Understanding the twin-image problem in phase retrieval

Manuel Guizar-Sicairos[1,2,*] and James R. Fienup[1]

[1]*The Institute of Optics, University of Rochester, Rochester, New York 14627, USA*
[2]*Presently at Paul Scherrer Institut, Villigen PSI 5232, Switzerland*
**Corresponding author: manuel.guizar-sicairos@psi.ch*

Received June 6, 2012; accepted September 6, 2012;
posted September 25, 2012 (Doc. ID 169975); published October 18, 2012

The twin-image problem in phase retrieval is characterized by the simultaneous occurrence of features from the original object and its inversion about the origin (twin image). This problem can occur in reconstructions for which the object support is centrosymmetric or loose, and in severe cases it can greatly hinder image quality. In this paper we examine this problem and find that it arises when the retrieved Fourier-domain phase is divided into sets of regions, some of which reconstruct the object while others the twin. We examine sample reconstructions that present the twin-image problem to different extents and find that, even when the twin-image problem is not visually evident, it can exist in small regions of the retrieved Fourier phase. The reduced-support constraint approach is shown to be effective in escaping stagnation caused by the twin-image problem.

OCIS codes: 100.5070, 100.3010, 110.7440, 070.0070.

1. INTRODUCTION

Phase retrieval algorithms are able to recover the phase of a coherent wave given an adequate measurement (or measurements) of its intensity distribution and some *a priori* known information about the wave. Compared to conventional imaging, phase retrieval techniques offer a great relaxation of experimental requirements at the expense of increased computations and have found important applications. For example, in optical wavefront measurement, the aberrations of a wavefront can be recovered without the need of a well-calibrated interferometer [1–4].

Phase retrieval also provides a route to perform coherent lensless imaging for cases where a suitable imaging optic or holographic reference wave cannot be used [5]. The object of interest is illuminated with a coherent wave, the transmitted (or backscattered) field propagates, often to the far field, and the intensity of the resulting field is measured with an intensity detector array. At optical wavelengths, this allows high-resolution imaging using a conformal, lightweight imaging system whose axial thickness is not increased with increasing numerical aperture [6,7], and at x-ray wavelengths it allows imaging small specimens with high resolution without the need of imaging optics [8–10]. Lensless imaging, by measuring the intensity of an x-ray diffraction pattern, is often referred to as coherent diffractive imaging (CDI). In both optical and x-ray applications, the intensity is typically measured in the far-field regime, and the object and measurement plane are related, to a good approximation, by a Fourier transform (FT) [11].

In practice we approximate the propagation to the far field by the discrete FT. The problem is then to recover the (in general complex-valued) object, $f(x,y)$, from its Fourier intensity, $|F(u,v)|^2$, where

$$F(u,v) = |F(u,v)|\exp[i\phi(u,v)] = \frac{1}{\sqrt{MN}}\sum_{x,y} f(x,y)\exp\left[-i2\pi\left(\frac{ux}{M}+\frac{vy}{N}\right)\right], \quad (1)$$

or equivalently to retrieve the Fourier phase, $\phi(u,v)$, from its magnitude, $|F(u,v)|$. For x-ray CDI, an object support constraint, i.e., a set of points, S, outside of which the object is known to be zero,

$$f(x,y) = 0, \quad \forall\ (x,y) \notin S, \quad (2)$$

is commonly used in the reconstruction. Phase retrieval algorithms that use this constraint and are reasonably robust to stagnation have been developed and successfully used [12–15]. If a nonnegativity constraint can be used, the latter-mentioned algorithms are also robust against having a loose support constraint, i.e., a support constraint larger than the true support. However, when the object is complex valued, reconstruction with a loose support is more difficult [6,16–18]. For cases when the support constraint is not known *a priori*, it can be estimated from the object's autocorrelation [19–21], which is obtained directly from the intensity measurement, or it can be iteratively refined [10]. When the support is estimated, it typically results in an upper bound on the support rather than the true object support; hence, in practical situations one must often deal with a loose support constraint.

Image reconstruction by phase retrieval using a support constraint suffers from a few inherent ambiguities that arise from the loss of the phase in the Fourier plane. For example, a change in the transverse position of the object would only cause the addition of a linear phase in the measurement (Fourier) plane. Because the phase in the Fourier domain is lost in the measured intensity, there is no information about the absolute transverse position of the object. The

reconstructed image will be bounded by the support constraint, and if we were to translate the support prior to reconstruction, the reconstructed image will also be translated.

Another inherent ambiguity arises if the object has a centrosymmetric support. Notice that the object, $f(x,y)$, and its twin, $f^*(-x,-y)$, have the same Fourier intensity, since the FT of $f^*(-x,-y)$ is

$$F^*(u,v) = |F(u,v)| \exp[-i\phi(u,v)]. \quad (3)$$

If the support is centrosymmetric, then both fit inside the support, and either one is a valid solution of the phase retrieval problem. This is true even for a noncentrosymmetric support if the object and its twin both fit inside a loose support constraint. Because of this inherent ambiguity, and because image quality is not hindered (they appear the same except for a 180° rotation), reconstruction of either $f(x,y)$ or $f^*(-x,-y)$ is considered successful for phase retrieval. Notice that this ambiguity is similar to that of off-axis holography [11,22]. Aside from the translation and twin-image ambiguity, it is thought that phase retrieval in two or higher dimensions is usually unique [23].

The twin-image problem in phase retrieval, earlier described by Fienup and Wackerman [24], is a stagnation mode that is characterized by the simultaneous appearance of features of the upright and twin images. This problem can occur when the object support is centrosymmetric or when both the object and twin can fit inside the support, and for severe cases it can significantly hinder image quality.

The twin-image problem is a persistent stagnation mode; i.e., continuing with iterations or randomly perturbing the retrieved phases rarely leads to escaping it, and the frequency with which it appears depends not only on the shape of the support but also on the object. It commonly appears when attempting reconstruction of objects that have an approximately centrosymmetric support. The latter can occur in x-ray CDI, for example, when reconstructing images of cells, viruses or other biological specimens [9,25,26]. It should be noted that a linear combination of the upright and twin images does not satisfy the Fourier intensity measurement [24] and until now, the nature of the twin-image problem remained not fully explained.

The twin-image problem can be avoided if one has the ability to fabricate an object having a noncentrosymmetric support. Alternatively, making a multiplicity of measurements after a known change in the experimental configuration (diversity of measurement) has also shown to reduce or eliminate this problem. Examples of the latter are focus diversity [1,3], piston diversity for segmented mirrors [4], and transverse translation diversity [27–29]. In many cases, however, it is not possible to control the object support, or it is impractical to physically obtain a diversity of measurements, or the object may move or change before diverse measurements can be made. Hence, it remains important to understand the twin-image problem.

In this paper we show that the twin-image problem arises because complementary regions in the Fourier domain reconstruct either the upright or twin image. As the iterations progress, the Fourier domain divides into two types of regions: in some, the phase $\phi(u,v)$ associated with $F(u,v)$ and $f(x,y)$ is approximately retrieved, whereas in others $-\phi(u,v)$, associated with $F^*(u,v)$ and $f^*(-x,-y)$. We illustrate this phenomenon and demonstrate the effectiveness of a technique to overcome it.

2. NUMERICAL SIMULATIONS

We have performed a numerical simulation to illustrate the twin-image problem. The amplitude and phase of a complex-valued object, contained in a circular support of 120 pixels diameter and embedded in a 256×256 array, are shown in Figs. 1(a) and 1(b), respectively. We purposely chose a complex-valued object having a centrosymmetric support, since that case is more prone to stagnation than a real-valued, nonnegative object or one having a nonsymmetric support. Making the complete extent of the support less than half of the computational window ensures that the Fourier intensity is at least Nyquist sampled. The magnitude and phase of the object FT are shown in Figs. 1(c) and 1(d), respectively. The Fourier magnitude shown in Fig. 1(c), along with knowledge of the object support (the circular shape), is the input to the phase retrieval algorithm. The solution sought by the phase retrieval algorithm is the object in Figs. 1(a) and 1(b) or equivalently the Fourier phase, $\phi(u,v)$, shown in Fig. 1(d).

For the reconstruction, we used a combination of the hybrid input–output (HIO) and error-reduction [12] versions of the iterative transform algorithm (ITA). For one series of iterations, we performed 45 iterations of HIO with $\beta = 0.7$ to explore the solution space followed by five iterations of error reduction to help the current reconstruction settle down. We performed 20 series of iterations (for a total of 1000 iterations) to give the algorithm plenty of time to either find a solution or to stagnate.

Because of its centrosymmetric support, this particular object was conductive to the twin-image problem. From 20 reconstructions performed from different random starting

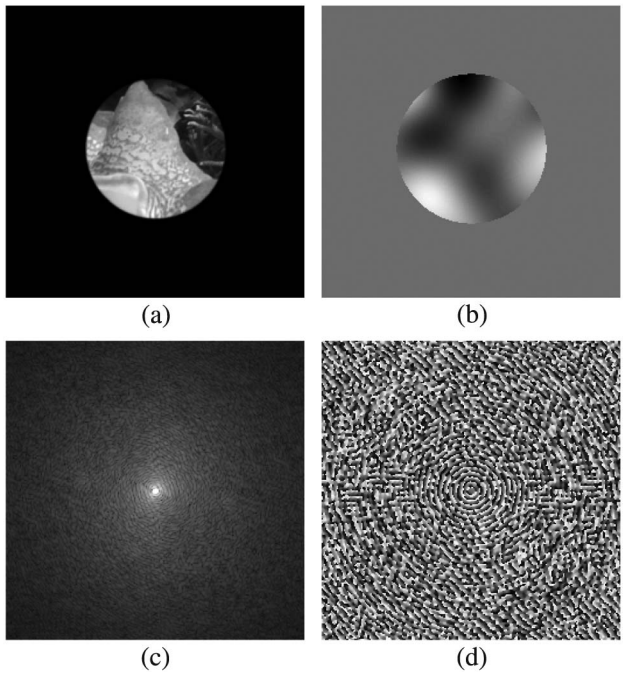

Fig. 1. Object (a) magnitude and (b) phase; Fourier (c) magnitude and (d) phase. Because of its large dynamic range, the magnitude is displayed in (c) raised to the 1/5 power. Phase is shown from −0.5 to 0.5 rad in (b) and from $-\pi$ to π in (d).

guesses, two were found to exhibit a pronounced twin-image problem, 15 showed a moderate problem (the object or its twin is substantially dominating over the other), and only three achieved successful reconstructions where the twin-image problem was not discernable. This classification was done by inspection of the reconstructions and the degree of symmetry they exhibited. Figures 2(a), 5(a), and 8(a) (below) show examples of what were classified as pronounced, moderate, and not visually discernable twin-image problems, respectively. Here, when we speak of a reconstruction (a reconstructed image), we mean the stagnated output image from the algorithm, which is known to be imperfect because there is imperfect agreement with the Fourier magnitude or with the object's support constraint.

A. Example with Pronounced Twin-Image Problem

The reconstruction shown in Fig. 2(a) exhibits a pronounced twin-image problem, having a symmetry that does not exist in the original object magnitude, Fig. 1(a). Not even the coarse features are easily recognizable.

The Fourier phase error, $\hat{\phi}(u,v) - \phi(u,v)$, of the reconstruction with respect to the ideal image is shown in Fig. 2(c). This was obtained by subpixel registering the reconstruction to the ideal image [30], removing a global phase difference, and subtracting their Fourier phases. The same procedure was repeated for the ideal twin image, $f^*(-x,-y)$, to obtain the Fourier phase error of the reconstruction with respect to the twin image, $\hat{\phi}(u,v) + \phi(u,v)$, shown in Fig. 2(d). Notice that this procedure requires knowledge of the ideal image and is thus only available for numerical simulations.

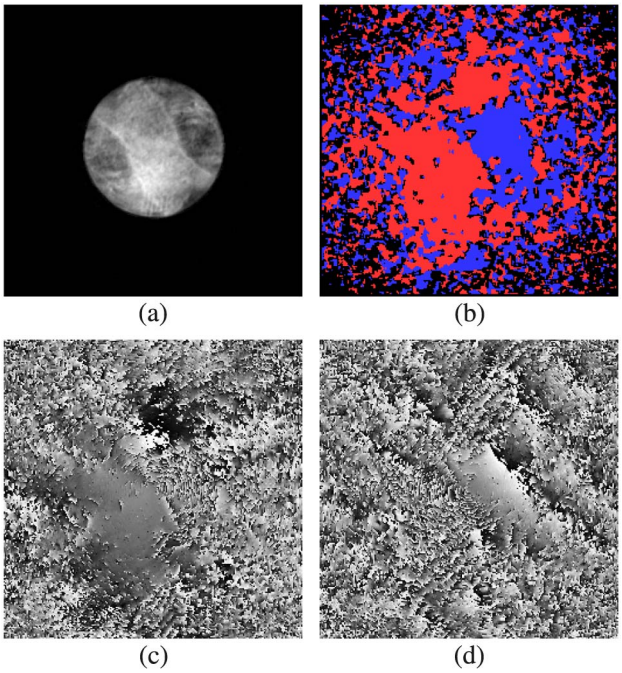

Fig. 2. (Color online) Image reconstruction exhibiting a pronounced twin-image problem. (a) The stagnated image magnitude; the Fourier phase errors with respect to (c) the upright ideal image, $f(x,y)$, and (d) the twin image, $f^*(-x,-y)$, and (b) regions of the Fourier domain that reconstruct either the upright image (bright red), the twin image (blue), or neither (black) after 1000 iterations. Phase error is shown from $-\pi$ to π radians in (c) and (d). The evolution of the regions with iteration number is shown in Media 1.

The Fourier phase error with respect to the upright image, Fig. 2(c), exhibits regions that are very smooth and regions where the phase error appears random from one pixel to the next. The smooth regions correspond to areas where the phase error is low, and thus a reasonable reconstruction of the Fourier phase, ϕ, was achieved. The Fourier phase error with respect to the twin image, Fig. 2(d), also exhibits both smooth regions, where the negative of the Fourier phase $-\phi$ is approximately recovered, and random regions. Notice that the smooth regions in Figs. 2(c) and 2(d) are approximately complementary. Also notice that disjoint smooth regions of the Fourier phase error can have a different global phase constant or a linear phase term between them.

As these results suggest, the twin-image problem is explained by the division of the Fourier data into regions, some of which reconstruct the upright image, whereas others reconstruct the twin image. Figure 2(b) shows the regions in the Fourier domain that reconstruct either the upright (bright red) or twin image (blue) for the reconstruction shown in Fig. 2(a). These areas were obtained by comparing the normalized root-mean-square error (NRMSE) of the reconstruction with respect to both $F(u,v)$ and $F^*(u,v)$ in a small pixel subset of the Fourier data. In general, a reconstructed FT $G(u,v)$ and the ideal $F(u,v)$ [or $F^*(u,v)$] may have a global (piston) or linear phase terms (corresponding to image translations) between them that should not be included as part of the error, so we computed a version of the NRMSE that is invariant to these terms:

$$\varepsilon^2 = \min_{\alpha,x_0,y_0} \frac{\sum_{u,v} \left|\alpha G(u,v) \exp\left[-i2\pi\left(\frac{ux_0}{M} + \frac{vy_0}{N}\right)\right] - F(u,v)\right|^2}{\sum_{u,v}|F(u,v)|^2}, \tag{4}$$

which, using Parseval's theorem, can be expressed in the image domain as [31]

$$\varepsilon^2 = \min_{\alpha,x_0,y_0} \frac{\sum_{x,y}|\alpha g(x-x_0,y-y_0) - f(x,y)|^2}{\sum_{x,y}|f(x,y)|^2}, \tag{5}$$

where α is an unknown complex-valued constant and (x_0,y_0) are unknown translations. This minimization can be solved by calculating the peak of a cross correlation [31], which can be done to subpixel accuracy with an efficient approach [30].

The NRMSE, as given in Eq. (4) but with the sum restricted to a 5×5 pixel subset of the FT, was computed for the reconstruction with respect to both $F(u,v)$ and $F^*(u,v)$. If the NRMSE with respect to the upright image is smaller than that computed with respect to the twin image, then the central pixel of the subset was considered to reconstruct the upright image. For cases where the NRMSE with respect to the twin image was smaller than that with respect to the upright image, the central pixel was considered to reconstruct the twin image. If neither of the NRMSEs was smaller than 0.5, then the recovered phase at the central pixel was considered erroneous—it reconstructs neither the upright nor the twin image [dark regions in Fig. 2(b)].

To verify that complementary regions of the Fourier domain are reconstructing either the upright image or the twin image, Fig. 3 shows the results of filtering the ideal image and the reconstruction in the Fourier domain using the regions shown in Fig. 2(b). Figure 3(b) was obtained by computing

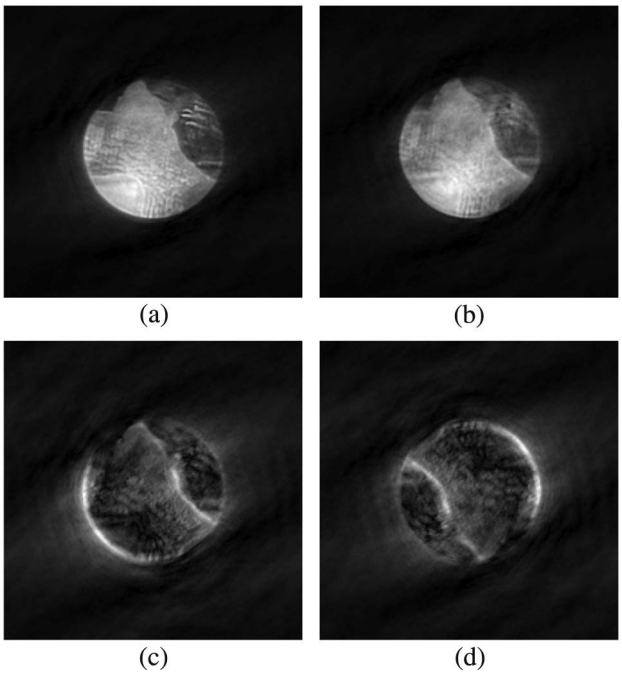

Fig. 3. Fourier filtering of the ideal image and the reconstruction shown in Fig. 2(a) using the Fourier regions shown in Fig. 2(b). Filtering of the (a) ideal image and (b) reconstruction using the bright red Fourier regions. Filtering of the (c) ideal image and (d) reconstruction using the blue Fourier regions.

the FT of the reconstruction, shown in Fig. 2(a), setting to zero all pixels that fall outside of the Fourier regions that reconstruct the upright image [bright red regions in Fig. 2(b)], and then computing the inverse FT. For comparison purposes, the ideal image was filtered the same way and is shown in Fig. 3(a). Notice that a good agreement is obtained and that, although the complete reconstruction shows a very pronounced twin-image problem, this is not the case where only the bright red regions in the FT are included in the filtered reconstruction.

Figure 3(d) shows the result obtained by filtering the reconstruction in a similar way, but only using the regions that reconstruct the twin image, i.e., the blue regions in Fig. 2(b). In this case the filtered version of the reconstruction, Fig. 3(d), is inverted with respect to the equivalently filtered version of the ideal image, Fig. 3(c). This supports our statement that the twin-image problem is explained by complementary regions in the Fourier domain reconstructing features of either the upright or twin image.

To monitor the progress of the reconstructions, we used three metrics: the support error E, the invariant NRMSE of the reconstruction versus the ideal image ε, and the fractional energy of upright and twin image present in the reconstruction. Of these metrics, only the support error, defined by

$$E^2 = \frac{\sum_{(x,y)\notin S}|g(x,y)|^2}{\sum_{x,y}|g(x,y)|^2}, \tag{6}$$

is available without prior knowledge of the solution.

Figure 4(a) shows the support error versus iteration number for the reconstruction shown in Fig. 2(a). Notice that, during the first 45 iterations, of each series of 50, the support error may increase on account of the ability of the HIO algorithm to escape local minima. In the last five iterations of each

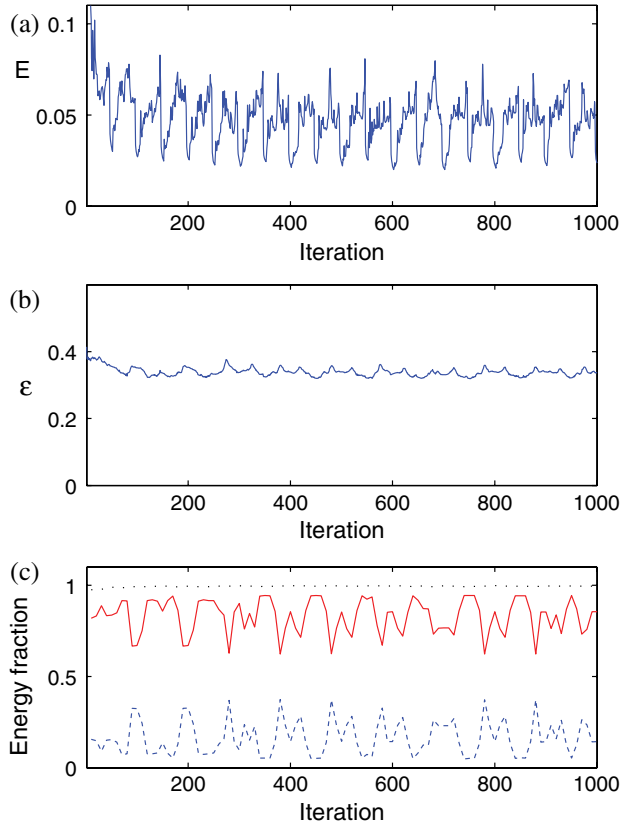

Fig. 4. (Color online) (a) Support error for the reconstruction shown in Fig. 2(a) versus iteration number. (b) NRMSE of the reconstruction with respect to the ideal image versus iteration number. (c) Energy fraction of the reconstruction versus iteration number for the Fourier regions that reconstruct the upright image (solid curve) and the twin image (dashed curve). The sum of the two energy fractions is indicated by a dotted curve in (c).

series of 50, the error-reduction algorithm rapidly decreases the support error, allowing the reconstruction to settle.

Figure 4(b) shows the NRMSE of the reconstruction versus iteration number. This was computed using Eq. (5) by registering the ideal image and the reconstruction to within 1/15 of a pixel [30]. Because either the upright or twin images are considered successful reconstructions, the NRMSE was computed between the reconstruction and both $f(x,y)$ and $f^*(-x,-y)$ and the minimum of these was considered to be the NRMSE of the reconstruction.

The Fourier regions were computed as indicated above every 10 iterations during the reconstruction. To monitor the twin-image problem, the sum of the object Fourier intensity (after normalizing it to have unit energy) was taken after filtering using the computed Fourier regions, thus obtaining the fraction of energy that is present in the upright and twin images. Figure 4(c) shows the fractional energy of the upright (solid curve) and twin image (dashed curve) versus iteration number. Notice that the fractional energy of the upright image remains in a somewhat narrow band throughout the reconstruction. This indicates that, from the early iterations, the reconstruction suffered from the twin-image problem to a considerable extent, and that, even after 1000 iterations, the ITA was unable to escape this very persistent stagnation mode.

Inspection of the regions versus iteration number for this reconstruction (Media 1) shows that the regions form early

in the reconstruction and remain somewhat fixed throughout the iterations. Every 50 iterations, the dark regions (regions of erroneous reconstruction) are reduced because of the few error-reduction iterations. The latter gives the video a characteristic pulsating appearance.

B. Example with a Moderate Twin-Image Problem

A sample reconstruction, from a different random starting guess, for which the twin-image problem was considered moderate, is shown in Fig. 5(a). Reconstructions with this visual quality were the most frequent (15 out of 20) for this particular object. In this particular reconstruction, the twin image visually dominates, and the twin-image problem is considered moderate because it appears less symmetric than that shown in Fig. 2(a) and the general shape of the object can be distinguished better. The Fourier phase errors with respect to the upright and twin images are shown in Figs. 5(c) and 5(d), respectively. The Fourier regions, computed as described above, are shown in Fig. 5(b). For this reconstruction, the support error, the NRMSE with respect to the ideal image, and the energy fraction in the Fourier regions are shown in Fig. 6. The error metrics shown in Figs. 6(a) and 6(b) are somewhat larger than those in Fig. 4, indicating that the reconstruction shown in Fig. 5(a) is worse than that shown in Fig. 2(a), even though the latter appears more symmetric. The relative energy fraction is also worse in Fig. 6(c) than in Fig. 4(c). All these metrics indicate that, although visually the reconstruction shown in Fig. 5(a) appears less symmetric than that shown in Fig. 2(a), the latter is a slightly better reconstruction from a squared-error point of view. Visual inspection of the apparent symmetry of the reconstruction is, in this case, not necessarily the best way to estimate or compare the severity of the twin-image problem. Conversely, the energy fraction of the reconstruction is an indicator of the presence of the twin-image problem but does not always correlate well with

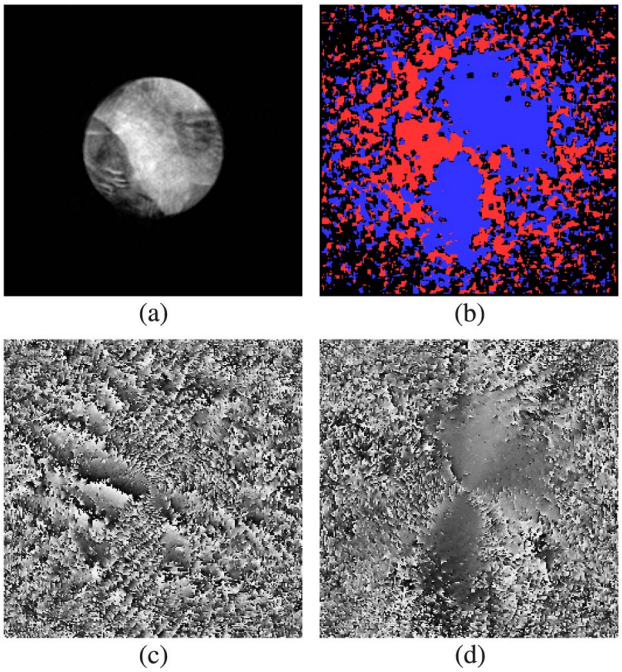

Fig. 5. (Color online) Same as Fig. 2 but for an image reconstruction exhibiting a moderate twin-image problem. The evolution of the regions with iteration number is shown in Media 2.

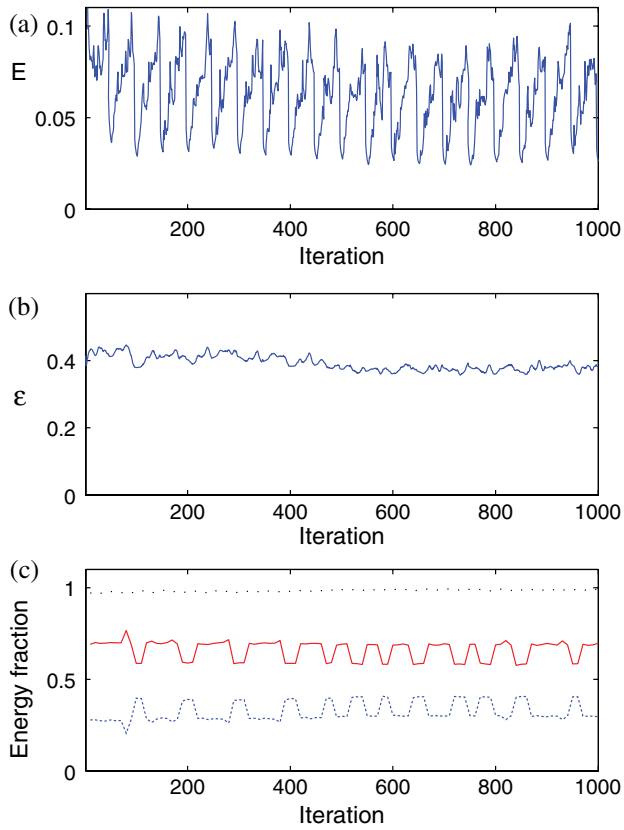

Fig. 6. (Color online) Same as Fig. 4, but for the reconstruction shown in Fig. 5(a).

perceived image quality. Exploring alternative metrics to quantify the image quality impact of the twin-image problem, for example by excluding the central pixels of the FT (which may skew the results as further discussed below), would be valuable for further understanding this problem. As discussed in Section 3, in the future it would also be useful to develop techniques that can identify the twin-image problem without requiring the true solution so that this problem can be diagnosed in real-world applications.

Notice that, while the twin image appears to visually dominate the reconstruction, the relative energy fraction [Fig. 6(c)] indicates that the upright image has more energy than the twin image. This counterintuitive result can be explained by two facts. (1) Although the area of the Fourier regions that reconstructs the twin image (blue) is evidently larger than the area of the regions that reconstruct the upright image [the ratio is about 3/2 judging from Fig. 5(b)], the very central region of the Fourier data is predominantly reconstructing the upright image. Because we are reconstructing a weak phase object, the intensity of the pixels near the central region of the FT is orders of magnitude greater than at the pixels that correspond to middle- and high-frequency components, thereby explaining the greater energy of the upright image for this reconstruction. (2) Notice in Fig. 5(b) that the Fourier regions that reconstruct the twin image have a very good coverage of the middle spatial frequencies. These frequencies are crucial for reconstructing visually appealing edges. Because of this, the sharp edge that is present in the object appears inverted in the reconstruction. This explains why from visual inspection we might conclude that the twin

image dominates the reconstruction in Fig. 5(a), whereas the energy ratio indicates the opposite.

Upon examination of the Fourier regions throughout the iteration process, for this particular random starting guess, we found that the central pixels consistently reconstruct the upright image (Media 2). The regions in the Fourier domain appear in the early iterations, and, although they remain relatively static, it is evident that some twin-image regions (blue) connect over time to form a larger region. As we will see later, this latter procedure can eventually lead to an acceptable solution of the phase retrieval problem. In this case, however, the reconstruction stagnated.

The reconstruction, shown in Fig. 5(a), was filtered using the Fourier regions shown in Fig. 5(b) in a form analogous to that shown in Fig. 3. The filtered reconstructions, using the Fourier regions that reconstructed the upright and twin images, are shown in Figs. 7(b) and 7(d), respectively. Notice that Fig. 7(d) has most of the object hard edges and details, while Fig. 7(b) is much smoother on account of the missing middle spatial frequencies.

C. Example with a Not-Discernable Twin-Image Problem

Finally, in Fig. 8(a) we show a reconstruction example where the twin-image problem could not be discerned upon visual inspection. Figures 9(a) and 9(b) show that the support error and the NRMSE are better than those of the previous reconstructions, and Fig. 9(c) shows that, from early in the iteration process, most of the energy in the reconstruction corresponds to the upright image. However, upon examination of the Fourier phase error with respect to the upright and twin images, shown in Figs. 8(c) and 8(d), respectively, we realize that this reconstruction is not exempt from the twin-image problem. Because the Fourier regions that reconstruct the twin image are confined to a portion of high spatial frequencies that have much less energy, the visual impact of the twin-image artifacts is very subtle.

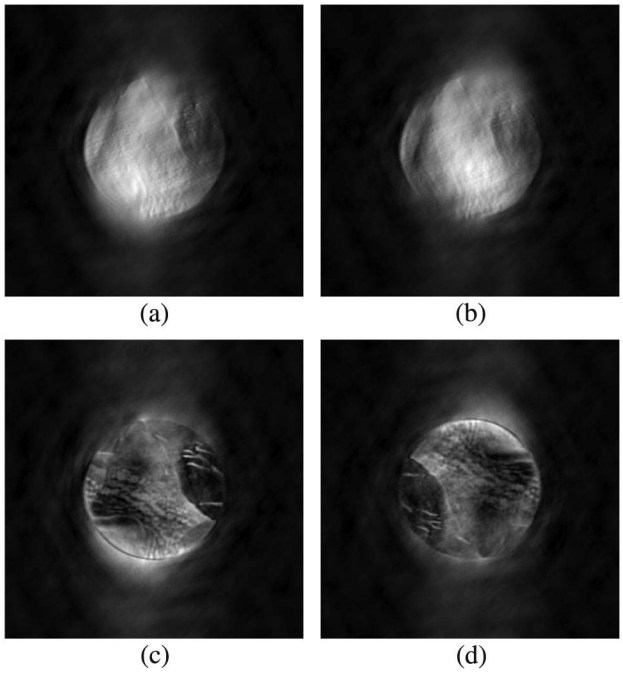

Fig. 7. Same as Fig. 3, but for the reconstructed image shown in Fig. 5(a).

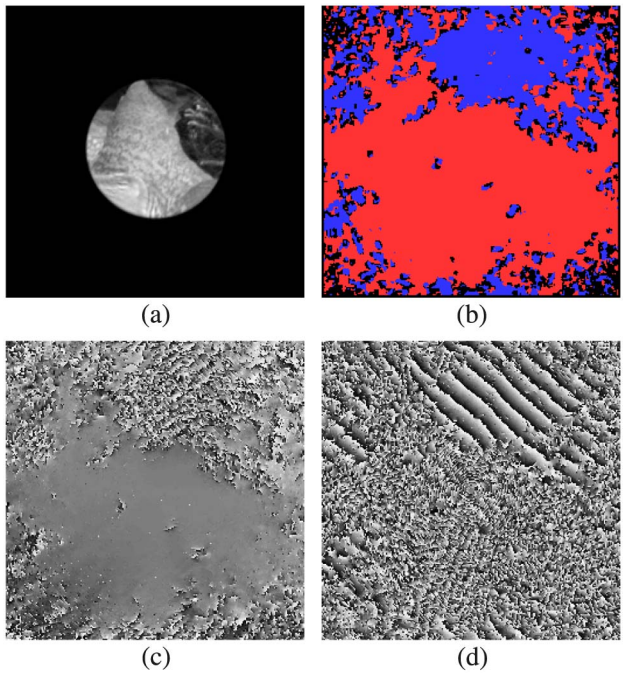

Fig. 8. (Color online) Same as Fig. 2 but for an image reconstruction where the twin-image problem cannot be detected visually. The evolution of the regions with iteration number is shown in Media 3.

Upon inspecting the evolution of the regions with respect to the iteration number, we notice that the regions that reconstruct either the upright or twin images are formed early in

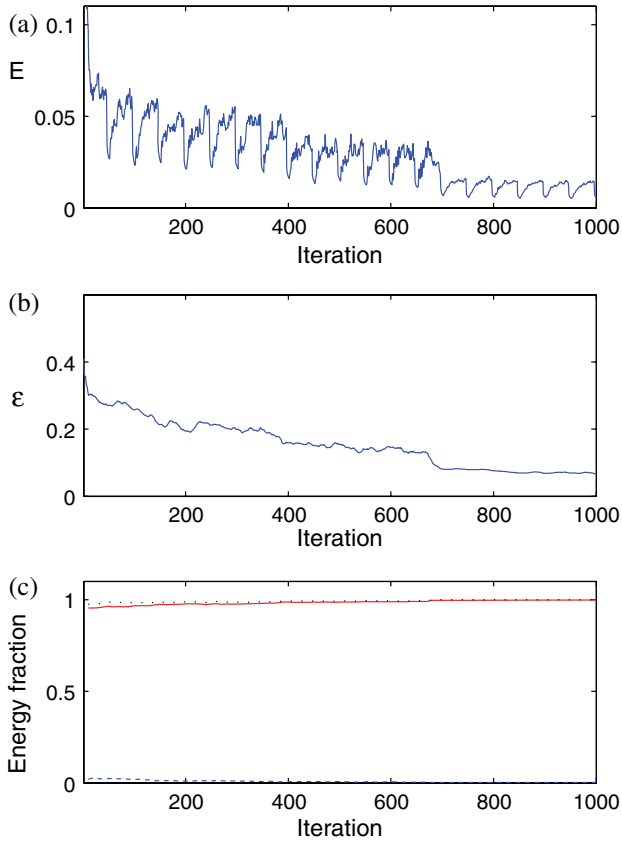

Fig. 9. (Color online) Same as Fig. 4 but for the reconstruction shown in Fig. 8(a).

the reconstruction process and that the success of the reconstruction relies heavily on the ability of these regions to connect to each other and form larger regions (see Media 3). Because disconnected regions can be out of phase with one another, the quality of the reconstruction is affected if the regions are separated.

Despite its good visual quality, for this reconstruction not all of the higher spatial frequencies were accurately recovered. The reconstruction is stagnated at a point that is much closer to the solution, and, although it cannot be seen from visual inspection of the reconstruction, this stagnation is caused by the twin-image problem as seen in Figs. 8(b)–8(d). Upon inspecting the other two reconstructions that visually appeared exempt from the twin-image problem, we found similar results; i.e., all of the reconstructions exhibited the twin-image problem to at least a small extent.

The Fourier regions that characterize this stagnation mode can be very persistent with iterations but change with different starting guesses. To assess the consistency of retrieved phases, it is then important to compare reconstructions from different random starting guesses, especially for objects that are conducive to the twin-image problem as described in Section 1.

3. IDENTIFYING THE TWIN-IMAGE PROBLEM

The methods for identifying the twin-image problem and finding the Fourier regions for the reconstructions described in the previous sections require knowledge of the solution to the phase retrieval problem and are thus only available for numerical simulations. In a real-world reconstruction, the metrics available for analysis are the support error and consistency between reconstructions from different starting guesses. The question arises of how to identify whether our solution suffers from the twin-image problem.

A possible method for identifying the twin-image problem is to take two reconstructions, generated from two different random starting guesses, and select the one with the lowest support error as a reference. Register this reference with the second reconstruction (to remove a linear phase between them in the Fourier domain) and remove a constant global phase factor between them; do the same with the twin image of the second reconstruction. The smoothness of the retrieved Fourier phase differences can give an indication of the twin-image problem. If either of the Fourier phase differences is smooth everywhere, this would be an indication that the reconstructions are free of the twin-image problem. In this test, the twin-image problem would be characterized by the appearance of smooth and random regions that appear complementary when comparing upright with twin images, much like the regions that appear in Figs. 2(c) and 2(d). It should be noted that the regions identified in this way will not directly separate regions that reconstruct upright or twin image, but rather will indicate the coincidence of regions of the reference and second reconstruction that are being compared. This test would fail if the Fourier regions of different reconstructions match by chance, but because the reconstructions are started from different random starting guesses, it is unlikely that they would have exactly the same Fourier regions. In the numerical example examined for this paper, we found different Fourier regions for each reconstruction from a different starting guess.

Furthermore, this test will only work with already good individual reconstructions. And although it could indicate the presence of the twin-image problem in at least one of the reconstructions, it cannot give direct information as to which of these regions reconstructs either the upright or twin images in either reconstruction. Further development of techniques to address this issue is still required.

4. ESCAPING THE STAGNATION

A method for overcoming the twin-image problem in phase retrieval is the reduced-area support constraint method, described in [24]. In this method, the object support mask is replaced with a noncentrosymmetric mask that is a subset of the original support. After a few iterations of the reconstruction algorithm, the temporary mask is replaced by the original mask and iterations continue.

Using the same random starting guess as for the reconstruction shown in Fig. 5(a), we performed 200 iterations of the ITA (series of 45 HIO followed by five of error reduction). The reconstruction after 200 iterations is shown in Fig. 10(a), and the Fourier regions that reconstruct either the upright or twin image are shown in Fig. 10(b). Notice that the twin-image problem is present in this reconstruction. After iteration 200, the circular support with 60 pixel radius was replaced by a semicircle obtained by halving the original support in the x direction. The reconstruction result and Fourier regions after 10 iterations of the HIO with this reduced support are shown in Figs. 10(c) and 10(d), respectively.

The reduced-support constraint introduces a strong perturbation on the estimation of the Fourier phase, which allows the algorithm to escape this very persistent stagnation mode. Notice how the Fourier regions after the perturbation show that the phase estimation is now inaccurate throughout the Fourier window. The success of the technique relies on the fact that the perturbation by the reduced-support method tends to favor either the upright or the twin image over the other. We applied the reduced-area method after only 200 iterations because stagnation by the twin-image problem was consistently found to be established early in the reconstruction. Even though for some cases further iterations gradually improved the results, this typically required many more iterations.

After the 10 HIO iterations, the reduced support was replaced with the original support constraint, and we performed 35 HIO iterations followed by five of error reduction. The reconstruction and Fourier regions after the first 10 HIO iterations with the original support are shown in Figs. 10(e) and 10(f), respectively. Notice that, after only these 10 iterations, the upright image heavily dominates over the twin image. Subsequent iterations were performed by alternating 45 iterations of HIO with five of error reduction. The final reconstruction, after a total of 1000 ITA iterations, shown in Fig. 11(a), has a considerably increased quality as compared to Fig. 5(a). Although no twin-image problem is noticeable by visual inspection of the reconstruction, the Fourier phase errors with respect to the ideal upright and twin images, shown in Figs. 11(c) and 11(d), respectively, show that the division into Fourier regions is still present. This problem, however, is even

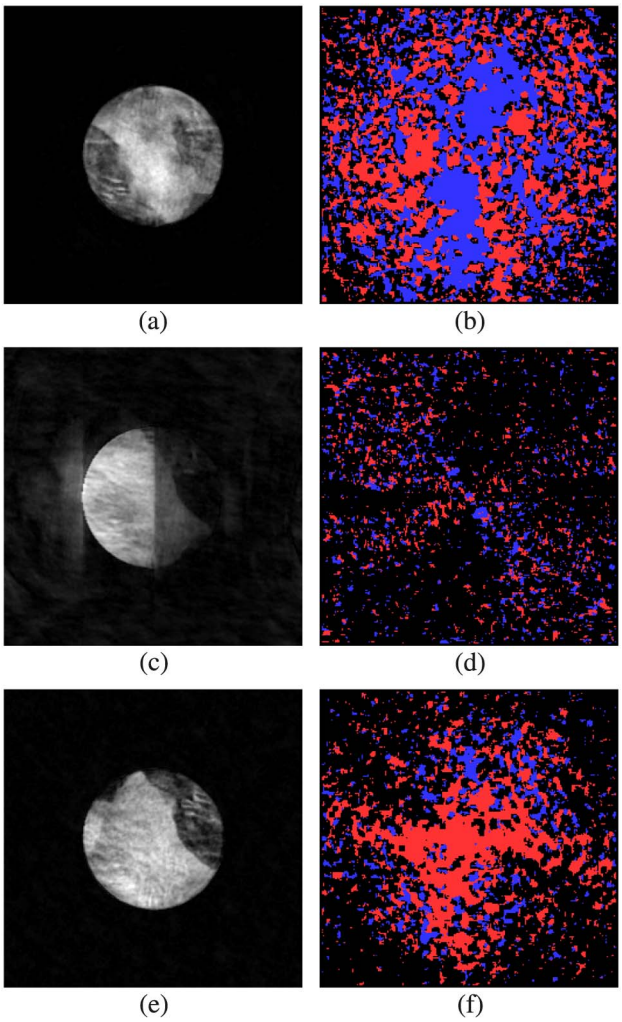

Fig. 10. (Color online) (a) Stagnated reconstruction after 200 iterations and (b) its Fourier regions. (c) Reconstruction after 10 iterations using the reduced-support method and (d) its Fourier regions. (e) Reconstruction after reinstating the original support constraint and performing 10 more iterations and (f) its Fourier regions. Regions of the far field that reconstruct either the upright or twin image are shown in bright red or blue, respectively, in (b), (d), and (f).

less than those obtained with the lucky starting guesses (one of which is shown in Fig. 8).

The support error, NRMSE, and energy fraction of the Fourier regions are shown in Figs. 12(a)–12(c), respectively. Notice that the reduced-support method (at iteration 200) considerably increases the support error and the NRMSE, but then after the original support is reinstated, these error metrics quickly drop and eventually settle. Subsequent iterations improve the results, but this reconstruction is still affected by a slight but persistent twin-image problem that is not visually discernable, keeping the NRMSE from going to zero.

We applied the reduced-support method, as described above, to reconstructions using the same 20 starting guesses (for the previously quoted 20 reconstructions). For all these reconstructions, we obtained good results, similar to those shown in Fig. 11, but with Fourier regions located in different portions of the Fourier window. Furthermore, they all converged to the upright image, which is a consequence of the choice of orientation of the reduced support. If the reduced support is inverted through the origin, then all the

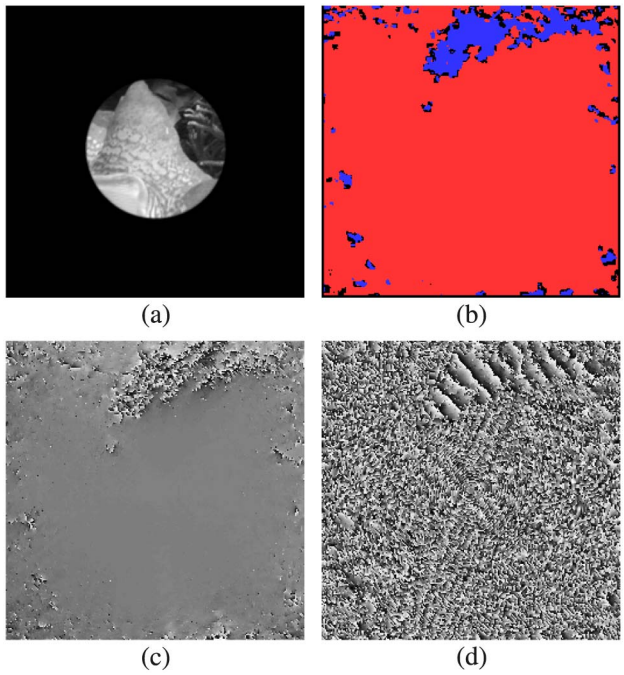

Fig. 11. (Color online) Same as Fig. 5, from same starting guess but using the reduced-support method in iteration 200. The evolution of the regions with iteration number is shown in Media 4.

reconstructions converge to the twin image, which is also considered a successful reconstruction. This was a great improvement over the conventional reconstruction with series of HIO

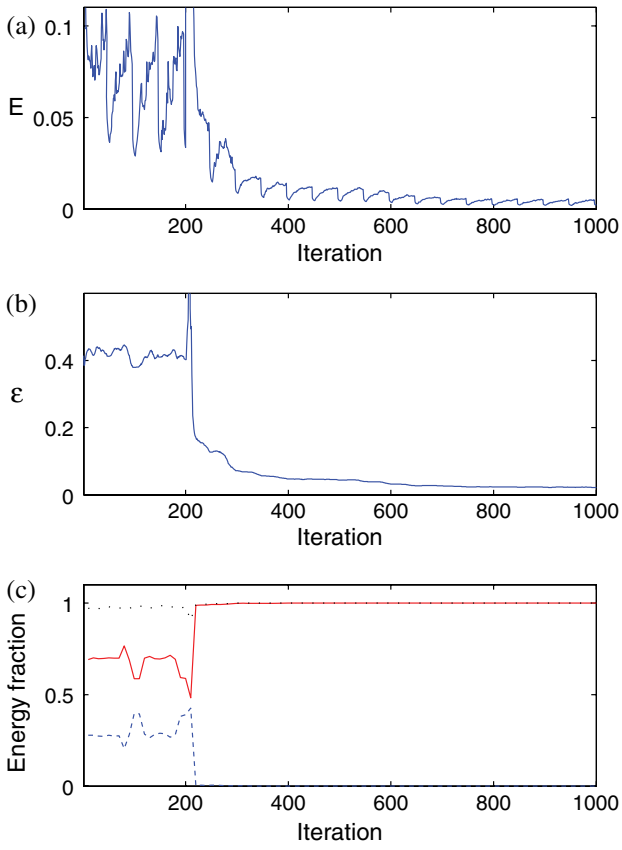

Fig. 12. (Color online) Same as Fig. 4, but for the reconstruction shown in Fig. 11, using the reduced-support method. The final support error was $E = 2.1 \times 10^{-3}$.

and error reduction alone, where (as previously quoted) 17 out of 20 reconstructions were found to have a moderate or severe twin-image problem.

We should note that the success of the reduced-support technique is highly dependent on the object, as discussed in [24]. In our case, the success was driven by the fact that the object had energy preferentially in one half of the support. The performance of this method will be decreased for objects that more uniformly fill the support.

Combining reconstructions from independent starting guesses can also be effective in suppressing or diminishing the twin-image problem as well as other stagnation modes. Registration and averaging of reconstructions is already widely used and has the advantage of suppressing spatial frequencies with large phase variations among reconstructions. Comparing the Fourier intensity of this averaged reconstruction to the measurement provides a tool to diagnose the repeatability of the phase recovery at each spatial frequency [9,10]. Other methods of selectively combining information in the Fourier domain from reconstructions of moderately good quality include the voting and the patching methods [6,24].

5. SUMMARY AND CONCLUSIONS

The twin-image problem is a persistent stagnation mode in phase retrieval, characterized by the simultaneous appearance of features of the image and its complex-conjugated centrosymmetrical inversion (twin image). This problem can occur when the image sought and its twin both fit inside the support constraint, and in severe cases it can significantly hinder image quality.

In this paper, we demonstrated, through numerical simulations, that this problem occurs when the reconstructed Fourier domain divides into regions in the early iterations of the reconstruction. Some of these Fourier regions reconstruct the upright image, while others reconstruct the twin image. The location and size of the regions varies with the random initial guess. We have also provided guidelines toward identifying the twin-image problem from reconstructions from different random starting guesses, and we have briefly described strategies to escape this stagnation.

Additionally, we have found that, for reconstructions of good visual quality for which the twin-image problem was not visually discernible, this problem may still be present to a small extent. Although for these cases where the impact on image quality is subtle, it still affects the retrieval of the phases at high spatial frequencies, thus affecting the reconstruction of fine-resolution features. The importance of understanding the twin-image problem comes from the natural symmetry exhibited by objects for many applications, including some samples of significant interest for x-ray CDI, such as biological specimens, which could be prone to this stagnation mode.

REFERENCES

1. R. A. Gonsalves and R. Childlaw, "Wavefront sensing by phase retrieval," Proc. SPIE **207**, 32–39 (1979).
2. J. R. Fienup, "Phase-retrieval algorithms for a complicated optical system," Appl. Opt. **32**, 1737–1746 (1993).
3. G. R. Brady and J. R. Fienup, "Nonlinear optimization algorithm for retrieving the full complex pupil function," Opt. Express **14**, 474–486 (2006).
4. M. R. Bolcar and J. R. Fienup, "Sub-aperture piston phase diversity for segmented and multi-aperture systems," Appl. Opt. **48**, A5–A12 (2009).
5. J. R. Fienup, "Reconstruction of an object from the modulus of its Fourier transform," Opt. Lett. **3**, 27–29 (1978).
6. J. R. Fienup, "Lensless coherent imaging by phase retrieval with an illumination pattern constraint," Opt. Express **14**, 498–508 (2006).
7. M. Guizar-Sicairos and J. R. Fienup, "Phase retrieval with Fourier-weighted projections," J. Opt. Soc. Am. A **25**, 701–709 (2008).
8. 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 **400**, 342–344 (1999).
9. D. Shapiro, P. Thibault, T. Beetz, V. Elser, M. Howells, C. Jacobsen, J. Kirz, E. Lima, H. Miao, A. M. Neiman, and D. Sayre, "Biological imaging by soft x-ray diffraction microscopy," Proc. Natl. Acad. Sci. USA **102**, 15343–15346 (2005).
10. H. N. Chapman, A. Barty, S. Marchesini, A. Noy, S. P. Hau-Riege, C. Cui, M. R. Howells, R. Rosen, H. He, J. C. H. Spence, U. Weierstall, T. Beetz, C. Jacobsen, and D. Shapiro, "High-resolution *ab initio* three-dimensional x-ray diffraction microscopy," J. Opt. Soc. Am. A **23**, 1179–1200 (2006).
11. J. W. Goodman, *Introduction to Fourier Optics*, 3rd ed. (Roberts, 2005).
12. J. R. Fienup, "Phase retrieval algorithms: a comparison," Appl. Opt. **21**, 2758–2769 (1982).
13. V. Elser, "Phase retrieval by iterated projections," J. Opt. Soc. Am. A **20**, 40–55 (2003).
14. H. H. Bauschke, P. L. Combettes, and D. R. Luke, "Hybrid projection-reflection method for phase retrieval," J. Opt. Soc. Am. A **20**, 1025–1034 (2003).
15. C. Chen, J. Miao, C. W. Wang, and T. K. Lee, "Application of optimization technique to noncrystalline x-ray diffraction microscopy: guided hybrid input—output method," Phys. Rev. B **76**, 064113 (2007).
16. R. H. T. Bates and D. G. H. Tan, "Fourier phase retrieval when the image is complex," Proc. SPIE **0558**, 54–59 (1985).
17. 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).
18. R. G. Paxman, J. R. Fienup, and J. T. Clinthorne, "Effect of tapered illumination and Fourier intensity errors on phase retrieval," Proc. SPIE **0828**, 184–189 (1987).
19. J. R. Fienup, T. R. Crimmins, and W. Holsztynski, "Reconstruction of the support of an object from the support of its autocorrelation," J. Opt. Soc. Am. **72**, 610–624 (1982).
20. T. R. Crimmins, J. R. 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).
21. J. R. Fienup, B. J. Thelen, M. F. Reiley, and R. G. Paxman, "3-D locator sets for opaque objects for phase retrieval," Proc. SPIE **3170**, 88–96 (1997).
22. M. Guizar-Sicairos and J. R. Fienup, "Holography with extended reference by autocorrelation linear differential operation," Opt. Express **15**, 17592–17612 (2007).
23. Y. M. Bruck and L. G. Sodin, "On the ambiguity of the image reconstruction problem," Opt. Commun. **30**, 304–308 (1979).
24. J. R. Fienup and C. C. Wackerman, "Phase-retrieval stagnation problems and solutions," J. Opt. Soc. Am. A **3**, 1897–1907 (1986).
25. E. Lima, L. Wiegart, P. Pernot, M. Howells, J. Timmins, F. Zontone, and A. Madsen, "Cryogenic x-ray diffraction microscopy for biological samples," Phys. Rev. Lett. **103**, 198102 (2009).
26. H. Chapman, "X-ray imaging beyond the limits," Nat. Mater. **8**, 299–301 (2009).
27. H. M. L. Faulkner and J. M. Rodenburg, "Movable aperture lensless transmission microscopy: a novel phase retrieval algorithm," Phys. Rev. Lett. **93**, 023903 (2004).
28. M. Guizar-Sicairos and J. R. Fienup, "Phase retrieval with transverse translation diversity: a nonlinear optimization approach," Opt. Express **16**, 7264–7278 (2008).
29. P. Thibault, M. Dierolf, A. Menzel, O. Bunk, C. David, and F. Pfeiffer, "High-resolution scanning x-ray diffraction microscopy," Science **321**, 379–382 (2008).
30. M. Guizar-Sicairos, S. T. Thurman, and J. R. Fienup, "Efficient subpixel image registration algorithms," Opt. Lett. **33**, 156–158 (2008).
31. J. R. Fienup, "Invariant error metrics for image reconstruction," Appl. Opt. **36**, 8352–8357 (1997).