

Effectiveness of iterative algorithms in recovering phase in the presence of noise

Garth Williams,*‡ Mark Pfeifer,§ Ivan Vartanyants¶ and Ian Robinson††

Department of Physics, University of Illinois, Urbana, IL 61801, USA. Correspondence e-mail: g.williams@physics.unimelb.edu.au

Inversion of coherent X-ray diffraction patterns allows the imaging of three-dimensional density distributions. The recovery of such shapes often requires application of iterative algorithms, such as Fienup's error reduction or hybrid input/output. Since the measurement of such a pattern records the intensity in reciprocal space, any errors due to noise will probably not have a straightforward impact on the final real-space result. In this paper, the effect of the types of noise common in coherent X-ray diffraction (CXD) experiments, counting statistics, scatter from alien particles and detector noise, on the recovered real-space density projection is explored by simulating a two-dimensional CXD pattern and adding noise. It is found that an R factor measuring the reproducibility between the best and second-best real-space result is a leading indicator of performance.

1. Introduction

X-ray diffraction is a powerful technique for resolving the structure of materials. In traditional techniques, a sample is illuminated by a beam containing many coherence volumes resulting in diffraction that corresponds to an ensemble average of domains within the sample. By contrast, if diffraction resulting from the interaction of a single coherence volume in the beam with the sample – so-called coherent X-ray diffraction (CXD) – can be sufficiently measured, the possibility to recover the three-dimensional shape of the sample arises. In this latter case, the diffracted intensity in the vicinity of a Bragg peak is related to the electron number density, $\rho(\mathbf{r})$, of the sample by

$$I(\mathbf{Q}) = \left| \int d\mathbf{r} \rho(\mathbf{r}) \exp(i\mathbf{Q} \cdot \mathbf{r}) \right|^2, \quad (1)$$

where \mathbf{Q} is the momentum transfer local to the Bragg point (Warren, 1990). Therefore, equation (1) presents us with an inverse problem: given the square modulus of the Fourier transform of a function, what is the function? Essentially, the problem reduces to one of finding a set of phases that are consistent with the measured Fourier modulus, defined to be the square root of the measured intensity, and some known real-space properties of the object, such as a finite support.

It is extremely important to explore the uniqueness of a solution to any inverse problem. In this particular problem,

the intensity is measured in one space and the object is known to occupy only a finite portion of the other. As is customary in the literature, we draw a distinction between trivial ambiguities and non-unique solutions. For example, if we measure the modulus of the Fourier transform of a real-space function $f(z)$, $|F(q_z)|$, the trivial ambiguities (Bates, 1982) to $f(z)$ are a translation, $f(z + z_0)$, a constant phase offset, $\exp(i\phi)f(z)$, and the 'twin' image, $f^*(-z)$. Bruck & Sodin (1979) provide the earliest work on the uniqueness of this problem. Their position was reinforced by the work of Bates (1982) and that of Barakat & Newsam (1984). The end result of this theoretical work is that the uniqueness of solutions to the inverse problem is highly dependent on the dimensionality of the problem. This dependence on dimensionality arises because non-unique solutions can occur when the diffracted amplitude is factorizable. Since a one-dimensional polynomial is always factorizable on the complex plane, the one-dimensional problem has no unique solution, although it may have a unique *physical* solution. In two dimensions and higher, factorizable polynomials become very rare and therefore the problem generally has a unique solution (Bates, 1982). One of the principal aims of this paper is to investigate whether the presence of noise in a measured diffraction pattern is likely to result in reconstructions that are not trivially related.

Iterative methods have been applied to this problem to recover two-dimensional and three-dimensional images of samples using coherent X-ray diffractive imaging (Sayre *et al.*, 1998; Miao *et al.*, 1999; Robinson *et al.*, 2001; Miao *et al.*, 2002; Williams *et al.*, 2003). An example is the CXD pattern in Fig. 1(a) and the corresponding reconstructed projection of the density of an Au microcrystal in Fig. 1(b). Reconstructions of

* Current affiliation: School of Physics, The University of Melbourne, Australia.

‡ Current affiliation: Department of Physics, La Trobe University, Australia.

§ Current affiliation: HASYLAB, DESY, Hamburg, Germany.

¶ Current affiliation: Department of Physics and Astronomy, University College London, London, England.

this kind are made possible by ‘oversampling’ the diffraction patterns. The oversampling ratio (Miao *et al.*, 1998) is generally defined as

$$\sigma = \frac{\text{size of the FFT array}}{\text{portion occupied by object}} \quad (2)$$

and must be larger than 2. This is essentially sampling the intensity in accordance with Shannon’s theorem. An important question that has not been completely answered is how noise in CXD data affects the final real-space solution. The effect of noise on stagnation and uniqueness has been addressed by Seldin & Fienup (1990) and Miao *et al.* (1998) for the case of Gaussian random noise. Here, we seek to discover if the specific kinds of noise generally present in a CXD experiment cause artifacts in the real-space objects that may be confused for a physical feature. Specifically, we consider noise from native – *i.e.* the Poisson distributed shot noise inherent in the diffracted signal – and alien – defined as any other source of X-ray scatter – scattering as well as noise generated in a typical CCD detector.

2. Algorithms

The first algorithm for the recovery of $\rho(\mathbf{r})$ in this type of experiment is due to Gerchberg & Saxton (1972). In their experiment, both the scattering and an image of the sample were available and the Gerchberg–Saxton (GS) algorithm finds a solution by Fourier transforming a possible solution between real and reciprocal space using the image in real space and the measured scattering in reciprocal space as constraints on the solution. This is a rather unusual situation and is markedly different from the more common case in which an image is typically not available. Fienup (1978) later generalized the GS algorithm by showing that the real-space constraint could be replaced by *a priori* physical knowledge – such as the finite size of the object or that the diffracting density is real positive. This new algorithm is called the generalized GS or error reduction (ER) algorithm, due to its equivalence – with a careful choice of real-space constraints – to gradient search methods.

ER possesses several undesirable features. Most notably, since it is similar to gradient search algorithms, it has a tendency to stagnate in local minima of the error metric, which is a measure of how well the reconstructed reciprocal-space object matches the measured CXD data. A very common stagnation arises from the tendency of the algorithm to reconstruct different elements of the image and its twin simultaneously. With this in mind, Fienup (1978) proposed the hybrid input/output (HIO) algorithm, which modifies the current iterate by using some fraction of the previous iterate. More recently, Elser (2003) proposed the ‘difference map’ (DM) of which HIO is a special case. It is advantageous to state DM in the context of projection operators acting on N -dimensional vectors in a Hilbert space (Bauschke *et al.*, 2002). With such notation, a projector, π_i , acts on an iterate, the N -dimensional ρ_n , in such a way that $\pi_i\rho_n$ is the closest point to ρ_n lying within the set of all vectors satisfying the

constraint i . A solution to the inverse problem is a vector that lies in the intersection of all of the constraint sets. In the case of the commonly used real-space constraints – positivity (π_+) and support (π_s) – the set of all such vectors is convex. On the other hand, the set comprised of all vectors that obey the

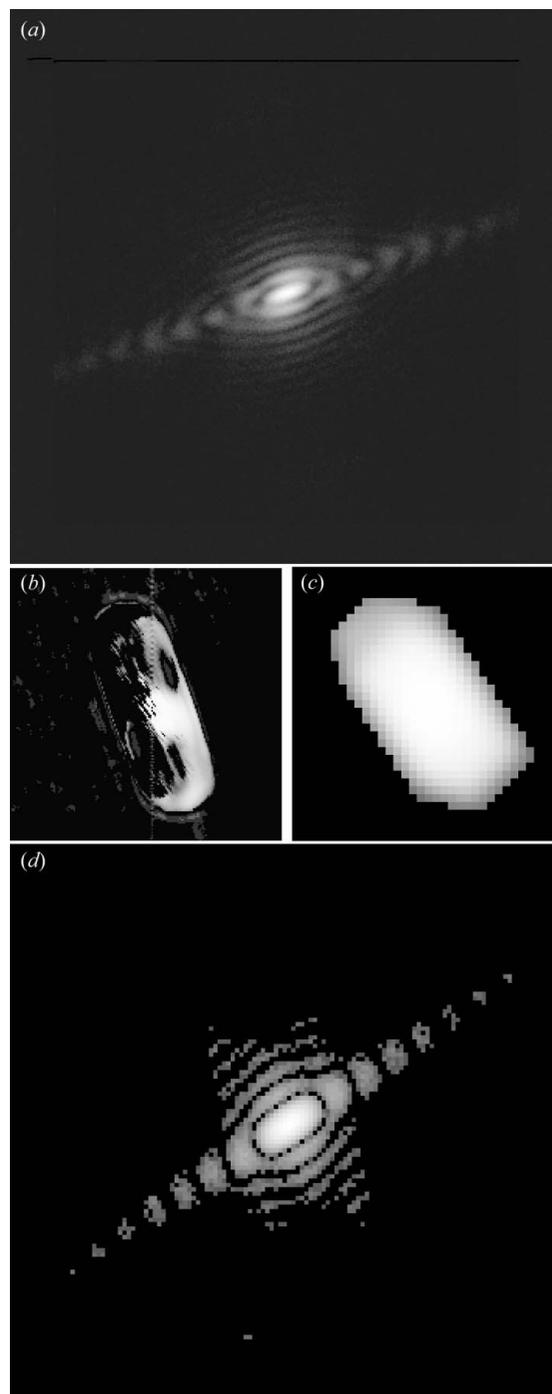


Figure 1

(a) An experimental CXD pattern collected from Au crystals. (b) The real-space density projection arrived at by reconstructing the complex amplitude whose magnitude is shown in (a) using the iterative methods described in the text. (c) A simulated density projection from a ‘crystal’. (d) The Fourier transform of (c). Note the similarities between (a) and (d). (a) and (d) are the logarithm of the ‘intensity’ while (b) and (c) are linearly scaled.

Fourier modulus (π_m) constraint is non-convex. Since this set is not convex, it is difficult to construct an algorithm that consistently finds the solution.

Traditionally, ER (Fienup, 1982) is stated

$$\rho_{n+1}(x) = \begin{cases} \rho_n(x) & \text{if } \rho_n(x) \in S \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

where the set S is all positive values inside the support constraint. Using the notation above, it becomes (Bauschke *et al.*, 2002)

$$\begin{aligned} \rho_{n+1} &= \pi_{s+} \pi_m \rho_n \\ &= \pi_s \pi_+ \pi_m \rho_n. \end{aligned} \quad (4)$$

DM is defined (Elser, 2003) to be

$$\rho_{n+1} = D\rho_n = (1 + \beta[\pi_i f_j - \pi_j f_i])\rho_n, \quad (5)$$

where 1 is the identity and $f_i = (1 + \gamma_i)\pi_i - \gamma_i$ maps ρ_n onto a point on the line connecting itself and its projection onto constraint set i . Three real scalar coefficients are introduced: γ_i , γ_j and β . In the context of the problem here, i would be s or $s+$ and j is m . For the choice of $\gamma_s = -1$ and $\gamma_m = 1/\beta$, equation (5) becomes HIO:

$$\rho_{n+1} = [1 + (1 + \beta)\pi_s \pi_m - \pi_s - \beta\pi_m]\rho_n. \quad (6)$$

In this paper, the behavior of DM with respect to the choice of the scalars β , γ_1 and γ_2 is not explored. It has been our experience that other choices of these parameters may also lead to successful reconstructions. In this paper, we will report the quantity $\pi_m \rho_n$ as the result of our simulated noisy experimental reconstruction. Generally, the progress of the iterative methods are monitored by means of an error metric, which is analogous to a crystallographic R factor (Prince, 2004; Stout & Jensen, 1968). In this paper, we define the metric *via* a sum over all elements in an array:

$$\chi^2 = \frac{\sum_{k=0}^{N-1} |\rho_n(k) - \pi_m \rho_n(k)|^2}{\sum_{k=0}^{N-1} I_k}, \quad (7)$$

where ρ_n is the real-space iterate on cycle n , $\pi_m \rho_n(k)$ is the iterate after application of the Fourier modulus constraint, and I_k is the measured intensity in pixel k . This metric is the traditional one and works very well for ER, where the Fourier modulus constraint is strictly enforced on the iterate. In the case of HIO-like algorithms, it is not appropriate to consider the iterate as an estimate of the solution to the inverse problem and so χ^2 will not reflect the error of the estimate as it does for ER. Nevertheless, for the purpose of classification, we will adhere to tradition and identify the ‘best’ reconstruction as the one with lowest χ^2 . To combat the propensity of the algorithms to become stuck in local minima, it is common to use multiple starting points. When utilizing this tactic, we may define a second error metric in analogy to equation (7):

$$\xi_a^b = \frac{\sum_{k=0}^{N-1} |[\pi_m \rho_n^{(a)}](k) - [\pi_m \rho_n^{(b)}](k)|^2}{\sum_{k=0}^N I_k}. \quad (8)$$

We will calculate two quantities in this way: the reproducibility of the two best reconstructions, in this case, the two best from

the five random phase sets generated, ξ_1^2 , and the fidelity of the best reconstruction to the truth image, ξ_1^o . Defined this way, a large ξ_a^b indicates a poor agreement between the two objects and $\xi_a^b = 0$ would be perfect agreement. The fidelity is the most interesting quantity, but with experimental data we are able to calculate only χ^2 and ξ_1^2 . ξ_1^o will be used to highlight the discrepancy in the behavior of χ^2 for ER and HIO discussed above and, as we will show, ξ_1^2 is more useful than χ^2 in predicting the success of the reconstruction.

3. Expected noise in CXD

The main question addressed in this paper is whether noise increases the likelihood of measuring a factorizable CXD pattern. This is particularly important since it is a necessary condition for non-unique solutions to this problem that the diffracted amplitude be factorizable, for example, given a measurement $|H(Q)|$, if $|H(Q)| = |F(Q)||G(Q)|$, more than one solution will exist corresponding to conjugation operations on the factors. Obviously, a noisy diffraction pattern is not a Fourier transform of a meaningful object, but we also seek to discover if the presence of noise is detrimental to the progression of the reconstruction, *i.e.* if noise encourages stagnation.

In an X-ray diffraction measurement, there is an inherent uncertainty in the measurement of the number of arriving photons governed by the Poisson distribution. Therefore, the measurement at each pixel has an uncertainty that goes as $M^{1/2}$, where M is the number of photons measured. Further noise is introduced by the detector and electronics necessary to measure the signal (see, for example, Janesick, 2001). Typically, a CCD detector is used for measurements of this kind, since two-dimensional arrays are relatively inexpensive, efficient and convenient. CCDs detect photon events by measuring the number of electrons freed in the depletion region of the chip by the collision of the photon with an electron–hole pair. This gives rise to an uncertainty in the number of electrons generated by a one-photon event given by the Fano factor. CCD detectors also possess a ‘dark current’ that is the result of electron–hole pairs splitting due to thermal fluctuations. Complicating matters further, CCDs tend to degrade and accumulate defects from radiation damage with time and the dark current generated in these regions may not be the same as in undamaged ones. There is also electronic noise, *e.g.* $1/f$ noise, present in the measurement. The final contribution to noise we consider is scattering from alien scatterers. This may be scatter from the substrate, incoherently illuminated sections of the sample, slit scatter or air scatter.

4. Test object and trial conditions

In this paper, we seek to investigate the effects of these experimental realities by means of a simulation. The procedure is as follows: (i) select the density projection of the object being ‘measured’; (ii) calculate the fast Fourier trans-

form (FFT) of the object; (iii) select the number of photons present in the measurement; (iv) using the FFT from (ii) calculate the intensity for use as a probability for determining the location of photons; (v) based on the integrated intensity in (iv) and the number of photons in (iii), calculate the average intensity per photon; (vi) at each pixel in the simulated CXD pattern, draw a number from a Poisson random distribution with mean given by the intensity in that pixel divided by the average intensity per photon. The result is an image with the inherent Poisson distributed noise that would be observed in a perfect measurement. This image is then converted from photons to an integer number of ‘analog to digital units’ (ADUs) mimicking what happens in an actual experiment. In the simulations here, the exact number of ADUs is drawn from a Gaussian random distribution with mean and width consistent with CXD data. The simulated pattern in Fig. 1(*d*) was generated in this way from the real-space density projection in Fig. 1(*c*). At this point, we are free to add alien scattering or bias to the pattern.

Alien scattering is assumed to be spatially uniform, so that no pixel is more likely to receive scatter than any other. This might correspond to air scatter in an experiment. The scatter is generated by assuming a mean number of photons per pixel, here between 0.005 and 0.5, which specifies the mean for a Poisson random distribution from which the actual simulated scatter is drawn. We then convert this integer number of photons to ADUs. A bias level would primarily arise from the dark current of the CCD. From the experimental data, we can derive the width of a Gaussian distribution from which we draw the actual bias contribution in ADUs. Here, a large mean is chosen, corresponding to a long exposure. A number is drawn from this distribution for each pixel and then a mock background subtraction is carried out. This consists of subtracting some constant from every element of the array. One departure from reality is that fractional ADUs from the bias and subtraction are allowed.

The object in Fig. 1(*c*) shows very good agreement between its simulated diffraction (Fig. 1*d*) and the diffraction measured from a small gold crystal (Fig. 1*a*); however, this shape is suboptimal because it is very symmetric, making it difficult to determine when stagnation due to mixing with its twin has occurred. So instead we use the object in Fig. 2(*a*), which is asymmetric and has very smooth edges. This latter feature reduces the aliasing common to FFTs. Fig. 2(*c*) is the FFT of the shape and is used to calculate Fig. 2(*d*), the simulated CXD pattern with 9×10^4 photons. Our experience is that CXD patterns with this number of photons can be collected in exposure times of tens of seconds for micrometre-sized objects. The utility of simulating a pattern from scratch, rather than adding noise to an existing data set, is that we can generate a ‘truth image’ to compare the reconstruction to. This is created by allocating a certain number of photons to a pattern while keeping the phase of the complex amplitude from the FFT for pixels that receive at least one photon. This complex finite-photon pattern is back transformed to generate the truth image. Fig. 2(*b*) is the truth image for those tests with 9×10^4 photons. As an illustration, Figs. 2(*e*) and (*f*) are

simulated patterns with high levels of alien scattering and added bias, respectively.

We assign a signal-to-noise ratio calculated as

$$\text{SNR} = \frac{\sum_{k=0}^{N-1} I_k}{\sum_{k=0}^{N-1} (I_k)^{1/2} + P_k + B_k}, \quad (9)$$

where I_k is the number of ADUs in a pixel generated by the six-step procedure above, P_k is the signal due to photons scattered into the detector by alien sources and B_k is the bias level arising from the detector and electronics. It is not expected that different kinds of noise will produce the same result in the reconstructed real-space image, so the SNR in one test may not yield equivalent results in another.

5. Tests

Here, we examine the effect of three scenarios on the quality of the reconstruction: finite photon number, alien scattering

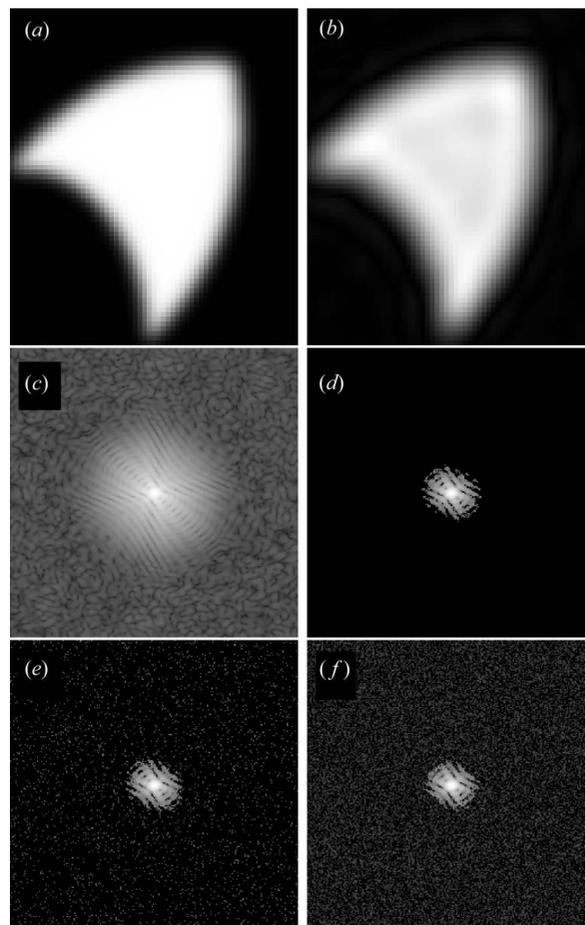


Figure 2

(*a*) A simulated density projection chosen for its smooth edges and non-symmetric shape. (*b*) The truth image when 9×10^4 ‘photons’ are present in the simulated CXD pattern. (*c*) The Fourier transform of (*a*) (log-scaled ‘intensity’ shown). (*d*) The simulated CXD pattern with 9×10^4 photons and the Fourier transform of (*b*) (log-scaled ‘intensity’ shown). (*e*) The simulated CXD after adding alien scattering noise. (*f*) The simulated CXD pattern adding bias, as from the detector background.

collected in the detector, and finite bias levels in the data due to detector noise. For each of the tests in this section, the standard procedure was to perform five reconstructions for each algorithm at each noise level. In §§5.1–5.3, 500 iterations were performed and, in §5.4, 300 were performed. Each of the five reconstructions starts with a different set of random phases. The HIO parameter β was held constant in §§5.1–5.3 at 1 and is varied in §5.4. For each reconstruction and algorithm, three values are reported: the error metric, χ^2 , the reproducibility, ξ_1^2 , and the fidelity to the truth image, ξ_1^0 . The simulated patterns were generated in a 256×256 pixel array. The reconstruction was carried out in the same array size. The real-space support constraint was 60×70 pixels and ER and HIO both required that the reconstructed density be real positive. Fig. 2(d) was used as the Fourier modulus constraint.

An interesting, albeit purely qualitative, result is that a noisy simulated CXD pattern often gives rise to a real-space density projection that also appears ‘noisy’. Fig. 3 is a collection of real-space best reconstructions resulting from such simulated patterns. In this context, we regard the smearing of boundaries and the apparently random internal density modulations as defining characteristics of a noisy real-space image. The first two panels, denoted (a) and (b), are the best result of an ER-only routine using alien-scatter-added CXD patterns as input. The next two panels are the result of reconstructing an imperfectly background subtracted simulation with ER only and HIO only, demonstrating that HIO also possesses this characteristic. The final two panels are the best reconstructions to a perfectly background subtracted simulation reconstructed with only HIO and different β s. An

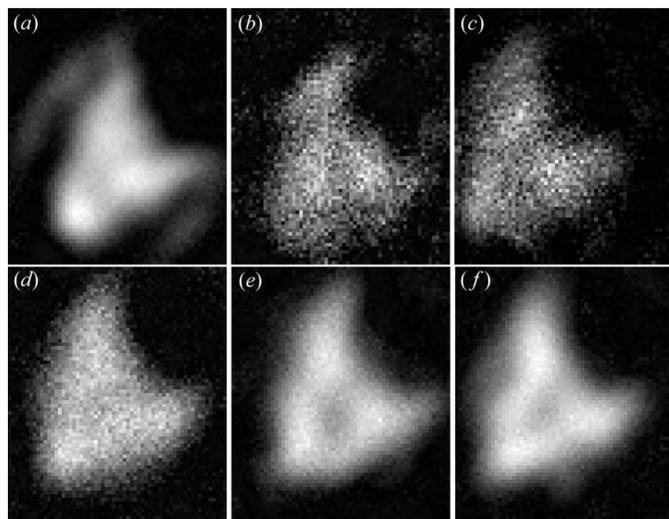


Figure 3 Real-space results from the reconstruction of noisy data demonstrating that the addition of noise in the simulated CXD pattern can lead to a ‘noisy’ appearance in the real-space result. Shown is the best result from each test where the support was the region displayed. (a) and (b) are the result of reconstructing data with alien scatter with Poisson means of 0.05 and 0.5 per pixel using only ER. (c) and (d) are the result of reconstructing the bias-added simulation with 95% background subtracted using only ER and only HIO. (e) and (f) are the result of reconstructing the bias-added simulation with 100% background subtraction using only HIO with $\beta = 0.2$ and $\beta = 0.5$.

important characteristic of these solutions is that comparison with the second-best result reveals that the pattern of internal contrast is not reproducible.

5.1. Photon number

Varying the number of photons in a simulated pattern between 9×10^4 and 10^9 gives SNR between 290 and 2.3×10^4 . Fig. 4(a) shows the trends of ξ_1^0 , ξ_1^2 and χ^2 with decreasing SNR. The graph shows a clear trend of increasing error metric with decreasing SNR as expected. The ξ_1^0 values are more interesting, as they remain more or less constant for reconstructions using HIO until $\text{SNR} = 10^3$, whereupon the fidelity rapidly degrades. The procedure using ER produces worse estimates of the truth image and, while the trend of ξ_1^0 with SNR is less clear, there is a modest increase as SNR decreases.

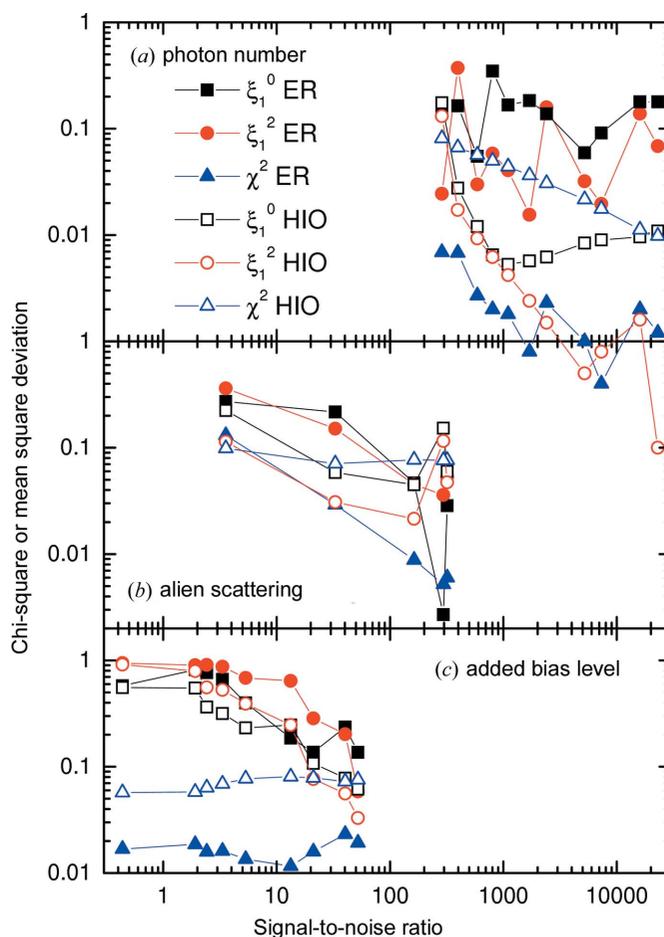


Figure 4 Scatter plots of reproducibility ξ_1^2 , fidelity ξ_1^0 and error metric against SNR for the three tests of noise level described in the text. Lower values of all three quantities defined in the text indicate better agreement. As expected, as photon number increases (a), the estimate of the truth object becomes more accurate. (b) shows that added alien scatter produces a more or less linear relationship between SNR and ξ_1^0 . Added bias, (c), also produces a linear dependence, but with a much smaller slope, indicating the algorithms do not perform as well in this case as in the case of alien scattering. Finally, we note that the behavior of ξ_1^2 more closely mirrors the behavior of ξ_1^0 than does χ^2 , revealing its importance as an indicator of a correct solution in the presence of noisy diffraction data.

For the lowest values of SNR, ER and HIO reconstructions seem to give roughly equivalent estimates of the simulated pattern. We also see that χ^2 is deceiving, inasmuch as its value is always better for reconstructions using ER than those using HIO, which is not the trend we see in the measure of fidelity to the truth image, ξ_1^o . In other words, this is a demonstration of the assertion in §2 that χ^2 is not an ideal metric for HIO. An encouraging result is that ξ_1^2 seems to be a much better indicator of ξ_1^o than is χ^2 . For the next two sections, we choose the worse case – the simulated pattern with SNR = 290 – for investigating the effect of alien scatterers and added bias levels. This case is similar to the experimental data in Fig. 1, where we estimate that 2×10^5 photons were captured.

5.2. Alien scatterer

Specifying a mean number of alien photons between 0.0005 and 0.5 per pixel, as described above, yields simulated patterns with SNR between 296 and 3. Fig. 4(b) shows the result from reconstructing these simulations. In this case, ER and HIO procedures both display a roughly linear dependence of increasing ξ_1^o and ξ_1^2 with decreasing SNR. It is interesting to note that, in this case, χ^2 for HIO is essentially unaffected by increasing noise while ξ_1^o grows steadily worse. χ^2 for ER in this case tracks with ξ_1^o and ξ_1^2 correctly indicating that the reconstruction quality is decreasing. In all, HIO arrives at a better estimate of the truth image and we again see that χ^2 does not predict the quality of the HIO reconstruction.

5.3. Added bias

By adding a bias level generated by drawing a different number from a Gaussian random distribution about a common mean for each pixel and then subtracting a common constant from every pixel, we mimic the noise inherent in a background subtraction from our CCD detector. Fig. 4(c) summarizes the result of reconstructing simulated patterns with background subtraction of 0, 10, 30, 50, 70, 90, 95, 99 and 100% of the mean of the Gaussian distribution, which yields SNR between 0.4 and 52. In this case, χ^2 is a very poor indicator of a good reconstruction. In fact, χ^2 is seen to decrease slightly with decreasing SNR, which is very misleading. By contrast, ξ_1^o and ξ_1^2 track together and increase more or less linearly with decreasing SNR, as expected. Again, HIO produces slightly better agreement with the truth image than ER. It is interesting to note that the linear dependence of increasing ξ_1^o with decreasing SNR is steeper for the case of added bias than for the case of alien scattering. This seems to indicate that the algorithms are less tolerant of a bias level than of random alien scattering.

5.4. Mixing algorithms

It has been reported previously that a combination of HIO and ER is particularly effective (Fienup & Wackerman, 1986) and this procedure has been used to reconstruct experimental data (Robinson *et al.*, 2001; Williams *et al.*, 2003). It is, therefore, interesting to investigate the effect of the β parameter in such a combination. To see if there exists a particular β for

which an ER/HIO combination outperforms ER and HIO separately, we choose a high SNR simulated pattern (with 5×10^6 photons) and conduct five reconstructions as above except that the reconstruction recipe is now 150 iterations of ER, 50 HIO and 100 ER rather than 500 of either of the algorithms alone. Fig. 5 displays the result of these trials.

For reference, the HIO-only reconstruction of this pattern resulted in: $(\chi^2)^{1/2} = 0.19$, $(\xi_1^2)^{1/2} = 0.04$, $(\xi_1^o)^{1/2} = 0.08$. The ER-only gave: $(\chi^2)^{1/2} = 0.03$, $(\xi_1^2)^{1/2} = 0.43$, $(\xi_1^o)^{1/2} = 0.13$. These values are in line with the trends above, χ^2 for HIO remains high but ξ_1^2 tracks ξ_1^o . That χ^2 for the best ER reconstruction is low while ξ_1^2 remains high is indicative of the propensity for ER to become trapped in local minima, that is, only one of the starting guesses was able to find a good estimate in this number of iterations. As the graph shows, a negative β , while permitted by DM, is not favorable under these conditions. For positive values of β , the combination algorithm outperforms ER- and HIO-only procedures for almost every β , including $\beta = 1$, the value used in §5.1. The results are not acutely sensitive to the value of β but appear to be best in the range $0.5 < \beta < 1.0$. This is consistent with our experience in reconstructing CXD data.

6. Conclusions

In using iterative methods to recover lost phase information from a CXD imaging experiment, it is imperative to consider the effect of noise on the resulting recovered phases before making quantitative statements about the sample. For this purpose, we have simulated a noisy CXD pattern to investigate the type and severity of artifacts introduced by the

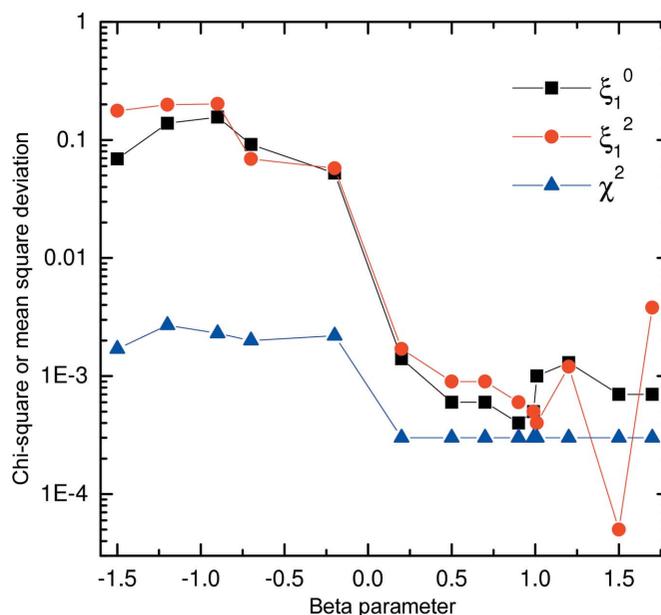


Figure 5 Scatter plot of reproducibility ξ_1^2 and fidelity ξ_1^o against SNR^{-1} for various values of the HIO β in a typical combination of ER and HIO, as described in the text. The combination performs best when β is chosen to be a little less than 1.

presence of noise. We have shown that, while χ^2 alone may not indicate a good estimate of the truth image, the combination of χ^2 and the calculable reproducibility ξ_1^o should allow the safe identification of a good estimate to the generally unknown truth image. We have also demonstrated that χ^2 is not, in general, a good metric for HIO. This is hardly surprising in that the iterate is not to be regarded as a potential solution in this scheme.

Of further interest is whether noise is likely to cause meaningful non-uniqueness of the final solution. Since ξ_1^o and ξ_1^2 tend to behave in the same way, it seems that noisy data increase the number of solutions that agree with the data equally well. However, these are not multiple solutions in the broadest sense; rather, they manifest themselves as fluctuation about or variations on the truth image not unlike a real-space 'noise'. Therefore, it may be reasonable to use multiple poor solutions to form a most likely solution, depending upon the situation. In our comparison of different kinds of noise, we see that scattering from alien sources is troublesome, but it may be ignored if its presence is not overwhelming. For example, a simulation in which on average 1 of every 20 pixels contains a photon scattered from an alien source induces RMS errors of about 10%. Additional background is more harmful. Not only does the fidelity grow worse more quickly with our measure of SNR, but ER – and to some extent HIO – will allow vortices (phase singularities) in the reciprocal-space amplitude or, equivalently, stripes (Fienup & Wackerman, 1986) in the real-space density projection. These stripes might be mistaken for actual physical phenomena, so care must be taken when a good background subtraction is not performed. Even so, assuming vortices are not formed, the resulting reconstruction is not materially different from the truth image when a bias level is present. Therefore, a good reconstruction with a bias level present is still a good reconstruction, albeit one with a large zero-order Fourier component arising from the high average value of the reciprocal-space pattern.

As expected, simulations with large numbers of photons allow a better estimate of the truth object. Interestingly, in the case of shot noise alone, a low χ^2 during HIO seems to imply a good reconstruction, but is deceiving in the presence of other contamination. Fortunately, at a third-generation synchrotron, count rates are typically very high and can be increased for small samples through the use of focusing optics. Therefore, we expect that the number of photons in the simulations presented here are realistic for a broad variety of samples that can be measured in modern synchrotron experiments.

It has been shown that combinations of ER and HIO are highly effective at increasing both the reproducibility of reconstructions and the fidelity to the truth object. Several strategies may be employed to further enhance the quality of the reconstructions presented here. For example, HIO is extremely sensitive to the tightness of the support constraint. In this paper, we have not altered the support, which has been reported to improve the result of the reconstruction (Marchesini *et al.*, 2003).

Lastly, it must be said that these results may depend to some extent on the kind of object or real-space density projection chosen. We chose a single compact object with smooth edges. Further, it possesses a highly asymmetric shape and is purely real. Nevertheless, it seems that realistic levels of noise in simulated data generated in agreement with experimental CXD data do not materially hamper the effectiveness of ER, HIO or ER/HIO combinations.

This research was supported by the NSF under grant DMR 03-08660. The UNICAT facility at the Advanced Photon Source (APS) is supported by the University of Illinois at Urbana-Champaign, Materials Research Laboratory (US DOE contract DEFG02-91ER45439, the State of Illinois–IBHE–HECA, and the NSF), the Oak Ridge National Laboratory (US DOE under contract with UT-Battelle LLC), the National Institute of Standards and Technology (US Department of Commerce) and UOP LLC. The APS is supported by the US DOE, Basic Energy Sciences, Office of Science under contract No. W-31-109-ENG-38. SEM work was carried out in the Center for Microanalysis of Materials, University of Illinois.

References

- Barakat, R. & Newsam, G. (1984). *J. Math. Phys.* **25**, 3190–3193.
- Bates, R. H. T. (1982). *Optik (Stuttgart)*, **61**, 247–262.
- Bauschke, H. H., Combettes, P. L. & Luke, D. R. (2002). *IEEE ICIP*, pp. II841–II844.
- Bruck, Y. M. & Sodin, L. G. (1979). *Optics Commun.* **30**, 304–308.
- Elser, V. (2003). *J. Opt. Soc. Am.* **A20**, 40–55.
- Fienup, J. R. (1978). *Optics Lett.* **3**, 27–29.
- Fienup, J. R. (1982). *Appl. Optics*, **21**, 2758–2769.
- Fienup, J. R. & Wackerman, C. C. (1986). *J. Opt. Soc. Am.* **A3**, 1897–1907.
- Gerchberg, R. W. & Saxton, W. O. (1972). *Optik (Stuttgart)*, **35**, 237–246.
- Janesick, J. R. (2001). *Scientific Charge-Coupled Devices*. New York: SPIE Press.
- Marchesini, S., He, H., Chapman, H. N., Hau-Riege, S. P., Noy, A., Howells, M. R., Weierstall, U. & Spence, J. C. H. (2003). *Phys. Rev. B*, **68**, 140101.
- Miao, J., Charalambous, P., Kirz, J. & Sayre, D. (1999). *Nature (London)*, **400**, 342–344.
- Miao, J., Ishikawa, T., Johnson, B., Anderson, E., Lai, B. & Hodgson, K. O. (2002). *Phys. Rev. Lett.* **89**, 088303.
- Miao, J., Sayre, D. & Chapman, H. N. (1998). *J. Opt. Soc. Am.* **A15**, 1662–1669.
- Prince, E. (2004). Editor. *International Tables for Crystallography*, Vol. C. Dordrecht: Kluwer Academic Publishers.
- Robinson, I. K., Vartanyants, I. A., Williams, G. J., Pfeifer, M. A. & Pitney, J. A. (2001). *Phys. Rev. Lett.* **87**, 195505.
- Sayre, D., Chapman, H. N. & Miao, J. (1998). *Acta Cryst.* **A54**, 232–239.
- Seldin, J. H. & Fienup, J. R. (1990). *J. Opt. Soc. Am.* **A7**, 412–427.
- Stout, G. H. & Jensen, L. H. (1968). *X-ray Structure Determination: a Practical Guide*. New York: Macmillan.
- Warren, B. E. (1990). *X-ray Diffraction*. New York: Dover.
- Williams, G. J., Pfeifer, M. A., Vartanyants, I. A. & Robinson, I. K. (2003). *Phys. Rev. Lett.* **90**, 175501.