Published in final edited form as:

J Opt Soc Am A Opt Image Sci Vis. 2010 September 1; 27(9): 1927–1937.

Phase Retrieval for Superposed Signals from Multiple Binary Objects

Andreas Alpers^{1,*}, Gabor T. Herman², Henning Friis Poulsen¹, and Søren Schmidt¹

¹ Center for Fundamental Research: 'Metal Structures in Four Dimensions,' Risø-DTU, Technical University of Denmark, DK-4000 Roskilde, Denmark

² Department of Computer Science, The Graduate Center, City University of New York, NY 10016, USA

Abstract

We introduce the binary superposed phase retrieval problem that aims at reconstructing multiple 0/1-valued functions with nonoverlapping bounded supports from moduli of superpositions of several displaced copies of their individual Fourier transforms. We discuss an application in coherent diffraction imaging of crystalline objects, propose two algorithms, and evaluate their performance by means of simulations.

1. Introduction

Reconstructing a function from the modulus (i.e., absolute magnitudes) of its Fourier transform (FT) constitutes the well-known *phase retrieval* problem: the phase of the complex valued FT has to be retrieved before it can be inverted. This nonlinear inverse problem arises in optics, astronomy, crystallography, and numerous other areas of physics and engineering [1–3], and is known as *binary phase retrieval* problem [4] if the function is 0/1-valued.

A particular imaging application in which objects have successfully been modeled by 0/1-functions is in *incoherent* X-ray diffraction imaging of small crystals (called *grains*) inside of polycrystalline materials [5–7]. *Coherent diffraction imaging* (CDI) of multiple grains however has not yet been accomplished, while non-crystalline nano-structures [8–10] and single nano-particles [11,12] have successfully been reconstructed. In CDI of single grains, there is a direct relation between the diffraction pattern and the FT modulus of the grain's shape function [13]. Experience with incoherent X-ray diffraction from multiple grains suggests that diffracted signals from different grains superpose, and that the modulus of this superposition can be measured on a detector. While the development of a detailed imaging model is left for a future paper, we show that – standard – phase retrieval algorithms (such as Algorithm II introduced below) are not well suited for reconstructing objects from highly superposed data. Our main contribution is Algorithm I, which yields reconstructions from superposed data of superior quality compared to our standard phase retrieval algorithm. Our presentation is deliberately kept rather general, because we feel that our algorithms are of independent interest in other (yet unexplored) application areas in which phase retrieval problems of superposed signals from binary objects occur.

We define the following generalization of the binary phase retrieval problem, to our knowledge it has not been considered in the literature before: the reconstruction of multiple 0/1-valued

*Corresponding author: awalpers@yahoo.de.

functions with nonoverlapping bounded supports from moduli of superpositions of several displaced copies of their individual FTs. In other words, the two-dimensional version of problem is the following (generalizations to higher dimensions are straightforward): for given $G, H \in \mathbb{N}$ (throughout this paper $\mathbb{N}$ denotes the set of positive integers) and *displacement vectors* $\mathbf{t}_{gh} \in \mathbb{R}^2$, $g = 1, \dots, G$, $h = 1, \dots, H$, the aim is to reconstruct the functions $f_1: \mathbb{R}^2 \rightarrow \{0, 1\}, \dots, f_G: \mathbb{R}^2 \rightarrow \{0, 1\}$ with nonoverlapping bounded supports from the knowledge of $\left| \sum_{g=1}^G \sum_{h=1}^H \widehat{f}_g(\cdot - \mathbf{t}_{gh}) \right|$, where $\widehat{f}_1, \dots, \widehat{f}_G$ denote the FTs of $f_1, \dots, f_G$, respectively. We call this variant the *binary superposed phase retrieval* problem. For $G = H = 1$, we obtain as special case the binary phase retrieval problem for a single object. However, we do not place any restrictions on the displacement vectors other than that they should be known, and thus the problem usually does not reduce to the binary phase retrieval problem.

Binary superposed phase retrieval poses major algorithmic challenges. Standard phase retrieval algorithms, in fact, aim at reconstructing a single function from its FT modulus. Thus only a *superposed* function would be returned as output from such an algorithm that is applied to superposed FT modulus data. We aim, however, at reconstructing the individual binary functions – and this is the major algorithmic challenge.

Algorithms I and II introduced in this paper are the first algorithms for binary superposed phase retrieval. Algorithm I is a Monte Carlo method, while Algorithm II is based on the *Hybrid-Input-Output* (HIO) algorithm [14]. Not surprisingly, Algorithm I is far superior than Algorithm II in terms of reconstruction quality (see Section 5). While Algorithm I aims at reconstructing the individual functions simultaneously, Algorithm II proceeds sequentially. We remark that Algorithm I is more time efficient than other comparable Monte Carlo methods, because we make use of the fact that the functions to be reconstructed are binary.

Details of Algorithm I are discussed in Section 2 and 3, here we remark only that Monte Carlo methods in imaging applications follow the paradigm of simulating Markov chains; usually only the type of prior information and the forward operator relating the image to the measurements differ. One ends up with an algorithm that is useless for practical applications if the priors and operators cannot be implemented effectively. Our particular forward operator, however, can be implemented efficiently, as we discuss in Section 3. For the (different) forward operator used for phase retrieval of non-crystalline objects, Monte Carlo methods have been developed in the 80s [15–17]. With the exception of [18], they seem to have been abandoned in favor of more effective methods, such as the HIO algorithm.

In the application area of CDI of multiple grains, several experimental imaging methods – such as holography or *ab initio* phase retrieval for the complete oversampled diffraction pattern [19] – are possible. However, experimental concerns (signal/noise and sampling issues) imply that robust reconstructions could be considerably facilitated by incorporating additional prior information. The approach that we propose incorporates diffraction data from both incoherent and coherent X-rays. Additional information in this case is that (1) the grains are binary objects (we consider the general grain shape and not the individual atom positions, as we discuss in Section 2.B) and (2) from existing x-ray methodology for incoherent beams we can obtain (partial) low-resolution maps [6,7,20–22], which we propose to use as prior knowledge for algorithms that aim at reconstructing objects from coherent X-ray diffraction data, see Fig. 1. Our algorithms are capable of including such prior knowledge.

We remark that grains are of central interest in many branches of materials science, because polycrystalline materials (such as metals, alloys, and ceramics) are composed of multiple grains, which in turn determine the material's physical, chemical and mechanical properties.

Our results show that, under simplifying assumptions, CDI of multiple grains is possible with Algorithm I. Thus, CDI is potentially a vital imaging tool for materials science.

2. Binary Superposed Phase Retrieval

We define the binary superposed phase retrieval problem and discuss a relevant application in simplified terms.

2.A. Our General Inverse Problem

We repeat from the introduction the definition of the two-dimensional *binary superposed phase retrieval problem*: For given $G, H \in \mathbb{N}$ and *displacement vectors* $\mathbf{t}_{gh} \in \mathbb{R}^2$, $g = 1, \dots, G$, $h = 1, \dots, H$, the aim is to reconstruct the functions $f_1: \mathbb{R}^2 \rightarrow \{0, 1\}, \dots, f_G: \mathbb{R}^2 \rightarrow \{0, 1\}$ with

nonoverlapping bounded supports from the knowledge of $\left| \sum_{g=1}^G \sum_{h=1}^H \widehat{f}_g(\cdot - \mathbf{t}_{gh}) \right|$. This is usually an ill-posed inverse problem, since the binary phase retrieval sub-case is already ill-posed. It is well known [23] that there are trivial ambiguities in the phase retrieval problem; for instance, it is impossible to determine the absolute position of the f_g , because $f_g(\cdot - \mathbf{t})$ with an arbitrary $\mathbf{t}$ has the same FT modulus as f_g . However, such ambiguities can often be resolved if prior information is available (see, e.g., [24]). The main prior information utilized in our case is provided by the availability of partial reconstructions obtained by other imaging methods and the knowledge that the function values are 0/1 (the effectiveness of such binary constraints has been demonstrated in [4]). In the next subsection we discuss an application for which our assumptions are largely justified.

We define the *continuous color map* $f: \mathbb{R}^2 \rightarrow \{0, 1, \dots, G\}$ with

$$f(\mathbf{x}) = \begin{cases} g, & \text{if there is a } g \in \{1, \dots, G\} \text{ with } f_g(\mathbf{x})=1, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Obviously, we can reverse this procedure by decomposing f into $f_1, \dots, f_G$, thus reconstructing the individual functions is equivalent to reconstructing the continuous color map.

In order to deal with the reconstruction problem on a computer, we need to discretize it. Following [25, Section 2] we define, for $N \in \mathbb{N}$ and positive *sampling distance* $d \in \mathbb{R}$, the discretization $f^{N,d}: [0, N-1]^2 \rightarrow \{0, 1, \dots, G\}$ of f by

$$f^{N,d}(y_1, y_2) = f\left(-\frac{N-2y_2-1}{2}d, \frac{N-2y_1-1}{2}d\right). \quad (2)$$

In this definition $[0, N-1]^2$ denotes the set of all *pixels* $\mathbf{y}^T = (y_1, y_2)$, where the positive integers y_1 and y_2 are less than N . Typically N and d would be chosen large enough so that $f(\mathbf{x}) = 0$ if $\mathbf{x}^T = (x_1, x_2)$ and $\max\{|x_1|, |x_2|\} \geq Nd/2$. We refer to such an $f^{N,d}$ in Eq. (2) as an $N \times N$ *color image*. Our algorithms aim at reconstructing such color images.

Since usually knowledge of N and d can be assumed or is not needed, we sometimes simplify the notation by using $\tilde{f}$ for a discretized version of a function f . Motivated by the application

in the next subsection, we refer to discretizations of the data $\left| \sum_{g=1}^G \sum_{h=1}^H \widehat{f}_g(\cdot - \mathbf{t}_{gh}) \right|$ as *detector data*.

To help the reader distinguish between continuous and discrete concepts, we use $\mathbf{x}$ (position in real space), $\mathbf{u}$ (vector in Fourier space), and $\mathbf{p}$ (position on a detector) throughout this paper to denote real vectors. The corresponding discrete counterparts are denoted by $\mathbf{y}$ (pixel in real space), $\mathbf{v}$ (vector in discretized Fourier space), and $\mathbf{q}$ (pixel on a detector). We note that, for

any $f_g: \mathbb{R}^2 \rightarrow \{0, 1\}$, the discrete Fourier transform (DFT) $\widehat{f_g^{N,d}}$ is an approximation of the discretization of the FT $\hat{f_g}$ with the same N but with a sampling distance $1/Nd$. Hence, for any fixed d , we can make the sampling distance in Fourier space as small as we wish by choosing N sufficiently large (this is usually referred to as *zero-padding*).

2.B. Application: Coherent Diffraction Imaging of Multiple Undeformed Grains

The aim in this application is to reconstruct collections of grains (small crystals) inside of bulk materials from (monochromatic) coherent X-ray diffraction images. The term “collection” means in our context that the relative positions of the grains are relevant. Typical collections are space-filling – for example, if the collection represents a sample from a *polycrystalline material* – but generally voids might be present, too. Grains are naturally 3D objects, but by considering 2D slices through the sample we restrict our attention to the two-dimensional case.

Initially, two aspects of a grain are relevant: (1) the positions in space occupied by the grain, and (2) the internal lattice structure specifying the positions of the grain’s unit cells. Signal/noise ratios and detector technology, however, imply that one cannot aim at atomic resolution for grain imaging of bulk materials (see, e.g., [26] for an unsuccessful attempt in electron microscopy). One thus needs to make assumptions about the imaged material. In our case, we assume that the grain lattice is perfect, which means that the grains are undeformed (see [6,7] for an imaging application in incoherent X-ray diffraction), deformed grains are outside the scope of this paper.

To date, collections of grains have been reconstructed only from incoherent X-ray diffraction data, but higher resolution images are expected to be obtained from coherent X-ray diffraction data. To facilitate robust reconstructions from coherent X-ray diffraction data, we propose to incorporate additional diffraction images from incoherent X-ray diffraction data into the reconstruction process, see Fig. 1. In fact, the lattice of undeformed grains can be inferred already from incoherent X-ray diffraction data [27,28], and the reconstruction task from coherent X-ray diffraction data thus reduces to the task of determining for each grain the positions it occupies in space. Fig. 4(e) shows a 2D collection of grains, the internal lattice structure is hidden.

Referring to the global sample coordinate system, we can describe each grain, labeled by a number $g \in \{1, \dots, G\}$, by two parameters f_g and $\mathcal{B}_g$, where $f_g: \mathbb{R}^2 \rightarrow \{0, 1\}$, the so-called *shape function*, is a function whose (necessarily bounded) support in the sample is occupied by the grain, and $\mathcal{B}_g \subseteq \mathbb{R}^2$ is a set of two basis vectors defining the internal 2D lattice structure. A collection of G grains is thus specified by $f_1, \dots, f_G$ and $\mathcal{B}_1, \dots, \mathcal{B}_G$. Positions in real-world samples are occupied by at most one grain, and thus f in Eq. (1) and its discretized version $f^{N,d}$ in Eq. (2), called the *grain map* in this application, are well-defined. For short, we refer to a grain labeled g as grain g ; note that this is also the color of the pixels occupied by that grain.

2.B.1. Forward Model—In the following we briefly describe our forward model, used to obtain from G , H , and $f_g, \mathcal{B}_g, \mathbf{t}_{gh}$, for $g = 1, \dots, G$, $h = 1, \dots, H$, the corresponding coherent X-ray diffraction image measured on a detector. For a single grain g , it is well-known [13] that measured detector intensities are proportional to the squared modulus of superpositions of FTs of shape functions f_g , and displacement vectors $\mathbf{t}_{gh}$ correspond to the reciprocal lattice positions (Bragg peaks) of the grain. The theory of coherent X-ray diffraction from multiple grains has not yet been developed in detail, but it is plausible that the signals (for Bragg diffraction) from

multiple grains superpose with a subsequent phase loss owing to the detector reading. It is well known that there is a one-to-one correspondence between grain lattices and their reciprocal lattices. The one-to-one correspondence is given – if we model both lattices by sums of delta distributions – by the FT [29,30]. Thus from $\mathcal{B}_g$ we can compute the $\mathbf{t}_{gh}$, for $h = 1, \dots, H$.

Fourier space can be probed by a properly aligned detector. If the detector had arbitrarily fine resolution, we could measure values proportional to $\left| \sum_{g=1}^G \sum_{h=1}^H \widehat{f}_g(\cdot - \mathbf{t}_{gh}) \right|^2$. We are free to set the proportionality factor involved to 1, because in practice one observes that the technical detector implementation introduces yet another factor. Taking square roots shows that the measured data are exactly of the form stated in the binary superposed phase retrieval problem. The detector, however, has a finite resolution. Thus an additional discretization step is involved, which is discussed in Subsection 3.B.

2.B.2. Inverse Problem—Assuming, as previously discussed, that G, H , and $\mathcal{B}_g, \mathbf{t}_{gh}$, for $g = 1, \dots, G, h = 1, \dots, H$, are known, the inverse problem in this application is to reconstruct the grain map $\tilde{f} (= f^{N,d})$, for suitable selected N and d). We have discussed in the previous section that this is a binary superposed phase retrieval problem.

3. Algorithm I: Monte Carlo Approach

Our first method is a Monte Carlo method. Similar algorithms have been developed for there construction of grain maps from to mographic diffraction data that have been acquired by *incoherent* X-rays [6,7,20–22]. The main difference from the algorithms presented in those papers is that here we have to use a totally dissimilar forward operator.

Let $\mathcal{I}$ denote the set of all possible $N \times N$ color images with at most $G + 1$ colors, and let $\mathcal{D}$ denote the set of all possible detector data. Elements of these sets can be rearranged in to vectors from finite-dimensional real vector spaces, and we thus have norms, such as the ℓ_1 -norm $\|\cdot\|_1$ at our disposal. For $\mathcal{L} \subseteq [0, N-1]^2$ and $f_{init} \in \mathcal{I}$, we define $\mathcal{I}(f_{init}, \mathcal{L})$ to be the set of color images $\tilde{f} \in \mathcal{I}$ whose pixel values are equal to those of f_{init} at each pixel in $[0, N-1]^2 \setminus \mathcal{L}$. The idea behind this notation is that we can think of f_{init} as a partial image reconstruction obtained by other methods, and of $\mathcal{L}$ as being the set of so-called *ambiguous pixels*, whose values are to be reconstructed with our method. We explicitly allow $\mathcal{L} = [0, N-1]^2$, which amounts to $\mathcal{I}(f_{init}, [0, N-1]^2) = \mathcal{I}$. Moreover, we assume that $\mathcal{L}$ is a subset of the union of the supports of $f_1, \dots, f_G$. It is straightforward to adapt our algorithm to the situation in which this assumption is violated.

Let $\mathbf{P}: \mathcal{I} \rightarrow \mathcal{D}$ denote the forward operator mapping each color image $\tilde{f}$ to its detector data, and let $\mathbf{P}_0 \in \mathcal{D}$ denote the measured detector data. As for many inverse problems, the general idea is to find

$$f^* := \operatorname{argmin}_{\tilde{f} \in \mathcal{I}(f_{init}, \mathcal{L})} \left\| \mathbf{P}(\tilde{f}) - \mathbf{P}_0 \right\|_1. \quad (3)$$

A technique borrowed from statistical mechanics is to introduce the family of probability distributions $\gamma_\beta: \mathcal{I}(f_{init}, \mathcal{L}) \rightarrow \mathbb{R}_+, \beta \in \mathbb{R}_+$, with

$$\gamma_\beta(\tilde{f}) = \frac{1}{Z_\beta} \exp(-\beta \left\| \mathbf{P}(\tilde{f}) - \mathbf{P}_0 \right\|_1), \quad (4)$$

where Z_β is a (usually unknown) normalizing constant. Distributions of the form in Eq. (4) are known in the literature as *Gibbs distributions* [31]. We “solve” Eq. (3) by sampling from the distribution γ_β . If $\mathcal{F}$ denotes the set of minimizers in Eq. (3), then one can prove (Proposition 5.2.1 in [31]) that

$$\gamma(\bar{f}) := \lim_{\beta \rightarrow \infty} \gamma_\beta(\bar{f}) = \begin{cases} 1/|\mathcal{F}|, & \text{if } \bar{f} \in \mathcal{F}, \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

In other words, sampling from γ_β at high β yields the minimizers with high probability. Estimates for β are often hard to obtain and often too weak for practical purposes. In practice, however, it often turns out that one can sample efficaciously if β is chosen reasonably large. We remark in passing that a common technique is to vary β according to some “cooling” scheme, and the resulting algorithm becomes a *simulated annealing algorithm* [32]. We do not follow this approach – in this paper β is fixed.

We have not yet discussed how to sample from γ_β . The *Metropolis algorithm* introduced in [33] achieves exactly this. It simulates a Markov chain and, after a suitably long “burn-in phase,” it samples from γ_β . Our Algorithm I, based on the Metropolis algorithm, is outlined in Fig. 2.

We highlight four aspects of Algorithm I:

- only ambiguous image pixels as specified in the input are processed;
- in theory, the initial assignment of colors to the ambiguous pixels in f_{init} is irrelevant – in our implementation we choose a random assignment;
- the algorithm reduces to the original Metropolis algorithm if $\bar{f}$ instead of f^* is returned at termination; and
- the terms $\gamma_\beta(\bar{f})/\gamma_\beta(\bar{f})$ and $\gamma_\beta(f^*)/\gamma_\beta(\bar{f})$ can be evaluated without knowledge of Z_β .

In summary, the core of Algorithm I is the Metropolis-based algorithm (Fig. 2) applied to γ_β as defined in Eq. (4). In Subsections 3.A and 3.B we discuss how $\mathbf{P}$ is implemented.

3.A. Fast Updates of the Discrete Fourier Transforms

The key step of the Metropolis algorithm is the proposal of changing the current color assignment to a pixel $\mathbf{y}^* \in [0, N-1]^2$. This affects only two discretized functions, i.e., the value of an $\bar{f}_a(\mathbf{y}^*)$ changes from 0 to 1, while the value of an $\bar{f}_b(\mathbf{y}^*)$ changes from 1 to 0. For the forward projection we need to update their 2-dimensional DFTs. However, we do not have to resort to a full re-computation using a fast Fourier transform algorithm (FFT). To see this let $\bar{f}_{a,1}$ (resp. $\bar{f}_{b,0}$) denote the function resulting from $\bar{f}_a$ (resp. $\bar{f}_b$) after the change, and let $\widehat{f}_{a,1}^{N,d}$ (resp. $\widehat{f}_{b,0}^{N,d}$) and $\widehat{f}_a^{N,d}$ (resp. $\widehat{f}_b^{N,d}$) denote the corresponding DFTs, respectively. We immediately obtain

$$\begin{aligned} \widehat{f}_{a,1}^{N,d}(\mathbf{v}) &= \sum_{\mathbf{y} \in [0, N-1]^2} \bar{f}_{a,1}(\mathbf{y}) \exp(-2\pi i \mathbf{v}^T \mathbf{y} / N) \\ &= \bar{f}_a(\mathbf{y}^*) \exp(-2\pi i \mathbf{v}^T \mathbf{y}^* / N) + \sum_{\mathbf{y} \in [0, N-1]^2 \setminus \{\mathbf{y}^*\}} \bar{f}_a(\mathbf{y}) \exp(-2\pi i \mathbf{v}^T \mathbf{y} / N) \\ &= \exp(-2\pi i \mathbf{v}^T \mathbf{y}^* / N) + \sum_{\mathbf{y} \in [0, N-1]^2} \bar{f}_a(\mathbf{y}) \exp(-2\pi i \mathbf{v}^T \mathbf{y} / N) \\ &= \exp(-2\pi i \mathbf{v}^T \mathbf{y}^* / N) + \widehat{f}_a^{N,d}(\mathbf{v}), \end{aligned} \quad (6)$$

for all $\mathbf{v} \in [0, N-1]^2$. In the same way, we obtain for the other change, that

$\widehat{f}_{b,0}^{N,d}(\mathbf{v}) = \widehat{f}_b^{N,d}(\mathbf{v}) - \exp(-2\pi i \mathbf{v}^T \mathbf{y}^*/N)$. We update the DFTs according to these formulas, i.e., we add(or subtract, respectively) the term $\exp(-2\pi i \mathbf{v}^T \mathbf{y}^*/N)$ from the previous DFT. Computing the DFTs in this way gives an $O(N^2)$ algorithm – FFT algorithms [34], on the other hand, would require a computation time of the order $O(N^2 \log(N^2))$.

To achieve an additional speed-up, one could compute and store the values of $\exp(-2\pi i \mathbf{v}^T \mathbf{y}^*/N)$ prior to the execution of the Metropolis algorithm. For memory reasons we decided against storing the values for all N^4 possible $(\mathbf{v}^T, \mathbf{y}^{*T}) = (v_1, v_2, y_1^*, y_2^*)$ -tuples. In our implementation we store the values of $\exp(-2\pi i v_1 y_1^*/N)$ for all (v_1, y_1^*) -pairs. The evaluation of $\exp(-2\pi i \mathbf{v}^T \mathbf{y}^*/N)$ requires then only one complex multiplication. A further speed-up is discussed in Subsection 3.C.

3.B. Computation of Simulated Detector Images

In order to calculate the ℓ_1 -norm in (4) we need to estimate $\mathbf{P}(\vec{f})$ for each pixel for which we have a detector reading in $\mathbf{P}_0$. Let us denote the spacing between detector pixel centers (both horizontally and vertically) by δ , this is the sampling distance of the detector data. As explained in the last paragraph of Subsection 2.A, for any $f_g: \mathbb{R}^2 \rightarrow \{0, 1\}$, the discretization of the FT $\hat{f}_g$ with a sampling distance $\delta' = 1/Nd$ can be estimated from $f_g^{N,d}$.

During the execution of the Metropolis Algorithm (Fig. 2), we need to calculate $\mathbf{P}(\vec{f})$ for various color images $\vec{f}$. Let us first look at the special simple case when $H = 1$ and $\mathbf{t}_{g1}$ is the zero vector, for $g = 1, \dots, G$. Then the detector data is a discretization of

$\left| \sum_{g=1}^G \sum_{h=1}^H \widehat{f}_g(\cdot - \mathbf{t}_{gh}) \right| = \left| \sum_{g=1}^G \widehat{f}_g(\cdot) \right|$. We see that, by an appropriate selection of N and d , we can make sure that $\delta/\delta' \in \mathbb{N}$. This means that in this case the values of $\mathbf{P}(\vec{f})$ at the detector pixels can be directly estimated by decomposing $\vec{f}$ and taking DFTs.

Consider now the general case. What we need is a method for estimating at the detector pixels the values of $\hat{f}_g(\cdot - \mathbf{p})$ for an arbitrary $\mathbf{p} \in \mathbb{R}^2$. As in the previous paragraph, we generate an approximation to a discretization of $\hat{f}_g(\cdot)$ with sampling distance δ' . By shifting the points at which $\hat{f}_g(\cdot)$ has been estimated by $\mathbf{p}$, we get estimates at certain points on the detector plane of $\hat{f}_g(\cdot - \mathbf{p})$. This situation is depicted in Fig. 3: the shaded dots indicate locations of the detector readings and the white dots indicate the places where we have estimates of $\hat{f}_g(\cdot - \mathbf{p})$. By selecting $\delta/\delta' \in \mathbb{N}$ large enough, we can ensure that for every shaded dot $\mathbf{p}'$ there will be an empty dot $\mathbf{p}^*$ that is near enough to it so that $\hat{f}_g(\mathbf{p}' - \mathbf{p})$ can be approximated by the known $\hat{f}_g(\mathbf{p}^* - \mathbf{p})$. Furthermore, and this is important for computational efficiency, due to the fact that $\delta/\delta' \in \mathbb{N}$, for any $\mathbf{p}$ we need to find only one such pair of points $\mathbf{p}'$ and $\mathbf{p}^*$; for all other detector pixels the contribution of $\hat{f}_g(\cdot - \mathbf{p})$ can be identified by a simple subsampling (using sampling distance δ) of the discretization of $\hat{f}_g(\cdot)$ with sampling distance δ' ; see Fig. 3. (This technique is similar to the so-called footprint-based texture mapping in computer graphics.) We remark that the process is repeated for several f_g s and $\mathbf{p}$ s. The detector values are not overwritten: all complex values at each detector pixel are summed up. The very last step to obtain the final detector image $\mathbf{P}(\vec{f})$ is to replace each complex valued detector pixel entry by its absolute value.

3.C. Additional Code Optimizations

Several other speed-ups have been implemented. The following list gives a short overview.

- *Updating of the simulated detector image:* We have discussed in Subsection 3.A how we update the DFTs of a single f_g . In a similar way we update the whole simulated detector image. We are not only keeping the simulated detector image, i.e., the

modulus of the superposed DFTs, but we also store the complex values in a different array. Now, if two DFTs are changed in a Metropolis proposal step, then we first subtract both unchanged DFTs (placed at positions as described in Subsection 3.B) from the complex valued array. In a second step, we add the DFTs of the changed DFTs to the complex valued array. Finally, we obtain the simulated detector image by taking the absolute values of these complex numbers. The described updating technique has the advantage that only two DFTs are involved when the simulated detector image is recalculated. The drawback, however, is that rounding errors might accumulate.

- *Avoiding unnecessary reassignments of large data arrays:* At each step of the Metropolis algorithm one needs to compute the simulated detector image for the proposed pixel change. We store these values in a different array, because this permits a quick of the old detector values in case the proposal is rejected. We only need to keep a pointer telling us which of the two arrays contains the current simulated detector image.
- *Look-up tables for frequent computations:* The idea of using look-up tables for frequently occurring computations has already been discussed for the DFT updates. We also use look-up tables in another situation: when the DFTs are copied onto the various places on the detector one needs to ensure that no pixels outside of the detector are considered. Instead of performing range checks in each iteration, we store the permissible ranges for all DFTs and displacement vectors in a look-up table. Thus, we need no range checks, we use only loops with variables in the range indicated by the look-up table.

4. Algorithm II: Adaptation of Phase Retrieval Algorithms

In the following we introduce a second method for binary superposed phase retrieval, called *Algorithm II*, that uses as a subroutine a standard phase retrieval method. In the next subsection we discuss the general framework of Algorithm II, details of the subroutine are given in Subsection 4.B.

4.A. Adaptation

Phase retrieval methods reconstruct single functions from single FTs. In certain cases we can reconstruct the $f_1, \dots, f_G$ independently. For example, if for each $g' \in \{1, \dots, G\}$ there is an $h' \in \{1, \dots, H\}$ such that, for all $\mathbf{u}$ in a neighborhood of $\mathbf{t}_{g'h'}$, we have

$$\left| \sum_{g=1}^G \sum_{h=1}^H \widehat{f}_g(\mathbf{u} - \mathbf{t}_{gh}) \right| \approx \left| \widehat{f}_{g'}(\mathbf{u} - \mathbf{t}_{g'h'}) \right|, \quad (7)$$

then we could try to reconstruct $f_{g'}$ from the data using a phase retrieval method. This motivates the following algorithm, which is our Algorithm II (with an integer parameter pycles): select for each $g' \in \{1, \dots, G\}$ the $h' \in \{1, \dots, H\}$ whose minimum distance to any other displacement vector $\mathbf{t}_{gh}$ is maximal and reconstruct $f_{g'}$ with pycles iterations of a phase retrieval algorithm

using the data $\left| \sum_{g=1}^G \sum_{h=1}^H \widehat{f}_g(\mathbf{u} - \mathbf{t}_{gh}) \right|$, with $\mathbf{u}$ in the neighborhood of $\mathbf{t}_{g'h'}$.

Equation (7) is a valid approximation only in special scattering situations in which superposition effects from different FTs are small. Our simulations in Section 5 include several

scattering situations, and the results can thus be interpreted as testing whether Algorithm II gives useful results in situations in which Eq. (7) does not strictly hold.

4.B. HIO

The phase retrieval algorithm that we use as subroutine in Algorithm II is basically the HIO algorithm introduced in [14]. Our exposition closely follows [35]. We remark that one could replace HIO with any other phase retrieval algorithm, e.g., with the *difference map* [36].

HIO follows the Gerchberg-Saxton [37] scheme for reconstructing a function f from its FT modulus $|\hat{f}|$ by generating a sequence $(s_k)_{k \in \mathbb{N}}$ of functions that approximate f . Starting with an initial guess s_0 (usually randomly chosen, in our case function values are randomly chosen from $\{0, 1\}$), the iteration step from s_k to s_{k+1} is as follows: (1) the amplitude of the FT of s_k is replaced by $|f|$; (2) the inverse FT of this function yields a function s'_k ; and finally (3) we obtain s_{k+1} from s_k , s'_k , and prior information.

Algorithms following the Gerchberg-Saxton scheme differ in the implementation of Step (3). For real, non-negative valued functions f – as in our case – one typically assumes that an estimate S_{object} of the support of f is available (details for our choice in the simulations are given in Section 5). Let c_k be defined as

$$c_k(\mathbf{x}) = \begin{cases} \Re(s'_k(\mathbf{x})), & \text{if } \mathbf{x} \in S_{object} \text{ and } \Re(s'_k(\mathbf{x})) > 0, \\ 0, & \text{otherwise,} \end{cases} \quad (8)$$

where $\Re(s'_k(\mathbf{x}))$ denotes the real part of the complex value $s'_k(\mathbf{x})$. The idea is that $s'_k(\mathbf{x})$ violates prior knowledge if $|c_k(\mathbf{x}) - s'_k(\mathbf{x})| > \varepsilon$, where ε denotes a prescribed tolerance level (see Subsection 5.B). In HIO we have

$$s_{k+1}(\mathbf{x}) = \begin{cases} s'_k(\mathbf{x}), & \text{if } |c_k(\mathbf{x}) - s'_k(\mathbf{x})| \leq \varepsilon, \\ s_k(\mathbf{x}) + \beta_{HIO}(c_k(\mathbf{x}) - s'_k(\mathbf{x})), & \text{otherwise,} \end{cases} \quad (9)$$

with a prescribed system parameter β_{HIO} (see Subsection 5.B).

Algorithm II thus performs pcycles iterations of HIO, thereafter the reconstructed function values are rounded to 0 or 1 depending on whether or not the real part is smaller than 0.5.

5. Simulations

In the following we describe five sets of simulations that we performed on a PC with a 2.66 GHz Intel Core2Duo processor and 3.23 GB RAM. The algorithms have been implemented as C programs running under Windows XP. Details of the five simulations are discussed below, a summary is given in Table 1.

5.A. Data Generation

For our tests we selected three grain maps of aluminum samples that were acquired by electron microscopy, more specifically, by electron backscattered diffraction (EBSD). In what follows we refer to them as Maps I, II and III, respectively. While we aim (in the next subsection) at

reconstructing maps $f^{N,d}$, we remark that each EBSD map is a higher-resolution discretization $f^{KN,d/K}$ of a continuous color map f . By introducing K , we ensure that the discretization for the simulated data generation is somewhat independent from and is more accurate than the discretization used for the reconstruction. This is a way of avoiding the commitment of the “inverse crime” of using the same model for generating the data for testing an algorithm as the one used in the design of the algorithm. These $f^{KN,d/K}$ are shown in the left column of Fig. 4. For Maps I and II (containing $G = 2$ grains) we have $N = 64$, and $K = 2$. For Map III (containing $G = 11$ grains) we have $N = 64$, and $K = 3$.

We generate the data (i.e., the moduli of the superposed FTs on the detector) from these maps at the sampling distance $1/Nd$, which is the same as the sampling distance when we approximate the FT $\hat{f}_g$ by the discrete Fourier transform (DFT) $f_g^{KN,d/K}$ (see the last paragraph of Subsection 2.A). However, we make the detector dimension $N_{det} \times N_{det}$ larger than $KN \times KN$, since the detector has to accommodate also the shifted versions of the $\hat{f}_g$. The actual values used in our five simulations are shown in Columns 2 – 5 of Table 1.

For each grain g we specify a lattice basis $\{\mathbf{a}_g, \mathbf{b}_g\}$, and define the set of $\mathbf{t}_{gh}$ s as $\{\lambda\mathbf{a}_g + \mu\mathbf{b}_g : \lambda, \mu \in \{-1, 0, 1\}\}$. We assume that $\mathbf{t}_{gh} = \mathbf{0}$ is in the center of the detector. Simulation 1 differs from Simulation 2 exactly in the choice of the $\mathbf{t}_{gh}$ s, while the $\mathbf{t}_{gh}$ s are equal in Simulation 1 (resp. 2) and Simulation 3 (resp. 4). Instead of listing the $\mathbf{t}_{gh}$ for each simulation, we choose to report two numbers “#Spots” and D_{best} , the former gives the total number of $\mathbf{t}_{gh}$ s whose discretized version $\bar{\mathbf{t}}_{gh}$ hits the detector, while the latter gives information about the general FT “overlap.” Since we choose for HIO the FT that is furthest apart from any other FT (see Section 4.A), it is natural to define D_{best} as $\min\{D_1, \dots, D_G\}$, where $D_g = \max_h \min_{(g',h') \neq (g,h)} \|\bar{\mathbf{t}}_{gh} - \bar{\mathbf{t}}_{g'h'}\|_\infty$ with $\bar{\mathbf{t}}_{gh}$ and $\bar{\mathbf{t}}_{g'h'}$ denoting the discretized version of $\mathbf{t}_{gh}$ and $\mathbf{t}_{g'h'}$, respectively, that appear on the detector. The values of “#Spots” and D_{best} in our simulations are given in Columns 6 and 7 of Table 1, respectively.

The detector data in our simulations are subjected to noise. The exact nature of the noise needs to be determined in future experiments, but at present it is clear that one needs to distinguish between random noise introduced by the detector reading and geometric noise that is due to uncertainties in the experimental geometry. To define these types of noise in precise terms, we need to introduce some notations. Let $\text{det\_value}(\mathbf{q})$ denote the value of the detector pixel $\mathbf{q}$ of the $N_{det} \times N_{det}$ detector image. A detector shift by, say s pixels, is obtained by replacing $\text{det\_value}(\mathbf{q})$ by $\text{det\_value}(\mathbf{q} + (s, 0)^T)$, with $\text{det\_value}(\mathbf{q} + (s, 0)^T) = 0$ if $\mathbf{q} + (s, 0)^T \notin [0, N_{det} - 1]^2$. In our simulations we distinguish four different noise scenarios:

- **No noise:** In this case no detector pixel value is altered;
- **Random noise:** The no noise detector image is altered by setting $\text{det\_value}(\mathbf{q}) \leftarrow \max\{0, \text{det\_value}(\mathbf{q}) + r\}$, where r is chosen uniformly at random from the interval $[-100 \cdot \text{det\_value}(\mathbf{q}), 100 \cdot \text{det\_value}(\mathbf{q})]$ for $\mathbf{q} \in [0, N_{det} - 1]^2$;
- **4 pixel shift:** The no noise detector image is shifted by 4 pixels;
- **8 pixel shift:** The no noise detector image is shifted by 8 pixels.

Detector data of all five no noise simulations are depicted in Fig. 5.

5.B. Recovery

As mentioned above that the aim is to reconstruct the color image $f^{N,d}$ for Maps I-III. Since we assume that prior information about approximate grain positions is available, we aim at reconstructing only “ambiguous pixels” and not the whole color image. We assume that the approximate grain position is given by means of the grain’s ($= f_g^{N,d}$) minimal bounding box that is extended by a number of N_{margin} pixels at each border, and by some partial reconstruction

that might have been obtained by other methods (e.g., by incoherent diffraction imaging). We simulate the partial reconstruction by eroding the grain boundaries. Those removed pixels that fall into the extended bounding boxes of at least two grains are defined to be the ambiguous pixels in our simulations. The right column of Fig. 4 shows the f^{N,d_s} together with the ambiguous pixels (colored white) and the total number of ambiguous pixels in the f^{N,d_s} is given in the second column of Table 2. Note that low-accuracy position information and lattice structure are implicitly incorporated into Algorithm II via the support constraints and the FT data that requires that the $\mathbf{t}_{ghs}$ (Bragg peak positions) be known.

We set up five simulations as summarized in Table 1. The parameter N_{margin} is set to 8, only in Simulation 5 we choose $N_{margin} = 0$. The only difference between Simulation 1–2 and Simulation 3–4 is that a different input map is used, while from Simulation 1 to Simulation 2 we only vary the set of $\mathbf{t}_{ghs}$. Simulation 5 tests the reconstruction of a more complex grain map.

A common technique adopted in many stochastic algorithms (and the HIO algorithm [35]) is to repeat reconstructions from random initial guesses and to return the solution with the best data fit (or objective function value) that occurred over these independent runs. We follow this approach and use for Algorithm I two independent reconstructions, while for Algorithm II we use five. Details about the initial guesses are discussed below. For each simulation 50 reconstruction results are obtained using Algorithm I and Algorithm II, respectively.

The number of iterations for all algorithms is $mc\_cycles = HIO\_cycles = 200$. Moreover, the following specific parameters are used.

Algorithm I—The parameter β in Eq. (4) is set to $\beta = 20$. The candidate colors for each ambiguous pixel are restricted to be the colors of the grains into whose extended bounding box it falls. Premature termination occurs if no change in the objective function γ_β is observed over 10 of the $mc\_cycles = 200$ iterations.

Algorithm II—In agreement with [14,35] we choose $\varepsilon = 0.01$ and $\beta_{HIO} = 0.9$. The support $S_{subject}$ of each grain is chosen as tight as possible, i.e., as the set of those pixels of the extended minimal bounding box that are either pixels in that grain or ambiguous. In each HIO iteration we enforce prescribed pixel values as additional object domain constraints. As initial guesses we randomly assign the values 0/1 to each ambiguous pixel in $S_{subject}$ (0 corresponding to background, 1 corresponding to grain pixel). Grains are processed in random order.

5.C. Results and Discussion

In Table 2 we report the results of the simulations. The reconstruction error is defined as the number of pixels in the reconstruction that differ from the pixels of the true color image. Since each simulation is based on 50 results obtained by each algorithm, we report mean reconstruction errors, median reconstruction errors, and the standard deviations (the latter two are given in parentheses). Additionally, we report mean running times (the standard deviation was comparatively small in all cases and is thus not reported). Each simulation is performed for each of the noise scenarios discussed above.

Several facts from Table 2 are worth noting:

- Algorithm II is usually dramatically faster than Algorithm I, but the quality of reconstruction is always better with Algorithm I. This can be explained by the fact that Algorithm I tries to fit simultaneously the complete data, while Algorithm II uses only one FT at a time for reconstruction.

- The random noise does not have an appreciable effect on the reconstruction quality of Algorithms I and II.

In Fig. 6 we show a typical difference image (black pixels correspond to pixels in the reconstruction that differ from the original image) obtained by Algorithm I in Simulation 5 (no noise). Reconstruction errors seemingly appear only at the grain boundaries, which is not surprising if one takes into account that data generation might introduce some uncertainties.

As in [6,7,20–22], one could think of improving Algorithm I by incorporating statistical information about grain shapes (via Gibbs priors) into the reconstruction process. However, we implemented this approach and did not find an appreciable improvement in our experiments.

One may argue that in Algorithm II the tight support constraints may not be available in practice. In this case one can use loose supports in the initial reconstructions, which could be tightened in subsequent reconstructions. Algorithm I, in contrast to Algorithm II, does not rely on support constraints, and thus its (already superior) performance will not be negatively affected by the weakening of support constraints.

Algorithm I and II, as presented, yield 2D images as output. Generalizations of the methods to 3D are straightforward. Algorithm I is particularly well-suited to be implemented in a parallel computing environment. Reconstructions in 3D can be obtained from 2D reconstructions if the data are acquired slice-by-slice. We remark, however, that CDI for multiple grains in 3D does not necessarily correspond to the 3D binary superposed phase retrieval problem as introduced in this paper if the data are acquired truly in 3D; in that case, the data would involve superposed FTs of *projections* of the $f_1, \dots, f_G$.

6. Conclusions

We have introduced the *binary superposed phase retrieval* problem, which generalizes the binary phase retrieval problem. As relevant application we discussed, under simplifying assumptions, the reconstruction of grain maps from coherent diffraction data. We proposed a Monte Carlo based algorithm and another algorithm based on standard phase retrieval methods. It appears that the Monte Carlo algorithm gives far superior results in cases where the superposition effects on the FTs cannot be ignored.

Acknowledgments

The authors would like to thank Ivan Vartanyants for helpful discussions and Fengxiang Lin for providing us with the EBSD maps. This work was partially supported by the EU program TotalCryst, the Danish National Research Foundation, the Danish XFEL initiative, and Award Number R01HL070472 from the National Heart, Lung, and Blood Institute. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Heart, Lung, And Blood Institute or the National Institutes of Health.

References

1. Millane RP. Phase retrieval in crystallography and optics. *J Opt Soc Am A* 1990;7:394–411.
2. Dainty, JC.; Fienup, JR. Phase retrieval and image reconstruction for astronomy. In: Stark, H., editor. *Image Recovery: Theory and Application*. Academic Press; Orlando, FL: 1987. p. 231–275.
3. Jaming P. Phase retrieval techniques for radar ambiguity problems. *J Fourier Anal Appl* 1999;5:309–329.
4. Lo VL, Millane RP. Reconstruction of compact binary images from limited Fourier amplitude data. *J Opt Soc Am A* 2008;25:2600–2607.
5. Poulsen, HF. *Three-Dimensional X-ray Diffraction Microscopy*. Springer-Verlag; Berlin: 2004.
6. Alpers A, Poulsen HF, Knudsen E, Herman GT. A discrete tomography algorithm for improving the quality of 3DXRD grain maps. *J Appl Cryst* 2006;39:582–588.

7. Alpers, A.; Rodek, L.; Poulsen, HF.; Knudsen, E.; Herman, GT. Discrete tomography algorithms for grain map reconstructions. In: Herman, GT.; Kuba, A., editors. *Advances in discrete tomography and its applications*. Birkhäuser; Boston, MA: 2007. p. 271-301.
8. Miao JW, Charalambous P, Kirz J, Sayre D. Extending the methodology of x-ray crystallography to allow imaging of micrometre-sized non-crystalline specimens. *Nature* 1999;400:342–344.
9. Miao J, Ishikawa T, Shen Q, Earnest T. Extending x-ray crystallography to allow the imaging of noncrystalline materials, cells, and single protein complexes. *Annu Rev Phys Chem* 2008;59:387–410. [PubMed: 18031219]
10. Zuo JM, Vartanyants I, Gao M, Zhang R, Nagahara LA. Atomic resolution imaging of a carbon nanotube from diffraction intensities. *Science* 2003;300:1419–1421. [PubMed: 12775837]
11. Vartanyants IA, Robinson IK, Onken JD, Pfeifer MA, Williams GJ, Pfeiffer F, Metzger H, Zhong Z, Bauer G. Coherent x-ray diffraction from quantum dots. *Phys Rev B* 2005;71:245302.
12. Pfeifer MA, Williams GJ, Vartanyants IA, Harder R, Robinson IK. Three-dimensional mapping of a deformation field inside a nanocrystal. *Nature* 2006;442:63–66. [PubMed: 16823449]
13. Ewald PP. X-ray diffraction by finite and imperfect crystal lattices. *Proc Phys Soc* 1940;52:167–174.
14. Fienup JR. Phase retrieval algorithms: A comparison. *J Opt Soc Am A* 1982;21:2758–2769.
15. Nietovesperinas M, Mendez JA. Phase retrieval by Monte-Carlo methods. *Opt Commun* 1986;59:249–254.
16. Nietovesperinas M, Navarro R, Fuentes FJ. Performance of a simulated-annealing algorithm for phase retrieval. *J Opt Soc Am A* 1988;5:30–38.
17. Navarro R, Fuentes FJ, Nietovesperinas M. Simulated annealing image-reconstruction in photon-limited stellar speckle interferometry. *Astron Astrophys* 1989;208:374–380.
18. Morris D. Simulated annealing applied to the Misell algorithm for phase retrieval. *IEE Proc H* 1996;143:298–303.
19. Miao J, Sayre D. On possible extensions of x-ray crystallography through diffraction-pattern oversampling. *Acta Cryst, Sect A* 2000;56:596–605. [PubMed: 11058847]
20. Alpers A, Knudsen E, Poulsen HF, Herman GT. Resolving ambiguities in reconstructed grain maps using discrete tomography. *Electron Notes Discrete Math* 2005;20:419–437.
21. Rodek L, Poulsen HF, Knudsen E, Herman GT. A stochastic algorithm for reconstruction of grain maps of moderately deformed specimens based on x-ray diffraction. *J Appl Cryst* 2007;40:313–321.
22. Kulshreshtha AK, Alpers A, Herman GT, Knudsen E, Rodek L, Poulsen HF. A greedy method for reconstructing polycrystals from three-dimensional x-ray diffraction data. *Inverse Probl Imaging* 2009;3:69–85.
23. Bates RHT. Uniqueness of solutions to two-dimensional Fourier phase problems for localized and positive images. *Comput Vis Graph Image Process* 1984;25:205–217.
24. Miao J, Sayre D, Chapman HN. Phase retrieval from the magnitude of the Fourier transforms of nonperiodic objects. *J Opt Soc Am A* 1998;15:1662–1669.
25. Herman GT, Davidi R. Image reconstruction from a small number of projections. *Inverse Problems* 2008;24:045011.
26. Wu J, Weierstall U, Spence JCH. Diffractive electron imaging of nanoparticles on a substrate. *Nat Mater* 2005;4:912–916. [PubMed: 16299507]
27. Lauridsen EM, Schmidt S, Suter RM, Poulsen HF. Tracking: A method for structural characterization of grains in powders or polycrystals. *J Appl Cryst* 2001;34:744–750.
28. Moscicki M, Kenesei P, Wright J, Pinto H, Lippmann T, Borbely A, Pyzalla AR. Friedel-pair based indexing method for characterization of single grains with hard x-rays. *Mat Sci Eng A-Struct* 2009;524:64–68.
29. Ashcroft, NW.; Mermin, DN. *Solid State Physics*. Thomson Learning; Toronto: 1976.
30. Warren, BE. *X-Ray Diffraction*. Dover Publications, Inc; New York: 1990.
31. Winkler, G. *Image Analysis, Random Fields and Markov Chain Monte Carlo Methods*. Springer-Verlag; Berlin: 2003.
32. Kirkpatrick S, CDG, Vecchi MP. Optimization by simulated annealing. *Science* 1983;220:671–680. [PubMed: 17813860]

33. Metropolis N, Rosenbluth AW, Rosenbluth MN, Teller AH, Teller E. Equation of state calculations by fast computing machines. *J Chem Phys* 1953;21:1087–1092.
34. Cooley JW, Tukey JW. An algorithm for the machine calculation of complex Fourier series. *Math Comp* 1965;19:297–301.
35. Vartanyants IA, Robinson IK. Partial coherence effects on the imaging of small crystals using coherent x-ray diffraction. *J Phys: Condens Matter* 2001;13:10593–10611.
36. Elser V, Rankenburg I, Thibault P. Searching with iterated maps. *Proc Natl Acad Sci USA* 2007;104:418–423. [PubMed: 17202267]
37. Gerchberg RW, Saxton WO. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik* 1972;35:237–246.

Figure 1.

Schematic of Grain Map Reconstruction from Coherent Diffraction Imaging (CDI) data. In addition to CDI data, we propose that incoherent X-ray diffraction images are acquired to provide partial grain maps as input for Algorithms I and II. Note, however, that Algorithms I and II do not strictly require such input data.

Input: A color image f_{init} with at most $G + 1$ colors and a set $\mathcal{L}$ of n ambiguous pixels defining the space $\mathcal{I}(f_{init}, \mathcal{L})$ of reconstructible color images.

Output: A color image $f^* \in \mathcal{I}(f_{init}, \mathcal{L})$ for which the value of γ is nearly maximal.

```

1: Choose  $\beta$  and mc_cycles sufficiently large (as discussed in the text);
2:  $f^* \leftarrow f_{init}$  and  $\bar{f} \leftarrow f_{init}$ ;
3: for  $i = 1$  to n·mc_cycles do
4:   Select randomly a candidate pixel  $\mathbf{q} \in \mathcal{L}$  and color  $c \in \{1, \dots, G\}$ ;
5:   Calculate  $r = \gamma_\beta(\bar{f}') / \gamma_\beta(\bar{f})$ , the ratio of the probabilities of the new color image  $\bar{f}'$ 
      (which we would obtain if we assigned color  $c$  to  $\mathbf{q}$ ) and the old color image  $\bar{f}$ ;
6:   Set  $\bar{f} \leftarrow \bar{f}'$  with probability  $\min(1, r)$ ;
7:   if  $\gamma_\beta(f^*) / \gamma_\beta(\bar{f}) < 1$  then
8:      $f^* \leftarrow \bar{f}$ ;
9:   end if
10: end for
11: return  $f^*$ ;

```

Figure 2.

Our general Metropolis-based algorithm for “solving” Eq. (3).

Figure 3.

Computation of $\mathbf{P}(\vec{f})$, see Subsection 3.B. Detector pixels are depicted as shaded dots. Locations at which estimates, provided by a shifted DFT, are available are depicted as white dots.

Figure 4.

(a) Map I comprising 2 grains, (b) corresponding $f^{N,d}$ with $N = 64$ and ambiguous pixels (white) for Simulations 1–2, (c) Map II comprising 2 grains, (d) corresponding $f^{N,d}$ with $N = 64$ and ambiguous pixels (white) for Simulations 3–4, (e) Map III comprising 11 grains, (f) corresponding $f^{N,d}$ with $N = 64$ and with ambiguous pixels (white) for Simulation 5.

Figure 5.
Detector data for (a) Simulations 1, (b) Simulations 2, (c) Simulations 3, (d) Simulations 4, and (e) Simulations 5.

Figure 6.
Typical reconstruction error of Algorithm I in Simulations 5 (no noise). Pixels of the reconstruction that differ from the original are depicted in black.

Table 1

Overview of the data generation in the simulations, see Section 5.

Simulations	Input Map	Size of Input Map ($KV \times KV$)	Detector Dimension ($N_{det} \times N_{det}$)	K	#Spots	D_{best}
1	Map I	128×128	200×200	2	14	35
2	Map I	128×128	200×200	2	16	25
3	Map II	128×128	200×200	2	14	35
4	Map II	128×128	200×200	2	16	25
5	Map III	192×192	500×500	3	69	21

Results of the simulations. Each entry is based on 50 simulations from random initial guesses. Errors (defined as number of pixels that differ from the corresponding pixels in the original) are reported as mean pixel errors, followed by the median and the standard deviation in parentheses. Median reconstruction times are given in the corresponding columns.

Table 2

	#Ambig. Pixels		Algorithm I		Algorithm II	
			error	time/s	error	time/s
Simul. 1:	103	No noise	0.1 (0.0/0.4)	21	5.5 (5.5/2.5)	4
		Random noise	0.0 (0.0/0.3)	21	6.4 (6.5/3.1)	4
		4 pixel shift	11.1 (8.0/7.8)	24	48.8 (49.0/5.5)	4
		8 pixel shift	17.3 (17.0/6.6)	24	45.4 (50.5/13.2)	4
Simul. 2:	103	No noise	4.2 (4.0/0.7)	23	24.1 (18.0/14.1)	4
		Random noise	4.4 (4.0/1.1)	24	23.9 (18.0/13.2)	4
		4 pixel shift	28.8 (29.0/6.4)	24	47.6 (54.0/17.5)	4
		8 pixel shift	44.4 (45.5/7.1)	23	54.0 (52.0/22.3)	4
Simul. 3:	311	No noise	0.0 (0.0/0.0)	61	8.4 (8.0/2.6)	4
		Random noise	1.4 (0.0/7.5)	66	9.6 (7.5/11.7)	4
		4 pixel shift	83.8 (81.0/19.9)	76	147.2 (144.5/10.2)	4
		8 pixel shift	112.4 (113.0/18.4)	81	146.8 (144.0/10.0)	4
Simul. 4:	311	No noise	2.9 (3.0/1.0)	68	135.5 (123.0/22.0)	4
		Random noise	2.8 (3.0/1.0)	68	133.3 (120.5/22.9)	4
		4 pixel shift	90.8 (89.5/20.2)	92	151.4 (142.5/15.0)	4
		8 pixel shift	128.7 (128.0/15.1)	93	153.0 (142.0/16.6)	4
Simul. 5:	1318	No noise	124.7 (124.5/17.3)	2; 265	582.0 (588.0/43.8)	106
		Random noise	123.3 (121.0/18.6)	2; 124	584.8 (583.0/45.5)	106
		4 pixel shift	478.6 (482.0/31.3)	2; 288	682.0 (682.5/58.5)	106
		8 pixel shift	574.5 (577.5/22.2)	2; 398	682.8 (679.5/61.6)	106