

DART: a robust algorithm for fast reconstruction of three-dimensional grain maps


K. J. Batenburg, J. Sijbers, H. F. Poulsen and E. Knudsen

J. Appl. Cryst. (2010). **43**, 1464–1473

Copyright © International Union of Crystallography

Author(s) of this paper may load this reprint on their own web site or institutional repository provided that this cover page is retained. Reproduction of this article or its storage in electronic databases other than as specified above is not permitted without prior permission in writing from the IUCr.

For further information see <http://journals.iucr.org/services/authorrights.html>



Many research topics in condensed matter research, materials science and the life sciences make use of crystallographic methods to study crystalline and non-crystalline matter with neutrons, X-rays and electrons. Articles published in the *Journal of Applied Crystallography* focus on these methods and their use in identifying structural and diffusion-controlled phase transformations, structure–property relationships, structural changes of defects, interfaces and surfaces, *etc.* Developments of instrumentation and crystallographic apparatus, theory and interpretation, numerical analysis and other related subjects are also covered. The journal is the primary place where crystallographic computer program information is published.

Crystallography Journals **Online** is available from journals.iucr.org

DART: a robust algorithm for fast reconstruction of three-dimensional grain maps

K. J. Batenburg,^{a,b*} J. Sijbers,^b H. F. Poulsen^c and E. Knudsen^c^aCentrum Wiskunde en Informatica, Science Park 123, NL-1098XG, Amsterdam, The Netherlands,^bUniversity of Antwerp, IBBT-Vision Laboratory, Universiteitsplein 1, B-2610, Wilrijk, Belgium, and^cCentre for Fundamental Research: 'Metal Structures in Four Dimensions', Risø National Laboratory, Technical University of Denmark, DK-4000 Roskilde, Denmark. Correspondence e-mail:

joost.batenburg@ua.ac.be

A novel algorithm is introduced for fast and nondestructive reconstruction of grain maps from X-ray diffraction data. The discrete algebraic reconstruction technique (DART) takes advantage of the intrinsic discrete nature of grain maps, while being based on iterative algebraic methods known from classical tomography. To test the properties of the algorithm, three-dimensional X-ray diffraction microscopy data are simulated and reconstructed with DART as well as by a conventional iterative technique, namely SIRT (simultaneous iterative reconstruction technique). For 100×100 pixel reconstructions and moderate noise levels, DART is shown to generate essentially perfect two-dimensional grain maps for as few as three projections per grain with running times on a PC in the range of less than a second. This is seen as opening up the possibility for fast reconstructions in connection with *in situ* studies.

© 2010 International Union of Crystallography
Printed in Singapore – all rights reserved

1. Introduction

During the past five years, hard X-ray imaging methods have appeared that enable nondestructive mapping in three-dimensions of the grains within undeformed polycrystals. Three-dimensional X-ray diffraction (3DXRD) microscopy has a setup and a reconstruction approach rather similar to those for classical parallel-beam tomography, but one probes the diffracted beam instead of the attenuation of the incoming beam (Poulsen, 2004). In a variant of 3DXRD microscopy known as diffraction contrast tomography (DCT), in addition one uses the fact that the diffracted beam from a given grain will give rise to an enhanced attenuation of the direct beam (Ludwig *et al.*, 2009). (In the language of electron microscopy, in DCT 'bright field' and 'dark field' imaging is combined.) In both cases, the reconstruction task is mathematically distinctly different from classical tomography, because projections of the grains are not apparent in each image but only at certain rotation angles, defined by Bragg's law. As such, typical transform methods such as 'filtered backprojection' give inferior performance. Instead, large-scale iterative algebraic methods have been used, such as the algebraic reconstruction technique (ART; Poulsen & Fu, 2003), Monte Carlo methods (Suter *et al.*, 2006) and pure forward-projection methods (<http://sourceforge.net/apps/trac/fable/wiki>).

Currently, both the 3DXRD and DCT approaches have generated grain maps comprising hundreds or thousands of grains, and the first *in situ* studies have been reported (Schmidt *et al.*, 2004, 2008; King *et al.*, 2008). Nevertheless, there are still some obstacles to overcome:

(a) To facilitate *in situ* studies, the tendency is to go towards illuminating more grains at the same time (to achieve the required grain statistics faster), to reduce the acquisition time of the individual images and to reduce the number of projections taken during each scan. For reasons such as spot overlap and signal-to-noise considerations, this implies that the number of useful projections per grain becomes very small, in the limit only 2–5.

(b) Ultimately, the aim is to generate a time series of hundreds of three-dimensional maps during, for example, an annealing of the sample. With maps comprising up to $500 \times 500 \times 500$ voxels, run time is a real concern. Furthermore, to take full advantage of the synchrotron facilities, the ability to compute accurate reconstructions within minutes would be a major step forward. This imposes significant constraints on the reconstruction algorithms that can be used, even with generous use of clusters.

A possible approach to deal with a small number of projections per grain is to construct new fast reconstruction algorithms that explicitly make use of the *a priori* knowledge we have of undeformed polycrystals, namely that

(1) the grains are approximately discrete objects in the sense that a given voxel in the sample will either belong fully to a given grain or not belong at all;

(2) the grains are simply connected three-dimensional space-filling objects within the borders given by the sample geometry, and hence it is sufficient to reconstruct the morphology of the grain boundaries;

(3) the grain boundaries are smooth and the grains are nearly convex.

It has been noted that the problem of grain reconstruction fits well within the emerging mathematical discipline of discrete tomography (Herman & Kuba, 1999). In particular, discrete tomography deals with the problem of reconstructing binary images, which corresponds to feature (1) listed above. By incorporating one or more of the listed features in the reconstruction algorithm as prior knowledge, the accuracy of the reconstruction can be vastly improved. The fact that the range of pixel values is not a convex set differentiates this problem from the approach of convex set theory to incorporate prior knowledge in tomography (Sezan & Stark, 1982). In a series of papers, Alpers, Rodek and co-authors (Alpers *et al.*, 2006; Rodek *et al.*, 2007) presented such an implementation for 3DXRD work based on a Monte Carlo scheme and the use of Gibbs priors. The superiority of the discrete methods was demonstrated. However, the Monte Carlo scheme is inherently slow.

In this paper, we present a grain-map reconstruction scheme based on the discrete algebraic reconstruction technique (DART). DART is a new discrete tomography algorithm (Batenburg & Sijbers, 2007) that is based on an underlying conventional iterative algebraic reconstruction method, such as ART (Gordon *et al.*, 1970) or SIRT (simultaneous iterative reconstruction technique; Gilbert, 1972; Gregor & Benson, 2008). The acronym 'ART' in DART therefore refers to a general iterative algebraic reconstruction technique. Iterations of the continuous algebraic technique are alternated with segmentation and fixation steps, which restrict the updates to the boundary of the reconstructed object. We apply this algorithm to the most basic 3DXRD-type analysis: that of reconstructing one internal two-dimensional section in a monophase undeformed polycrystal. Similar to the previous work by Poulsen and Fu based on ART, the

new approach involves three steps: identification of grains by an existing polycrystalline indexing program, reconstruction of the boundary of each grain separately and the stitching together of the boundary maps for the grains.

Reconstructing a single grain at a time allows for the application of more general reconstruction algorithms used in other tomography applications, which do not have to take into account the constraints with respect to the connectivity and overlap between grains. The stitching phase that follows the reconstruction ensures that these constraints are satisfied in the final reconstructed grain map.

Following an outline of the suggested approach, in §2, 3DXRD is modelled as an algebraic reconstruction problem. In §3, the quality of a set of 3DXRD grain-map simulations is reported, comparing at all times the use of DART with that of a conventional iterative algorithm, SIRT. The robustness of the two algorithms towards the number of projections and noise is surveyed. Next, in §4, the performance is discussed with respect to the Monte Carlo-based algorithms, and the prospect of generalizations is outlined.

2. Approach

A full description of 3DXRD microscopy is beyond the scope of this article and can be found elsewhere (Poulsen, 2004), yet a brief explanation of the relevant operational mode will be given here.

3DXRD microscopy can be performed in various modes (Poulsen, 2004). For simplicity we will in the following assume the setup shown in Fig. 1. Firstly, the incoming beam is focused in one direction: the direction parallel to the axis of rotation. Secondly, focusing is moderate so that the beam may be considered collimated, and thirdly, two area detectors are used: a low-resolution large effective area detector positioned far (~ 30 cm) from the sample and a high-resolution area detector close to the sample (~ 5 mm).

Thus, a collimated monochromatic X-ray beam impinges on a polycrystalline sample. Depending on the X-ray wavelength and the orientation of the crystallites, or grains, in the sample, a fraction of the beam is diffracted according to Bragg's law. Assuming the composition and space group of the crystallite is known, the intensity of one diffracted beam relative to another originating from the same crystallite is known. Further, in the limit of kinematical scattering, the overall intensity diffracted by each crystallite is directly proportional (Warren, 1990) to the number of unit cells in the crystallite, and therefore proportional to the size of the crystallite. Modelling the polycrystalline sample by volume elements (voxels) where each voxel is considered to have one crystallographic orientation only, we may apply inverse methods, making it possible to reconstruct three-dimensional orientation maps from the detected diffracted intensity.

When the sample is undeformed, it generally consists of crystallites with negligible spread in orientation within each crystallite and no voids in between. In such cases it is often convenient to consider the crystallites as independent objects.

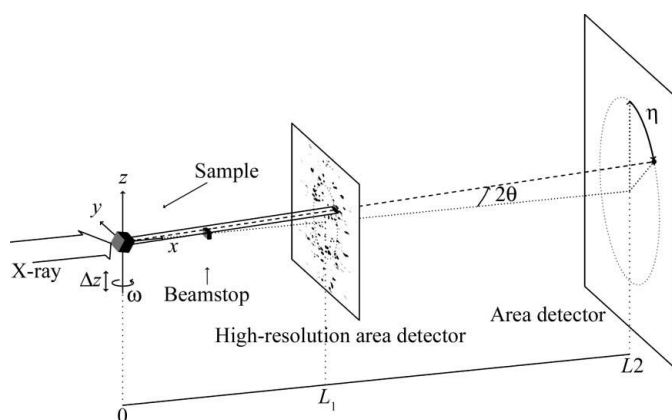


Figure 1

Schematic drawing of the 3DXRD setup. An X-ray beam, focused in the vertical direction, impinges on the sample. The sample is mounted on a rotation stage, which is used to bring different lattice planes into the diffracting condition. Diffracted intensity is detected on one or more semi-transparent area detectors. In the present case, a large-area low-resolution detector is positioned further away from the sample and a high-resolution detector close to the sample. The low-resolution detector data are used for indexing purposes, whereas the high-resolution data are used for reconstruction.

Each object may then be described by its orientation, centre-of-mass position and boundary morphology and may be handled separately. In the particular case presented here where the beam is focused in one direction, the boundary morphology can be reduced to a two-dimensional map. Assuming that the boundaries of the individual crystallites have been reconstructed correctly, these boundary maps can be assembled (stitched) to form a two-dimensional orientation map (grain map) of a slice through the complete sample. By displacing the sample vertically and repeating the procedure, a set of grain maps can be reconstructed to form a three-dimensional grain map of the sample.

As mentioned, the proposed approach for constructing the two-dimensional grain map for a given layer comprises three steps: identifying (indexing) grains, reconstructing the shape for each grain and stitching the grains together. The first task is handled by a polycrystalline indexing algorithm such as *GRAINDEX* (Lauridsen *et al.*, 2001), *ImageD11* (<http://sourceforge.net/apps/trac/fable/wiki>) or *GrainSpotter* (<http://sourceforge.net/apps/trac/fable/wiki>) operating on data from the low-resolution area detector. Since the low-resolution detector has a larger effective area, the indexing algorithm may include more orders of diffraction, and therefore more spots, than the reconstruction algorithm. The lower resolution is acceptable in this step since only the positions of the diffraction spots are used. For each grain, these routines will output the orientation, approximate size and a list of associated diffraction spots. For each spot, we may compute its position and identify it with a spot found on the high-resolution detector. Furthermore, we can identify a region-of-interest (ROI) around each identified spot, within which the spot is confined. It is important to keep in mind that, although a spot may be used for indexing and a finite ROI may be defined on the high-resolution detector, the spot may not be useful for reconstruction. There may, for instance, be imperfections in the detector in that region. Further, as diffraction is governed by Bragg's law, all spots on the far-field detector appear on Debye–Scherrer rings, at fixed scattering angles (Warren, 1990). Because of this, the full two-dimensional-phase space of the detector may not be filled, even with an infinite number of grains. Instead, an increased number of grains simply increases the likelihood of overlapping spots. Such overlapping spots can be allowed in the indexing algorithms, provided that a full set of spots, matching the crystallographic structure of the sample material, can be identified. On the other hand such spots have to be excluded from an analysis based on separate grains (such as the one presented below), or preprocessed by specialized algorithms. Further, the integrated diffracted intensity in each spot originating from one grain should, when corrected for crystallographic structure factors and Lorentz factors, match up. Spots associated with more than one grain and spots with integrated intensities that do not match are considered overlapping.

As it cannot be known *a priori* to what extent spots from different grains will overlap, a practical reconstruction algorithm should therefore be flexible with respect to both quality and the number of spots used for reconstruction.

2.1. Algebraic formulation

We proceed by discretizing the layer of interest in the sample into voxels by imposing a regular grid. Next, for each grain and each reflection r , the basic assumption of kinematical scattering implies that one can formulate a linear relationship between the density of the grain in the various voxels in the layer and the integrated intensity in the detector pixels within the ROI. We may write this as $\mathbf{A}^r \mathbf{x} = \mathbf{b}^r$, where \mathbf{A}^r is a matrix comprising information on the geometry associated with the projection from the grain into the reflection spot r , \mathbf{x} is a column vector where element j contains the density of the grain in the j th voxel and \mathbf{b}^r is a column vector consisting of recorded pixel intensities on the area detector. More precisely, the value of element a_{ij}^r of \mathbf{A}^r indicates to what extent the projection of the j th sample voxel contributes to the i th detector pixel given the projection geometry associated with diffraction spot r . In other words a_{ij}^r indicates the fraction of the X-ray intensity diffracted from the j th voxel into reflection r that is deposited in the i th detector pixel. Herein lies a subtle difference from transmission tomography, where the geometry, and therefore \mathbf{A}^r , is completely determined by the setup and does not depend on the sample data. In our case, as implied earlier, \mathbf{A}^r depends on the chosen reflection r and on the crystallographic orientation of the grain, and therefore on the previous indexing of diffraction spots.

Next, we form a compound matrix $\mathbf{A} = (a_{ij})$, by stacking the \mathbf{A}^r matrices for all the reflections belonging to a given grain, and similarly define the compound vector \mathbf{b} :

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}^1 \\ \mathbf{A}^2 \\ \vdots \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} \mathbf{b}^1 \\ \mathbf{b}^2 \\ \vdots \end{pmatrix}. \quad (1)$$

With these definitions, the basic equation for reconstruction of the grain shape is

$$\sum_{j=1}^n a_{ij} x_j = b_i, \quad i = 1 \dots m. \quad (2)$$

In general, m and n will be large, the rank of $\mathbf{A} = (a_{ij})$ will be *a priori* unknown, and the set of equations may be under-determined as well as over-determined. For such situations, it is not possible to solve the linear system directly. Instead, iterative methods can be used that compute successive approximations of its solution. Poulsen & Fu (2003) solved equation (2) by applying the continuous ART routine with the following *ad-hoc* implementation of the constraint on the admissible intensity values (ρ is a normalized density of the material):

$$0 < x_j < \rho, \quad \forall j. \quad (3)$$

2.2. Reconstruction of one grain

In this paper, we will present results for two alternative ways to solve equation (2). The first one is SIRT – a classical iterative algorithm for continuous tomography (see ch. 7 of Kak & Slaney, 2001). In contrast to other methods such as

ART or SART (simultaneous algebraic reconstruction technique), where only some of the projection data are used to compute each update step, SIRT uses all available spot projections simultaneously to compute the update of the current grain image. This results in a longer computation time but reduces the influence of noise in the projection data on the final reconstruction.

The second algorithm, called DART, is a new approach, which combines elements of discrete and continuous tomography. It alternates between steps of a continuous reconstruction algorithm, such as SIRT, and discretization steps that involve tracking the grain boundary. In the following, we describe the flow of the new algorithm in detail.

First, we briefly review the continuous SIRT algorithm. Each iteration of the SIRT algorithm consists of three basic steps. First, the vector \mathbf{b}' of projections of the current reconstruction \mathbf{x} is computed: $\mathbf{b}' = \mathbf{A}\mathbf{x}$. Next, the difference \mathbf{d} between the projections of \mathbf{x} and the measured projections \mathbf{b} is computed: $\mathbf{d} = \mathbf{b}' - \mathbf{b}$. The difference for each detector pixel is then distributed among all image pixels that contribute to this detector value. The update for each pixel is computed by averaging the computed updates for all detector pixels onto which this pixel projects. To be more exact, the value $x_j^{(k+1)}$ of voxel j in iteration $k + 1$ is computed from the current reconstruction $x^{(k)}$ by the following update equation:

$$x_j^{(k+1)} = x_j^{(k)} + \frac{\sum_i \left[a_{ij} \left(b_i - \sum_h a_{ih} x_h^{(k)} \right) / \sum_h a_{ih} \right]}{\sum_i a_{ij}}. \quad (4)$$

This iterative scheme converges to a weighted least-squares solution of the reconstruction problem [equation (2)] (Gregor & Benson, 2008). The second term in the update equation (4) above is often scaled by a relaxation parameter, $0 < \lambda < 2$, which may improve the rate of convergence, depending on the conditioning of the matrix \mathbf{A} . We used $\lambda = 1$ in all experiments reported below.

Fig. 2(a) shows a grain phantom image for which three projection spots have been simulated. Fig. 2(b) shows the resulting SIRT reconstructions from the simulated noiseless projections. Although the original grain image contains only two grey levels, representing the interior and background (exterior) of the grain, the SIRT reconstruction contains a broad spectrum of grey levels. Of course, the SIRT reconstruction can also be converted into a binary image by thresholding it. Fig. 2(c) shows the result of this step, where the threshold has been set halfway between the grey levels for the interior and the background. These grey levels are ρ and 0 for the interior and exterior, respectively, where ρ is determined by the material density and experimental conditions such as background and impinging X-ray intensity.

This thresholded SIRT reconstruction forms the starting point for DART. The main idea of DART is the observation that, for homogeneous objects (such as the grains in our case), most of the errors in the reconstruction will be near the boundary of the object. The boundary of the current approximate reconstruction (*i.e.* the thresholded SIRT

reconstruction) can be computed automatically. The result is shown in Fig. 2(d). The set of pixels that are at this boundary are called boundary pixels, whereas the remaining pixels are called nonboundary pixels. We now assign all nonboundary pixels the discrete grey levels for the background (always 0) and grain interior, respectively. Next, a new series of SIRT iterations is performed, where only the boundary pixels are allowed to change their values. The nonboundary pixels are all kept fixed at their two discrete grey levels. The resulting system of linear equations has far fewer variables (as most of the pixels are kept fixed), yet it has the same number of equations as the original system. If the nonboundary pixels have been fixed at their correct values, the new solution for the boundary pixels will be much more accurate than the previous solution. Fig. 2(e) shows the result of applying two SIRT iterations in this manner.

On the other hand, suppose that a nonboundary pixel has been assigned the wrong discrete grey level (interior, instead of background). In that case, the SIRT iterations typically steer nearby boundary pixels into the opposite grey-value direction (*i.e.* towards the background), in order to compensate for this. This means that in the next step of DART, where a new boundary is formed, the erroneous pixel will now be at the boundary and will no longer be fixed. Although this is a rather heuristic argument, this mechanism can be clearly observed in experiments with DART.




Figure 2
Illustration of steps in the proposed reconstruction algorithm for one grain (§2.2).

The SIRT iterations result in a new, partially continuous, reconstruction. A smoothing filter is applied to the boundary pixels, after which a new segmented reconstruction is computed by thresholding. Fig. 2(f) shows the new boundary computed from this thresholded reconstruction. It can be clearly seen that, after a single DART iteration, the boundary matches the original phantom much more accurately than the original boundary computed from the SIRT reconstruction. The same DART procedure is now repeated several times. Fig. 2(g) shows the SIRT result after the second DART iteration. Fig. 2(h) shows the final segmented reconstruction, obtained by thresholding the SIRT result of the final DART iteration.

Fig. 3 provides an algorithmic description of DART. Every run of DART is characterized by three parameters: N_S , N_D and N_B . The parameter N_S corresponds to the number of SIRT iterations performed before entering DART. The parameter N_D corresponds to the total number of DART iterations, *i.e.* the total number of segmentation steps performed. Finally, the parameter N_B corresponds to the number of SIRT iterations performed on the boundary pixels within each DART iteration. Values for these parameters must be chosen by the user.

After computing the initial reconstruction \mathbf{x}^0 , the algorithm enters a loop. The threshold operation $r(x)$, which is applied to each component of the current reconstruction \mathbf{x}^{t-1} , is defined by

$$r(x) = 0 \text{ if } x \leq \rho/2 \text{ and } r(x) = \rho \text{ otherwise.} \quad (5)$$

It yields an image \mathbf{s}^t that only contains grey levels 0 and ρ . Based on this image, the set I^t of nonboundary pixels is computed, where a nonboundary pixel is defined as a pixel for which all of its horizontal or vertical neighbours have the same value as that pixel. Based on \mathbf{x}^{t-1} and \mathbf{s}^t , a new image \mathbf{y}^t is computed that is equal to \mathbf{x}^{t-1} for the boundary pixels and equal to \mathbf{s}^t for the nonboundary pixels. In other words, the interior and exterior pixels of the grain have been segmented, while the boundary pixels still retain their continuous spectrum of grey levels. Next, a series of SIRT iterations is applied to the boundary pixels while keeping the interior and exterior pixels fixed at their segmented values, yielding \mathbf{x}^t . To apply SIRT while keeping a set of pixels fixed, the projections of these pixels are first subtracted from the projection data \mathbf{b} and

Compute a start reconstruction \mathbf{x}^0 by applying N_S iterations of SIRT;

for $t := 1$ to N_D **do**
begin

 Compute the segmented image $\mathbf{s}^t = r(\mathbf{x}^{t-1})$;

 Compute the set I^t of nonboundary pixels of \mathbf{s}^t ;

 Compute the image \mathbf{y}^t from \mathbf{x}^{t-1} and \mathbf{s}^t , setting
 $y_i^t := s_i^t$ if $i \in I^t$ and $y_i^t := x_i^{t-1}$ otherwise;

 Using \mathbf{y}^t as the start solution, apply N_B iterations of SIRT
 \mathbf{x}^t , while keeping the pixels in I^t fixed;

 Apply a smoothing operation to the pixels that are not in I^t ;

end

Figure 3

Pseudo-code description of the DART algorithm.

the corresponding columns are removed from the projection matrix \mathbf{A} . The SIRT algorithm is then carried out using the modified system of equations. In the SIRT algorithm, all pixels at the boundary are updated independently, without taking spatial smoothness into account. This can result in significant noise in the reconstructed boundary. A Gaussian smoothing operator is applied to the pixels at the boundary to compensate for this effect. This completes an iteration of the DART algorithm, which is then repeated iteratively.

2.3. Stitching grains together

A full two-dimensional grain map may be obtained by superposing the solutions (the shapes) of the individual grains. However, such a map may not be space filling as boundaries from neighbouring grains may overlap or leave ‘voids’ in the map. To overcome this problem, we use the continuous DART result, after the last series of SIRT iterations (see Fig. 2g). Assume that a pixel has a nonzero value in several of the grain reconstructions. Then, similar to the approach of Alpers *et al.* (2006), such a pixel will be associated with the grain for which the grey level is the highest. If, on the other hand, the pixel value is below a certain threshold in all grain reconstructions (or has a negative value), the pixel is not assigned to any grain. In §3, experimental results are presented for both SIRT and DART. For the SIRT reconstructions, the grain maps are formed using the same procedure as for DART.

If the projection data are perturbed by a significant amount of noise, the above procedure is too simplistic – it may lead to a grain map with noisy grain boundaries and spurious holes. A postprocessing step is then applied to the resulting grain map, using the following heuristic but fast and efficient filtering algorithm:

(a) Grain-boundary smoothing

In this step, grain-boundary pixels are iteratively processed using a 3×3 pixel neighbourhood. In a first phase, isolated pixels are removed by replacing the central voxel by its 4-connected neighbourhood value if all 4-connected neighbourhood pixels are equal.

Next, if the central pixel value differs from one or more of the pixel values in the 4-connected neighbourhood, it is tentatively replaced by the most frequent pixel value in an 8-connected neighbourhood (where the weight of each pixel is taken inversely proportional to the distance to the central pixel). Then, subsequently, the new 3×3 neighbour configuration is compared with the three cases shown in Fig. 4. If it is identical to one of them, the central pixel is replaced by the value of the respective configuration.

(b) Filling of spurious holes

Spurious holes (*i.e.* small background areas within the grain map) are removed by processing boundary background pixels employing the following heuristic rules:

(1) If at least two neighbouring 4-connected foreground pixels are identical, the central background pixel is replaced by the respective pixel value.

(2) If (1) does not hold, the central background boundary pixel value is replaced by the pixel value that occurs most

frequently in a 4-connected neighbourhood. If all 4-connected pixel values have equal frequency of occurrence, one of the pixel values is randomly selected.

This process is repeated until all background pixels have been processed.

3. Simulations

Two series of simulation experiments have been performed to compare the reconstruction quality of DART and SIRT, based on two different phantom maps.

The first series of experiments is based on the phantom 128×128 pixel orientation map shown in Fig. 5(a). The aim of these simulations was to quantify the reconstruction quality as a function of experimental noise and the number of reflections available per grain. In both cases, the more robust the algorithm, the faster the data acquisition can be and the more grains can be studied simultaneously. A comparison was made between the DART and SIRT solutions throughout.

The phantom map comprises 44 grains and originates from electron microscopy data acquired on an annealed Al sample. This map is seen as a standard test map, also used by Alpers *et al.* (2006) and Rodek *et al.* (2007). The simulated geometry of the 3DXRD microscope and the measuring sequence was also identical to the setting in these two previous papers. We refer to these papers for details and mention here only that the energy was set to 50 keV, and that 91 images were generated corresponding to equally spaced rotation angles over a rotation range of 90° . The {111}, {200}, {222} and {311} reflection families were included in the analysis.

Although the 128×128 phantom map in Fig. 5(a) is well suited for performing large numbers of computational experiments, because of its limited size it comprises a smaller number of grains than typically encountered in practical experiments. As a more complex test case, we performed a series of simulation experiments with the 400×400 pixel orientation map shown in Fig. 6. This larger map is the direct result of electron backscatter diffraction (EBSD) measurements on a pure copper sample.

This grain map is significantly more challenging for any reconstruction algorithm, as many grains have complex, nonconvex shapes and some grains contains inclusions of other grains. The experiments for this map aim to compare the reconstruction quality of SIRT and DART for the practically




Figure 4
Identification of certain neighbourhoods of relevance to the postprocessing step (§2.3). In each case all possible rotations are applied to find possible matches. If the neighbourhood matches that of the map, the central pixel will take the value of the 'black' pixels.

relevant case of moderate noise, combined with a small number of available projections per grain.

Throughout both series of experiments, a run of SIRT is characterized by the single parameter N_S , corresponding to




Figure 5
Example of simulated data. Each grey level represents one grain with one crystallographic orientation; black is background. (a) The phantom image obtained by electron microscopy on an annealed aluminium sample; (b) DART reconstruction with $N_S = 3$, $N_D = 3$, $N_B = 3$ and no noise; (c) SIRT reconstruction with $N_S = 15$ and no noise. The symbols N_S , N_D and N_B all refer to the number of iterations of various parts of the algorithm (§3).

the total number of SIRT iterations performed, whereas every run of DART is characterized by three parameters: N_S , N_D and N_B . As an example, Fig. 5 shows the phantom grain map, its DART reconstruction for $N_S = 3$, $N_D = 3$ and $N_B = 3$, and a SIRT reconstruction for $N_S = 15$. The DART reconstruction had a running time of around 0.5 s on a desktop PC, *versus* 0.7 s for the SIRT reconstruction.

In the first series of experiments (§§3.1–3.3), the algorithm parameters for SIRT and DART are kept fixed at $N_S = 10$ and $(N_S, N_D, N_B) = (3, 3, 3)$, respectively. For these settings, both reconstruction methods have approximately the same running time and yield acceptable reconstruction quality for a wide range of noise levels. Increasing the number of iterations typically improves the reconstruction quality for both methods if the noise level is low, while it makes the reconstruction worse at high noise levels. The effect of increasing the number of iterations at a moderate noise level is investigated for the 400×400 grain map in §3.4.

In all experiments, the total number K of pixel errors between the phantom map and the reconstructed map was determined to obtain a quantitative measure of the reconstruction quality.

3.1. Variation with noise for many reflections per grain

The simulated reflection images for the phantom in Fig. 5(a) were subject to varying amounts of noise. As an approximation of Poisson noise, the noise was implemented by adding to each pixel intensity I_0 (*i.e.* the photon count) a random number following a Gaussian distribution with mean 0 and standard deviation $\sigma = cI_0$, where the parameter c determines the signal-to-noise ratio. Applying Gaussian noise in this




Figure 6
Phantom image representing a complex 400×400 map, obtained by electron microscopy on an annealed, pure copper sample.

manner may result in some negative detector values, which were set to 0. In the following, the noise levels will be indicated in per cent, where a noise level of 100% corresponds to the case $c = 1$. For noise levels varying from 0 to 100%, grain-map reconstructions have been computed using both SIRT with $N_S = 10$ and DART with $(N_S, N_D, N_B) = (3, 3, 3)$. For each test case, 100 independent experiments were run. The results are shown in Fig. 7, both before and after applying the filter postprocessing step.

The results demonstrate that both SIRT and DART are capable of reconstructing the simulated two-dimensional grain map fairly accurately. For $c < 0.1$ – which is typical of experimental data with a long exposure time – DART clearly outperforms SIRT (the number of wrongly assigned pixels at $c = 0$ is 14 for DART and 65 for SIRT). This is also the case for $c > 0.5$, but not by a large margin. At high noise levels, the filter operation is observed to be highly effective in decreasing the number of errors in both the SIRT and the DART reconstructions. To further illustrate the effect of the filtering procedure, maps reconstructed from noisy data, before and after filtering, are shown in Fig. 8.

3.2. Variation with number of reflections

As noted in §1, only a limited subset of the reflection spots used for indexing a grain may be suitable for reconstruction purposes. The robustness of an algorithm with respect to the number of reflections is therefore an important aspect. To investigate this, reconstructions of the phantom in Fig. 5(a), based on noiseless simulations, were computed using a varying number of reflections per grain. The number of reflections per grain was varied from 12 down to one. For every test case, 100 independent experiments were run, where the reflections for each grain were chosen randomly from the set of all available reflections. Fig. 9 shows the averaged reconstruction results for

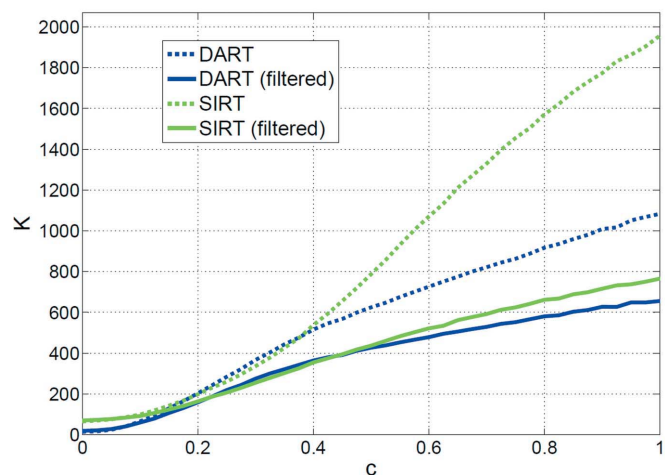


Figure 7
Number of pixel errors K as a function of the noise level for SIRT ($N_S = 10$) and DART ($N_S = 3$, $N_D = 3$, $N_B = 3$), before and after applying the filter postprocessing step. The parameter c determines the noise level by $\sigma = cI_0$. On average, 12 spots per grain were used for reconstruction.

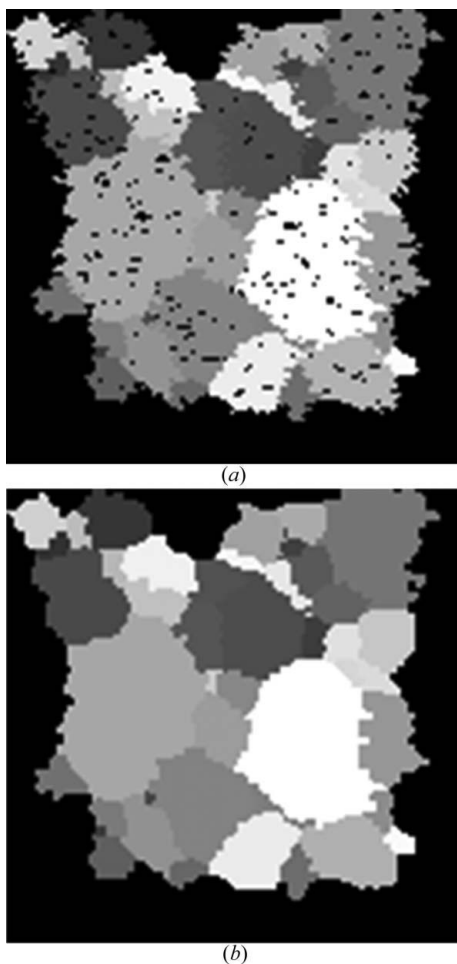


Figure 8
Result of the filtering process applied to a reconstruction using DART ($N_S = 3$, $N_D = 3$, $N_B = 3$) from noisy data ($c = 0.8$) (§3.1). Each grey level represents one grain with one crystallographic orientation; black represents background or an undetermined pixel. (a) Before filtering; (b) after filtering.

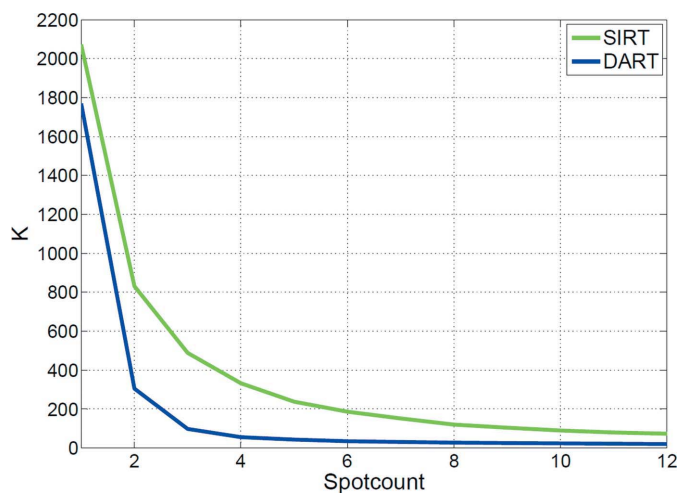


Figure 9
Number of pixel errors K as a function of the number of projections per grain for SIRT ($N_S = 10$) and DART ($N_S = 3$, $N_D = 3$, $N_B = 3$), after applying the filter postprocessing step.

both SIRT and DART, after applying the filter postprocessing step.

The results clearly demonstrate a principal advantage of using DART. Accepting an error level of, say, $K = 100$, we see that ten reflections are needed for SIRT, but only three for DART.

3.3. Variation with noise for four reflections per grain

From the results in the previous subsection, it is clear that DART yields more accurate reconstructions than SIRT if only a few reflections are available for each grain. However, these experiments were performed using simulated projection data without noise. To determine the effect of noise on the reconstruction from a small number of reflections, the experiments from §3.1 were repeated using a fixed number of four reflections per grain. For every test case, 100 independent experiments were run, where four reflections for each grain were chosen randomly from the set of all available reflections. The results are shown in Fig. 10 and confirm the conclusions reached above.

3.4. Reconstruction of a large complex two-dimensional orientation map

In the second series of simulation experiments, we investigated the reconstruction quality for both SIRT and DART for the more complex 400×400 phantom in Fig. 6. In each experiment, three reflections for each grain were selected randomly from the set of all available reflections, and a noise level of $c = 0.1$ (see §3.1) was applied to all reflections. This constitutes a conservative estimate of the noise usually seen in experimental data. The noise may be reduced by longer exposure times, but the dynamical range of the detector and stability impose an upper limit on the exposure times possible. In the same manner, three projections per grain is a conservative estimate, but one that enables reconstruction of large maps where spot overlap is considerable.

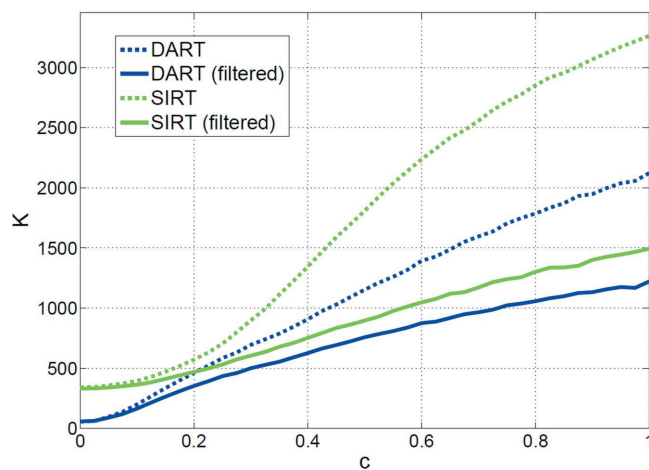


Figure 10
Number of pixel errors K as a function of the noise level for SIRT ($N_S = 10$) and DART ($N_S = 3$, $N_D = 3$, $N_B = 3$), using only four reflection spots per grain. The parameter c determines the noise level by $\sigma = c I_0$.

In the experiments of §§3.1–3.3, the number of iterations for both SIRT and DART was kept fixed. Here, we investigate how the reconstruction quality varies with the number of iterations. The parameter settings for each DART experiment correspond to a triple of the form $(N_S, N_D, N_B) = (M, 20, M)$, i.e. M initial SIRT iterations followed by 20 DART iterations, each performing M SIRT iterations on the boundary pixels. For comparison in a single graph, we compare the result of DART with SIRT, using $N_S = M + 20M = 21M$ iterations. Note that the actual workload performed in DART is much smaller than for the corresponding SIRT experiment, as only the boundary pixels are updated within the DART iterations. For each value of M , ten independent experiments have been performed.

Fig. 11 shows the average reconstruction error, as a function of the number N_S of SIRT iterations, where a value of $M = N_S/21$ has been used for the DART experiments. The results before and after filtering were highly similar; therefore we report the results for the unfiltered reconstructions.

The results show that, for this rather complex test set, SIRT yields superior reconstruction accuracy if only a few iterations are applied, but for a large number of iterations, DART attains a low number of pixel errors that is never reached by SIRT, regardless of the number of SIRT iterations. For the case $(M, N_S) = (10, 210)$, computing the SIRT reconstruction requires approximately 20 s of computing time on a standard desktop PC, whereas DART takes about 10 s.

4. Discussion and outlook

Several algorithms have been proposed for solving the reconstruction task of this paper. Compared to the original ART approach from Poulsen & Fu (2003), it is clear that the reconstructions computed by both SIRT and DART result in far less ambiguity in the reconstructed grain map. For moderate noise levels, and using more than ten reflections per

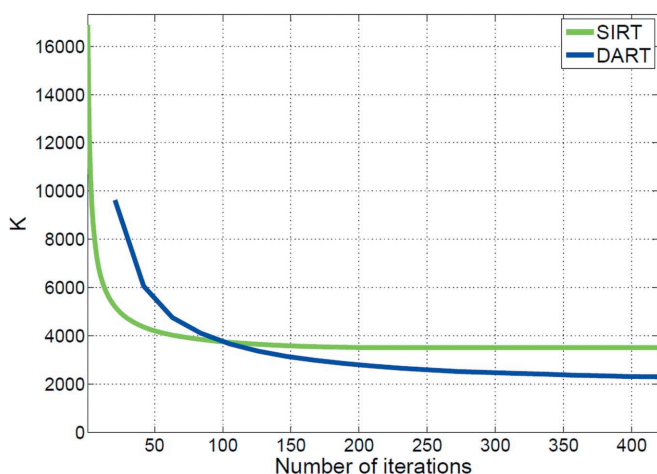


Figure 11 Number of pixel errors K for the phantom in Fig. 6 as a function of the number N_S of SIRT iterations. The DART reconstructions have been computed using $(N_S, N_D, N_B) = (M, 20, M)$, where $M = N_S/21$. In all experiments, three reflections per grain were used and a noise level $c = 0.1$ was applied to the projection data.

grain, very few errors can be seen in the reconstruction, even before the filtering operation is applied. In contrast, in the ART reconstruction presented by Poulsen & Fu (2003), the reconstructed grain map still contained large areas that could not be matched with one of the grains.

The Monte Carlo-based approach proposed by Alpers *et al.* (2006) resulted in a substantial improvement in reconstruction quality compared to the ART algorithm. The Metropolis algorithm, which uses prior knowledge of the frequency distribution of certain local image features in the reconstruction, can be seen as an advanced version of the filtering step applied in the present paper. The main problem with the Metropolis algorithm is its running time. As reported by Alpers *et al.* (2006), reconstruction of a grain map of size 128×128 , as used in our simulation, can take up to 1 min of computation time. As real 3DXRD maps may have a size of up to 500×500 , the running time is prohibitive for many *in situ* studies. In comparison, both SIRT and DART are extremely fast, requiring less than 1 s for reconstructing a 128×128 grain map and about 10–20 s for reconstructing the complex 400×400 map. Combined with the benefits of DART with respect to the small number of required reflections, this allows for a combination of fast image acquisition and fast reconstruction. We emphasize that the applicability of the approach presented here is confined to the reconstruction of undeformed grains, which allow for the representation of a grain by a binary image. This restriction also holds for the approaches of Poulsen & Fu (2003) and Alpers *et al.* (2006). For the more complex case of moderately deformed grain maps, we refer to Rodek *et al.* (2007).

Our simple model for the postprocessing filter does not capture some aspects of the nature of real microstructures. However, the simulations above demonstrate the validity of the heuristic approach, as the simple filter results in a significant improvement in accuracy over the unfiltered reconstructed grain maps. In this context, we note that one important piece of prior information, which is used in the stochastic approach of Alpers *et al.* (2006) but not by DART, is the knowledge that grains may not overlap, and that there cannot be empty space between the grains. This knowledge could be incorporated by simultaneously reconstructing several grains, while applying the constraints on the connections between grains (see Poulsen & Fu, 2003). Although such a procedure can potentially lead to improved reconstruction quality, the simultaneous reconstruction will likely lead to longer reconstruction times, as a result of the increased complexity of the reconstruction problem that now must incorporate constraints on the overlap and connectivity of grains. For cases where fast reconstruction is important, we feel that DART, when combined with the postprocessing filter, provides a good trade-off between reconstruction quality and reconstruction time.

Obviously, convergence is a relevant concern for any reconstruction algorithm. Although the specific variant of DART presented here is based on the SIRT algorithms, for which convergence can be proved formally, the complete DART algorithm does not have guaranteed convergence.

Even after a large number of iterations, the reconstruction may still change between subsequent iterations. To obtain a variant of DART with guaranteed convergence properties, it can be mixed with a convergent algorithm such as SIRT, where the new image computed in each iteration consists of a linear combination of an update term computed by SIRT and an update term computed by DART. If the contributions of the DART terms can be bounded and decrease exponentially as a function of iteration number, the resulting combined algorithm will formally converge (Butnariu *et al.*, 2007). However, the resulting reconstruction is not guaranteed to be discrete in the grey-level domain.

We remark that the convergence of DART, as well as the accuracy of the reconstructed image with respect to the ground truth, depends on many factors. Firstly, there is a dependency with respect to the data: the number of required projection images will depend on the complexity of the shapes of the grains and the size of the boundary. The characteristics of experimental noise will differ, at least to some extent, from the simulated noise used in our experiments. Secondly, the choice of algorithm parameters, such as the number of DART iterations, will influence the result. Finding the optimal parameters for a given data set is a difficult task, just as with classical continuous iterative methods.

The main application of the DART methodology for 3DXRD studies may be for a different mode, where the beam illuminates the full sample instead of only a layer. In this mode, the three-dimensional grain map is then reconstructed without the need for stacking. The use of iterative methods – ART – for this mode has been demonstrated by Markussen *et al.* (2004). This development is of major interest for *in situ* studies as the total data acquisition time is substantially reduced. However, it was also shown that the number of diffraction spots per grain needed for a good reconstruction in three dimensions is twice the number needed for two dimensions. This requirement is likely to be an issue. Fortunately, the generalization of both SIRT and DART to three dimensions is straightforward. This is also believed to be the case for applying the method to DCT.

The case for using DART for multiphase mapping by means of X-ray absorption or phase tomography is seen as equally strong. The difference is that, in this case, individual components cannot be reconstructed independently and then stitched together. Nevertheless, we anticipate – based also on first results from electron tomography (Bals *et al.*, 2007, 2009) – that DART can be successfully applied for such data as well.

5. Conclusion

Materials associated with a few fixed contrast levels are common in X-ray imaging. The DART algorithm makes it possible to use such *a priori* knowledge directly in the reconstruction. As illustrated by the case of grain maps based on 3DXRD microscopy, the algorithm enables quality reconstructions based on very few projections and is associated with very short run times. While DART is not substantially faster

than SIRT (or indeed similar tomography algorithms) it combines speed with the high reconstruction quality of discrete tomography. This combination of speed and quality is seen as the enabling factor for on-line 3DXRD mapping of undeformed samples using DART.

The authors would like to thank J. Bowen and L. Fengxiang for providing the EBSD micrographs as well as A. Alpers, H. O. Sørensen and S. Schmidt for helpful discussions. This work was partially supported by the Flemish Fund for Scientific Research (FWO), by the Danish National Research Foundation and by the EU 6th framework NEST/ADVENTURE project TotalCryst.

References

- Alpers, A., Poulsen, H. F., Knudsen, E. & Herman, G. T. (2006). *J. Appl. Cryst.* **39**, 582–588.
- Bals, S., Batenburg, K. J., Liang, D., Lebedev, O., Van Tendeloo, G., Aerts, A., Martens, J. A. & Kirschhock, C. E. A. (2009). *J. Am. Chem. Soc.* **131**, 4769–4773.
- Bals, S., Batenburg, K. J., Verbeeck, J., Sijbers, J. & Van Tendeloo, G. (2007). *Nano Lett.* **7**, 3669–3674.
- Batenburg, K. J. & Sijbers, J. (2007). Proceedings of the IEEE Conference on Image Processing (ICIP), San Antonio, USA, Vol. IV, pp. 133–136.
- Butnariu, D., Davidi, R., Herman, G. T. & Kazantsev, I. G. (2007). *IEEE J. Sel. Top. Signal Process.* **1**, 540–547.
- Gilbert, P. (1972). *J. Theor. Biol.* **36**, 105–117.
- Gordon, R., Bender, R. & Herman, G. T. (1970). *J. Theor. Biol.* **29**, 471–482.
- Gregor, J. & Benson, T. (2008). *IEEE Trans. Med. Imaging*, **27**, 918–924.
- Herman, G. T. & Kuba, A. (1999). Editors. *Discrete Tomography: Foundations, Algorithms and Applications*. Boston: Birkhäuser.
- Kak, A. C. & Slaney, M. (2001). *Principles of Computerized Tomographic Imaging*. Philadelphia: Society for Industrial and Applied Mathematics.
- King, A., Johnson, G., Engelberg, D., Ludwig, W. & Marrow, J. (2008). *Science*, **321**, 382–385.
- Lauridsen, E. M., Schmidt, S., Suter, R. M. & Poulsen, H. F. (2001). *J. Appl. Cryst.* **34**, 744–750.
- Ludwig, W., King, A., Reischig, P., Herbig, M., Lauridsen, E. M., Schmidt, S., Proudhon, H., Forest, S., Cloetens, P., du Roscoat, S., Rolland, S., Buffiere, J. Y., Marrow, T. J. & Poulsen, H. F. (2009). *Mater. Sci. Eng. A*, **524**, 69–76.
- Markussen, T., Fu, X., Margulies, L., Lauridsen, E. M., Nielsen, S. F., Schmidt, S. & Poulsen, H. F. (2004). *J. Appl. Cryst.* **37**, 96–102.
- Poulsen, H. F. (2004). *Three-Dimensional X-ray Diffraction Microscopy: Mapping Polycrystals and their Dynamics*. Berlin: Springer.
- Poulsen, H. F. & Fu, X. (2003). *J. Appl. Cryst.* **36**, 1062–1068.
- Rodek, L., Poulsen, H. F., Knudsen, E. & Herman, G. T. (2007). *J. Appl. Cryst.* **40**, 313–321.
- Schmidt, S., Nielsen, S. F., Gundlach, C., Margulies, L., Huang, X. & Juul Jensen, D. (2004). *Science*, **305**, 229–232.
- Schmidt, S., Olsen, U. L., Poulsen, H. F., Sørensen, H. O., Lauridsen, E. M., Margulies, L., Maurice, C. & Juul Jensen, D. (2008). *Scr. Mater.* **59**, 491–494.
- Sezan, M. I. & Stark, H. (1982). *IEEE Trans. Med. Imaging*, **MI-1**, 95–101.
- Suter, R. M., Hennesy, D., Xiao, C. & Lienert, U. (2006). *Rev. Sci. Instrum.* **77**, 123905.
- Warren, B. E. (1990). *X-ray Diffraction*. Mineola: Dover Publications.