research papers

SPEDEN: reconstructing single particles from their diffraction patterns

Stefan P. Hau-Riege,[a]* Hanna Szoke,[a] Henry N. Chapman,[a] Abraham Szoke,[a] Stefano Marchesini,[a] Alexander Noy,[a] Haifeng He,[b] Malcolm Howells,[b] Uwe Weierstall[c] and John C. H. Spence[c]

[a]Physics and Advance Technology, Lawrence Livermore National Laboratory, PO Box 808, Livermore, CA 94539, USA, [b]Advanced Light Source, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA, and [c]Department of Physics and Astronomy, Arizona State University, Tempe, AZ 85287-1504, USA. Correspondence e-mail: hauriege1@llnl.gov

SPEDEN is a computer program that reconstructs the electron density of single particles from their X-ray diffraction patterns, using a single-particle adaptation of the holographic method in crystallography [Szöke, Szöke & Somoza (1997). Acta Cryst. A53, 291–313]. The method, like its parent, is unique because it does not rely on 'back' transformation from the diffraction pattern into real space and on interpolation within measured data. It is designed to deal successfully with sparse, irregular, incomplete and noisy data. It is also designed to use prior information for ensuring sensible results and for reliable convergence. This article describes the theoretical basis for the reconstruction algorithm, its implementation, and quantitative results of tests on synthetic and experimentally obtained data. The program could be used for determining the structures of radiation-tolerant samples and, eventually, of large biological molecular structures without the need for crystallization.

1. Introduction

This paper describes our computer program SPEDEN, which reconstructs the density from the diffraction patterns of individual particles. SPEDEN is of interest for three reasons. Diffractive imaging promises to improve the resolution, sensitivity and practical wavelength range in X-ray microscopy for three-dimensional objects that are tolerant to X-rays. A few examples are defects in semiconductor structures, phase separation in alloys, nanoscale machines and laser fusion targets. A long-term vision is the possibility of high-resolution reconstruction of diffraction patterns of single biomolecules. Of broad theoretical interest is SPEDEN's unique approach to the reconstruction of scatterers – a difficult mathematical problem. In the rest of this section, we expand on these three topics.

Reconstruction of the electron density from non-uniformly sampled three-dimensional diffraction patterns is of wide interest and applicability with present-day sources. In radiation-tolerant samples, X-ray diffraction and diffraction tomography are capable of higher resolution than (straight- or cone-beam) tomography alone. In tomography, resolution is limited by the quality of the incident beam and by the spatial resolution of the detector; in diffraction, the resolution can be as fine as the wavelength of the incident radiation. Experi-

mentally, diffraction imaging has already produced X-ray images at higher resolution than possible with available X-ray optics (Miao et al., 1999; He et al., 2003). The price to be paid for these benefits is the intrinsic difficulty of the reconstruction. Nevertheless, several successful reconstructions from experimental X-ray data, using the iterative hybrid input–output version of the Gerchberg–Saxton–Fienup (GSF) algorithm, have been reported recently (Miao et al., 2002; Marchesini et al., 2003). The first successful application of this algorithm to electron diffraction data was reported in 2001 (Weierstall et al., 2002), and it has been used more recently to produce the first atomic resolution image of a single carbon nanotube (Zuo et al., 2003). In biology, the use of the GSF has recently been shown to dramatically reduce the number of images needed for tomographic cryo-electron-microscopy of protein monolayer crystals, so that phasing can be based mainly on the three-dimensional diffraction data (Spence et al., 2003).

The development of SPEDEN was also prompted by the promise of new ways to image biomolecules. Free-electron lasers can, in principle, provide X-ray pulses of tens to hundreds of femtoseconds in length and brightness up to ten orders of magnitude greater than synchrotron radiation. It was predicted that, under such circumstances, it should be possible to dispense with crystals and reconstruct the electron density

of single biological particles from their diffraction pattern (Neutze *et al.*, 2000). In proposed experiments, a large number of single particles will be injected into the X-ray beam in random orientation and their diffraction patterns will be recorded, each in a single shot of the free-electron laser. Such diffraction patterns will be very noisy and their resolution will be limited by the signal-to-noise (S/N) ratio. The measured diffraction patterns that correspond to different orientations of the particle will be classified into a number of mutually exclusive classes. The images within each class will then be averaged and the class averages assembled into a three-dimensional diffraction pattern by finding the mutual orientation of the individual diffraction patterns. Finally, the three-dimensional diffraction pattern will be reconstructed to yield the electron density of the molecule.

We have worked on the analysis of all three steps of such an experiment. The essence of the first analysis is that the maximum X-ray intensity at a given pulse length is limited by the requirement that the molecule stay intact during the pulse, even though it eventually disintegrates (Hau-Riege *et al.*, 2004). The second analysis discusses the division of noisy diffraction patterns into a number of distinct classes. If the images are divided into too few classes, the available resolution is not realized. If the patterns are divided into too many classes, the class averages will be poor and the pattern quality suffers. The individual class averages, each corresponding to a well defined orientation of the particle, will be assembled into a three-dimensional diffraction pattern. The result will be a three-dimensional diffraction pattern that is measured at a limited number of orientations. It will be, therefore, sparse and irregular and will have a limited signal-to-noise ratio (Huldt *et al.*, 2003).

The program *SPEDEN*, described in this paper, provides a way to optimally determine the electron density from such a three-dimensional ensemble of continuous diffraction patterns. We first give an analysis of their properties and discuss the methods and the expected difficulties of reconstructing a 'sensible' electron density from them. We then describe how *SPEDEN* adapts the holographic method (Szöke, 1997) in crystallography to deal with continuous diffraction patterns as opposed to discrete Bragg spots; this will be discussed in the next section. We then report quantitative results of preliminary tests for verifying the correctness of our method. These tests use computed and measured diffraction patterns from samples of inorganic particles.[1]

2. *SPEDEN*: the method

2.1. Theoretical considerations

2.1.1. Mathematical background. The reconstruction of the density of scatterers from its diffraction pattern is an 'inverse problem'. Other well studied inverse problems are those of computed tomography, image deblurring, phase recovery in astronomy, and crystallography. In tomography, for example,

an inversion algorithm (*e.g.* filtered back projection) is used to recover the density of scatterers from the measured tomograms. It is widely recognized that the reconstructed density is very sensitive to inaccuracies in the measurement. Small errors in the diffraction pattern cause large errors in the reconstruction. This property is called ill posedness or ill conditioning.

The reconstruction of the electron density from X-ray diffraction patterns is indeed ill conditioned. It also has two additional difficulties. First, in contrast to tomography, there are no direct inversion algorithms – not even approximate ones. Second, the reconstructed electron density at any sample point is influenced strongly by the electron densities of all sample points, as opposed to a limited number of them. Therefore, errors in density are non-local and 'propagate' far.

Fortunately, very good fundamental discussions of these subjects are provided in the books of Daubechies (1992), Bertero and Boccacci (Bertero, 1989; Bertero & Boccacci, 1998) and Natterer (Natterer, 1996; Natterer & Wübbeling, 2001). In somewhat simplified terms, the reconstruction of the electron density is similar to finding the inverse of an ill conditioned non-square matrix, a subject thoroughly discussed in Golub & van Loan (1996). We consider these mathematical properties to be essential for understanding the successes and limitations of reconstruction algorithms; we will try to be fully cognizant of them in the discussion that follows.

2.1.2. The phase problem of crystallography, oversampling. The crystallographic phase problem is a good starting point for further discussion. It was first realized by Sayre (1952) that the number of observable complex structure factors, limited by the Bragg condition, is equivalent to a critical sampling of the electron density in the unit cell of the crystal. The sampling theorem of Whittaker and Shannon teaches us that, if the amplitudes and phases of all the diffraction peaks were accurately measured, the electron density could, in principle, be reconstructed everywhere (Bricogne, 1992). Unfortunately, only the amplitudes of the Bragg reflections are measured, not their phases. Therefore there is not enough information in the diffraction pattern for a unique reconstruction of the electron density. Sayre (1980) proposed that if we could measure the diffraction amplitudes 'in between' the Bragg peaks, we should have enough information to reconstruct the electron density, or to 'phase' the diffraction pattern. This is exactly the situation in diffraction from a single particle.

Nevertheless, it is still difficult to reconstruct the electron density, even from a well 'oversampled' diffraction pattern. One corollary of critical sampling is that the amplitudes and phases of the Bragg reflections of a crystal are independent of one other, but any structure factor in between them depends on the surrounding ones to some extent. Therefore, too much oversampling does not help to obtain independent data, although it does improve the S/N ratio by reducing the noise. Ill posedness is still with us, although with oversampling the error propagates less. An additional difficulty with diffraction patterns from a set of discrete orientations of a particle is that at low resolution the diffraction pattern is well oversampled while at high resolution the sampling is sparse. A fundamental

---

property of diffraction is that the position and the handedness of the electron density are undetermined, resulting sometimes in stagnation of the algorithm and drift in the position of the results (Stark, 1987).

There are two well known necessary remedies for the lack of information and for the ill posedness of the reconstruction problem. The more important one is the need for more information. For example, one way to include *a priori* knowledge is to accept reconstructed electron densities only if they are 'reasonable'. The second remedy is to use 'stabilized' or pseudo-inversion algorithms. In the next section, we introduce our version of a real-space reconstruction algorithm; we will argue that our algorithm deals with all these problems optimally, at least in some sense. We return to the comparison of our algorithm with other methods for phase recovery in §2.3.

2.2. Details of SPEDEN, a real-space algorithm

In this section, we outline the workings of our reconstruction program, *SPEDEN*. *SPEDEN* uses a real-space method for reconstruction; its acronym stands for single-particle electron density. For computational efficiency, the particle to be recovered is put into a fictitious unit cell that is several times larger than the particle itself. All reconstruction algorithms use this artifice in order to be able to calculate structure factors by fast Fourier transform (FFT) techniques. The resulting similarity with crystallography enables the use of many crystallographic concepts. In fact, the recognition of this similarity enabled us to write *SPEDEN* based on our crystallographic program, *EDEN*, with relatively small modifications.

The most significant difference between the two programs is that in crystallography the Bragg condition restricts the reciprocal-lattice vectors to integer values, while the continuous diffraction pattern can be – and usually is – measured at arbitrary non-integer values of the reciprocal-lattice vectors.

In *SPEDEN*, in common with *EDEN*, the (unknown) electron density is represented by a set of Gaussian basis functions, with unknown amplitudes, that fill the fictitious unit cell uniformly. This way the recovery is reduced to the solution of a large set of quadratic equations. The program 'solves' these equations by finding the number of electrons in each basis function so as to agree optimally with the measured diffraction intensities as well as with other 'prior knowledge'. Prior knowledge includes the emptiness of the unit cell outside the molecule, the positivity of the electron density, possibly some low-resolution image of the object *etc*. Each one of those conditions is described by a cost function that measures the deviation of the calculated data from the observed data. One of the cost functions describes the deviation of the calculated diffraction pattern from the measured one; others depend on the deviation of the recovered density from prior knowledge. Measured data are weighted by their certainty (inverse uncertainty), other prior knowledge is weighted by its 'reliability'. The mathematical method used is (constrained) conjugate-gradient optimization of the sum of cost functions.

At each step of the optimization, there is a set of amplitudes available that describe the current electron density in the full unit cell. A full set of structure factors is calculated by Fourier transforming the current electron density. When the unit cell is larger than the particle, the structure factors can be stably interpolated to compare them with measured structure-factor amplitudes.

We refer to the cited literature that shows that the procedure we outlined is equivalent to a stabilized (quasi) solution of the inverse problem (Daubechies, 1992; Bertero & Boccacci, 1998; Natterer & Wubbeling, 2001). As such, it is optimally suited for sparse, irregular, incomplete and noisy data.

In the following subsections, we describe very briefly the common features of *EDEN* and *SPEDEN* as well as their differences. A more complete description of *EDEN* can be found in previous papers (Szöke, 1997, 1998).

2.2.1. Representation of the electron density. The electron density is represented as a sum of basis functions, adapted to the resolution of the data. Specifically, we take little Gaussian 'blobs' of width comparable to the resolution, and put their centers on a regular grid that fills the 'unit cell' and whose grid spacing is comparable to the resolution. The amplitudes of the Gaussians are proportional to the local electron density. In fact, the number of electrons contained in each Gaussian constitutes the set of our basic unknowns. The above is identical to the representation of the electron density in *EDEN*.

Some mathematical details follow. The actual formula for the representation of the electron density as a sum of Gaussians is

$$\rho_{\text{unknown}}(\mathbf{r}) \cong \frac{1}{(\pi\eta\Delta r^2)^{3/2}} \sum_{p=1}^{P} n(p) \exp\left[\frac{-|\mathbf{r} - \mathbf{r}(p)|^2}{\eta\Delta r^2}\right]. \quad (1)$$

The centers of the Gaussians are positioned at grid points, $\mathbf{r}(p)$, where p is a counting index. In our fictitious unit cell, the grid is orthogonal, the grid spacing is Δr and the centers of the Gaussians are usually on two intercalating (body-centered) grids for best representation of the electron density. The number η, of the order unity, determines the width of the Gaussians relative to their spacing, Δr. Finally and most importantly, $n(p)$ is the unknown number of electrons in the vicinity of the grid point $\mathbf{r}(p)$. The values of $n(p)$ are real, and in future may also be complex valued to allow for photoabsorption in addition to scattering. (Photoabsorption can be significant when diffraction measurements are made at longer X-ray wavelengths.)

Given $n(p)$, the structure factors can be calculated by

$$F_{\text{calc}}(\mathbf{h}) = \exp[-\eta(\pi\Delta r|\mathcal{F}^T\mathbf{h}|)^2] \sum_{p=1}^{P} n(p) \exp[2\pi i\mathbf{h} \cdot \mathcal{F}\mathbf{r}(p)],$$

$$(2)$$

using a fast Fourier transform. The vector $\mathbf{h}$, a triplet of integers, denotes the reciprocal-lattice vector, the operator $\mathcal{F}$ transforms from real-space (Cartesian) coordinates to fractional coordinates, and $\mathcal{F}^T$ denotes the dual transformation.

The constants appearing in equations (1) and (2) were discussed in some detail previously (Szöke *et al.*, 1997). For completeness, we define them here. The crystallographic *B* factor is given by $B = (2\pi\Delta r)^2\eta$. The 'crystallographic resolution', *d*, determines the grid spacing, Δr, by the relation $\Delta r \approx f_{space}d$, where f_{space} is a constant of the order unity. For a body-centered lattice, we set $f_{space} = 0.7$ and $\eta = 0.6$. For a simple lattice, we use $f_{space} = 0.6$ and $\eta = 0.8$.

Note that the Gaussian basis functions are not used in a one-to-one correspondence with single atoms but are simply used to describe the three-dimensional electron density at the resolution that is appropriate to the data resolution. In the special case that the resolution was about the size of an atom and an atom happened to be sitting exactly on a grid point, that atom would be represented by a single basis function. If the atom is not on a grid point, or if the atom happens to be smeared out because of thermal motion, that same atom would be represented by many basis functions. Similarly, at lower resolution, a single basis function represents assemblies of atoms.

2.2.2. Reciprocal-space knowledge. The measured diffraction pattern of the molecule is proportional to the absolute square of the structure factors. In *SPEDEN*, we do account for the curvature of the Ewald sphere. There are two subtle points: the diffraction pattern is measured only in a finite number of directions, $\mathbf{H}(i)$; and, as a rule, those directions are not along the reciprocal-lattice vectors of the (fictitious) unit cell. In other words, the measurement directions, $\mathbf{H}(i)$, are usually not integers and they are not uniformly distributed in reciprocal space. This is the main difference between crystallography and single-particle diffraction and, therefore, between *EDEN* and *SPEDEN*. The essence of any reconstruction algorithm is to try to find an electron-density distribution such that the calculated diffraction pattern matches the observed one. In our representation, we try to find a set of $n(p)$ such that

$$|F_{obs}(\mathbf{H}(i))|^2 = |F_{calc}(\mathbf{H}(i))|^2 \qquad (3)$$

for each measurement direction $\mathbf{H}(i)$. Let us assume for a moment that $\mathbf{H}(i)$ are integers. When the representation of the unknown density is substituted from (2), for each measured value of $\mathbf{H}(i)$, (3) becomes a quadratic equation in the unknowns, $n(p)$. The number of equations is the number of measured diffraction intensities. It is usually not equal to the number of independent unknowns that are the number of grid points in the unit cell. The equations usually contain inconsistent information owing to experimental errors. The equations are also ill conditioned and therefore their solutions are extremely sensitive to noise in the data. Under these conditions, the equations may have many solutions or, more usually, no solution at all. Our way of circumventing these problems is to obtain a 'quasi-solution' of (3) by minimizing the discrepancy or cost function (see *e.g.* Stark, 1987; Bertero & Boccacci, 1998):

$$\chi_{Speden} = \sum_i w'(\mathbf{H}(i))^2[|F_{obs}(\mathbf{H}(i))| - |F_{calc}(\mathbf{H}(i))|]^2. \qquad (4)$$

The weights, $w'(\mathbf{H}(i))^2$, are usually set to be proportional to the inverse square of the uncertainty of the measured structure factors, $1/\sigma(\mathbf{H}(i))^2$. As discussed by Szöke (1999), this is equivalent to a maximum-likelihood solution of the equations.

Let us now discuss the first, previously ignored, difficulty in the reconstruction. When we try to reconstruct the electron density from real experimental data, we have to compare the set of measured $|F_{obs}(\mathbf{H}(i))|$, where $\mathbf{H}(i)$ are not necessarily integers, with the calculated structure-factor amplitudes, $|F_{calc}(\mathbf{h})|$, that are on a regular grid, *i.e.* have integer $\mathbf{h}$. In principle, given an electron density of the molecule, one could calculate the structure factors in the experimental directions. Nevertheless, for computational efficiency, we put the (unknown) molecule or particle into a fictitious unit cell that is larger than the molecule. We will also demand that the Gaussians outside the molecular envelope be empty. (In practice, sizes of molecules are known from their composition; particle sizes and shapes may also be known from lower-resolution imaging.) As long as the distances of the Gaussian basis functions are determined by the experimental resolution, the number of 'independent' unknowns neither increases nor decreases, in principle, by this computational device. The structure factors are calculated on an integer grid in the large unit cell, so they are essentially oversampled in each dimension by the ratio of the size of the large cell to the size of the molecule. The oversampling allows stable interpolation of the calculated structure-factor amplitudes from integer $\mathbf{h}$ to the fractional $\mathbf{H}(i)$ everywhere, independent of the density of the actual measurements. Note that interpolation from fractional $\mathbf{H}(i)$ to integer $\mathbf{h}$ is not always a stable procedure!

In the present implementation of *SPEDEN*, we get sufficient accuracy with the simplest trilinear interpolation in the amplitudes of $|F_{calc}(\mathbf{h})|$ if we choose the fictitious unit cell to be three times larger than the molecule in each dimension. Now, some mathematical details: the reciprocal-space vector $\mathbf{H}(i)$ is within a cube, bounded by eight corners $\mathbf{h}(i,j)$, $\{j = 1, \ldots, 8\}$ with integer values. Let us denote the fractional parts of the components of $\mathbf{H}(i)$ as (H, K, L). We define weights for the eight corners, $w(i,j)$, by taking the products of the fractional parts of H or $(1-H)$ with those of K or $(1-K)$ and L or $(1-L)$. The cost function to be minimized now becomes

$$\chi_{Speden} = \sum_i w'(\mathbf{H}(i))^2\left[|F_{obs}(\mathbf{H}(i))| - \sum_{j=1}^8 w(i,j)|F_{calc}(\mathbf{h}(i,j))|\right]^2. \qquad (5)$$

A similar approach of applying crystallographic algorithms to continuous diffraction data has been done with direct methods (Spence *et al.*, 2003). In that case, however, the Ewald sphere was approximated by a plane.

2.2.3. Real-space knowledge (targets). Let us assume that we have some, possibly uncertain, knowledge of the electron density in parts of the unit cell from an independent source, *i.e.* one that does not come from the X-ray measurement itself. This is the kind of knowledge present when the unknown molecule is placed into a larger unit cell and we demand that

the unit cell be empty outside the molecule. This kind of knowledge was also referred to as a 'sensible' electron density in the introduction. We represent this knowledge by a target electron density $n_{target}(p)$ and by a real-space weight function $w'(p)^2$. It will be desirable that the actual electron density of the molecule, $\rho(\mathbf{r})$, as represented by $n(p)$, be close to the target electron density; the weight function $w'(p)^2$ expresses the strength of our belief in the suggested value of the electron density. Note that target densities can be assigned in any region of the unit cell independently of those in any other region. The simplest way to express the above statement mathematically is to minimize the value of the cost function

$$\chi_{space} = \frac{A}{2} \lambda_{space} \sum_{p=1}^{P} w'(p)^2 [n(p) - n_{target}(p)]^2, \qquad (6)$$

where A is a normalizing constant, described in Somoza *et al.* (1995). The relative weight, λ_{space}, is a scale factor that determines the strength of the real-space target relative to the reciprocal-space target. It expresses a value judgement about the certainty of the constraints. In practice, it is determined by trial and error. (In the absence of information at and around the molecule, weights are generally unity where it is known that there is no molecule and zero elsewhere.) The same procedure is used in *EDEN*.

2.2.4. Low-resolution target (phase extension). The knowledge of the electron density at low resolution can be expressed by a low-resolution spatial target. Crystallographers call this phase extension. The essence is that, during the process of the search for an optimal electron density, we try to keep its low-resolution component as close to the known density as possible. Actually, it is easier to carry out the computation in reciprocal space. The low-resolution target, $F_{target}(\mathbf{h})$, is prepared from the (presumably) known electron density, $n_{target}(p)$. Then, given the current solution $n(p)$, we smear out its Gaussian representation and compare it to the equally smeared out target. We define

$$\chi_{phasext} = \lambda_{phasext} \sum_{i} w'(\mathbf{h})^2 |F_{smear}(\mathbf{h}) - F_{target}(\mathbf{h})|^2, \qquad (7)$$

where the current 'smeared' structure factors are calculated using the low resolution, ΔR,

$$F_{smear}(\mathbf{h}) = \exp[-\eta(\pi \Delta R |\mathcal{F}^T \mathbf{h}|)^2] \sum_{p=1}^{P} n(p) \exp[2\pi i \mathbf{h} \cdot \mathcal{F} \mathbf{r}(p)]. \qquad (8)$$

The value of $\lambda_{phasext}$ is determined similarly to λ_{space}. The same procedure is used in *EDEN*.

2.2.5. Minimization of the cost function. In the presence of a target density, the actual cost function used in the computer program is the sum of χ_{Speden} (5), χ_{space} (6) and $\chi_{phasext}$ (7):

$$\chi_{total} = \chi_{Speden} + \chi_{space} + \chi_{phasext}. \qquad (9)$$

The fast algorithm described in Somoza *et al.* (1995) and Szöke *et al.* (1997) is always applicable to the calculation of the full cost function, (9). There is a clear possibility of defining more

target functions. They are all added together to form χ_{total}, which is minimized to find the optimum electron density.

The minimization of the cost function (9) is carried out in *SPEDEN* (as in *EDEN*) by D. Goodman's conjugate gradient algorithm (Goodman, 1991). It has proven to be very robust and efficient in years of use in *EDEN*. The essential properties of the algorithm that make it so advantageous for our application is that the positivity of the electron density, $n(p) \geq 0$, is always enforced and that the gradient vector in real space can be calculated by fast Fourier transform. The gradient calculation needed only a very simple modification for the interpolation in reciprocal space, equation (5). The line search algorithm does not use the Hessian, so matrices are never calculated.

As with any local minimization, global convergence is not achieved. We discussed this problem in our previous papers and came to the conclusion that, usually, the minimum surface of the cost function (7) is so complicated that finding a global minimum would take more computer time than the age of the universe!

2.3. Comparison to iterative algorithms

Reconstruction of the scatterer from a continuous diffraction pattern has a tangled history replete with repeated discoveries. Some of the present authors are also guilty of ignorance of prior work. We referred to the pioneering insights in §2.1.2.

The 'recent' period of algorithms started with the work of Miao *et al.* (1998), who pointed out that the fraction of the unit cell where the density is known is an important parameter for convergence. In somewhat later work, with oversampled structure factors calculated on a regular grid, the crystallographic program *EDEN* successfully demonstrated the recovery of the electron density using a simulated data set from the photoactive yellow protein (Szöke, 1999). The protein was put into a fictitious unit cell, twice the size of the original one, and a target with zero density was used outside the original unit cell of the protein. Similarly, Miao & Sayre (2000) have studied empirically how much oversampling is required in two- and three-dimensional reconstructions of a simulated data set using a version of the Gerchberg–Saxton–Fienup (GSF) algorithm. Among recent articles, we mention Robinson *et al.* (2001), Williams *et al.* (2003), Marchesini *et al.* (2003), and references therein, in addition to those mentioned in the *Introduction*.

All reconstruction algorithms of oversampled diffraction patterns use *a priori* information on the shape and size of the particle. In our previous studies in crystals, we found that such information is very valuable. For example, *EDEN* converges surprisingly well for proteins at low resolution where the only information used is that the molecule is a single 'blob'. *EDEN* also converges for synthetic problems with a good knowledge of the solvent volume if it is greater than 50% (Béran & Szöke, 1995) or 60% (*EDEN*). A similar conclusion was reached in

Miao *et al.* (1998). In comparison, when a molecule is embedded in a three times larger fictitious unit cell, the empty 'solvent' occupies ~96% of the cell volume.

As discussed previously, the reconstruction of scatterers from their diffraction pattern is a difficult mathematical problem. In many cases, the indeterminacy of the absolute position of the object and of its handedness causes difficulties in convergence. That is definitely the case with *SPEDEN* so, in that sense, *SPEDEN* is not a good algorithm. Empirically, the GSF algorithm has a larger radius of convergence and deals better with stagnation (Marchesini *et al.* 2003).

Another family of difficulties arises when there is *a priori* information available, but there is only incomplete and noisy data. Under such conditions, the main questions are how to find a solution that optimally takes into account the available information and that is the best 'sensible' one that reproduces the noisy and incomplete data to its limited accuracy. It is this second set of conditions for which *SPEDEN* was written. Although, in this paper, we show only its performance for artificial and 'easy' but incomplete data, *SPEDEN*'s older sister *EDEN* has been shown to have those properties on a large range of crystallographic data sets, ranging from CuO_2 to the ribosome. We expect that such properties of *EDEN* will be inherited by *SPEDEN*, considering that their fundamental mathematical properties are sufficiently similar.

The best known and successful class of algorithms is the group of iterative transform algorithms that we refer to as Gerchberg–Saxton (GS) (Gerchberg & Saxton, 1972), and its development, in which support constraints and feedback are added, the GS–Fienup (GSF) or hybrid input–output algorithm (Fienup, 1982; Aldroubi & Grochenig, 2001; Bauschke *et al.*, 2002, 2003). The essence of the GSF algorithms is that they iterate the N-pixel data between real and reciprocal spaces *via* FFTs and enforce the known constraints in each of these spaces.

Depending on the degree of noise in the data, these algorithms usually converge in about a hundred to a thousand iterations. The weak convergence (non-divergence) of the GS algorithm has been proven in the absence of noise (Fienup, 1982). There is no mathematical proof that these algorithms will converge in general, but it is reasonable that, by sequentially projecting onto the set that satisfies the real-space constraints and the set that satisfies the reciprocal-space constraints, the intersection (corresponding to a valid solution) should be approached. This is definitely true for projections onto convex sets, but unfortunately these sets are not convex (Stark, 1987). In practice, despite the fact that the modulus constraint is non-convex, the algorithms often converge even in the presence of noise in several hundred iterations.

The main difference between the GSF algorithms and *SPEDEN* is that *SPEDEN* does not iteratively project onto the sets of solutions that satisfy the real-space or reciprocal-space constraints separately, but rather it minimizes a cost function that includes all the constraints of both spaces. It does this by varying quantities in real space only [the $n(p)$'s] and the cost evaluation only requires a forward transform from real space to reciprocal space. As the cost function never increases, *SPEDEN* reaches only a local minimum.

In spite of its 'small' radius of convergence, there are some expected advantages to *SPEDEN*'s algorithm. In *SPEDEN*, we compare the F_{calc} to the F_{obs} by interpolating from the samples of $F_{calc}(\mathbf{h})$, calculated on a regular grid, to the measured sample vectors $\mathbf{H}$. Since the gridded $|F_{calc}(\mathbf{h})|$ are a complete set (because they are sampled above the Nyquist frequency), the interpolation is stable and performed with little error. Since an inverse transform is required in the GSF algorithms, the measured diffraction data $|F_{obs}(\mathbf{H}(i))|$ (recorded on Ewald spheres in reciprocal space) must be interpolated to the gridded data points $\mathbf{h}$. The observed data might not be a complete set, especially at high resolution where the density of samples is sparser: this may lead to error. An additional possible difficulty with the GSF algorithms is that the effective number of unknowns may increase with the size of the fictitious unit cell, while in *SPEDEN* the effective number of unknowns is constant. Finally, in *SPEDEN*, weightings can be properly applied to all data and knowledge. The measured data are inversely weighted by their uncertainties; it is a procedure equivalent to maximum-likelihood methods and it should be optimal, at least in theory. As a side effect, when the constraints are inconsistent, *SPEDEN* still converges to a well defined and correct solution. (Note that for noisy data the constraints are almost always inconsistent.) Weightings are also applied to reflect our confidence in real-space constraints, and these weightings are consistently used in real space.

3. Tests of the program

SPEDEN has certain built-in limitations. In particular, of course, reconstruction can only be as good as the diffraction measurements and the structure-factor amplitudes derived from them. There are also other less obvious limitations. For example, there are inherent inaccuracies due to the Gaussian representation of the electron density in real space (Szöke, 1997). Also, the trilinear interpolation from integer (hkl) structure-factor amplitudes to a non-integer position is only approximate. Finally, as a fundamental limitation of any reconstruction method, both the absolute position and the handedness of the molecule are undefined.

We performed preliminary tests to verify the capabilities and limitations of our reconstruction method using computed and experimentally obtained diffraction patterns. In this section, we first describe the reconstruction of simple 'molecules' from synthetic diffraction patterns with *SPEDEN*. Specifically, we discuss how the convergence of *SPEDEN* is affected by the errors due to interpolation in reciprocal space, by the quantity of observed structure factors, $|F_{obs}|$, by the extent of the 'known' starting model, and by the uniformity of sampling in reciprocal space. We then describe the reconstruction of simple two-dimensional objects from synthetic and measured diffraction patterns with *SPEDEN*.

research papers

3.1. Generation of synthetic test data

We created synthetic 'molecules' in the format of the Protein Data Bank (pdb) files (Berman *et al.*, 2000). Each molecule was composed of 15 point-like C atoms, placed at random positions within a cube measuring 16.8 Å in each dimension and 'measured' to 4 Å resolution; these values correspond to crystallographic B factors of 185.7 Å^2. The molecule was then shifted so that its center of mass was at the center of the cube. All our simulations were repeated using molecules with several different random arrangements. Initially, the 'unit cell' coincided with the dimensions of the cube in which the atoms were placed; later, larger cells were used and the atoms were positioned in their center. We also generated 'starting models' by removing atoms from the full 'molecule'.

Both full and partial molecules served to generate structure factors (F_{calc}), using the atomic positions and the B factors. Starting models were generated from the F_{calc}, using *SPEDEN*'s preprocessor, *BACK*, which finds the optimal real-space representation for an input F_{calc}. Initially, sets of 'measurements' ($|F_{obs}|$) were generated by deleting the phases of the calculated structure factors of the full or partial molecule. The F_{obs} files so generated all had integer $\mathbf{H}(i)$. The uncertainty of the measured structure factors, $\sigma(\mathbf{h})$, were chosen to be

$$\sigma(\mathbf{h}) = \alpha[\langle F_{obs}\rangle/10 + F_{obs}(\mathbf{h})]^{1/2}, \tag{10}$$

with $\alpha = 0.1$ for $\mathbf{h} \neq (0,0,0)$ and $\alpha = 0.01$ for $\mathbf{h} = (0,0,0)$. $\langle F_{obs}\rangle$ is the average over all the structure factors. We used two alternative methods to generate F_{obs} files for non-integer $\mathbf{H}(i)$. In the first method, we used a unit cell whose dimensions were incommensurate with one another and with the edge of the cube, but whose volume equaled that of the cube. The resulting F_{obs} file, generated again from F_{calc} files by deleting the phases, then had its indices scaled back appropriately, yielding fractional $\mathbf{H}(i)$. A different second method was also used and will be explained in §3.6.

For each of these F_{obs} files, some constraining information is required in order to find the atom positions. We used two types of constraints: one of them identified the (approximate) empty region; the other one used F_{calc} at a considerably lower resolution. We call them the empty target and the low-resolution target, respectively. Both are based on the assumption that, at a considerably lower resolution, the general position of atoms as one or more 'blobs' in empty space is known. The low-resolution FS_{calc} was prepared by smearing the full F_{calc} file to 10 Å. The empty target identified the empty points in terms of a mask. Then, throughout *SPEDEN*'s iteration process, using the solver *SOLVE*, the program attempted to match the current electron voxel^{-1} values in masked-in regions to empty $[n(p) = 0]$ values. The phase-extension target used the same FS_{calc} file; during the iteration process, at each step, the current real-space solution was smeared to that low resolution in reciprocal space and restrained to agree with that target.

3.2. Assessment of quality of reconstruction

Besides inspecting the reconstructed electron density visually, we used four quantitative measures to compare the reconstructed image with the electron density from the full molecule at 4 Å.

(*a*) Real-space root-mean-square (RMS) error: we calculated the real-space electron density from the electron voxel^{-1} files, a process we call regridding. We then calculated the RMS error of the electron densities, ρ_1 and ρ_2, defined as

$$\mathrm{RMS} = \frac{\{\sum_r[\rho_1(r) - \rho_2(r)]^2\}^{1/2}}{\{\sum_{r'}\rho_1^2(r') + \sum_r\rho_2^2(r)\}^{1/2}}. \tag{11}$$

We permitted one file to be shifted with respect to the other file in order to minimize the distance.

(*b*) Phase difference: we calculated the average (amplitude-weighted) phase difference between the F_{calc} at the end of the run and the corresponding F_{calc} generated from the full molecule.

(*c*) Final R factor: we calculated the crystallographic R factor (Giacovazzo *et al.*, 2002) at the end of the run.

(*d*) Count error: we compared the integrated real-space electron density generated from a run result with the true number of electrons as identified in the pdb file on an atom-by-atom basis. The integration was performed around each atom over a sphere with a radius that was 1.5 times the grid spacing. The figure of merit is the RMS error.

Of all these measures, the final R factor was the least useful to assess the quality of the reconstructed image. The R factor tends to be lower for a small number of entries in the observation file since there are fewer equations to satisfy during the reconstruction. In such a case, a visual inspection shows that the reconstructed electron density may have little resemblance to the 15 C atoms. However, there was a good correlation among the other three measures. The solutions looked correct when the phase difference between solution and full F_{calc} was less than 20°, the count error between solution and pdb model was 0.2 or lower (out of 6), and the RMS distance measure was less than 0.2.

3.3. Overcoming inaccuracies due to trilinear interpolation through oversampling

For computational purposes, we placed the (unknown) molecule or particle into a fictitious unit cell that is larger than the molecule and calculated the structure factors on an integer grid in the large unit cell, oversampled by the same ratio: the size of the large cell to the size of the molecule. We then calculated the structure factor at fractional $\mathbf{H}$ from the structure factor at integer $\mathbf{h}$ using trilinear interpolation. The interpolation error becomes smaller when larger unit cells are used (at the expense of computation time). We studied the effect of the unit-cell size on convergence of *SPEDEN*.

In the initial tests, the unit cell was the same size as the original molecule, and the starting (known) part of the molecule consisted of the full molecule. When the molecule consisted of atoms on grid points and the F_{obs} files had integer $\mathbf{H}(i)$, unsurprisingly *SPEDEN* converged, as did *EDEN* on the

same data set. However, we found that *SPEDEN* did not converge to a unique solution if either the atoms were not on grid points or the F_{obs} file had non-integer $\mathbf{H}(i)$, or both, since the solution meandered in real space.

In subsequent tests, we generated larger unit cells and applied a target over the empty part of the unit cell, in an attempt to restrain the meandering problem. We embedded the molecule in a cell that was two or three times greater in each dimension. An empty target was used that essentially covered the empty 7/8th or 26/27th of the unit cell, respectively. We applied a high relative weight for this empty target, and we still used the full molecule as a starting position. We found that both the larger cell and the empty target are of critical value in enabling *SPEDEN* to converge to the correct solution. Comparing the twofold larger unit cell *versus* the threefold unit cell, there was a great improvement in the latter case. These results show that, for trilinear interpolation, it is adequate to use a unit cell that is three times greater in each dimension. We expect that more sophisticated interpolation algorithms should allow use of smaller unit cells.

3.4. Dependence of reconstruction on the quantity of input data

A threefold enlarged unit cell increases the number of unknown amplitudes of the Gaussians, $n(p)$, by a factor of 27. In principle, the emptiness of the volume around the molecule restrains the effective number of independent unknowns. Nevertheless, if the number of equations, which is given by the number of entries in the F_{obs} file, is not increased, it is easy for the solver to 'hide' electrons among the large number of unknowns in the system, even when the empty target constraint is used. In fact, we found that when we compared the final F_{calc} from *SOLVE* against the starting F_{calc}, on the one hand, and the correct F_{calc} on the other, *SOLVE*'s final F_{calc} was closer to the starting one than to the correct one. In other words, the cost function in reciprocal space was not a sufficiently strong constraint in *SPEDEN*'s algorithm, for this synthetic problem. In a similar real case, more experimental diffraction patterns need to be collected in order that *SPEDEN* would be able to find the corresponding image without other information.

3.5. Recovery of missing atoms

In the next set of synthetic tests, we attempted to recover missing information by starting from a partial model that contained less than the full complement of 15 atoms. In these simulations, we used F_{obs} files with non-integer $\mathbf{H}(i)$, a three-fold enlarged unit cell, an empty target or a phase-extension target, and randomly positioned atoms.

We found that a low-resolution spatial target significantly helps *SPEDEN* to converge. Fig. 1(a) shows the results of the comparison of the reconstructed image with the electron density from the full pdb file when a phase-extension target is used. The phase-extension target was calculated at a resolution of 10 Å. We found that it was generally possible to recover 5, 10 or even all 15 of the atoms. Please note that the

amount of information in such a phase-extension target is $(4\,\text{Å}/10\,\text{Å})^3 \approx 6\%$ of the information in the perfect solution.

It was more difficult to reconstruct the original electron density when we used an empty target, as shown in Fig. 1(b). *SPEDEN* was able to recover 5 out of the 15 atoms, but did not converge when 10 atoms were unknown. Perhaps surprisingly, the case where there was no starting model at all (0 atoms known) did consistently better than those cases where a partial model was given as a starting condition.

3.6. Effect of non-uniform sampling on recovery

In this set of synthetic tests, we addressed the question of how difficult it is to recover the molecule from a non-uniform set of samples in reciprocal space, similar to real data sets, and how the results compare with the reconstruction from a uniformly sampled data set.

We generated two-dimensional diffraction patterns of the synthetic carbon molecule for different particle orientations, corresponding to recorded diffraction patterns in a real experiment. The two-dimensional diffraction patterns were linearly interpolated from a three-dimensional diffraction pattern; the latter was calculated on an additionally double-fine grid over the already triple-sized unit cells, *i.e.* using a unit
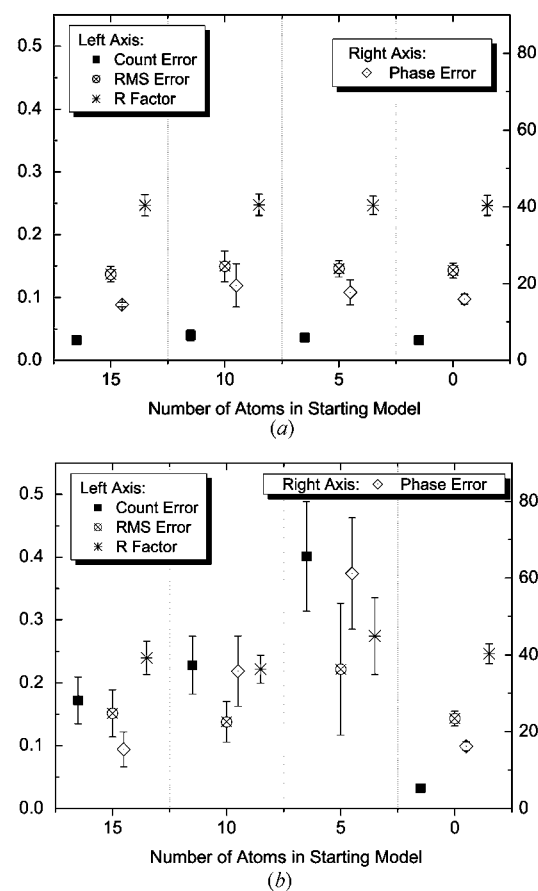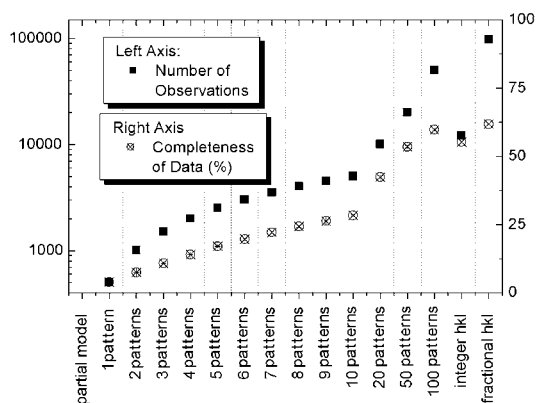

Figure 1
Results of the comparison of the reconstructed image with the electron density from the full pdb file in the cases when (a) a phase-extension target and (b) a low-resolution target are used. The phase error is given in units of degrees.

cell that was a total of six times larger than the molecule in each direction. Further refining the grid of the three-dimensional diffraction pattern did not alter the results significantly. The three-dimensional diffraction pattern, in turn, was the Fourier transform of the gridded electron density of the synthetic molecule.

Two-dimensional patterns do not sample the diffraction space uniformly. The sampling density near the center of the diffraction space is much larger than the sampling density further away. We used a completeness measure to characterize the sampling uniformity. The reciprocal space is divided into cells that are $4\pi/a$ by $4\pi/b$ by $4\pi/c$ in size, where a, b and c are the molecule sizes in each dimension. The completeness then is the ratio of cells in reciprocal space that contain at least one measurement over the total number of cells. Fig. 2 shows the completeness of the input observation files as a function of the number of diffraction patterns. Also shown in Fig. 2 is the number of calculated diffraction intensities.

We then used *SPEDEN* to recover 7 out of the 15 atoms. The molecule was embedded in a unit cell that was three times larger in each dimension, and we used an empty solvent target. Fig. 3 shows the errors of the reconstructed electron density as a function of the number of two-dimensional diffraction patterns. The orientations of the diffraction patterns were chosen at random, and the calculations were repeated for four different molecules. For comparison, also shown in Fig. 3 are the errors of the electron density of the eight known atoms ('partial model') and the errors of the reconstructed electron densities when a three-dimensional diffraction pattern is used, which was oversampled three times ('integer *hkl*') or six times ('fractional *hkl*').

As discussed above, the R factor is not a useful measure to assess the quality of the reconstructed image, but the RMS and count errors are better measures for the reconstruction quality. Surprisingly, we found that four two-dimensional
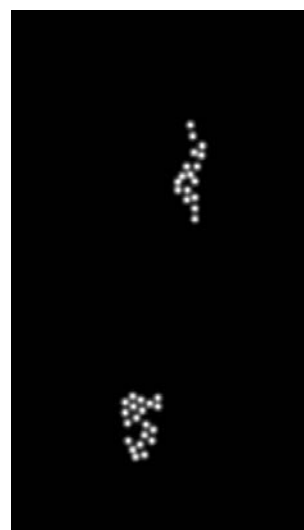

Figure 2
The completeness and length of the input observation files used for the calculations shown in Fig. 3.
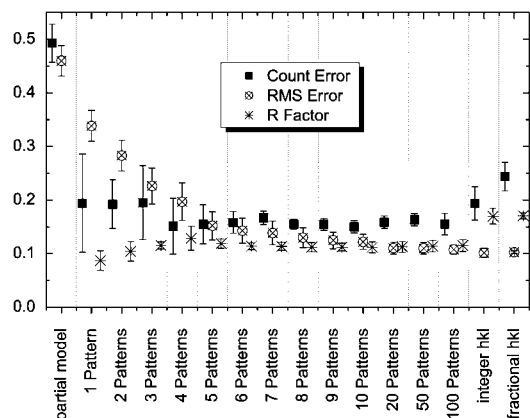
Figure 3
The error in the reconstructed electron density as a function of the number of two-dimensional diffraction patterns. The orientation of the diffraction patterns was chosen at random and the results were repeated for four different molecules. Also shown are the error of the electron density of the eight known atoms ('partial model') and the error of the reconstructed electron densities for a three-times ('integer *hkl*') and six-times ('fractional *hkl*') in-each-dimension-oversampled three-dimensional diffraction pattern.

Figure 4
Planar arrangement of 37 Au balls for two-dimensional reconstruction.

Figure 5
Initial starting model used for two-dimensional reconstruction from synthetic data.

patterns are sufficient to reconstruct the electron density as well as in the case when the full three-dimensional diffraction pattern is given. Four two-dimensional patterns have a remarkably low sampling completeness of only 14%. Further increasing the completeness or the number of observations does not improve the quality of the reconstruction. We would like to note, however, that these results could be dependent on the sparsity of the test model where positivity, enforced by *SPEDEN*, is a powerful constraint. Therefore, for less-sparse test models the number of required two-dimensional patterns may be larger.

3.7. Recovery of two-dimensional data

In the final set of tests, we demonstrate *SPEDEN*'s ability to recover missing information for a two-dimensional configuration of 37 Au balls in a plane. We reconstructed the Au

balls using (i) a synthetic diffraction pattern and (ii) an experimentally obtained diffraction pattern as discussed by He *et al.* (2003). In the following, we will discuss both cases, starting with the synthetic diffraction data.

The 37 Au balls are arranged in a plane as shown in Fig. 4. The arrangement of the balls is similar to the experimental case discussed by He *et al.* (2003). The Au balls were 50 nm in diameter. We generated an artificial set of 'measurements' ($|F_{obs}|$) by calculating the structure factors (F_{calc}) to 30 nm

Figure 8
Initial starting model for two-dimensional reconstruction from experimental data.

Figure 6
Weight function used for two-dimensional reconstruction from synthetic data.

Figure 9
Weight function used for two-dimensional reconstruction from experimental data.

Figure 7
Reconstructed electron density using synthetic data.
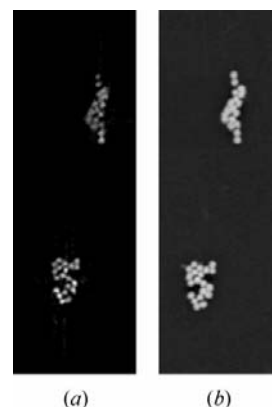
Figure 10
(*a*) Reconstructed electron density using reseperimental data. (*b*) SEM Image of Au balls.

resolution and deleting the phases. The uncertainty of the measured structure factors, $\sigma(\mathbf{h})$, were chosen according to equation (10). We generated an initial model by smearing the full F_{calc} file to 90 nm, and running *BACK* on it. The initial model is shown in Fig. 5. In Figs. 5–7, we only show one plane. We also used this smeared F_{calc} to generate a low-resolution spatial target as well as an empty target outside the molecule. The corresponding weight function is shown in Fig. 6.

We then used *SPEDEN* to reconstruct the Au balls. As shown in Fig. 7, *SPEDEN* reconstructed the electron density successfully. However, we further found that, if we use an empty starting model, *SPEDEN* has difficulties converging to the correct electron density. There are two reasons for this behavior. First, without an initial model, the symmetry of the system is not broken and *SPEDEN* stagnates since the support does not distinguish between the object and its centrosymmetric copy. Second, the mask and the reconstructed electron density are possibly shifted with respect to each other. If the initial model is empty, the position of the reconstructed electron density is mostly determined in the early iteration of the *SOLVE* algorithm and can be partially cut off by the solvent. The algorithm has difficulty shifting the result. It is necessary to provide information about the location of the electron density to a certain degree, for example in the form of a smeared model. Note that the GSF algorithms are designed to overcome these problems when there is abundant and accurate data available.

We will now discuss the reconstruction of the Au balls using experimental data. To generate a starting model, we took the experimental $|F_{obs}|$ data along with the phases obtained by He *et al.* (2003) using a version of the GSF algorithm, and smeared this data to 80 nm. The starting model is shown in Fig. 8. We used the same data to generate a real-space target with a target fraction of 99.7%, shown in Fig. 9. Similar to the case of the synthetic test data, we chose $\sigma(\mathbf{h})$ according to equation (10). We then used *SPEDEN* to reconstruct the Au balls. Fig. 10(*a*) shows the reconstructed electron density and Fig. 10(*b*) shows a scanning-electron-microscope (SEM) picture of the sample. We found that *SPEDEN* reconstructed the electron density from the experimental data successfully.

4. Summary and conclusions

In this paper, we presented *SPEDEN*, a method to reconstruct the electron density of single particles from their X-ray diffraction patterns, using an adaptation of the holographic method in crystallography. Unlike existing GSF algorithms, *SPEDEN* minimizes a cost function that includes all the constraints of both real space and reciprocal space, by varying quantities in real space only, so that the cost evaluation requires only a forward transform from real space to reciprocal space. *SPEDEN* finds a local minimum of the cost function using the conjugate gradient algorithm. We implemented *SPEDEN* as a computer program and tested it on synthetic and experimental data. Our initial results indicate that *SPEDEN* works well.

References

Aldroubi, A. & Grochenig, K. (2001). *SIAM (Soc. Ind. Appl. Math.) Rev.* **43**, 585–620.

Bauschke, H., Combettes, P. & Luke, D. R. (2002). *J. Opt. Soc. Am.* **19**, 1334–1345.

Bauschke, H. H., Combettes, P. L. & Luke, D. R. (2003). *J. Opt. Soc. Am.* **20**, 1025–1034.

Béran, P. & Szöke, A. (1995). *Acta Cryst.* A**51**, 20–27.

Berman, H. M., Westbrook, J., Feng, Z., Gilliland, G., Bhat, T. N., Weissig, H., Shindyalov, I. N. & Bourne, P. E. (2000). *Nucleic Acids Res.* **28**, 235–242.

Bertero, M. (1989). *Advances in Electronics and Electron Physics 75*. New York: Academic Press.

Bertero, M. & Boccacci, P. (1998). *Introduction to Inverse Problems in Imaging.* Bristol: Institute of Physics.

Bricogne, G. (1992). *International Tables for Crystallography*, Vol. B, edited by U. Shmueli. Dordrecht: Kluwer Academic Publishers.

Daubechies, I. (1992). *Ten Lectures on Wavelets. CBMS-NSF Regional Conference Series in Applied Mathematics* 61. Philadelphia, PA: Society for Industrial and Applied Mathematics.

Fienup, J. R. (1982). *Appl. Opt.* **21**, 2758–2769.

Gerchberg, R. W. & Saxton, W. O. (1972). *Optik (Stuttgart)*, **35**, 237–246.

Giacovazzo, C., Monaco, H. L., Artioli, G., Viterbo, D., Ferraris, G., Gilli, G., Zanotti, G. & Carri, M. (2002). *Fundamentals in Crystallography*, 2nd ed. Oxford Science Publications.

Golub, G. H. & van Loan, C. F. (1996). *Matrix Computations*, 3rd ed. Johns Hopkins University Press.

Goodman, D. M. (1991). In *Proceedings of 7th Workshop on Multi-Dimensional Signal Processing, Sept. 1991, Lake Placid, New York.* Piscataway, NJ: IEEE Signal Processing Society.

Hau-Riege, S. P., London, R. A. & Szöke, A. (2004). *Phys. Rev. E.* To be published.

He, H., Marchesini, S., Howells, M., Weierstall, U., Chapman, H., Hau-Riege, S., Noy, A. & Spence, J. C. H. (2003). *Phys. Rev. B*, **67**, 174114–174117.

Huldt, G., Szöke, A. & Hajdu, J. (2003). *J. Struct. Biol.* **144.** 219–227.

Marchesini, S., He, H., Howells, M., Weierstall, U., Chapman, H., Hau-Riege, S. & Spence, J. C. H. (2003). *Opt. Expr.* **11**, 2344–2353.

Miao, J., Charalambous, P., Kirz, J. & Sayre, D. (1999). *Nature (London)*, **400**, 342–344.

Miao, J., Ishikawa, T., Johnson, B., Anderson, E. H., Lai, B. & Hodgson, K. O. (2002). *Phys. Rev. Lett.* **89**, 088303-1–4.

Miao, J. & Sayre, D. (2000). *Acta Cryst.* A**56**, 596–605.

Miao, J., Sayre, D. & Chapman, H. N. (1998). *J. Opt. Soc. Am.* A**15**, 1662–1669.

Natterer, F. (1996). *The Mathematics of Computerized Tomography.* Chichester: Wiley.

Natterer, F. & Wubbeling, F. (2001). *Mathematical Methods in Image Reconstruction.* Philadelphia, PA: Society for Industrial and Applied Mathematics.

Neutze, R., Wouts, R., van der Spoel, D., Weckert, E. & Hajdu, J. (2000). *Nature (London)*, **406**, 752–757.

Robinson, I. K., Vartanyants, I. A., Williams, G. J., Pfeifer, M. A. & Pitney, J. A. (2001). *Phys. Rev. Lett.* **87**, 195505-1–4.

Sayre, D. (1952). *Acta Cryst.* **5**, 843.

Sayre, D. (1980). *Image Processing and Coherence in Physics. Springer Lecture Notes in Physics* 112, edited by M. Schlenker, pp. 229–235. Berlin: Springer.

Somoza, J. R., Szöke, H., Goodman, D. M., Beran, P., Truckses, D., Kim, S.-H. & Szöke, A. (1995). *Acta Cryst.* A**51**, 691–708.

Spence, J., Weierstall, U., Fricke, T., Glaeser, R. & Downing, K. (2003). *J. Struct. Biol.* **144**, 209–218.

Stark, H. (1987). *Image Recovery: Theory and Applications.* New York: Academic Press.

Szöke, A. (1997). *J. Imaging Sci. Tech.* **41**, 332–341.

Szöke, A. (1998). *Acta Cryst.* A**54**, 543–562.

Szöke, A. (1999). *Chem. Phys. Lett.* **313**, 777–788.

Szöke, A., Szöke, H. & Somoza, J. R. (1997). *Acta Cryst.* A**53**, 291–313.

Weierstall, U., Chen, Q., Spence, J. C. H., Howells, M. R., Isaacson, M. & Panepucci, R. R. (2002). *Ultramicroscopy*, **90**, 171–195.

Williams, G. J., Pfeifer, M. A., Vartanyants, I. A. & Robinson, I. K. (2003). *Phys. Rev. Lett.* **90**, 175501-1–4.

Zuo, J., Vartanyants, I., Gao, M., Zhang, M. & Nagahara, L. (2003). *Science*, **300**, 1419–1421.