $|\mathbf{k}_{\text{in}}|$ are constant and equal to $1/\lambda$, these spatial frequencies \mathbf{u} lie on the Ewald sphere of radius $1/\lambda$,^{22,23} where λ is the x-ray wavelength. This construction is equivalent to the condition that to scatter light by an angle 2θ from the forward direction (the z axis), the volume grating must be tilted by an angle θ from perpendicular to the forward direction (Bragg's law). With the convention used here we have $|\mathbf{q}| = q = (2/\lambda) \sin \theta$. The diffraction amplitudes in the direction \mathbf{k}_{out} are proportional to $O(\mathbf{q})$, and in diffraction imaging we measure the intensities, proportional to $|O(\mathbf{q})|^2$. In particular, in the Born approximation (which can be thought of in this context as single scattering), the number of photons per second measured in a CCD pixel, with solid angle Ω , is given by

$$I(\mathbf{q}; \Omega) = I_0 \Omega P |O(\mathbf{q})|^2, \quad (2)$$

where I_0 is the flux (photons per second per unit area) of the incident plane wave on the sample and P is the polar-

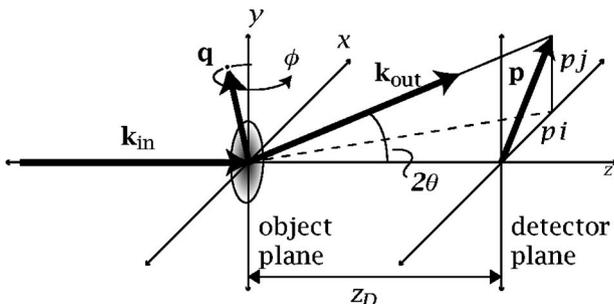


Fig. 1. Scattering geometry for coherent x-ray diffraction imaging. The sample is rotated about the y axis by an angle ϕ .

ization factor; $P = (1 + \cos^2 \psi)/2$ for unpolarized light, with $\psi = 2\theta$.²²

The complex scattering potential $o(\mathbf{x})$ that we aim to recover from measurements of $I(\mathbf{q})$ is related to the complex refractive index $n(\mathbf{x})$ of the object by (Ref. 21, Sect. 13.1)²⁴

$$o(\mathbf{x}) = r_e \rho(\mathbf{x}) = \frac{\pi}{\lambda^2} [1 - n^2(\mathbf{x})]. \quad (3)$$

In the soft x-ray region the complex refractive index is usually written in terms of the optical constants as $n(\mathbf{x}) = 1 - \delta(\mathbf{x}) - i\beta(\mathbf{x})$. For optical constants much less than unity, which is generally the case for soft x rays, Eq. (3) can then be well approximated by

$$o(\mathbf{x}) \approx \frac{2\pi}{\lambda^2} [\delta(\mathbf{x}) + i\beta(\mathbf{x})] = \frac{2\pi}{\lambda^2} \Delta n(\mathbf{x}). \quad (4)$$

The limit of validity of Eq. (2) under the Born approximation requires that $D|\Delta n(\mathbf{x})| < 2\pi\lambda C$, where D is the thickness of the object and $C \approx 0.2$.²⁵

A. Experimental Requirements

The recovery of the 3D image $o(\mathbf{x})$ from $O(\mathbf{u})$ requires the phases of $O(\mathbf{u})$ to be recovered and combined with the square root of the measured intensities. Both phase retrieval and image recovery place requirements on the 3D sampling of the diffraction intensities. Image recovery requires that the object be adequately sampled in real space to resolve the finest desired feature size over a given field of view. The requirements of the phase retrieval step are more demanding, in particular because our phase retrieval algorithm has the additional task of overcoming gaps and missing regions in the 3D sampled data by performing the interpolation tasks of recovering intensities that were blocked by a beam stop or that were missed due to a limited angular range of measurements. The 3D image recovery requires reconstructing the complex amplitudes $O(\mathbf{u})$ throughout a volume of reciprocal space. Since a single diffraction pattern is limited to frequencies $\mathbf{u} = \mathbf{q}$ on the Ewald sphere, diffraction data must be collected for various orientations of the sample.

In this work we perform phase retrieval and image recovery by full 3D Fourier synthesis, which requires interpolating the measured intensities from the Ewald sphere onto a uniform 3D Cartesian grid. In reciprocal space the grid has a width of N samples, spaced by Δu , and is centered at the zero spatial frequency. For N even, the spatial frequencies along each grid axis run from $-(N/2 - 1)\Delta u$ to $(N/2)\Delta u$. In real space we characterize the grid by a spacing Δx and a field width $w = N\Delta x$. Since $\Delta x \Delta u = 1/N$, we have the relationship $\Delta u = 1/w$; thus the largest spatial frequency component along a grid axis is given by $u_{x,\text{max}} = N\Delta u/2 = 1/(2\Delta x)$.

From Eqs. (1) and (2) the inverse Fourier transform of the intensity diffraction pattern is proportional to the autocorrelation function of the image that would be recovered when the phases are known:

$$i(\mathbf{x}) = \mathcal{F}^{-1}\{I(\mathbf{q})\} \propto o(\mathbf{x}) \otimes o^*(-\mathbf{x}), \quad (5)$$

where \otimes denotes the convolution operator. Consider an object of finite extent with a maximum width D along any axis of the real-space grid. The autocorrelation image $i(\mathbf{x})$ in that direction has a maximum width of $2D$, and hence the diffraction intensities are band limited. That is, the smallest grid spacing required to record all information present in the diffraction intensities is $\Delta u = 1/(2D)$, the Shannon critical sampling frequency, and the function can be recovered everywhere by a sinc-series expansion of the measured samples.²⁶ We define the sampling ratio s per dimension with $w = sD$ (Shannon sampling of intensities occurs for $s=2$). The oversampling of data relative to the sampling of $\Delta u = 1/D$ is what enables the phase retrieval techniques to be employed. In practice we may collect data on a finer grid than that strictly required to completely specify the diffraction intensities, as a way to increase detector dynamic range (sampling with $s > 2$ does not increase the information content of the measurement). However, it is possible to successfully retrieve phases from measurements with less than this factor of 2 in each dimension.²⁷ The minimum required sampling ratio has been shown by simulation to be $\sqrt{2}$ for 2D images and $\sqrt[3]{2}$ for 3D images.²⁸

The CCD must be placed so that it intersects a large enough range of scattering angles up to the desired spatial resolution. Usually the CCD is far enough from the sample to be in the diffraction far field, in which the angularly resolved diffraction pattern does not vary with propagation distance. For an object of width D the far field exists beyond distances of $z_F = 2D^2/\lambda$ from the object.²⁹ For a detector with pixels of width p placed a distance z_D from the object, we have, for small scattering angles, $\Delta q = p/(z_D\lambda)$. That is, to sample a field width of $w = sD$, the detector must be placed a distance of $z_D = spD/\lambda$. This will be in the far field if $z_D > z_F$, which can be satisfied if $D < sp/2$, or the condition that the sample must be smaller than the CCD pixel spacing when $s=2$. If the CCD is closer to the sample than z_F , then the sample and diffraction planes are related by a Fresnel, rather than a Fourier, transform, and the reconstruction algorithms must be appropriately modified.

Experimental requirements are placed on the transverse and longitudinal coherence of the incident beam. The transverse spatial coherence length of the incident beam must be at least as large as the entire field width $w = sD$.³⁰ The effect of partial coherence may be modeled as an incoherent source located some distance from the sample, whereby the diffraction intensity pattern is convolved with a demagnified intensity image of the source. In real space this convolution modulates the autocorrelation of the object with an envelope function that is proportional to the modulus of the transform of the source function. By the Van Cittert-Zernike theorem²¹ this envelope function is the mutual coherence of the source. The measured diffraction intensity is also convolved with the pixel response function of the detector, which modulates the autocorrelation image with an envelope proportional to the modulation transfer function of the detector. The spectral bandwidth $\Delta\lambda/\lambda$ of the incident light should be narrower than $2/N$, since we require $\Delta q/q = \Delta\lambda/(2\lambda)$ so

that the range of angles diffracted from a single spatial frequency by a range of wavelengths spreads by no more than half a detector pixel. This is equivalent to a minimum required longitudinal coherence length of $wq_{\max}\lambda = 2w \sin \theta$, which will be the maximum path length for light scattering by 2θ to the edge of the detector from points spaced transversely by w , or by the same angle from points spaced longitudinally by $w/(\tan 2\theta)$.

In our experiments we rotate the sample about an axis perpendicular to the incident beam direction to build up the 3D data set. At the highest spatial frequencies recorded, an angular increment of $\Delta\phi$ leads to a spacing between the Ewald sphere surfaces of $\Delta q = q_{\max}\Delta\phi$. That is, the Crowther resolution³¹ matches the critical sampling of the diffraction intensities ($s=2$) when

$$\Delta\phi = \Delta q/q_{\max} = \Delta x/D, \quad (6)$$

with the real-space sampling interval Δx . Note that this angular increment leads to a higher than necessary sampling at the lower spatial frequencies. For the examples in this paper we collected diffraction data with angular increments that are 2–4 times larger than those given by Eq. (6). In the process of phase retrieval we additionally recover both the amplitudes and phases of the missing data between the Ewald surfaces, including those in a large gap resulting from a limited range (usually $\pm 70^\circ$) of rotation angles, data blocked by a beam stop, and the missing “cone” of data resulting from rotating the sample about a single axis. This amplitude and phase retrieval of the missing data is essentially a superresolution technique and is achieved with the same real-space image constraints that we use for phase retrieval.³² Recovery of unmeasured data behind a beam stop has been demonstrated previously in 2D coherent x-ray diffraction imaging using this technique,^{6,33} and data missing due to a limited number of views have been recovered in the context of computed tomography using iterative algorithms³⁴ similar to those that we use (described in Subsection 2.C). Depending on how much amplitude information is missing, there will be a null space of images that are not constrained by the real- or reciprocal-space constraints^{32,35} and that may need to be regularized in the reconstruction.³⁶

B. Interpolation of the Diffraction Intensities

We interpolate the diffraction intensities measured on the CCD detector onto a regular Cartesian grid. The location of the CCD pixel indexed by integers (i, j) is given by the vector $\mathbf{p}_{i,j} = p(\hat{i}\mathbf{i} + \hat{j}\mathbf{j})$, relative to the pixel intersecting the direct beam, as shown in Fig. 1. We have then, with $\mathbf{k}_{\text{in}} = (1/\lambda)\hat{\mathbf{k}}$,

$$\mathbf{q}_{i,j} = \mathbf{k}_{\text{out}} - \mathbf{k}_{\text{in}} = \frac{1}{\lambda} \left(\frac{\mathbf{p}_{i,j} + z_D\hat{\mathbf{k}}}{\sqrt{p^2(i^2 + j^2) + z_D^2}} - \hat{\mathbf{k}} \right), \quad (7)$$

where z_D is the distance from the sample to the detector. Hence, for example,

$$q_{i,j}^x = \mathbf{q}_{i,j} \cdot \hat{\mathbf{i}} = \frac{1}{\lambda} \frac{pi}{\sqrt{p^2(i^2 + j^2) + z_D^2}}, \quad (8)$$

$$q_{i,j}^z = \mathbf{q}_{i,j} \cdot \hat{\mathbf{k}} = \frac{1}{\lambda} \left(\frac{z_D}{\sqrt{p^2(i^2 + j^2) + z_D^2}} - 1 \right). \quad (9)$$

In practice each diffraction pattern in our data set has a different rotation in the 3D Fourier space of the object, and the coordinate of each pixel in this space is given by

$$\mathbf{u}_{i,j,\phi} = R_\phi \mathbf{q}_{i,j}, \quad (10)$$

where R_ϕ is the 3D rotation matrix derived from the known object orientation. The coordinates $\mathbf{u}_{i,j,\phi}$ are then mapped onto a uniform Cartesian grid by nearest-neighbor sampling. Where more than one pixel from the set of intensity measurements contribute to a given voxel, the pixel values are averaged to determine the appropriate intensity value at that point. We find that for the data sets that we have collected, where the angular increment $\Delta\phi$ is at least four times larger than the condition specified by Eq. (6), nearest-neighbor interpolation gives slightly better reconstructions than does bilinear interpolation. At the high resolutions, where the reciprocal-space distance between data from adjacent object orientations is several pixels, the latter method leads to many interpolated values that are simply some fraction of their neighboring value. We find it better in this sparse case to have the intensity value concentrated at a single \mathbf{u} position (even if that position is slightly wrong) than distributing the intensity value over a cluster of voxels.

We note that there are alternatives to the nearest-neighbor interpolation onto a regular grid that we use in this work. The interpolation could be avoided by solving the inverse transform by inverse methods, such as is performed in the computer program BACK (Ref. 37, Sect. A5.3), which utilizes a constrained conjugate gradient solver and which is used in the computer program SPEDEN³⁸ (SPEDEN is a program to perform 3D phase retrieval, designed specifically to optimally include prior data and avoid instabilities caused by interpolation). Alternatively, it should be possible to employ fast algorithms to compute discrete Fourier transforms of nonequidistant data.³⁹ In the nondiffracting limit (i.e., computed tomography, CT), the reconstruction method of filtered backprojection can be shown to be a computationally efficient method that is equivalent to Fourier synthesis via a polar-to-Cartesian interpolation.^{40,41} A corresponding algorithm, called filtered backpropagation,⁴² has been developed for the diffracting case where the diffraction amplitude and phase are measured, which again can be shown to be equivalent to Fourier synthesis after interpolation.⁴³

C. Phase Retrieval

Our phase retrieval method follows from previous work on 2D diffraction imaging.^{6,7,8,44} In particular we have extended the Shrinkwrap algorithm⁶ to operate on 3D arrays. This algorithm is based on an iterative transform algorithm (ITA), which cycles between real and reciprocal space, respectively enforcing the constraints of known object support and known diffraction modulus. Usually an ITA requires knowledge about the shape of the object to set the support constraint. This support is usually larger than the actual boundary of the object; what is termed a

loose support. For general complex-valued objects (e.g., consisting of more than one element or compound), where a positivity constraint cannot be applied, the ITA gives higher-quality reconstructions when the support constraint more closely and tightly matches the object's boundary.⁴⁵ The reason for this is explained in Subsection 4.B. The Shrinkwrap algorithm periodically updates the estimate of the support based on the current object estimate. The updated support is chosen by low-pass filtering the current estimate and setting the support to be the region for which the intensity is above a certain threshold (usually a prescribed fraction of the maximum image intensity). The method can be started from an estimate of a very loose support, from a threshold of the object's autocorrelation function, or even from the entire array. A method that exists for finding an estimate of the object support from the autocorrelation function's support could also be used.^{46,47} While the Shrinkwrap method can be used with any ITA, such as the hybrid input-output (HIO)⁴⁸ or difference map⁴⁹ algorithms, we used the HIO and relaxed averaged alternating reflection (RAAR)⁵⁰ algorithms for this work.

Many of the phase retrieval ITAs can be written as fixed point iterative equations, which can be written generally in the form $g_{n+1} = \mathcal{T}g_n$ for a generic operator \mathcal{T} . The RAAR algorithm can be represented by the fixed point iterative equation of the complex-valued real-space image iterate g (Ref. 50):

$$\begin{aligned} g_{n+1} &= \left[\frac{1}{2}\beta(R_S R_M + I) + (1 - \beta)P_M \right] g_n \\ &= [2\beta P_S P_M + (1 - 2\beta)P_M + \beta(P_S - I)] g_n, \end{aligned} \quad (11)$$

where the operator $R = 2P - I$ is the reflector corresponding to the projector P , I is the identity operator, and β is a feedback parameter, which we usually set to $\beta = 0.9$. The two operators P_M and P_S are the projections onto the modulus constraint and support constraint, respectively. We apply the modulus according to

$P_M g$

$$= \mathcal{F}^{-1} \begin{cases} \frac{G(\mathbf{u})}{|G(\mathbf{u})| + \epsilon} [\sqrt{I(\mathbf{u})} + \sigma_{\sqrt{I}(\mathbf{u})}] & \text{if } G(\mathbf{u}) > \sqrt{I(\mathbf{u})} \\ & + \sigma_{\sqrt{I}(\mathbf{u})} \\ \frac{G(\mathbf{u})}{|G(\mathbf{u})| + \epsilon} [\sqrt{I(\mathbf{u})} - \sigma_{\sqrt{I}(\mathbf{u})}] & \text{if } G(\mathbf{u}) < \sqrt{I(\mathbf{u})} \\ & - \sigma_{\sqrt{I}(\mathbf{u})} \\ G(\mathbf{u}) & \text{otherwise, or } \mathbf{u} \notin M \end{cases}, \quad (12)$$

where $\sigma_{\sqrt{I}}$ is the estimated standard deviation of the measured diffraction amplitudes, $G(\mathbf{u}) = \mathcal{F}\{g(\mathbf{x})\}$, and ϵ is a small number. M is the set of \mathbf{u} where $I(\mathbf{u})$ has been measured. For example, $\mathbf{u} \notin M$ in the missing sector of data present when diffraction is recorded over a limited range of angles. The operator P_M of Eq. (12) retains the phase of the complex Fourier amplitude $G(\mathbf{u})$ and projects its modulus $|G(\mathbf{u})|$ to the nearest measured diffraction ampli-

tude, in the interval $[\sqrt{I(\mathbf{u})} - \sigma_{\sqrt{I}(\mathbf{u})}, \sqrt{I(\mathbf{u})} + \sigma_{\sqrt{I}(\mathbf{u})}]$ [or does nothing if the modulus already lies within that range or if $I(\mathbf{u})$ has not been measured]. Given the support S from Shrinkwrap, we apply the support constraint when retrieving the phase of a complex image, using

$$P_S g = \begin{cases} g(\mathbf{x}) & \text{if } \mathbf{x} \in S \\ 0 & \text{otherwise} \end{cases} \quad (13)$$

We also perform phase retrieval where we impose real and positive constraints on the image amplitudes, replacing P_S with

$$P_{S+g} = \begin{cases} \Re\{g(\mathbf{x})\} & \text{if } \mathbf{x} \in S \text{ and } \Re\{g(\mathbf{x})\} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

The HIO algorithm can be written in terms of a fixed point iterative equation only when applying the support constraint P_S , not when applying positivity constraints.⁵⁰ In general the HIO algorithm is given by

$$g_{n+1} = \begin{cases} P_M g_n & \text{if } \mathbf{x} \in S' \\ (I - \beta P_M) g_n & \text{otherwise} \end{cases} \quad (15)$$

where S' is the set of elements where $P_M g_n$ satisfies the support and (if desired) the reality and positivity constraints. As with the RAAR algorithm we use a value of the feedback parameter of $\beta=0.9$.

Regardless of algorithm, we monitor the reconstruction with the real-space image error

$$E_S^2 \equiv \frac{\sum |g_n - P_S g_n|^2}{\sum |P_S g_n|^2} = \frac{\sum_{\mathbf{x} \notin S} |g_n(\mathbf{x})|^2}{\sum_{\mathbf{x} \in S} |g_n(\mathbf{x})|^2} \quad (16)$$

This metric is a measure of the total power in the image that remains outside the support and is zero for the case of perfectly satisfying the real-space constraints. We define, in a completely analogous way to Eq. (16), the error E_M corresponding to the modulus constraint defined by P_M :

$$E_M^2 \equiv \frac{\sum |g_n - P_M g_n|^2}{\sum |P_M g_n|^2} = \frac{\sum \| |G_n| - \sqrt{I} \|^2}{\sum I} \quad (17)$$

where the equality follows from Parseval's theorem and is true only for $\sigma_{\sqrt{I}}=0$. The error metrics E_S and E_M are the normalized distances between the current iterate g_n and the support and modulus constraint sets, respectively. The reconstructed image from a reconstruction run (from a particular set of starting phases) is given by

$$\gamma_M = P_M g_n \quad (18)$$

for the final iterate g_n of both the RAAR and HIO algorithms.

The Shrinkwrap algorithm has been used previously to reconstruct 2D images of thin objects at a resolution of about 20 nm.⁶ We have found in subsequent studies that the step of updating the support would sometimes shrink the support to a shape smaller than the actual boundary of the object. To counter this effect, we have improved the Shrinkwrap algorithm to prevent it from overshrinking

the support. Depending on the parameters of the low-pass filter and the threshold level, the support may start to cut off the extremities of the object. At this point the support constraint error E_S^2 increases rapidly, and the reconstructions rapidly degrade with further iteration. This error is thus a good indicator of when to halt the support refinement. We simply monitor the error metric, and when it increases above a set point, we choose the support saved from ten iterations prior. This then becomes our best estimate of the support and is used as a fixed support in combination with the RAAR algorithm for many more (typically 100–1000) iterations. We further decrease the uncertainty of the retrieved phases by averaging the retrieved complex images from independent and random starting diffraction phases using the Shrinkwrap-derived support constraint³⁶ as described in Eq. (22) of Subsection 5.A. If the phase at a particular spatial frequency is randomly recovered from trial to trial, the average modulus will average to zero and hence be filtered out of the recovered image.

The 2D reconstructions shown in this paper were performed by using the RAAR algorithm [Eq. (11)], and the 3D reconstructions were performed by using a combination of HIO [Eq. (15)] and RAAR algorithm. A typical reconstruction process proceeds as follows. First we define the initial object support mask by applying a 2% intensity threshold to the object autocorrelation, obtained by Fourier transforming the measured diffraction pattern. The support constraint, defined by the current object mask, is applied to the solution in real space once per iteration. We typically use a feedback parameter of $\beta=0.9$ in the RAAR or HIO algorithm. The object support S is recomputed every 30 iterations by convolving the absolute value of the current reconstruction γ_M with a Gaussian of FWHM of initially 3 pixels in all dimensions and applying a threshold to the resultant image at 15% of the maximum value. As the iterations progress, we reduce the width of the Gaussian blurring function from 3 pixels to 1 pixel, following the prescription $w_S = 1 + 2 \exp(-n^2/n_w^2)$, with n_w regulating the speed at which w_S decreases with iteration number n . The reduction in the blurring width enables the support to better conform to the solution as the quality of the reconstruction increases. We perform this Shrinkwrap support determination without applying any real-space positivity or reality constraint on the image amplitudes (that is, we use the constraint P_S in the RAAR algorithm or $S'=S$ in the HIO algorithm). The final support is usually obtained after 300–600 iterations, with a stopping criterion that the support constraint error E_S^2 does not exceed 0.2. Once the support is determined, we carry out many iterations of the RAAR algorithm, starting from random phases, using a feedback parameter of $\beta=0.9$. In some cases additional real-space constraints, such as positivity or reality of the image amplitudes, are also applied.

As shown in Eq. (12), in diffraction space the amplitudes of the object guess are matched in magnitude to the measured diffraction pattern amplitude over those parts of 3D diffraction space where the measured intensity is defined. Those parts of 3D diffraction space where there are no measured data are allowed to float and are not constrained. This includes the regions between the measured

Ewald spheres, the missing wedge of data from the finite range of rotation angles, the central beam-stop region, and those parts of the diffraction pattern where the measured intensity is sufficiently low to be regarded as noise. An additional, optional Fourier space constraint is to set those pixels beyond the radius of the spatial frequencies measured by the CCD chip to zero. This asserts lack of knowledge of spatial frequencies higher than those measured by the CCD camera and effectively provides a pupil function for the imaging system in 3D space.

Providing an initial guess for the 3D object support is not typically necessary but speeds the reconstruction process and helps break inversion symmetry present in the object autocorrelation. An initial 3D support estimate can be obtained from the diffraction data by first performing Shrinkwrap phase retrieval on a 2D central section, as described in Subsection 4.C. We then extrude the support mask into a 3D shape that is invariant in the projection direction and that matches the 2D mask that was generated in all planes normal to this direction, to provide an initial 3D support estimate. If several 2D reconstructions are available from a range of views, the intersection of these support functions in three dimensions can be used to provide a more detailed initial support estimate. Experience has shown that even a low-resolution or comparatively poor support estimate is sufficient to almost immediately break any inversion symmetry in the reconstruction and hasten convergence of the 3D solution. If three or more extruded support functions are used to form an intersection, then the support functions may need to be initially aligned to form an intersection, although the Shrinkwrap algorithm will usually improve the support even if misalignment yields a too small intersecting volume. A better method for providing an initial support estimate appears to be obtainable from intersection operations of the support of the autocorrelation function,^{46,47} which is readily computed as described in Subsection 3.C. Although we have not yet attempted it, we believe that estimating the initial support from the support of the autocorrelation, followed by the dynamic support constraint of Shrinkwrap, should provide a very robust reconstruction methodology.

3. METHODS

A. Sample Preparation

A goal of this study was to be able to unambiguously compare reconstructed x-ray images of a 3D object with images obtained by another high-resolution method, such as a scanning electron microscope (SEM). To accomplish this, we fabricated a test object that consists of a silicon nitride membrane with a 3D pyramid shape that is decorated with 50-nm-diameter colloidal gold spheres, similar to that previously described.⁷ The object is 3D and has a comparable width, height, and depth, measuring $2.5 \mu\text{m} \times 2.5 \mu\text{m} \times 1.8 \mu\text{m}$.

The pyramid-shaped membrane was fabricated by lithography with methods similar to those used to make silicon nitride windows and silicon nitride atomic-force microscope tips. The starting material was a double-side polished 200- μm -thick wafer of silicon crystal with the crystal 100 axis oriented normal to the surface. Pits with

an inverted pyramid shape were etched into one side of the wafer by anisotropic etching through a pattern of 2.5- μm -wide square holes, lithographically printed and developed in photoresist. The anisotropic etch leaves the 111 crystal planes exposed, so that the surface normal of any of the four faces of the pyramid makes an angle of 54.7° to the window normal and the ratio of the depth of the pit to its base width is $1/\sqrt{2}$. After removal of the photoresist a low-stress silicon nitride film of 100-nm thickness was grown on the surface by chemical vapor deposition. Window openings were then etched from the other side of the wafer after first resist coating and patterning that side, making sure to align to marks etched in the front surface. The etch from the back removes silicon but leaves a free-standing membrane of silicon nitride, which in this case had one pyramid-shaped indentation per window. The windows were made with a slotted shape about 2-mm-wide by 50- μm high. With the 200- μm thickness of the silicon frame and the pyramid positioned in the center of the window, this allows a line of sight through the window at a maximum rotation angle (about an axis in the plane of the window, parallel to the short window dimension) of 78° .

The gold-sphere test object was made by dragging a small drop of solution of gold balls in water, suspended from a micropipette, across the silicon nitride window so that it intersected with the pyramid indentation. Best success was achieved with a slightly hydrophilic silicon nitride surface, which could be obtained by cleaning the surface in an oxygen plasma. As the drop was moved over and away from the indentation, a smaller drop broke away from the main drop and was captured in the pyramid. This drop quickly evaporated and left the gold balls in a characteristic pattern where the gold tended to fill in the edges of the pyramid. The main drop was completely dragged away from the window, so the only gold balls on the window were those in the pyramid. A plan-view SEM image (membrane and wafer perpendicular to the electron beam) of the object is shown in Fig. 2. The SEM is, however, sensitive only to the surface of the object—the electrons do not penetrate the gold spheres or the membrane. The depth of focus of the SEM was larger than the thickness of the object, and from the plan view we can determine the lateral coordinates of the topmost balls and infer the third coordinate from the known geometry of the pyramid.

The silicon nitride window was glued to a pin so that the pyramid was close (within about 20 μm) to the rotation axis of the pin. The pin was mounted in a collar that was attached to a JEOL electron microscope sample holder. This assembly was mounted into the modified goniometer holder of the diffraction apparatus.¹⁷

B. Data Acquisition

Experiments were carried out at an undulator source at beamline 9.0.1¹⁸ of the Advanced Light Source (ALS) using the Stony Brook University diffraction apparatus.¹⁷ Pertinent to this experiment, 750-eV (1.65-nm-wavelength) x rays were selected from the undulator radiation by a zone-plate monochromator with a spectral resolution of $\lambda/\Delta\lambda=1000$. The 5- μm -diameter monochromator exit pinhole also selects a transversely spatial co-

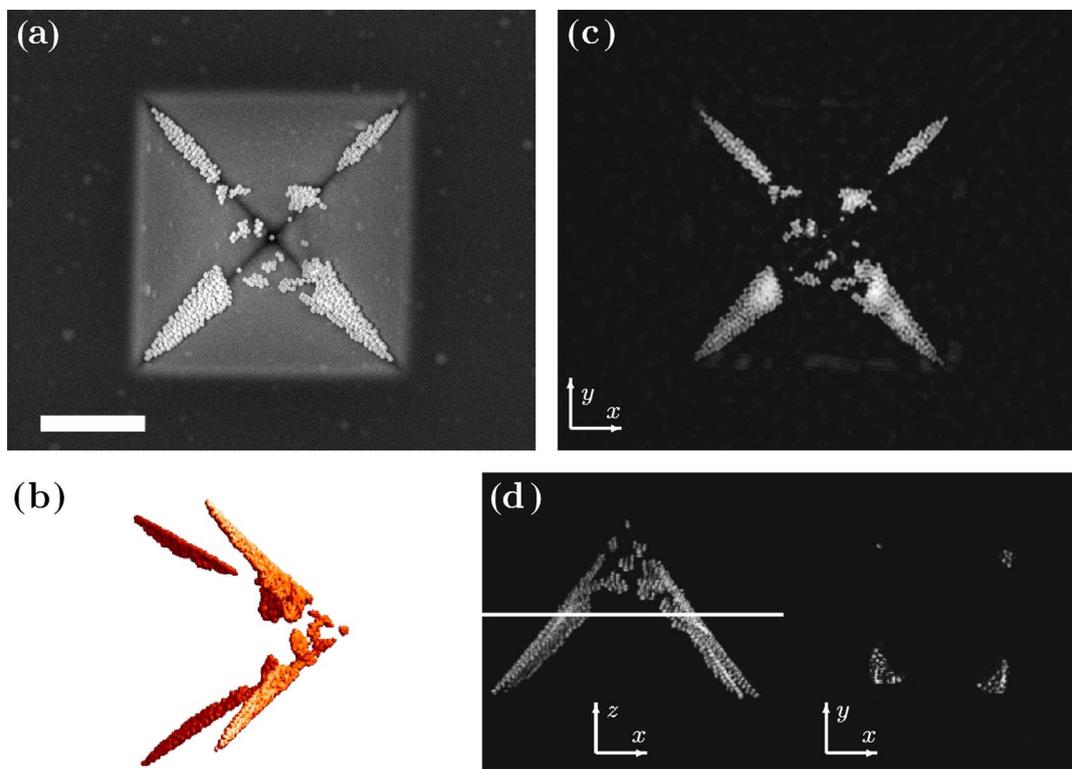


Fig. 2. (Color online) (a) SEM image of the pyramid test object, consisting of 50-nm-diameter gold spheres lining the inside of a pyramid-shaped indentation in a 100-nm-thick silicon nitride membrane. The membrane extends over a window of size $50\ \mu\text{m} \times 1.7\ \text{mm}$, the pyramid base width is $2.5\ \mu\text{m}$, and the height is $1.8\ \mu\text{m}$. (b) Isosurface rendering of the reconstructed 3D image. (c) Extremely large depth-of-field x-ray projection image from a central section of the 3D diffraction data set, reconstructed with the Shrink-wrap algorithm. (d) Maximum value projection of the 3D reconstructed image (left) with a horizontal white line indicating the location of a tomographic slice (right). The scale-bar length is $1\ \mu\text{m}$ and applies to all images.

herent patch of the beam. The sample was located 20 mm from this pinhole. A direct-detection bare CCD detector, with $20\text{-}\mu\text{m}$ pixel spacing, 1340×1300 pixels, was located 142 mm behind the sample. We selected subarrays of 1200×1200 elements, centered on the location of the zero spatial frequency (direct beam). At these CCD and wavelength settings we have a real-space sampling interval in x and y of $\Delta x = 9.8\ \text{nm}$ (in the small-angle approximation) and a field width of $w = N\Delta x = 11.7\ \mu\text{m}$. With these settings the $2.5\text{-}\mu\text{m}$ -wide pyramid object satisfies the far-field and sampling conditions discussed in Subsection 2.A. The diffraction from the pyramid object is more than $4 \times$ oversampled in each dimension ($s = 4.6$).

The frame of the slotted window in which the pyramid is formed blocks most of the high-angle scatter from the pinhole that would otherwise illuminate the CCD. This scatter reveals a projection shadow image of the slotted window, useful for aligning the pyramid to the beam. The diffraction pattern of the pyramid measured by the CCD is shielded from this remaining pinhole scatter with a small aperture placed 6 mm upstream of the sample (a distance at which the sample can be rotated without interference). A beam stop blocks the direct undiffracted beam from impinging on the CCD. More details are given by Beetz *et al.*¹⁷

Diffraction patterns were collected with the sample oriented at rotation angles of -57° to $+72^\circ$ at 1° intervals [compared with 0.27° angular increments required for full sampling according to Eq. (6)]. The shadow of the sample

support frame limited useful data to -57° to $+66^\circ$. We additionally collected data at 0.5° increments for a range of 19° centered at an object orientation of $\phi = -26^\circ$ from the head-on ($\phi = 0^\circ$) orientation. To keep the sample centered in the $5\text{-}\mu\text{m}$ beam, we calibrated the position of the sample by performing a 2D raster scan of the rotation and y goniometer motors. The total scattered counts (not including those blocked by the beam stop) were collected for each motor position and the optimum y position (a translation motion perpendicular to the rotation axis) was then computed for each rotation angle, and these were fitted to a smooth curve as a function of rotation angle. To collect the 3D data set, at each rotation angle we took several exposures to accommodate the large dynamic range of 10^5 of the diffraction pattern and to reduce the area occluded by the beam stop (by setting the beam stop to different positions). After subtraction of dark noise, pixel data that were not saturated and not masked by the beam stop were summed over these exposures and then normalized by the accumulated incident flux corresponding to that sum. A typical diffraction pattern is shown in Fig. 3(a), which was composed of ten or more individual exposures of 0.1-, 1-, 10-, and 60-s duration, for a cumulative exposure of 73 s. The diffraction pattern intensities are displayed on a logarithmic gray scale in Fig. 3. At the highest angles of the pattern (highest resolution, at $0.07\ \text{nm}^{-1}$, along the diagonal), the mean photon count is 1.9 photons/pixel for this sample orientation. The maximum normalized photon count, which occurs in a pixel

near the center of the pattern, is 109,000 photons. The estimated incident flux was 8×10^9 photons $s^{-1} \mu m^{-2}$ (per 400 mA of storage-ring current), and the normalized incident fluence for the accumulated sum of Fig. 3(a) was 3×10^{11} photons/ μm^2 . The total normalized scattered counts at the CCD over the accumulated exposure time for the pattern in Fig. 3(a) was 1.6×10^8 photons (equal to the total counts that would be recorded if the detector had infinite dynamic range and did not saturate).

Views of the diffraction data cube are shown in Figs. 4(a)–4(c) below and discussed in Subsection 4.A. This cube was assembled from the 123 diffraction patterns at 1° sample orientation increments, and 32 patterns at half-degree intervals, by interpolating onto $\mathbf{u}_{i,j,\phi}$. The total integrated exposure time for the complete data set was 3.2 h, with a total incident fluence of 5×10^{13} photons/ μm^2 .

C. Data Diagnostics

As seen in relation (5), the autocorrelation of the object can be determined from a single inverse Fourier trans-

form of the measured data without having to recover the diffraction phases. We find that computing the autocorrelation image from the measured diffraction data is a useful diagnostic to determine if the measurement fulfills the sampling requirements, to help identify the object, and to assess the measurement quality. The inverse Fourier transform of the diffraction pattern shown in Fig. 3(a) is given in Fig. 3(b), after first applying the high-pass filter described below in Eq. (19) to remove effects due to the abrupt intensity change caused by the beam stop. The displayed autocorrelation image has been cropped by half in width and height from the inverse Fourier transform of the diffraction pattern, since, with the linear gray scale displayed, the rest of the field was black. This autocorrelation image has a well-defined support that is confined within the image field, showing that we are indeed oversampling the diffraction intensities.

The Fourier relationship of expression (5) offers a further method to diagnose the diffraction pattern $I(\mathbf{q})$ as a function of the spatial frequency \mathbf{q} across the pattern. A property of the Fourier transform of a function, often dis-

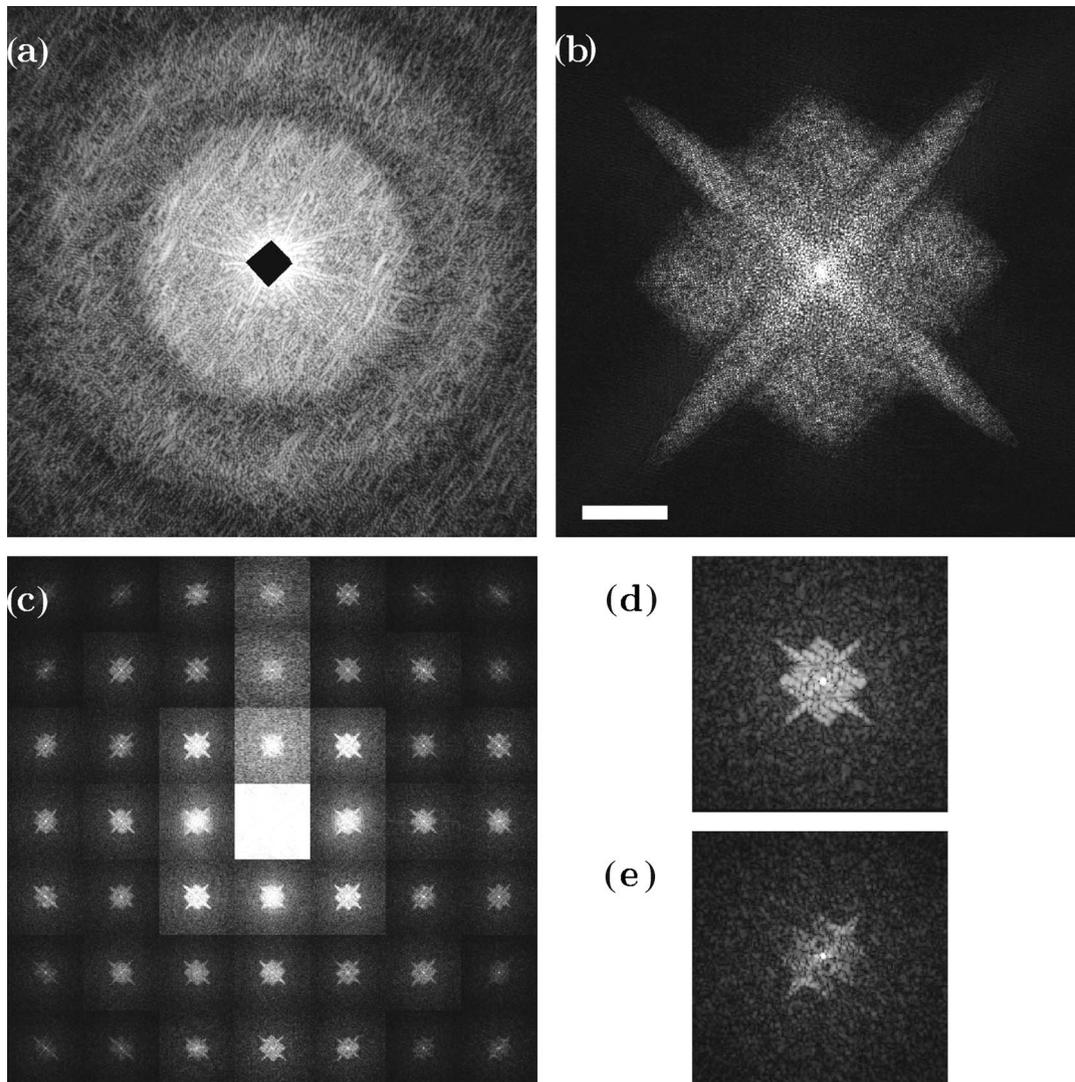


Fig. 3. (a) Diffraction pattern for the $\phi=0^\circ$ orientation of the pyramid. (b) Autocorrelation image $i(\mathbf{x})$, formed by Fourier transforming the diffraction intensity pattern of (a) after filtering, displayed with a linear gray scale (white corresponds to highest intensity). The scale bar is $1 \mu m$. (c) Autocorrelation spectrogram of the same single-view diffraction pattern of the pyramid, displayed with a logarithmic gray scale. The central and rightmost images of the bottom row of (c) are redisplayed in (d) and (e), respectively.

cussed in the context of holography, is that a subregion of the transform (or hologram) can be inverse Fourier transformed to give a spatially filtered image of the original function. The field of that image is the full field of the original function. The filtered image may differ from subregion to subregion, since each subregion contains information pertaining to particular spatial frequencies of the original object function. Nevertheless, for nonperiodic object functions, these images should be consistent and not vary too dramatically between neighboring regions. Large inconsistencies in images formed in neighboring regions point to inconsistencies in the measured diffraction data. This gives a convenient way to qualitatively check the diffraction intensities (and also reconstructed phases) across the diffraction pattern by forming an array of subimages, each corresponding to a particular subregion. We term this array a “spectrogram,” since it displays the image information as a function of spatial frequency, much like the time–frequency spectrograms used in audio analysis. We apply the spectrogram analysis both to the autocorrelation image and to reconstructed images to determine the consistency of the data or reconstructed phases.

An example of an autocorrelation spectrogram is shown in Fig. 3(c), where each image is formed by applying a 2D Gaussian window function to the diffraction pattern, centered at a location indicated by the image position in the array. One effect that is immediately noticeable to the eye is that, at the higher-resolution positions, the images vary according to their azimuthal position in the spectrogram. In particular, features that are tangential to circles of constant resolution have greater contrast than features that vary in a radial direction. The effect gives the appearance that the spectrogram is rendered onto the surface of a sphere, but it is in fact a signature of limited longitudinal coherence of the incident beam. For a given \mathbf{q} , represented in the spectrogram by the position of the image, pairs of points of the object that are perpendicularly bisected by the vector \mathbf{q} will interfere with zero path difference. These points will interfere with the same phase for all wavelengths (assuming no dispersion of the scattering factors). The path difference at \mathbf{q} of other pairs of points along this line perpendicular to \mathbf{q} depend only in second order on their mean displacement from the direction of \mathbf{q} . The path differences of rays scattering from pairs of points separated in the radial direction, however, vary to first order with their radial separation. Therefore a limited longitudinal coherence, which limits the path difference over which interference occurs, will reduce interference for points separated by a direction parallel to the \mathbf{q} vector by a much greater extent than that for the perpendicular direction. The spectrogram gives a good visual determination of the degree of longitudinal coherence, and we see from Fig. 3(c) that the longitudinal coherence is adequate for diffraction out to the center edge of the pattern [Fig. 3(d)] but not quite adequate for the corner of the pattern [Fig. 3(e)]. By comparison with spectrograms of simulated diffraction patterns, we estimate that the relative bandwidth of the incident beam in this case is approximately $\Delta\lambda/\lambda=1/600$.

It is also clear from Fig. 3(c) that the data are inconsistent at the low spatial frequencies, since at those \mathbf{q} positions there is noticeable intensity outside the support of

the pyramid object autocorrelation. This is due to the fact that low-frequency data are missing because of the beam stop, and also to a lesser degree due to scattering from the sample substrate or the coherence-defining pinhole. The data are visibly noisier in windowed regions located in the first three rows of the fourth column of Fig. 3(c), due to the stalk that holds the beam stop and that was moved over several positions in this region for the cumulative exposure. The noise and inconsistency can cause the Shrinkwrap algorithm to fail (in which it keeps shrinking beyond the object’s boundary), especially when applied without an image-space constraint such as positivity. We find that the Shrinkwrap algorithm consistently converges to a stable object support when we apply a high-pass filter to the diffraction intensities before interpolation. This filter has the form

$$f(q) = \begin{cases} (q/2a)^4 \exp(2 - q^2/2a^2) & \text{if } q < 2a \\ 1 & \text{otherwise} \end{cases}, \quad (19)$$

where $q=|\mathbf{q}|$ and the filter radius $a=100$ pixels, or less than 10% of the array width. The image thus formed will be a high-pass filtered image, equivalent to the coherent image formed by a lens with a central obscuration. The filter was applied to the data of Fig. 3(a), before Fourier transformation, to diminish the effects of the beam stop. This filter also regularizes the inverse transform, which is ill posed in part due to the missing data behind the beam stop,^{32,35} by simply setting the unknown intensities to be zero. The effect of this high-pass filter is to convolve the coherent image with the Fourier transform of the filter. This causes a ringing of the image, which gives rise to negative amplitudes in the image, and a slightly larger image support. We also zero the diffraction intensities of the bright cross streaks seen in the x - z central section, to suppress artifacts that they may cause.

D. Computational Implementation

The two key computational challenges in implementing high-resolution 3D phase retrieval at the time of writing are performing the numerous 3D FFTs required in a reasonable period of time and managing the memory requirements of the large 3D data arrays.

Memory requirements are dictated by the size of the data sets acquired and by the phase retrieval algorithms used. For the iterative transform phase retrieval methods described in Subsection 2.C, we require four or more 3D arrays with the same number of elements as that in the interpolated input diffraction volume. Specifically, the arrays required are the input diffraction modulus data ($\sqrt{I(\mathbf{u})}$, floating point), the current and previous iterates [$g_n(\mathbf{x})$ and $g_{n-1}(\mathbf{x})$, complex floating-point data], and the support constraint (S , byte data). The estimated modulus standard deviation $\sigma_{\sqrt{I}}$ requires another floating-point array, but in the 3D reconstructions we set $\sigma_{\sqrt{I}}=0$ to reduce the memory requirement and speed up the reconstructions. In principle FFTs can be performed on arbitrary sized arrays; however, it is advantageous to perform reconstructions on a regular grid with 2^n grid points on a side. Our interpolated data array is a gigavoxel data cube containing $1024^3=2^{30}$ elements, which requires a total of 8 GB of computer memory per array for single-precision

Table 1. Minimum Memory Footprint Required for Iterative 3D Phase Retrieval for Various Array Sizes^a

Array Size	Single Precision	Double Precision
256 ³	336 MB	592 MB
512 ³	2.6 GB	4.6 GB
1024 ³	21 GB	37 GB
2048 ³	168 GB	296 GB

^aThe arrays required are the input diffraction data (floating point), the current and previous iterates (complex single- or double-precision floating-point data), and the support constraint (byte data).

Table 2. Computing Times Using a Cluster-Based Fourier Transform and Reconstruction Code on 16 G5 Dual-Processor Xserve Compute Nodes^a

Array Size	Time per 3D Fourier Transform	Time per 3D Reconstruction
256 ³	73 ms	10 min
512 ³	850 ms	1.5 h
1024 ³	7.9 s	14 h

^aFourier transform timings are wall time per individual FFT. Reconstruction timings are for a complete 3D reconstruction consisting of 2000 iterations of HIO phase retrieval complete with two FFTs per iteration plus other operations required to calculate the reconstruction.

complex data. The minimum memory footprint for single-precision iterative object reconstruction using either the HIO or RAAR algorithm on a 1024³ data cube is therefore 2 × 8 GB complex arrays, plus 1 × 4 GB diffraction cube data and 1 × 1 GB support array, giving a total memory footprint of 21 GB of data, where we use the definition of 1 GB = 2³⁰ bytes. The minimum memory footprint for performing basic HIO and RAAR reconstruction on 3D arrays of different sizes is given in Table 1. Note that this is the minimum memory footprint needed to perform a HIO reconstruction and that more memory may be required depending on the specific implementation. For example, FFT speed can be increased through use of temporary “workspace” arrays that require additional memory, and maintaining a running sum of successive images γ_M requires an additional complex-valued array to be retained in memory. The memory calculations above include only the data arrays and do not take account of operating system requirements and the executable code itself.

The second computational challenge is efficient evaluation of the numerous 3D Fourier transforms required for 3D phase retrieval. The Fourier transform of a single large data set is not trivially parallelizable, in that the problem can be easily broken into separate parallel tasks and distributed over many computer processors, as is the case, for example, with ray tracing and partially coherent imaging, where each CPU can work on a subset of the entire problem without the need for intensive internode communication during execution. The nature of the Fourier transform means that any element of the input array affects all elements of the output, requiring internode exchange of array data at each Fourier transform step to ensure that all CPUs work together to solve the one large FFT.

We overcome the problem of efficiently calculating distributed Fourier transforms by using the `dist_fft` distributed gigaelement FFT library from Apple Computer specifically written for this project by the Apple Advanced Computation Group.²⁰ This FFT library distributes the Fourier transform calculation load efficiently over many processors and has been hand-optimized to take advantage of the G5 architecture used in the Apple Macintosh line of computers and the “AltiVec” single-instruction multiple-data floating-point vector processing unit. Distributed FFT libraries are also available elsewhere, for example in version 2 of the FFTW libraries,⁵¹ but at this time these do not support single-instruction multiple-data vector processing extensions and proved to be slower on our platform. `dist_fft` decomposes the input 3D data set into n_{proc} discrete data slabs consisting of an $n \times n \times (n/n_{\text{proc}})$ voxel subportion of the original data array. Only a distinct portion of the array resides on each CPU at any given time, enabling data sets much larger than the memory of each individual node to be computed, and the distributed memory nature of the FFT is exploited through parallelization of all steps in the reconstruction code. Standard message passing interface⁵² (MPI) commands are used to communicate data between processes.

We ran fully parallelized reconstruction code on a 16-node 2.0-GHz dual-processor (32 processors total) Macintosh Xserve G5 cluster with 4-GB RAM per node. To maximize interprocess communication speed, we used high-speed, low-latency Mellanox Infiniband interconnects to carry message passing interface traffic between compute nodes. If this cluster is used, the processing time on a 512³ array is 2.2 s/iteration with the HIO phase retrieval algorithm, and an acceptable 3D reconstruction can be produced in under 2500 iterations for a total computation time of 2.5 h on a 512³ grid. The individual FFT timing and total reconstruction time for typical array sizes on this cluster are given in Table 2.

4. IMAGE RECONSTRUCTION

A. Three-Dimensional Images

A full 3D image is obtained by performing phase retrieval on the entire 3D diffraction data set. The resulting volume image reveals the structure of the object in all three dimensions and can be visualized in many ways including forming projections through the data or slices (tomographs) of the data. Specific segmentation analyses can be carried out on the volume image to determine properties such as strength of materials.⁵³ 3D reconstructions were performed by interpolating the diffraction intensities at $\mathbf{u}_{i,j,\phi}$ onto a 1024³ grid. Representations of the interpolated diffraction intensities are given in Fig. 4. Note that the 1° angular increments of the object rotation are just less than four times larger than the 0.27° requirement of Eq. (6) for this object and that we have a 40° sector of missing data due to our limited range of object orientations, as well as data lost to the beam stop. The effect of the 1° rotation increment is apparent in Fig. 4(b), where the gaps between the measured Ewald spheres are seen in the u_x - u_y plane (referred to as a central section) extracted from the data cube. The limited range of views is readily apparent in Fig. 4(c), which shows the u_x - u_z central section.

The 3D phase retrieval code described above in Subsection 3.D was applied to the assembled 3D data to produce a full 3D reconstruction from the diffraction cube. We applied the Shrinkwrap algorithm, as described in Subsection 2.C, to determine the 3D support mask and the diffraction phases. We performed phase retrieval by using either the real-positive real-space constraint P_{S+} or the support constraint P_S . For the complex image reconstruction, as with the case of reconstruction from central sections discussed below in Subsection 4.C, the solution was regularized by first applying the high-pass filter of Eq.

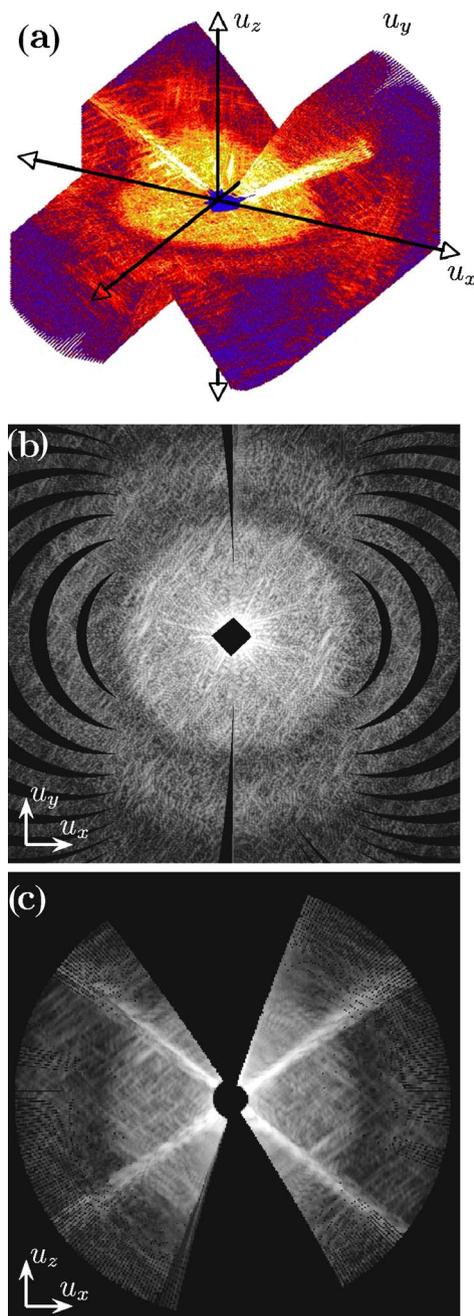


Fig. 4. (Color online) (a) Rendering of the entire 3D diffraction data set. (b) Central slice of the data set in the plane rotated by -24° about the y axis from the u_x - u_y plane. (c) Central slice of the data set in the u_x - u_z plane. All patterns are displayed on a logarithmic gray scale (white corresponds to highest intensity). The half-width of each pattern is $u_{x,\max}=0.048 \text{ nm}^{-1}$.

(19) to the diffraction intensities. For the real-positive reconstruction the missing amplitudes were unconstrained and were allowed to be recovered by the algorithm. The reconstruction success with the sparsity of data that we have in this case is undoubtedly due to the sparseness of the object itself. In essence the object is a membrane, and the 3D speckles are elongated by up to 50 pixels in directions perpendicular to the pyramid faces, as can clearly be discerned in Fig. 4(c).

A 3D isosurface rendering of the real-positive constrained reconstructed pyramid is shown in Fig. 2(c), where we also display a slice from the volume image in Fig. 2(d). Three images from the 3D pyramid image are shown in Fig. 5, showing the maximum value projection, along the three orthogonal axes, of the real part of the 3D image that was reconstructed by using the support constraint with real positivity, P_{S+} . Each pixel of this image is given by the maximum value encountered along the path that projects onto that pixel, illustrating a visualization method available only with the full 3D image array. In initial reconstruction trials using only a support constraint P_S , we observed a linear ramp in the imaginary component. This was essentially a phase ramp, due to a shift of the diffraction data in Fourier space caused by an inaccurate measurement of the location of the zero spatial frequency (the direct beam) on the CCD. We shifted the diffraction data by an amount that minimized the real-space phase shift, which required shifting the data by half-pixel amounts. This recentering of the data was necessary before we could apply the real-positive constraint on the image. Further analysis of the images is carried out in Section 5. Movies of the reconstructed image and 3D diffraction intensities are available on the Web.⁵⁴

B. Two-Dimensional Images

2D images are useful for visualizing and quantifying objects, and most diffraction imaging experiments performed to date have been 2D. However, if the object is thick, then the interpretation of an image reconstructed from a single Ewald sphere is not trivial. Most notably, as compared with our full reconstructed 3D image, the 2D image will have defocus artifacts that do not diminish in power with displacement of the object along the beam axis. However, in some cases, obtaining a full 3D reconstruction may not be achievable, for example when imaging nonreproducible objects with single pulses of an x-ray free-electron laser. It is thus instructive to compare 2D images reconstructed from single-view diffraction patterns with the 3D image.

We first consider how thin an object must be to be considered 2D. In a 2D reconstruction from a single diffraction pattern, the spatial frequencies that are passed to the diffraction pattern are constrained on the Ewald sphere according to

$$q_z = 1/\lambda - \sqrt{1/\lambda^2 - q_x^2 - q_y^2} \approx -\frac{\lambda}{2}(q_x^2 + q_y^2), \quad (20)$$

where the approximation is for small scattering angles, or $q_x \ll 1/\lambda$. We can define the numerical aperture of the diffraction pattern as $\text{NA}=q_{x,\max}\lambda$, in analogy with imaging

with a lens (of square aperture for the case of a square detector, with the NA defined here along the half-width of the square rather than the diagonal), which gives the expression of maximum longitudinal distance of the Ewald surface, $q_{z,\max} \approx -\text{NA}^2/(2\lambda)$. For a 2D object of thickness $D \rightarrow 0$, the 3D transform will be independent of the longitudinal frequency q_z (rods in the z direction), and so the measurement on the Ewald sphere is equivalent to measurement of the 2D diffraction intensities $|O(q_x, q_y, 0)|^2$. In such a case there will be no artifact in the image caused by not actually measuring intensities on $q_z=0$. An object of finite thickness D will have a coherent diffraction pattern with speckles of width $1/D$ in the longitudinal direction. If, at the highest transverse frequencies, the Ewald sphere substantially cuts through a speckle that is centered at $q_z=0$ then the measurement will again be equivalent to the 2D diffraction intensities on the $q_z=0$ plane. That is, we can consider an object to be thin or 2D if the Ewald departure is no more than $1/(4D)$, or half the speckle half-width, which corresponds to

$$D < \frac{\lambda}{2\text{NA}^2}, \quad (21)$$

or, equivalently, the thickness D must be less than a depth of focus. For the experiments with the pyramid object at $\lambda=1.65$ nm and $\text{NA}=0.084$, this thickness limit is $D=120$ nm, which is considerably smaller than the 1.8 μm thickness of the pyramid.

Relation (21) does not imply, however, that diffraction imaging performs optical sectioning where only the parts of the object located within the depth of focus are imaged. The thickness limit simply implies that the 2D single-pattern image of an object thicker than D will contain artifacts due to the information that is cut off by the transfer function. Consider an object containing two parts (e.g., screens) that are separated by more than a depth of focus. As with coherent imaging with the equivalent aberration-free thin lens, partial information from both screens of that object will be transferred in the imaging process. In fact, in diffraction imaging, there is not necessarily any preferred image plane, since, by the Fourier shift theorem, a shift δz of an object along the beam axis z will cause only a phase shift given by $-2\pi\delta z q_z$ and hence there will be no change to the diffracted intensities (for small enough δz that the change in distance to the detector does not change the effective NA and scale the pattern on the detector). Note that from relation (20) the phase shifts of the 2D spatial frequencies of the image, due to the defocus δz , will be $\pi\delta z\lambda(q_x^2+q_y^2)$, as expected from the Fresnel propagator.⁵⁵ The position of the focal plane can be chosen in the phase retrieval step, a fact that was demonstrated computationally and experimentally by Spence *et al.*⁵⁶ In that work the focus of the retrieved image of an object of two screens separated by some depth could be chosen by setting a tight support for the features in one screen or the other. As shown by Spence *et al.*, once the phases of the diffraction intensities have been retrieved,

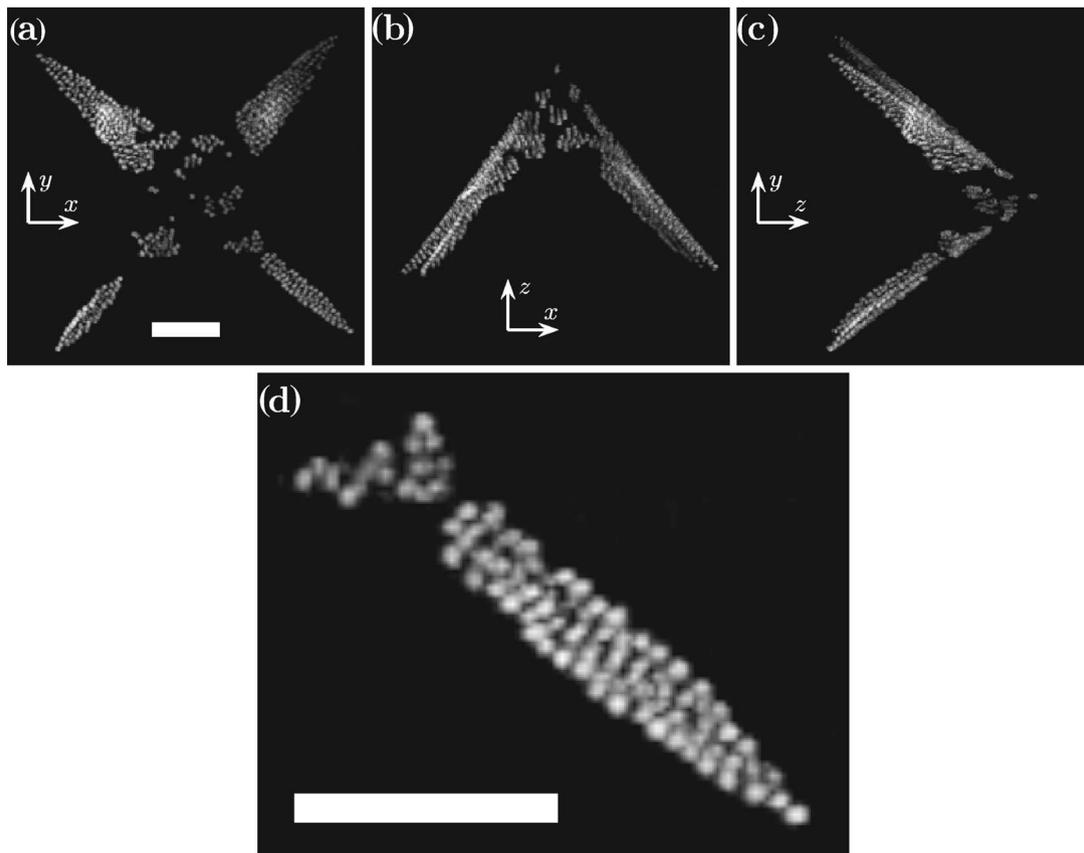


Fig. 5. Maximum value projections along three orthogonal directions of the reconstructed 3D image. Projections were performed along (a) z , (b) x , and (c) y directions. (d) Enlarged region of (a) for comparison with Fig. 8 below. The 3D image was reconstructed with reality and positivity constraints. The scale bars are 500 nm.

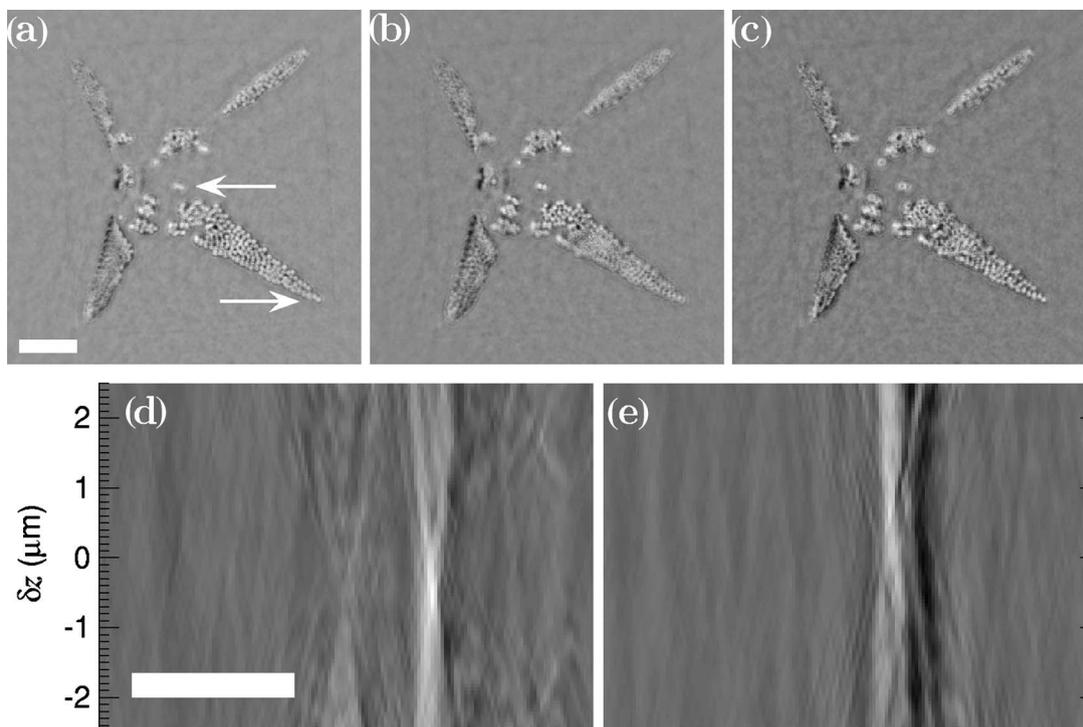


Fig. 6. Real part of the image reconstructed from (a) a single-view diffraction pattern, and real part of the image formed by numerically propagating (a) by (b) $-0.5 \mu\text{m}$ and (c) $+0.7 \mu\text{m}$. Lineouts from the image near (d) the pyramid center and (e) the arm extremity for a range of propagation from -2.5 to $+2.5 \mu\text{m}$. The locations of these lineouts are indicated by arrows in (a). The difference in the plane of best focus for these two image locations is apparent. The scale bars are 500 nm .

images can be generated at any position through focus by Fresnel propagating the image wave field (equivalent to applying the appropriate quadratic phase term to the diffraction phases).

The defocus effects of a single view are illustrated in Fig. 6, where we show 2D images of the wave field at the pyramid object, reconstructed from a single-view diffraction pattern. In this example we use the diffractogram for the object rotated by $\phi=24^\circ$ from the head-on (plan-view) orientation. The image γ_M reconstructed by Shrinkwrap, from the single-view diffraction pattern, is shown in Fig. 6(a). No real-space reality or positivity constraint was applied, and the reconstructed image is complex. For this object and view the edges of the object (its support) are at a range of heights along the beam axis. In this case the end-point support that the Shrinkwrap algorithm arrived at was tightest around the balls halfway along the arms of the pyramid, and consequently this is the plane of best focus. This focal plane gives the greatest overall image contrast, which explains why Shrinkwrap converges to it. The complex image can be numerically propagated, by convolution with the Fresnel propagator, by any arbitrary amount δz . We generated a series of numerically refocused images, where δz varies between -2.5 and $+2.5$ in 50 nm steps. As the refocus distance is increased, the best focus moves along the pyramid arms to their extremities. The difference in focus of balls near the vertex and arm extremities can be seen in Figs. 6(d) and 6(e), which show $x-\delta z$ lineouts of the real part of the complex image. The difference between the best focus for these two cases is $1.2 \mu\text{m}$, which agrees with the 3D image (Subsection 4.A) and the known geometry of the pyramid. It should be

noted that this computational focusing does not constitute 3D imaging but is simply the propagation of a 2D coherent field. The coherent transfer function (CTF) for the 2D image formed from a single-view diffraction pattern is confined to the Ewald surface, and in this situation with coherent illumination the integrated intensity of this 2D image does not change with defocus (a consequence of Parseval's theorem and the invariance of the diffraction intensities with defocus). That is, it is unlikely that numerical defocusing of a complicated object gives results that could be as easily interpreted as those for the pyramid-membrane test object used here. This situation is unlike partially coherent imaging in a microscope, where out-of-focus objects of compact support contribute less power to the image and some optical sectioning can be carried out.¹⁶

Another consequence of the "defocus artifact" of 2D images is that the 2D image of a thick real object is complex, which means that a real-space positivity constraint cannot be applied during the phase retrieval process. A positivity constraint, when valid, is known to be very effective in deriving the diffraction phases and is important in direct methods in crystallography and a strong constraint in diffraction imaging. Here a real object is one in which the object's 3D complex transmission function $o(\mathbf{u})$ is real, to a multiplicative complex constant. Propagation of the out-of-focus parts of the object to the selected image plane will give rise to a large variation in the complex values of image, as demonstrated in Fig. 7. Here we show the complex amplitudes of images recovered from calculated diffraction patterns of simulated objects. The simulated objects consisted of gold balls of equal size, distributed in a

similar way to that for the pyramid test object. In the first case [Fig. 7(a)] the z coordinate of all ball centers was set to zero to construct a quasi-2D object. Ignoring the arbitrary phase shift, the reconstructed image is real although not strictly positive (the negativity of the image is due to the truncation of the diffraction pattern). The calculated image values are complex for the 3D object [Fig. 7(b)], and there is a rough correlation between absolute value and phase of the values. This nonreality can also be explained by the curvature of the Ewald sphere. The 3D diffraction magnitudes of a real object are centrosymmetric, whereas the Ewald sphere does not cut through both $O(\mathbf{u})$ and $O(-\mathbf{u})$.⁵⁷ In general a positivity constraint will be applicable only for the full 3D image, 2D projections (discussed in Subsection 4.C), 2D images of thin objects, and 2D images of objects with a mirror-plane symmetry.

If the object can be considered 2D and positive, a positivity constraint will have the effect of focusing the image. Usually the support constraint is loose, and even if the shape of the object is well-known or determined from the Shrinkwrap algorithm, for example, there may still be room for a defocused image to be contained within the support. The degree of defocus allowed by the support depends on how tight it is. The defocused image of a real 2D object with sharp edges or high frequencies will be real but include negative values. The focused image will be that which is most positive, and hence a positivity constraint will select that image. This is true not only for defocus but for any other low-order aberration (astigmatism, coma, etc.). That is, without the positivity constraint, there are slowly varying phase modes that cannot be determined, and the number of these modes de-

pends on how tight the support is. The same argument applies for 3D reconstruction of positive 3D objects. In this case the phase errors will be low-order 3D modes, which cannot be thought of as focus or other aberrations of an optical system but are simply unconstrained phase errors in Fourier space.

C. Two-Dimensional Projection Images

Defocus in a 2D image formed from a single diffraction pattern is a consequence of the Ewald sphere CTF, as described in Subsection 4.B. The focal plane of the image may be varied by multiplying the Fourier transform of the 2D image by a quadratic phase term. In a full 3D reconstruction, there is no concept of defocus. A shift of the object by δz along the beam axis causes the phase ramp $-2\pi\delta z u_z$ across the 3D transform. This causes a shift of the image, no different from shifts δx or δy in the other dimensions. There is no optical axis in the reconstructed 3D image, so there is no defocus. Similarly there is no defocus in a 2D projection image formed by integrating the 3D image along a specific direction. A 2D projection may be recovered from the diffraction intensities without having to first undergo a full 3D reconstruction, and we find that this is a useful step in quickly examining our 3D datasets. By the Fourier projection theorem the projection image is formed from a central section in reciprocal space; e.g., the plane $u_z=0$ gives the projection along the z axis. We have performed phase retrieval on central sections of the pyramid diffraction data by first extracting the appropriate diffraction intensities from all views. One example of a central section is shown in Fig. 4(b), which was generated by linear interpolation of measured intensities at

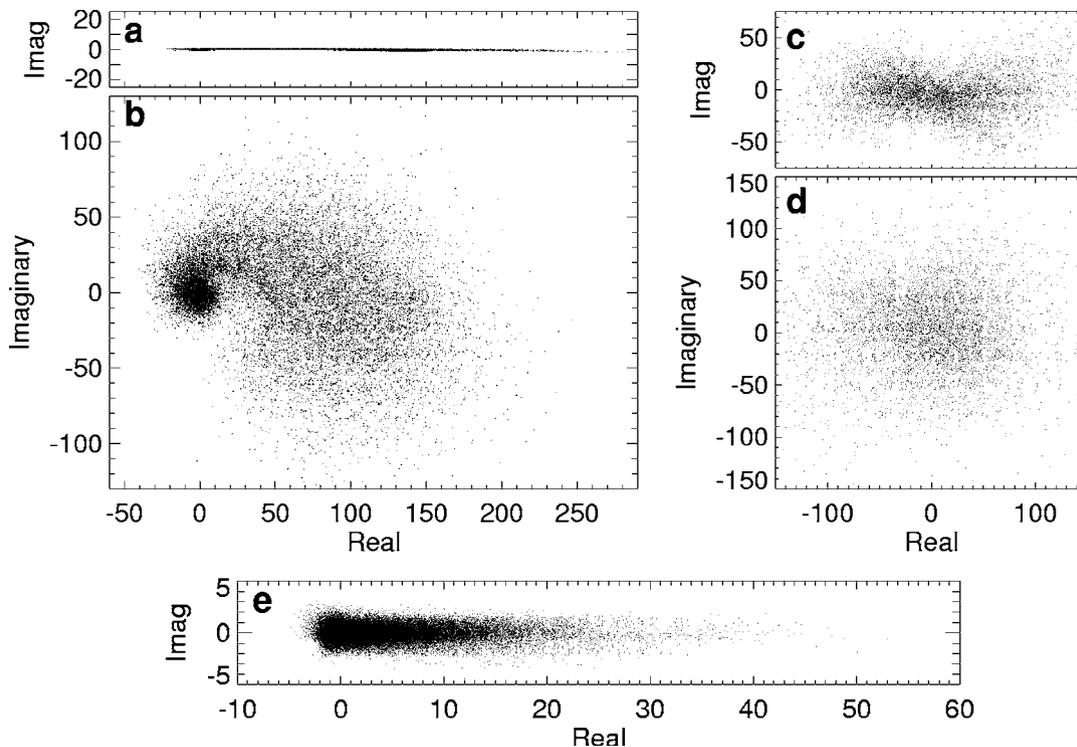


Fig. 7. Distributions of the real-space complex amplitudes $\bar{\gamma}_M$, in the Argand plane, of simulated single-view coherent images for (a) a 2D and (b) a 3D object consisting of 50-nm-diameter gold balls for an x-ray wavelength of 1.6 nm. Distributions of complex amplitudes of images reconstructed from experimental data are given for (c) the 2D projection image shown in Fig. 8, (d) the single-view 2D image of Fig. 6, and (e) the full 3D image. Cases (c) and (d) were reconstructed by using P_S , and (e) by using P_{S+} .

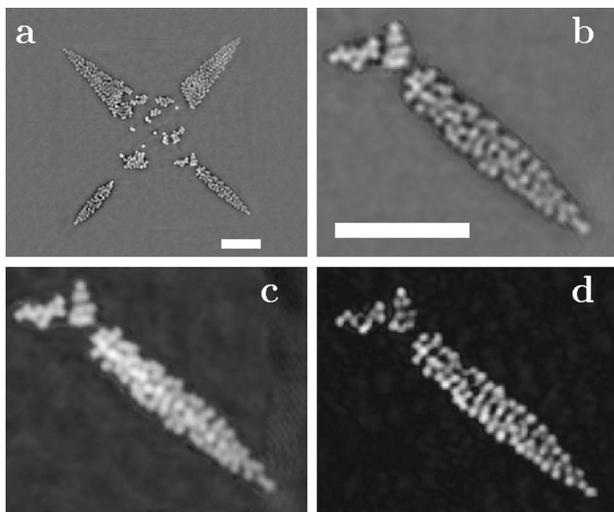


Fig. 8. Extremely large depth-of-focus projection images for the object orientation $\phi=0^\circ$. (a) Reconstruction from a 2D central section interpolated from the 3D diffraction data set. The reconstruction was performed without a positivity constraint ($E_S^2=0.167$). (b) Enlargement of the lower right arm of (a). (c) [and also Fig. 2(c)] Reconstruction from the 2D central section using a positivity constraint ($E_S^2=0.072$). (d) Projected image formed by integrating the full 3D reconstructed image ($E_S^2=0.113$). The scale bars are 500 nm.

$\mathbf{u}_{i,j,\phi}$ onto the $u_z=0$ plane. The projection images that we reconstruct from experimental data are superior to the reconstruction on a single Ewald sphere. One example is shown in Figs. 8(a) and 8(b), which can be compared with Fig. 6. In the projection images, balls at the apex of the pyramid are similar to the balls at the base, whereas in the single-view image the balls at the apex appear out of focus. The image of Figs. 8(a) and 8(b) was obtained with the Shrinkwrap algorithm (parameters given in Subsection 2.C) after first regularizing by filtering the diffraction intensities according to Eq. (19). The missing data in the arc-shaped regions seen in Fig. 4(b) were allowed to float in the reconstruction of the complex image, according to Eq. (12).

Figures 8(a) and 8(b) depict the real part of the complex image, and the distribution of complex values of the reconstructed image is given in Fig. 7(c). As compared with the distribution of complex values of a reconstructed image from a single-view diffraction pattern, given in Fig. 7(d), the values of the projection image are clustered closer to the real axis. This is as expected, since there are no defocus artifacts, and the object consists mostly of a single material (gold) as was simulated in Figs. 7(a) and 7(b). Thus we should be able to apply the reality and positivity constraints of Eq. (14) to the projection image, to further improve it and allow this extra information to help reconstruct the spatial frequency amplitudes and phases that are missing behind the beam stop. This was the case for Figs. 2(b) and 8(c), which is the real-positive constrained reconstruction from the same central section as that for Fig. 8(a). In this case the diffraction data were not filtered. This reconstruction was carried out with the same support mask derived by Shrinkwrap and used for the reconstruction of the complex image. Since they were constrained, the complex amplitudes of the image were

distributed along the real axis, with some deviation from real for smaller amplitudes that could be attributed to noise and scattering material other than gold (e.g., the silicon nitride pyramid).

5. IMAGE ANALYSIS

Both the reconstructed x-ray 2D projection image described in Subsection 4.C and the 3D image described in Subsection 4.A clearly show good agreement with the SEM image of Fig. 2(a). When we overlay a semitransparent version of the projection image of Fig. 2(b) on the SEM image of Fig. 2(a), we see that the locations of all balls visible in the SEM match with the balls visible in the x-ray image to within a pixel. In the x-ray volume image, however, we can locate more balls than are visible in the SEM image. The slice image of Fig. 2(d) reveals that the gold balls of the object are not entirely a single layer; rather, the arms of the structure are several balls deep in places. The balls were deposited on the inside surface of the silicon nitride pyramid, and it is clearly seen that these balls are indeed flush with the intersecting edges of the pyramid. The regions where the balls are layered give rise to a higher projected image intensity, which shows up as brighter regions in the projection image of Fig. 2(b). We confirm that the 3D pyramid geometry determined from the reconstructed volume image is consistent with the manufacture of the pyramid. We measure an included angle between opposite faces of the pyramid of $70^\circ \pm 1^\circ$, compared with the included angle of 70.5° between the 111 and $\bar{1}\bar{1}1$ crystal planes of silicon. The 3D image does not appear to contain the pyramid-shaped silicon nitride membrane itself, which in principle should be reconstructed as a predominantly phase object (with a phase of 1.3° per voxel). However, due to the missing data in the beam stop, we observe only the high-frequency edges of the pyramid, which are faintly visible in projection images.

The volume images display good spatial resolution in the x , y , and z dimensions. Quantifying resolution is not straightforward, since we do not have an exactly known 3D standard—the SEM shows only the surface of the object, for example, and this method cannot reveal the 3D structure. We estimate the resolution of our images by examining both their Fourier and real-space representations. In Fourier space we base measures of resolution on the signal-to-noise ratio (SNR) of measured diffraction intensities and the consistency of recovered phases, whereas in real space we fit models to 3D images of isolated balls.

A. Reconstruction Consistency and Resolution

The performance of our imaging technique could be quantified in Fourier space, in principle, by measuring the CTF. For the numerical reconstruction technique used here, this CTF would encapsulate resolution limits due to SNR, data alignment, and regions of missing data, as well as algorithm stability and uniqueness. The direct computation of the CTF would require computing the ratio of the image Fourier amplitudes to the Fourier amplitudes of the actual object, which again requires an accurate quantitative model of the actual object structure at high

resolution. Without such a model we can base an estimate of the upper limit of the CTF cutoff on the SNR of the measured diffraction data plotted in Fig. 9. The largest spatial frequency used in the interpolated 3D diffraction data set (recorded near the corner of the CCD) is at $u_{\max} = \sqrt{2N}\Delta q = 0.068 \text{ nm}^{-1}$. At this resolution shell we recorded an average of <1 photon/pixel, and obtained a SNR of 1 photon/pixel at $u = 0.062 \text{ nm}^{-1}$. (Since the noise level of our camera is considerably less than 1 photon, we assume that the noise in our diffraction patterns is determined by photon shot noise.) If we assume hypothetically that the diffraction phases are known, then the image can be fully represented, without loss of information, with a pixel sampling of $\Delta q = 1/D$, where D is the width of the object, corresponding to $s=1$, and so we could rebin our oversampled data into larger pixels with a correspondingly higher photon count. Summing in this way over pixels (referred to as pixel binning) is not the same as resampling, however, and such an operation would multiply the autocorrelation image with the Fourier transform of the summed pixel shape, which will be a function that falls from unity at the image center to $2/\pi$ at the edge of the autocorrelation image. The effect could be deconvolved from the pattern, but we avoid that by binning to a pixel sampling of $\Delta q = 1/(sD)$, with $s=2$, which is the critical sampling interval of the object's autocorrelation function. The measured data were collected at $s=4.6$, so resampling to $s=2$ gives an average of 1 photon/pixel (SNR=1) at $u = 0.066 \text{ nm}^{-1}$. If we take a measure of resolution as the frequency at which the SNR of the rebinned data is unity, then we find that the average 3D cutoff is 0.066 nm^{-1} , or a smallest resolvable half-period of 7.5 nm. This is very close to the smallest half-period of 7.3 nm limited by the detector NA.

The phase retrieval process recovers the diffraction phases with a limited accuracy, due to factors including SNR of the diffraction amplitudes, missing data, the inconsistency of constraints, and systematic errors in the data (such as errors in interpolation). These errors in phase reduce the resolution of the synthesized image.

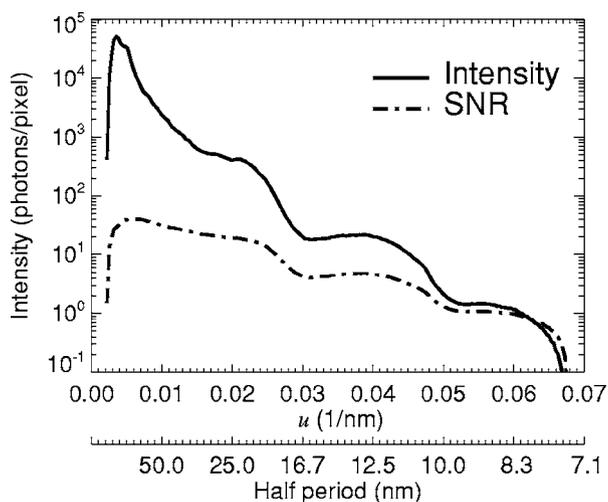


Fig. 9. 3D diffraction intensities $I(\mathbf{u})$, averaged over shells of constant u , in units of average photon count per CCD pixel. The average over constant u of the 3D SNR of the measured intensities is shown as a dotted's dashed curve.

With a complex image a loose support constraint will lead to unconstrained low-order aberrations, for example, as was discussed in Subsection 4.B. In our case of reconstructing complex 2D images, with low frequencies missing due to the beam stop, we have observed that phase retrieval from independent random starts may differ by a phase vortex (right or left handed) centered at the zero spatial frequency. This too has the effect of reducing the image resolution. One way to quantify the effect of these phase variations is to determine the correlation between phases retrieved from independent random starts of the phase retrieval algorithm. For example, we could compute the differential phase residual of these two solutions in the same way that independent images are compared in cryoelectron microscopy (Ref. 58, Chap. 3, Sect. B). Since we have the ability to compute an unlimited number of reconstructions from independent random starts, a more appropriate choice is to determine the variation in retrieved phases as a function of resolution, as suggested by Shapiro *et al.*³⁶ More specifically, the average of the independent complex reconstructions is computed, and the square of the Fourier amplitudes of this average are compared with the measured diffraction intensities. Where the phases are consistently retrieved to the same value, the squared modulus of the average will be equal to the constrained modulus and the ratio will be unity. Where the phases are random and completely uncorrelated, the average will approach zero. Thus the ratio is effectively a transfer function for the phase retrieval process, and the average image is the best estimate of the image: Spatial frequencies are weighted by the confidence with which their phases are known.³⁶ It should be possible to obtain even better estimates of the complex amplitudes by a more detailed analysis of the distributions of retrieved phases over independent trials, for example by applying patching and voting methods as described by Fienup and Wackerman.⁵⁹

All 2D and 3D images displayed in this paper are averages of more than 300 independent phase retrieval trials. That is, the best estimate of the image is given by

$$\bar{\gamma}_M = \langle \gamma_M \exp(i\phi_0) \rangle, \quad (22)$$

where $\langle \rangle$ denotes an average over independent reconstructions. Analogous to the modulus of the CTF of a coherent imaging system, we define the phase retrieval transfer function (PRTF) as

$$\text{PRTF}(\mathbf{u}) = \frac{|\mathcal{F}_{\mathbf{u}}\{\bar{\gamma}_M\}|}{\sqrt{I(\mathbf{u})}} = \frac{|\langle \Gamma_M(\mathbf{u}) \exp(i\phi_0) \rangle|}{\sqrt{I(\mathbf{u})}}, \quad (23)$$

where Γ_M is the diffraction amplitude with retrieved phases, the Fourier transform of Eq. (18). Plots of the PRTF, averaged over shells of constant u and where $I(\mathbf{u})$ are nonzero, are shown in Fig. 10(a) for the 3D image of Fig. 5 and for the 2D projection image of Fig. 8(a).

When we compute the average image $\bar{\gamma}_M$, the arbitrary multiplicative phase constant ϕ_0 of each image must be adjusted to a common value so that the random variation of this constant does not reduce the average, which would result in a low value of the transfer function. We do this for the first reconstructed image $\gamma_M^{(0)}$ by finding the con-

stant phase that maximizes the real part of that image, which can be achieved by finding the value ϕ_0 that maximizes⁶⁰

$$\begin{aligned}\alpha &= \sum_k [\Re\{\gamma_M^{(0)}(k)\exp(i\phi_0)\}]^2 \\ &= \sum_k \{2|\gamma_M^{(0)}(k)|^2 + [\gamma_M^{(0)}(k)]^2 \exp(2i\phi_0) \\ &\quad + [\gamma_M^{(0)}(k)]^{*2} \exp(-2i\phi_0)\}/4\end{aligned}\quad (24)$$

for an image with pixels (or voxels) indexed by integers k and complex values $\gamma_M(k)$. We maximize the square of the real part to allow for positive and negative real values. The value α can be maximized by maximizing either the second or third term of Eq. (24), and we do so by finding the phase ϕ of the complex value $\sum_k [\gamma_M^{(0)}(k)]^2$ and setting $\phi_0 = -\phi/2$. The subsequent images γ_M are adjusted by finding the constant phase ϕ_1 that minimizes $\sum_k |\gamma_M^{(0)}(k) - \gamma_M(k)|^2$. This phase is that which maximizes $\sum_k \Re\{\gamma_M^{(0)}(k) \times (k) \gamma_M(k) \exp(i\phi_1)\}$, which is simply the phase of the complex value $\sum_k [\gamma_M^{(0)}(k)]^* \gamma_M(k)$.

In the case of 2D images we also improve the average by separating out the vortex modes mentioned above. This was achieved by using the voting method,⁵⁹ simply by correlating each phase retrieval solution with the previous solutions and separating the solutions into three classes (which were found to differ by left- and right-handed phase vortices) based on the value of the correlation. We found that the class with the most frequent members (60% of trials) gave rise to the best image, whereas the other two classes were equally frequent (20% each) and gave rise to images for which the balls were larger, had bright edges, and showed reduced intensity at their centers. Based on the appearance of the balls, we assumed that the most frequent class was that which did not have a vortex mode. The effect of removing the vortex modes from the average image is shown in Fig. 10(a). As is seen in that figure, the PRTF is uniformly increased across all frequencies. This is due to the fact that the left- and right-handed vortex modes sum together to give a variation of the modulus, which varies as the cosine of the azimuthal angle in the diffraction pattern and which averages to zero in the average around this complete circle for each u .

The resolution cutoff of the phase retrieval process can be given by the spatial frequency at which the PRTF extrapolates to zero. For all cases here, this frequency is greater than the measurement cutoff of $u_{\max} = 0.068 \text{ nm}^{-1}$, or half-period resolution length of 7.4 nm. A more conservative estimate of the resolution is given by the frequency at which the PRTF reaches a value of 0.5. For the vortex-corrected 2D reconstruction this occurs just at u_{\max} , but for the 3D image this corresponds to 0.048 nm^{-1} , or a half-period resolution length of 10.4 nm. In this case the average resolution cutoff is worse than in the 2D case because the 3D PRTF is diminished along the u_z direction where the diffraction data are missing (which reduces the average over the u shell). This is illustrated in Fig. 10(b), where we display the 3D PRTF as a surface for which it has a value of 0.5. The PRTF is not defined in the regions of missing diffraction data, which are seen as

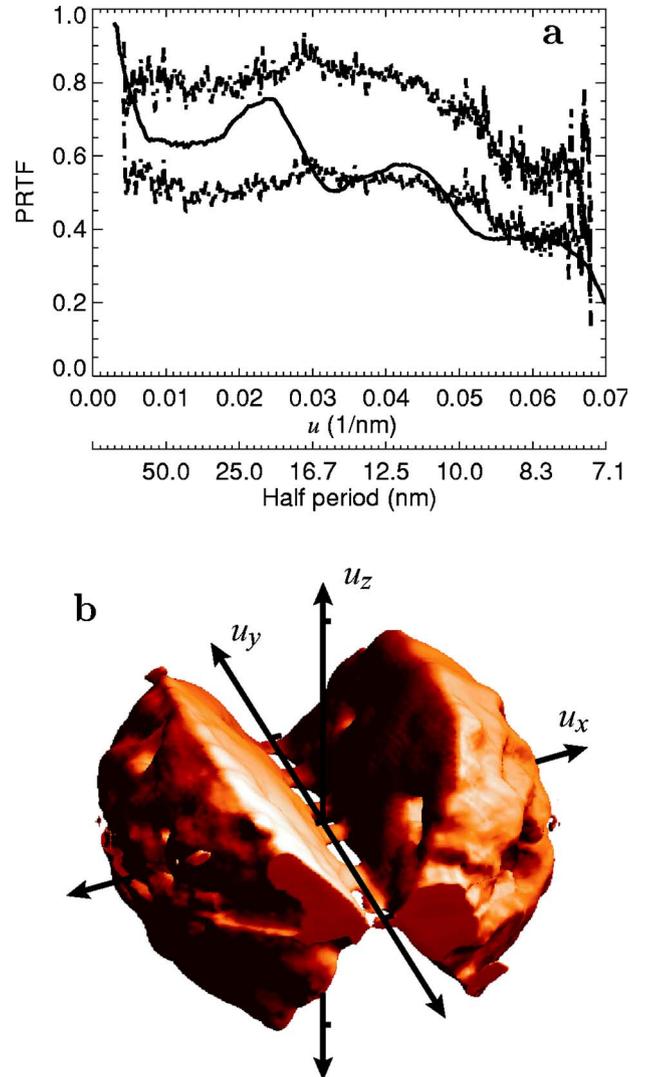


Fig. 10. (Color online) (a) Phase retrieval transfer function (PRTF), averaged over shells of constant u , for the real-positive 3D projection image (solid curve) and averaged over circles of constant u for the complex 2D image (dotted-dashed curves). The dotted-dashed curve with lower values is for the 2D projection image without correction of vortex phase modes. (b) Isosurface rendering of the 3D PRTF at a threshold level of 0.5. The tick marks on the u_y and u_z axes indicate 0.05 nm^{-1} .

the missing wedges in the surface. It is seen that the resolution is approximately the same in all directions of \mathbf{u} where intensities were measured.

When applied to the average image $\bar{\gamma}_M$, the modulus constraint error E_M^2 of Eq. (17) is equal to the intensity-weighted integral over \mathbf{u} of $|1 - \text{PRTF}(\mathbf{u})|^2$. That is, it gives a single measure of how well diffraction intensities of the average image agree with the measurement. This is generally higher than the metric E_M^2 applied to the iterate g_n , which gives an estimate for how well the algorithm fits the intensity data. The value of E_M^2 applied to the average 3D image is 0.368, and 0.059 for the average 2D projection image that was corrected for vortex phase errors (0.312 without vortex correction). We expect that a similar correction of low-order phase modes in the 3D image would lead to a similar improvement in the error metric, and the relatively high value of E_M^2 for the average 3D im-

age is a result of the overall filtering due to the variation of these low-order phase modes.

We can also compute the agreement of the average image $\bar{\gamma}_M$ with the real-space support constraint E_S^2 of Eq. (16). We find a value of 0.228 when the computation is applied to the average 3D image and 0.167 for the average 2D complex-valued projection image reconstructed from the central section. Note, however, that in the 3D image the support S accounts for 0.10% of the image voxels whereas S covers 4.1% of the pixels in the projection 2D images, and so the average error per pixel outside the support is much less for the 3D than for the 2D reconstruction. We find that, with the addition of the real-space positivity constraint, E_S^2 of the average 2D projection image decreases from 0.167 to 0.072. However, in this case the modulus constraint error E_M^2 increases from 0.059 to 0.172.

B. Real-Space Resolution

The measures of resolution from the SNR and PRTF reveal the effects of noise, consistency of the diffraction data, and how well the image obeys the imposed constraints. These measures are contributors to the overall image resolution. A direct measure of a lower limit of resolution can be obtained by examining the images of isolated and closely spaced gold balls. Lineouts of the isolated ball located on the lower left arm of Fig. 5(a) are shown in Fig. 11 for all three orthogonal directions. The ball image has full widths at half-maximum (FWHM) of 30, 35, and 70 nm in the x , y , and z directions, respectively. Images of other isolated balls in the object are very similar to those shown in Fig. 11. Assuming that the balls are 50 nm in diameter, we obtain a good fit to the images by modeling a coherent imaging system with a CTF that is unity within a cube of half-width 0.05 nm^{-1} (centered at the zero frequency) and that is zero within a sector of 60° as rotated about the y axis and centered about the z axis. Lineouts of the modeled coherent images, computed by convolving an isolated 50-nm sphere with the Fourier transform of the CTF (that is, the point spread function, or PSF), are shown as dotted curves in Fig. 11. The FWHMs of the modeled coherent image are 36, 40, and 64 nm in the x , y , and z directions, respectively, in good agreement with the reconstructed image. We do not ex-

pect the model to be an exact fit to the data, since the actual PSF is more complicated and depends on the details of the phase retrieval, which is better characterized by the PRTF in Fig. 10. However, the fits are reasonable, and the widths of the modeled PSF are in good agreement with the measures of resolution obtained from analysis of the diffraction intensities and recovered phases. The modeled PSF, given by the Fourier transform of the CTF, has a half-width of $10 \text{ nm} \times 10 \text{ nm} \times 40 \text{ nm}$. Here the half-width is defined as the distance from the central maximum of the PSF to the first zero. Since the imaging process is coherent, the image width depends on the phase of the PSF, which has different distributions for the x and y directions. This explains the variations of image width in the x and y directions and why the image FWHMs in these directions are in fact smaller than the ball diameter. As expected, the resolution is much worse in the z direction than in the x and y directions, due to the missing sector of data that arises from recording diffraction over a limited range of angles.

6. SUMMARY AND CONCLUSIONS

We have presented 3D images reconstructed *ab initio* from coherent x-ray diffraction, which exhibit high resolution in all three dimensions. These images are the highest-resolution 3D images of nonperiodic objects where the resolution is comparable in all three dimensions. The work presented here marks an important advance in that we have fully demonstrated the experimental methods to collect 3D coherent x-ray diffraction and the computational tools to manage the data, and to reconstruct 3D images of more than 10^9 voxels.

The coherent x-ray diffraction recorded from our 3D test object comprised 140 views, at 1° intervals, and extended to a maximum spatial frequency of 0.068 nm^{-1} , or a smallest reconstructible half-period of 7.4 nm. Although we cannot exactly quantify the resolution of the image, which would require knowing the object's 3D structure, we have determined the consistency of the retrieved phases, which gives us an estimate of an upper bound of the CTF of the imaging process. Our analysis shows that we can consistently retrieve phases out to the maximum spatial frequency recorded. This consistency measure

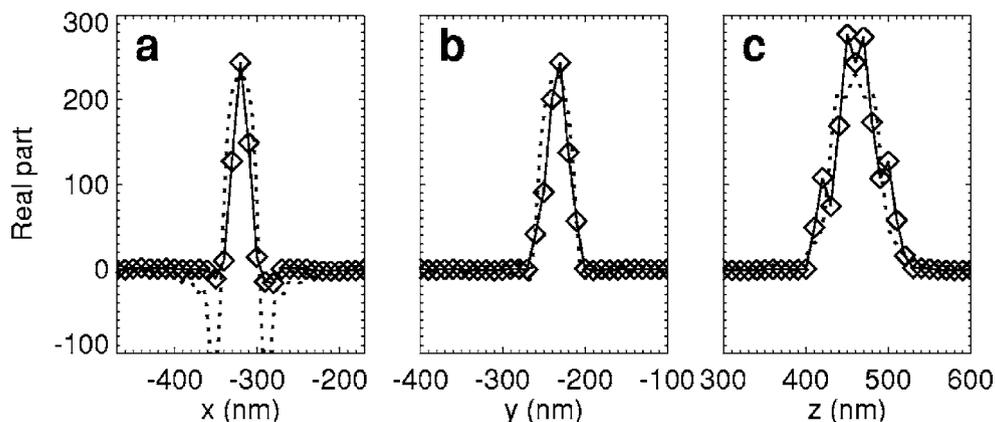


Fig. 11. Lineouts of the real part of the reconstructed complex amplitude 3D image, for three orthogonal directions (a) x , (b) y , and (c) z , through the isolated single ball at the pyramid apex. Coordinates are relative to the center of the 3D image array. Dotted curves show lineouts from a simulated 3D coherent image with a cube CTF with a 60° missing sector.

does not tell us anything about systematic errors, such as interpolation of the data, errors in assigning spatial frequency \mathbf{u} to the intensities (imperfect knowledge of the beam center), and missing data due to the beam stop or limited range of object orientations. However, we easily resolve 50-nm spheres that are touching each other, and from such image lineouts, and comparisons of reconstructed x-ray images with the SEM image, we have confidence that our achieved image resolution is close to our upper estimate.

We have found that our Shrinkwrap algorithm,⁶ which determines the object support *ab initio*, is robust and works well even with missing Fourier-space data due to limited object orientations or the beam stop. The phase retrieval process can be essentially characterized by a 3D transfer function (the PRTF), which is influenced by the noise of the measured diffraction intensities. While the algorithm also allows the amplitudes at the locations of missing data to be recovered, these values are not consistently reconstructed and are averaged to zero, leaving worse resolution in the depth (z) direction. We expect that with a data set collected over the full range of sample orientation angles we would achieve equal resolution in all three dimensions. As it is, we obtained an estimate of 10 nm in x and y and 50 nm in z .

We have shown that high-NA x-ray coherent imaging of thick objects can properly be carried out only in the context of three dimensions. Here we define high-NA imaging of thick objects to be imaging under conditions that lead to a depth of focus less than the depth of the object in any of its orientations. Since the imaging is coherent, a 2D image of a thick object in any view will exhibit defocus artifacts that do not diminish in overall power with the degree of defocus and that lead to difficulties in the interpretation of the image. In addition, these artifacts cause the image of a real-positive object, for example, to be complex, hence hampering quantitative evaluation of the image. 2D images free from defocus artifacts can be quickly generated from central sections extracted from the diffraction data. 3D images are synthesized from the entire 3D diffraction data set. The tools are now in place to perform full 3D reconstructions of thick samples. Currently we have reconstructed arrays with almost 2×10^9 elements. If the minimum oversampling of $\sqrt[3]{2}$ is used in each dimension, then this would correspond to objects of width $9.5 \mu\text{m}$ at a pixel spacing of 10 nm, or a half-period resolution length of 7 nm along the diagonal. When single-particle x-ray free-electron laser imaging at atomic resolution becomes feasible, then these demonstrated computational capabilities could be used to reconstruct objects of 480-nm width at 0.7-nm resolution, for example. This would correspond to a large virus, or a large protein complex such as the ribosome.

ACKNOWLEDGMENTS

We thank Ray Mariella (LLNL) for the idea of using a silicon nitride pyramid as a test object and Jackie Crawford and Dino Ciarlo (LLNL) for determining its manufacturing process. We thank Janos Kirz (LBNL and Stony Brook) for technical advice and extensive discussions about our experiments. We acknowledge stimulating dis-

cussions with Abraham Szöke (LLNL), Gösta Hultdt (Uppsala Universitet), and Eugene Ingerman (CBST). We gratefully acknowledge Richard Crandall and the Advanced Computations Group (Apple Computer, Inc.) for the development of the `dist_fft` software. We acknowledge the support of the Advanced Light Source (ALS) staff. This work was performed in part under the auspices of the U.S. Department of Energy by University of California, Lawrence Livermore National Laboratory, under contract W-7405-Eng-48. The work of the Lawrence Berkeley National Laboratory participants and the operation of the ALS facility was supported by the Director, Office of Energy Research, Office of Basic Energy Sciences, Materials Sciences Division, U.S. Department of Energy, under contract DE-AC03-76SF00098. This work has been supported by funding from the National Science Foundation. The Center for Biophotonics, an NSF Science and Technology Center, is managed by the University of California, Davis, under cooperative agreement PHY 0120999. This work was also partly supported by the Department of Energy under grant DEFG0202ER45996.

Corresponding author e-mail address:
henry.chapman@llnl.gov.

REFERENCES

1. D. Sayre and H. N. Chapman, "X-ray microscopy," *Acta Crystallogr., Sect. A: Found. Crystallogr.* **51**, 237–252 (1995).
2. D. Sayre, H. N. Chapman, and J. Miao, "On the extendibility of x-ray crystallography to noncrystals," *Acta Crystallogr., Sect. A: Found. Crystallogr.* **54**, 232–239 (1998).
3. J. Miao, P. Charalambous, J. Kirz, and D. Sayre, "Extending the methodology of x-ray crystallography to allow imaging of micrometre-sized non-crystalline specimens," *Nature (London)* **400**, 342–344 (1999).
4. I. K. Robinson, I. A. Vartanyants, G. J. Williams, M. A. Pfeifer, and J. A. Pitney, "Reconstruction of the shapes of gold nanocrystals using coherent x-ray diffraction," *Phys. Rev. Lett.* **87**, 195505 (2001).
5. G. J. Williams, M. A. Pfeifer, I. A. Vartanyants, and I. K. Robinson, "Three-dimensional imaging of microstructure in Au nanocrystals," *Phys. Rev. Lett.* **90**, 175501 (2003).
6. S. Marchesini, H. He, H. N. Chapman, S. P. Hau-Riege, A. Noy, M. R. Howells, U. Weierstall, and J. C. H. Spence, "X-ray image reconstruction from a diffraction pattern alone," *Phys. Rev. B* **68**, 140101 (2003).
7. S. Marchesini, H. N. Chapman, S. P. Hau-Riege, R. A. London, A. Szöke, H. He, M. R. Howells, H. Padmore, R. Rosen, J. C. H. Spence, and U. Weierstall, "Coherent x-ray diffractive imaging: applications and limitations," *Opt. Express* **11**, 2344–2353 (2003).
8. H. He, S. Marchesini, M. Howells, U. Weierstall, H. Chapman, S. Hau-Riege, A. Noy, and J. C. H. Spence, "Inversion of x-ray diffuse scattering to images using prepared objects," *Phys. Rev. B* **67**, 174114 (2003).
9. M. R. Howells, T. Beetz, H. N. Chapman, C. Cui, J. M. Holton, C. J. Jacobsen, J. Kirz, E. Lima, S. Marchesini, H. Miao, D. Sayre, D. A. Shapiro, and J. C. H. Spence, "An assessment of the resolution limitation due to radiation-damage in x-ray diffraction microscopy," arxiv.org e-print archive, physics/0502059, February 11, 2005, <http://arxiv.org/pdf/physics/0502059>.
10. R. Neutze, R. Wouts, D. van der Spoel, E. Weckert, and J. Hajdu, "Potential for biomolecular imaging with femtosecond x-ray pulses," *Nature (London)* **406**, 753–757 (2000).
11. J. Miao, K. O. Hodgson, and D. Sayre, "An approach to three-dimensional structures of biomolecules by using

- single-molecule diffraction images," *Proc. Natl. Acad. Sci. U.S.A.* **98**, 6641–6645 (2001).
12. J. C. H. Spence and R. B. Doak, "Single molecule diffraction," *Phys. Rev. Lett.* **92**, 198102 (2004).
 13. W. S. Haddad, I. McNulty, J. Trebes, E. Anderson, R. Levesque, and L. Yang, "Ultra-high-resolution x-ray tomography," *Science* **266**, 1213–1215 (1994).
 14. D. Weiss, G. Schneider, B. Niemann, P. Guttman, D. Rudolph, and G. Schmahl, "Computed tomography of cryogenic biological specimens based on x-ray microscopic images," *Ultramicroscopy* **84**, 185–197 (2000).
 15. C. A. Larabell and M. A. Le Gros, "X-ray tomography generates 3-D reconstructions of the yeast, *Saccharomyces cerevisiae*, at 60-nm resolution," *Mol. Biol. Cell* **15**, 957–962 (2004).
 16. N. Streibl, "Three-dimensional imaging by a microscope," *J. Opt. Soc. Am. A* **2**, 121–127 (1985).
 17. T. Beetz, M. Howells, C. Jacobsen, C. Kao, J. Kirz, E. Lima, T. Mentess, H. Miao, C. Sanchez-Hanke, D. Sayre, and D. Shapiro, "Apparatus for x-ray diffraction microscopy and tomography of cryo specimens," *Nucl. Instrum. Methods Phys. Res. A* **545**, 459–468 (2005).
 18. M. R. Howells, P. Charalambous, H. He, S. Marchesini, and J. C. H. Spence, "An off-axis zone-plate monochromator for high-power undulator radiation," in *Design and Microfabrication of Novel X-Ray Optics*, D. C. Mancini, ed., *Proc. SPIE* **4783**, 65–73 (2002).
 19. C. Giacovazzo, *Direct Phasing in Crystallography* (Oxford U. Press, 1998), p. 468.
 20. R. Crandall, E. Jones, J. Klivington, and D. Kramer, "Gigaelement FFTs on Apple G5 clusters," Advanced Computation Group, Apple Computer (2004), http://images.apple.com/acg/pdf/20040827_GigaFFT.pdf.
 21. M. Born and E. Wolf, *Principles of Optics*, 7th ed. (Cambridge U. Press, 2002).
 22. R. W. James, *The Optical Principles of the Diffraction of X-Rays* (Bell, 1962).
 23. E. Wolf, "Three-dimensional structure determination of semi-transparent objects from holographic data," *Opt. Commun.* **1**, 153–156 (1969).
 24. J. Kirz, C. Jacobsen, and M. Howells, "Soft x-ray microscopes and their biological applications," *Q. Rev. Biophys.* **28**, 33–130 (1995).
 25. F. Natterer, "An error bound for the Born approximation," *Inverse Probl.* **20**, 447–452 (2004).
 26. R. N. Bracewell, *The Fourier Transform and Its Applications*, 2nd ed. (McGraw-Hill, 1986).
 27. J. Miao, T. Ishikawa, E. H. Anderson, and K. O. Hodgson, "Phase retrieval of diffraction patterns from noncrystalline samples using the oversampling method," *Phys. Rev. B* **67**, 174104 (2003).
 28. J. Miao, D. Sayre, and H. N. Chapman, "Phase retrieval from the magnitude of the Fourier transforms of nonperiodic objects," *J. Opt. Soc. Am. A* **15**, 1662–1669 (1998).
 29. J. W. Goodman, *Introduction to Fourier Optics*, 2nd ed. (McGraw-Hill, 1996).
 30. J. C. H. Spence, U. Weierstall, and M. Howells, "Coherence and sampling requirements for diffractive imaging," *Ultramicroscopy* **101**, 149–152 (2004).
 31. R. Crowther, D. DeRosier, and A. Klug, "The reconstruction of a three-dimensional structure from its projections and its applications to electron microscopy," *Proc. R. Soc. London* **A317**, 319–340 (1970).
 32. E. Salerno, "Superresolution capabilities of the Gerchberg method in the band-pass case: an eigenvalue analysis," *Int. J. Imaging Syst. Technol.* **90**, 181–188 (1998).
 33. Y. Nishino, J. Miao, and T. Ishikawa, "Image reconstruction of nanostructured nonperiodic objects only from oversampled hard x-ray diffraction intensities," *Phys. Rev. B* **68**, 220101 (2003).
 34. T. Sato, S. J. Norton, M. Linzer, O. Ikeda, and M. Hiramama, "Tomographic image reconstruction from limited projections using iterative revisions in image and transform spaces," *Appl. Opt.* **20**, 395–399 (1981).
 35. M. Bertero and E. R. Pike, "Resolution in diffraction-limited imaging, a singular value analysis. I. The case of coherent illumination," *Opt. Acta* **29**, 727–746 (1982).
 36. D. Shapiro, P. Thibault, T. Beetz, V. Elser, M. Howells, C. Jacobsen, J. Kirz, E. Lima, H. Miao, A. M. Neimann, and D. Sayre, "Biological imaging by soft x-ray diffraction microscopy," *Proc. Natl. Acad. Sci. U.S.A.* **102**, 15343–15346 (2005).
 37. A. Szoke, H. Szoke, and J. R. Somoza, "Holographic methods in x-ray crystallography. V. Multiple isomorphous replacement, multiple anomalous dispersion and noncrystallographic symmetry," *Acta Crystallogr., Sect. A: Found. Crystallogr.* **53**, 291–313 (1997).
 38. S. P. Hau-Riege, H. Szoke, H. N. Chapman, A. Szoke, S. Marchesini, A. Noy, H. He, M. Howells, U. Weierstall, and J. C. H. Spence, "SPEDEN: reconstructing single particles from their diffraction patterns," *Acta Crystallogr., Sect. A: Found. Crystallogr.* **60**, 294–305 (2004).
 39. D. Potts, G. Steidl, and M. Tasche, "Fast Fourier transforms for nonequispaced data: a tutorial," in *Modern Sampling Theory: Mathematics and Applications*, J. J. Benedetto and P. Ferreira, eds. (Springer, 2001), Chap. 12, pp. 249–274.
 40. H. Choi and D. C. Munson, Jr., "Direct-Fourier reconstruction in tomography and synthetic aperture radar," *Int. J. Imaging Syst. Technol.* **9**, 1–13 (1998).
 41. F. Natterer, *The Mathematics of Computerized Tomography* (SIAM, 2001).
 42. A. J. Devaney, "A filtered backpropagation algorithm for diffraction tomography," *Ultrason. Imaging* **4**, 336–350 (1982).
 43. S. Pan and A. Kak, "A computational study of reconstruction algorithms for diffraction tomography: interpolation versus filtered-backpropagation," *IEEE Trans. Acoust., Speech, Signal Process.* **SP-31**, 1262–1275 (1983).
 44. H. He, S. Marchesini, M. Howells, U. Weierstall, G. Hembree, and J. C. H. Spence, "Experimental lensless soft-x-ray imaging using iterative algorithms: phasing diffuse scattering," *Acta Crystallogr., Sect. A: Found. Crystallogr.* **59**, 143–152 (2003).
 45. J. R. Fienup, "Reconstruction of a complex-valued object from the modulus of its Fourier transform using a support constraint," *J. Opt. Soc. Am. A* **4**, 118–123 (1987).
 46. J. R. Fienup, "Phase retrieval and support estimation in x-ray diffraction," in *Coherence 2005: International Workshop on Phase Retrieval and Coherent Scattering* (ESRF, 2005), <http://www.esrf.fr/NewsAndEvents/Conferences/Coherence2005/Proceedings/files/Talks/Fienup.pdf>.
 47. T. R. Crimmins, J. Fienup, and B. J. Thelen, "Improved bounds on object support from autocorrelation support and application to phase retrieval," *J. Opt. Soc. Am. A* **7**, 3–13 (1990).
 48. J. R. Fienup, "Phase retrieval algorithms: a comparison," *Appl. Opt.* **21**, 2758–2769 (1982).
 49. V. Elser, "Phase retrieval by iterated projections," *J. Opt. Soc. Am. A* **20**, 40–55 (2003).
 50. D. R. Luke, "Relaxed averaged alternating reflections for diffraction imaging," *Inverse Probl.* **21**, 37–50 (2005).
 51. M. Frigo and S. G. Johnson, "The design and implementation of FFTW3," *Proc. IEEE* **93**, 216–231 (2005). (Special issue on "Program Generation, Optimization, and Platform Adaptation").
 52. "The message passing interface (MPI) standard," <http://www-unix.mcs.anl.gov/mpl/>.
 53. A. J. Ladd, J. H. Kinney, D. L. Haupt, and S. A. Goldstein, "Finite-element modeling of trabecular bone: comparison with mechanical testing and determination of tissue modulus," *J. Orthop. Res.* **16**, 622–628 (1998).
 54. "Movies of three-dimensional diffraction data and reconstruction of the pyramid test object," http://als.lbl.gov/esg_beamlines/cxdi.
 55. J. M. Cowley, *Diffraction Physics* (North-Holland, 1981).
 56. J. C. H. Spence, U. Weierstall, and M. Howells, "Phase recovery and lensless imaging by iterative methods in

- optical, x-ray and electron diffraction," *Philos. Trans. R. Soc. London, Ser. A* **360**, 875–895 (2002).
57. G. Huldt, Biomedical Centre, Uppsala Universitet (personal communication, 2005).
58. J. Frank, *Three-Dimensional Electron Microscopy of Macromolecular Assemblies* (Academic, 1996).
59. J. R. Fienup and C. Wackerman, "Phase-retrieval stagnation problems and solutions," *J. Opt. Soc. Am. A* **3**, 1897–1907 (1986).
60. J. R. Fienup, "Invariant error metrics for image reconstruction," *Appl. Opt.* **36**, 8352–8357 (1997).