

A discrete tomography algorithm for improving the quality of three-dimensional X-ray diffraction grain maps

A. Alpers,^a H. F. Poulsen,^{b*} E. Knudsen^b and G. T. Herman^a

^aDepartment of Computer Science, The Graduate Center, City University of New York, NY 10016, USA, and ^bCenter for Fundamental Research: 'Metal Structures in Four Dimensions', Risø National Laboratory, Dk-4000 Roskilde, Denmark. Correspondence e-mail: henning.friis.poulsen@risoe.dk

A discrete tomography algorithm is presented for the reconstruction of grain maps based on X-ray diffraction data. This is the first algorithm for this task, inherently exploiting the discrete structure of grain maps. Gibbs potentials serve to characterize the statistics of the local morphology of the grain boundaries. A Monte Carlo based algorithm is applied as a restoration method for improving the quality of grain maps produced by a classical (non-discrete) tomography algorithm (ART). The quality of the restored maps is demonstrated and quantified by simulation studies. The robustness of the algorithm with respect to the choice of Gibbs potentials is investigated.

© 2006 International Union of Crystallography
Printed in Great Britain – all rights reserved

1. Introduction

Three-dimensional X-ray diffraction (3DXRD) microscopy is an emerging methodology for the structural characterization of polycrystals (Poulsen, 2004). Based on diffraction by monochromatic hard X-rays, area detectors, and reconstruction principles, three-dimensional maps can be generated, capturing the morphology of the individual grains within millimetre-thick specimens (Poulsen *et al.*, 2001; Poulsen & Fu, 2003; Markussen *et al.*, 2004). Furthermore, in favourable cases, three-dimensional movies can be produced (Schmidt *et al.*, 2004). Hence, for the first time the local dynamics of processes such as recrystallization and grain growth can be observed directly under realistic bulk conditions. This is seen as a pre-requisite to establishing first-principles models for these industrially very important processes.

More specifically, a tomographic-type reconstruction procedure, 2D-ART (two-dimensional algebraic reconstruction technique), has been demonstrated for undeformed specimens (Poulsen & Fu, 2003). In this procedure, layers (two-dimensional sections) within the sample are illuminated consecutively. To generate three-dimensional maps, the results for the layers are stacked. Furthermore, as the diffraction spots associated with different grains typically do not overlap on the detector, the spots can be sorted with respect to the grain of origin by a polycrystal indexing routine (Lauridsen *et al.*, 2001). As a result, the reconstruction can be performed for each grain separately, based on the intensity distributions within the associated spots. For a given discretized layer in a specimen, one can assign a 0/1 value to each cell/pixel, indicating whether this pixel is occupied by the grain in question. All the occupied pixels contribute linearly to the signal intensity measured on the detector. Thus, we suppose that the collection of pixels is associated with a 0/1 valued vector \mathbf{x} ,

which characterizes the morphology of the grain. In a forward-projection model, we can then compute the intensities within the individual pixels on the detector (collected in a vector \mathbf{b}), by a linear function, \mathbf{Ax} , which follows from the experimental geometry. In this way we can represent the process by a linear system of equations: $\mathbf{Ax} = \mathbf{b}$.

The problem at hand is the inverse problem: given \mathbf{A} and \mathbf{b} , compute \mathbf{x} . The algorithm previously used is a variant of the algebraic reconstruction technique (ART) used in biomedical imaging (Gordon *et al.*, 1970), which is an iterative method for solving systems of linear equations, especially tailored to tomographical data acquisition settings. Notably ART is a classical/non-discrete algorithm, which generates real-valued results. However, the condition $0 \leq x_j \leq 1$ can be imposed for all elements j in the vector \mathbf{x} . The values x_j are interpreted as grey levels, giving an indication whether or not a pixel belongs to a grain. In the 2D-ART approach, the solution is binarized (pixels belong to the grain or not) by setting a threshold at 0.5.

Presently, the spatial resolution of 2D-ART maps is mainly limited by the number of useful diffraction spots per grain. Due to variation in structure and Lorentz factors with Bragg angle, geometry of available detectors and the wish to minimize acquisition times for grain maps (in order to be able to perform *in situ* annealing studies), the number of spots per grain may be as few as five (Poulsen & Fu, 2003). As a result, the system of equations, $\mathbf{Ax} = \mathbf{b}$, is very underdetermined. Due to this, along with experimental noise and the fact that grains are reconstructed independently, the association of some pixels at the boundary between grains will be ambiguous. In particular, 2D-ART may associate some pixels with two or more grains and some with no grains at all.

In this paper, we aim at improving the quality of grain maps generated by 2D-ART and making them space-filling by applying a subsequent 'restoration algorithm' to the ambigu-

uous parts of the map. The basis of the restoration is twofold. First, a stochastic approach is used that simultaneously operates on all grains. Second, grain morphologies are not random and, as such, some local boundary configurations are highly unlikely. Making use of this, the restoration algorithm is an application (and, in part, extension) of the ‘discrete tomography using Gibbs priors’ formalism presented by Herman & Kuba (1999) and Carvalho *et al.* (1999). To our knowledge, it is the first time that this class of algorithms, which inherently takes the discrete nature of grain maps into account, has been applied in materials science.

In the next section we outline the suggested approach. In §3, the quality of resulting grain maps is characterized by means of simulations, and restorations with and without the use of Gibbs priors are compared with pure 2D-ART reconstructions. Furthermore, the robustness of the algorithm with respect to the choice of values for the Gibbs priors and system parameters is investigated. Finally, the applicability of the method and its extensions to full-beam 3DXRD using 3D-ART (Markussen *et al.*, 2004) is discussed. A formal definition of the Gibbs priors used in this work has been provided in a preliminary paper with focus on the mathematical aspects (Alpers *et al.*, 2005).

2. Approach

The method is based on statistics over the local morphology. It is designed to be simple and computationally efficient, but nevertheless efficacious. A (discretized) grain map can be interpreted as a multi-colour image, where different colours correspond to different grains. For each position in a grain map, we define its local configuration as a 3×3 array of black and white pixels as follows: the central pixel is always white and any other pixel is white if, and only if, it belongs to the same grain as the central pixel. These configurations are partitioned into seven classes, G_0, G_1, \dots, G_6 , each containing configurations of similar morphology, such as ‘grain interior’, ‘grain edge’, *etc.* For $1 \leq i \leq 6$, the class G_i consists of the configuration illustrated in Fig. 1 and all the configurations that can be obtained from it by a sequence of 90° rotations around the centre and mirror reflections about the central vertical line. Configurations not in any of G_1, \dots, G_6 are put into G_0 . [The physical assumptions inherent in this model can be roughly stated as: (i) at the scale of 3×3 configurations the grain’s morphology does not depend on the orientation of its neighbours, (ii) at the same scale, triple junctions play a minor role (that is, the grain boundary characteristics are statistically independent of whether there are two or three neighbours), and (iii) the grain microstructure is isotropic. These assumptions are at least partially justified by the success, as reported below, of the method based on them.]

Next, we approximate the distribution of grain maps by a Gibbs distribution, *i.e.* we assume that the distribution π of all grain maps f that have to be restored can be approximated by

$$\pi(f) = \frac{1}{Z} \exp\left(\beta \sum_{i=1}^6 N_i U_i\right), \quad (1)$$

where U_i is the Gibbs potential associated with the class G_i , N_i is the number of configurations of class G_i in this map, β is the inverse temperature and Z is a normalization constant. We used an approach based on counting 3×3 configurations in a typical microstructure (the training set) for determining the Gibbs potentials. Since the distribution π does not change if we add a constant to all Gibbs potentials, we put $U_0 = 0$.

The restoration algorithm seeks to maximize the expression

$$\gamma(f) = \frac{1}{Z} \exp\left[\beta\left(\sum_{i=1}^6 N_i U_i - \alpha \|\mathbf{P}_f - \mathbf{P}\|_1\right)\right], \quad (2)$$

where \mathbf{P} is a vector representing the measured pixel intensities on the detector, while \mathbf{P}_f represents the simulated (projected) pixel intensities given the map f . The scale parameter α determines how much weight is put on correspondence with data as compared with obeying the Gibbs distribution. Both α and β are system parameters to be optimized by simulations. The maximization is performed by the Metropolis algorithm (Metropolis *et al.*, 1953; Brémaud, 1999), which is directly applicable to expression (2). Following Vardi *et al.* (2001), a ‘look-up’ table is used to speed up the maximization.

Evidently, different samples (and sample regions) will be associated with different Gibbs potentials. It would be cumbersome, and in some cases impossible, to establish these potentials prior to every 3DXRD experiment. However, grain microstructures are broadly speaking similar (*e.g.* grain growth is often associated with self-similar patterns). Hence, we predict that the algorithm is sufficiently robust with respect to choice of Gibbs potentials, and so one set of parameters can be of use for a large set of samples. This prediction will be tested below.

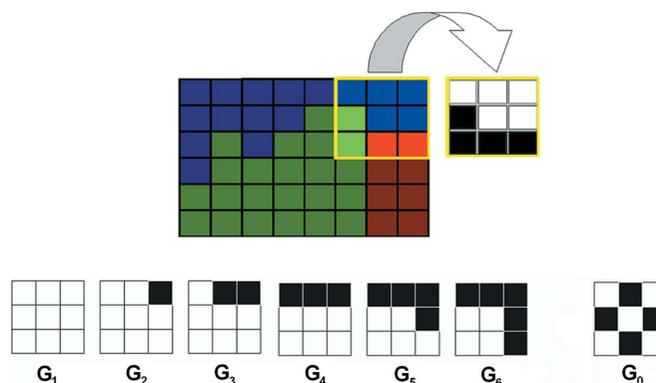


Figure 1 Top: example of producing a local configuration. A 3×3 pixel region (that overlaps three different grains) produces the local configuration for its central pixel by assigning white to those pixels that belong to the same grain as the central one, and black to the others. Bottom: representative elements of the six classes G_1, \dots, G_6 of 3×3 local configurations used in the model. (The other elements of a class can be obtained from the representative element by a sequence of rotations and mirroring.) The class G_0 contains all other (‘random’) configurations; we also show one element of this class.

3. Simulations

Simulations were performed to optimize the parameters of the algorithm and to characterize quantitatively the quality of the reconstructions as functions of the experimental noise and the number of spots per grain.

The simulations were based on a 128×128 pixel orientation map of an annealed Al sample, shown in Fig. 2 (left). This map was acquired by electron microscopy using the EBSD (electron back-scatter diffraction) technique (Alam *et al.*, 1954; Lassen *et al.*, 1992). The orientation map was transformed into a grain map in which pixels are associated with grain number rather than orientation. By sweeping the 3×3 grid over the image, the numbers K_i of configurations belonging to class G_i , $i = 0, \dots, 6$ were determined. The appropriate Gibbs potentials can then be estimated as in the work of Carvalho *et al.* (1999):

$$U_i = \kappa \left[\ln \left(\frac{K_i}{|G_i|} + 1 \right) - \ln \left(\frac{K_0}{|G_0|} + 1 \right) \right], \quad (3)$$

where $|G_i|$ denotes the number of configurations in class G_i , *i.e.* $(|G_0|, |G_1|, |G_2|, |G_3|, |G_4|, |G_5|, |G_6|) = (227, 1, 4, 8, 4, 8, 4)$.

Applying this equation with κ such that the largest number is 1.4 (which is an arbitrary choice), the following Gibbs potentials were derived: $U^1 = (U_1, U_2, U_3, U_4, U_5, U_6) = (1.4, 0.71, 0.61, 0.79, 0.5, 0.61)$. That the choice of κ can indeed be arbitrarily adjusted follows from the observation that the same values are obtained for γ in equation (2) if we change the κ and adjust the values of α and β accordingly.

Next, the 3DXRD diffraction patterns associated with the orientation map were simulated. We generated 91 images corresponding to equally spaced rotation angles over a rotation range of 90° . The $\{111\}$, $\{200\}$, $\{222\}$ and $\{311\}$ reflection families were included. The simulated patterns were generated from the EBSD grain map. The EBSD data also provide the orientation of each grain. For each grain and each reflection within a given reflection family, the diffraction vector was computed, and it was determined at which rotation step the associated diffraction spot would appear. Successively, each pixel within the grain was projected along this diffraction

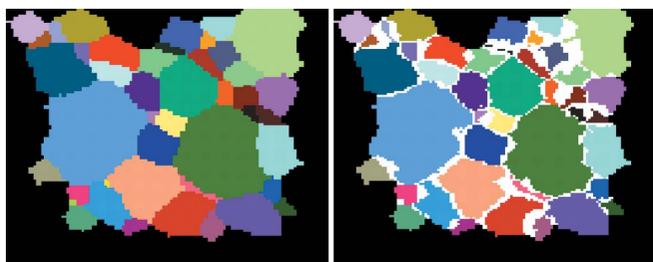


Figure 2 Left: the grain map (experimentally determined by electron microscopy) upon which the simulations are based. Colours represent orientations according to a scheme based on Euler angles. Right: a simulated result of grain-by-grain reconstruction by 2D-ART from 3DXRD data of the grain map shown on the left. Due to the limited number of diffraction spots included in the 3DXRD analysis, there are areas of ambiguity (shown in white). The aim of the restoration is to resolve the ambiguities in the white areas.

vector onto the detector and intensities were summed accordingly. The detector geometry mimicked the experimental setup at the 3DXRD microscope at the European Synchrotron Radiation Facility (ESRF) at the time, implying that on average only eight diffraction spots per grain were acquired. The details of the setup are as follows. The X-ray energy was 50 keV. The sample area to be reconstructed comprised 128×128 pixels and had dimensions of $294.4 \times 294.4 \mu\text{m}$. The two-dimensional detector had 1024×1536 pixels; positioned at a distance of 4.186 mm to the centre of the sample, its field-of-view was 2.4×3.5 mm. The typical pixel intensity within a diffraction spot was 4.3 counts.

To evaluate the quality of the reconstruction that was computed by our algorithm when applied to grain-by-grain reconstructions by 2D-ART (in which the resulting composed grain map has overlapping and missing pixels, which form areas of ambiguity), we performed 2D-ART reconstructions of the individual grains at a fixed threshold. The ambiguous areas were defined as the set of pixels which either became associated with zero grains or with several grains. These areas are shown as white areas in Fig. 2 (right). The total number of pixels in these ambiguity areas was 1490. Note that 682 of the non-white pixels were assigned to the wrong grains by 2D-ART (mainly due to calibration problems), and since the restoring algorithm processes only white pixels, these incorrect assignments cannot be undone.

Based on the set of Gibbs potentials U^1 , the diffraction images and the ambiguity areas defined above, a series of restorations were made with varying system parameters α , β , the number of Monte Carlo cycles (MCC, where one MCC is defined to consist of as many steps in the Metropolis algorithm as the number of initially ambiguous pixels), as well as the noise. Noise was implemented by adding a value of 1 to randomly chosen detector pixels. In the following, noise levels will be indicated in percent as the ratio of the total number of added 1's to all the detector pixels. The quality of the restorations is determined by the 'number of errors', which is the number of pixels in the white areas that have been assigned to a grain in the restored map that is different from the grain in the original map. To improve the statistics, each simulation was repeated ten times with different initial seeds to the Metropolis algorithm.

The restorations were found to converge rapidly. The variations in number of errors as function of (α, β) at $\text{MCC} = 1000$ are shown in Fig. 3. For noise levels clearly below 100%, a range that is expected to include the level of noise in real 3DXRD experiments, a broad optimum is roughly found around the values $(\alpha_0, \beta_0) = (1, 1)$. The exact values showing the quality of reconstruction are listed in Table 1. The dependency of α_0 and β_0 on the noise level is discussed in more detail by Alpers *et al.* (2005), and is shown to be quite robust, as can also be seen from Fig. 3. For the rest of this paper, the system parameters (α, β) are fixed at $(1, 1)$.

The variation with noise of the restoration based on equation (2) is compared in Fig. 4 with results for projections only. 'Projections only' means no use of Gibbs priors, *i.e.* based on equation (2), $\gamma(f) = (1/Z) \exp(\beta\alpha \| \mathbf{P}_f - \mathbf{P} \|_1)$, with $\alpha = 1$ and

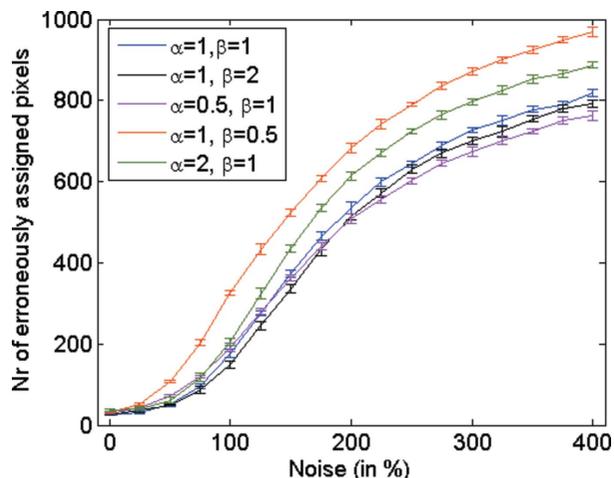


Figure 3 Optimization of system variables α and β . For various (α, β) , the variation in restoration quality, defined as the total number of pixels within the ambiguous areas of the reconstructed map (the white areas in Fig. 2, right) which are assigned to the wrong grain, is shown as a function of the ‘additive noise’ in the simulated detector images. The error bars mark the standard deviation of the ten Monte Carlo simulations.

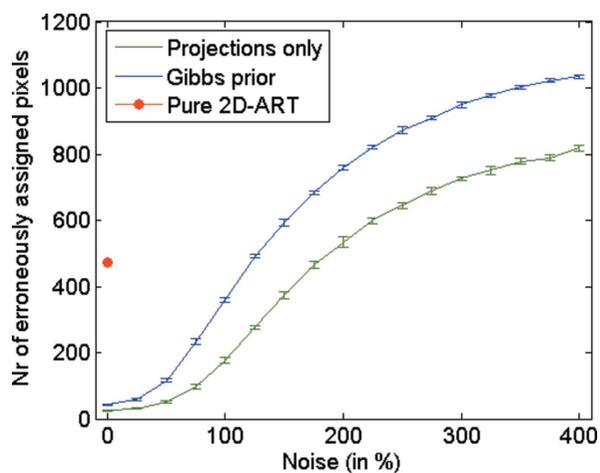


Figure 4 Reconstruction quality as a function of the amount of additive noise in the simulated detector images. Two map restoration methods are compared: restoration by projection only and restoration by projections plus the Gibbs prior based on U^1 (see text). In addition, the result for reconstruction based purely on 2D-ART for a noise level of 0% is indicated.

$\beta = 1$. Also shown is the result of a pure 2D-ART reconstruction based on the simulated diffraction images. Ambiguous pixels were in this case allocated to the grain yielding the highest value in the single-grain reconstructions. Pixels with value 0 in all individual reconstructions can, in this case, leave holes inside the map. Apart from differences that may be seen directly by inspecting sample reconstructions (Fig. 5), a more detailed study shows that there are significant differences in performance between the three methods. These results are interpreted as follows.

The ‘projections only’ restoration and pure 2D-ART both optimize grain maps based only on the correspondence between the projection of the map and the diffraction data. Previous work has demonstrated that 2D-ART can provide

Table 1 Values of reconstruction quality for low noise levels as depicted in Figs. 3, 4 and 8.

Wherever α and β are not specified, they are both equal to 1, and wherever the choice of Gibbs potential is not specified, it is U^1 .

Method	Noise				
	0%	25%	50%	75%	100%
$\alpha = 1, \beta = 1$	24.8	31.7	50.3	96.7	176.8
$\alpha = 1, \beta = 2$	26.5	35.4	49.9	86.5	148.5
$\alpha = 0.5, \beta = 1$	27.9	41.4	72.6	120.7	190.9
$\alpha = 1, \beta = 0.5$	30.2	51.8	107.8	203.4	324.8
$\alpha = 2, \beta = 1$	35.8	40.8	61.4	117.3	205.4
Projections only	43.5	58.5	116.8	232.6	359.6
U^2	25.8	34.1	52.7	96.3	174.5
U^3	25.9	32.9	51.7	102.5	175.9
U^4	27.0	33.3	55.4	97.7	182.5
U^5	25.7	32.5	51.1	93.2	171.9
U^6	32.7	42.7	78.3	159.6	274.9

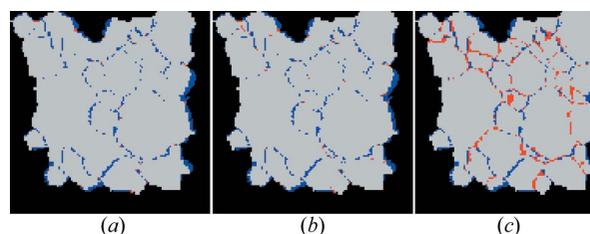


Figure 5 Error in the reconstruction of grain maps by three different methods: (a) restoration by projections and Gibbs priors; (b) restoration by projections only; (c) reconstruction based purely on 2D-ART where ambiguous pixels have been assigned to the grain yielding the highest rational number in the reconstruction. Erroneously assigned pixels in the defined ambiguous areas (the white areas in Fig. 2, right) are indicated in red; incorrect assignments in other parts, due to errors in the initial map, are indicated in blue. Grey pixels indicate correct assignments. In all cases, on average, five useful spots per grain were used, and the noise level was 0%. The number of erroneously assigned pixels in the ambiguous areas is (a) 27, (b) 49 and (c) 471 pixels.

high-quality reconstructions of one single grain (Poulsen & Fu, 2003). Hence, it appears that the difference in performance between ‘projections only’ and 2D-ART is (at least partly) caused by the restoration algorithm allocating pixels based on an optimization with respect to all neighbouring grains simultaneously, while 2D-ART treats grains separately. An alternative route to the same effect has been presented by Poulsen & Fu (2003). However, this is the first time that the superior effect of a simultaneous reconstruction (as compared with a single grain reconstruction) has been quantified. The difference between restorations based on Gibbs priors plus projections and projections only is evidently due to the use of the additional information in the priors.

Similar results are shown in Fig. 6 for the variation with average number of useful spots per grain (at noise levels of 0% and 100% for the figures on the left and right, respectively). The number of spots was varied by removing spots from the simulated detector images. The removal was carried out arbitrarily, by terminating the spot simulation (favouring no particular kind of spots) after a prescribed amount of spots occurred on the detector. Again the superior quality of the

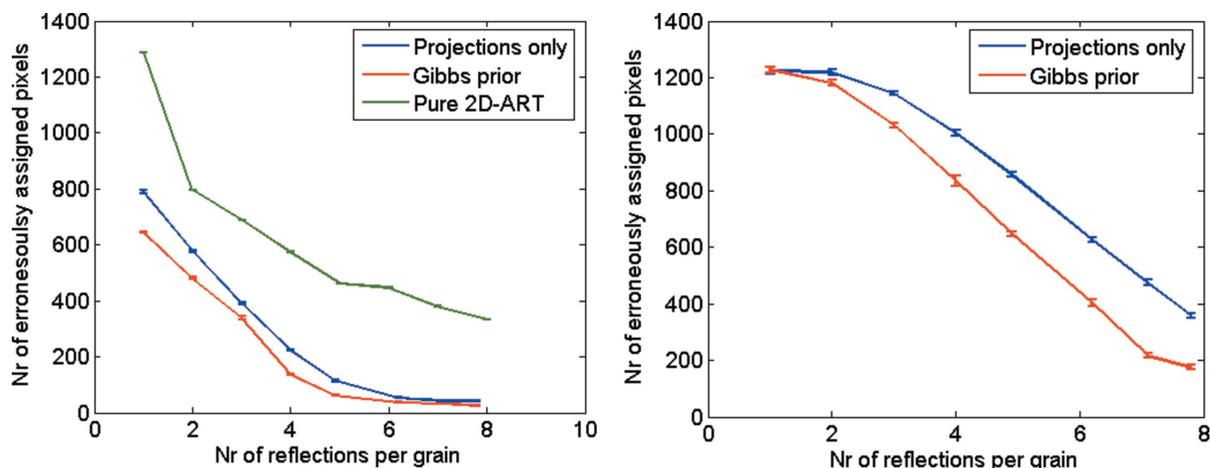


Figure 6 Reconstruction quality as a function of the average number of useful spots per grain. Left: at a noise level of 0%, two map restoration methods are compared: restoration by projection only and restoration by projections plus the optimized Gibbs priors. Right: similar results for a noise level of 100%. In addition, for the noise-free case, results for reconstruction based purely on 2D-ART are also indicated.

mapping based on the restoration algorithm with Gibbs priors over the pure 2D-ART approach is evident.

Finally, the robustness of the method with respect to the choice of Gibbs potentials was tested. In order to do so, five other sets of potentials were defined. Samples of the six associated grain map distributions are shown in Fig. 7. Four of the sets were based on counting configurations in existing EBSD maps of an aluminium alloy with 0.13% Mg.

$U^1 = (1.4, 0.71, 0.61, 0.79, 0.5, 0.61)$: a medium-scale grain structure (deformed at a strain of 10 and annealed for 1 h at 573 K).

$U^2 = (1.4, 0.8, 0.8, 0.94, 0.19, 0.82)$: a coarse-grained structure (deformed at a strain of 10 and annealed for 1 h at 673 K).

$U^3 = (1.4, 0.94, 0.94, 1.09, 0.12, 1.01)$: a fine-grained structure (deformed at a strain of 10 and annealed for 1 h at 498 K).

$U^4 = (1.4, 0.91, 0.91, 1.08, 0.02, 1.0)$: a partially recrystallized sample (deformed at a strain of 3 and annealed for 1 h at 498 K).

The following two sets of potentials (U^5 and U^6) were determined based on visual inspection of Monte Carlo simulations. For various choices of the potentials, the grain maps were allowed to develop from random noise for 1000 MCC using no projection data (*i.e.* $\alpha = 0$) and $\beta = 1.5$.

$U^5 = (1.5, 1.2, 0.84, 1, 1.25, 0.6)$ displayed grain-like features.

$U^6 = (0.5, 0.4, 1.0, 0.8, 0.1, 0.6)$ did not produce grain-like features.

Based on these six sets of potentials, but otherwise with the same settings of system parameters ($\alpha = \beta = 1$), restorations were made to the grain map shown on the right in Fig. 2. Similar to the results shown in Fig. 4, the variation in performance with noise was simulated at an average of eight spots per grain. The results are shown in Fig. 8. (For clarity of display, error-bars have been omitted; their magnitudes are similar to those in Figs. 3 and 4.) Evidently, the quality of reconstruction based on the five sets of potentials related to grain-like features is nearly identical. This is seen as a strong indication that reconstructions of reasonably similar samples with unknown microstructures can be based on a set of Gibbs

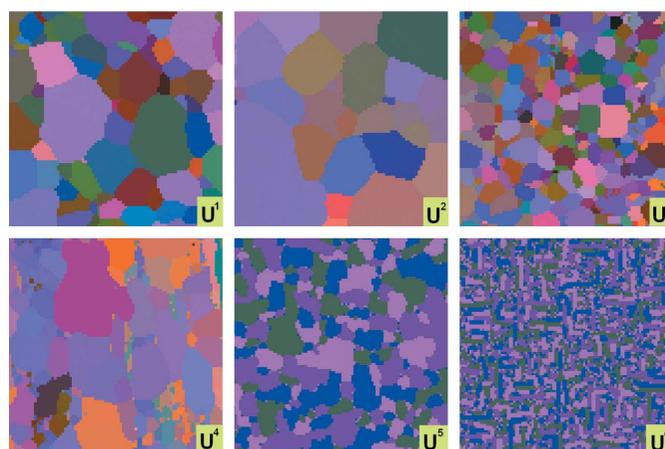


Figure 7 Samples of the grain map distributions defined by the six sets of Gibbs potentials U^1 , U^2 , U^3 , U^4 , U^5 and U^6 (see text).

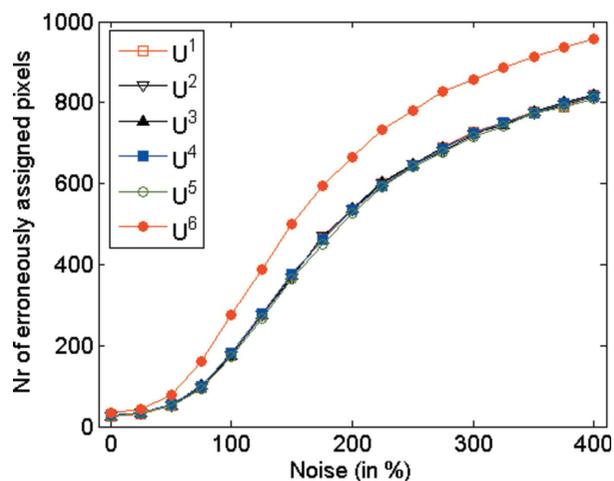


Figure 8 Test of robustness of the method. The quality of restoration is shown as a function of noise for six sets of Gibbs potentials U^1 , U^2 , U^3 , U^4 , U^5 and U^6 . (Samples of the associated grain map distributions are shown in Fig. 7.)

potentials derived from electron microscopy investigations of one representative.

4. Discussion and outlook

The reported performance obtained using Gibbs priors is statistical. Relying on the restoration may be dangerous, *e.g.* for determining the local grain boundary curvature near a specific triple junction. However, the method is applicable to any 3DXRD study in which the focus is on deriving statistical properties.

Likewise, we are not claiming that distributions of physical grain maps can be precisely represented by Gibbs distributions. In fact, typical grain maps are not representative of steady-state processes, but snapshots of grain growth processes, and as such Gibbs distributions may not apply. Hence, the Gibbs distributed configurations are introduced here as a convenient approximate representation of the microstructure. The simulations described above demonstrate the validity of the approach. Alternative representations of the local microstructure may be possible, and for all we know these may give rise to similar or even better restorations.

Similarly, we are aware that our simple model does not capture some aspects of the nature of real microstructures, *e.g.* the fact that the local curvature of grain boundaries is correlated to the orientation relationships between the associated neighbouring grains. Also, most microstructures are not completely isotropic. Again, the simulations above demonstrate the validity of these approximations.

One important factor affecting the reconstruction quality is the quality of the initial grain map, which determines the ambiguous (white) areas. In the simulations presented in §3, 2D-ART was used to generate this map. Notably, one could be more conservative with the results from 2D-ART (dilating the white regions by some pixels) or other algorithms could be used to generate the initial map, such as the program *GRAINSWEEPER* (Schmidt, 2006), to be presented elsewhere.

An extreme approach is to rely only on the outcome of *GRAINDEX* (see Poulsen, 2004); that is, the information that there are M grains in the map and a list of their orientations. To test this approach, simulations were repeated with an initial grain map where all the 10299 non-black pixels of the grain map in Fig. 2 were assigned to be white points. Reconstructions (including Gibbs priors) for 0%, 50% and 100% noise are shown in Fig. 9. The running time was six times longer than for the results presented in §3. The total numbers of incorrectly assigned pixels in the three cases are 120, 586 and 5482. Notably, for 0% noise the total number of errors is smaller than the corresponding result with an initial map provided by 2D-ART, while the number of errors in the ambiguous parts is larger. The results for the 100% noise case are much worse than those obtained using an initial map, due to the fact that the Metropolis algorithm is stuck in a local minimum, which might be resolved by vastly increasing the number of MCC. These results indicate that an optimal trade-off between the accuracy of the initial map and the presented method with

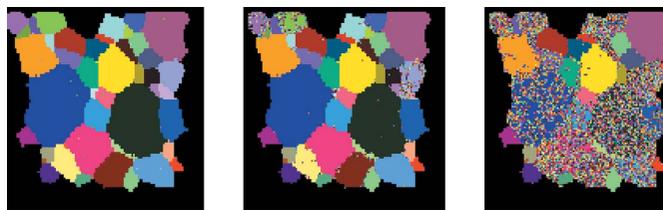


Figure 9
Reconstructed grain maps based on an initial map where the ambiguity area of Fig. 2 (right) is extended to include all non-black points of that map. Optimizing with respect to both projections and Gibbs priors, reconstructions are shown for 0% (left), 50% (middle) and 100% (right) noise. The numbers of erroneously assigned pixels are 120, 586 and 5482 pixels, respectively.

respect to running time, noise and reconstruction quality has to be found. This optimum may depend on the specifics of the grain map and is a topic of current research.

The method presented was implemented as a C program on a state-of-the-art PC. The running time for each restoration in §3 (based on 1000 MCC and 1490 ambiguous pixels) was 10 s. Real 3DXRD maps may have a size of up to 500×500 or even 1000×1000 . It is difficult to predict exactly how the running time scales with the size of the map, since other issues such as spot overlap, size of grains, spatial resolution and quality of initial map may play a role. In the case of constant grain size, fixed spatial resolution and negligible spot overlap, a near linear increase may be expected, as the reconstruction in this case can be divided into a set of nearly independent smaller reconstructions. Fortunately, the algorithm can be run in parallel in several ways, *e.g.* each layer can be processed separately, and even within one layer one can simultaneously restore disjoint ambiguity areas.

Next we comment on some details of the approach. First, the Gibbs potentials were determined from experimental data by the so-called ‘heuristic method’ [equation (3)] (Carvalho *et al.*, 1999). In the work of Liao & Herman (2004) two alternative expressions are presented for deriving the potentials by means of counting. Under idealized conditions, these are shown to lead to more accurate potentials and ultimately to superior restorations. We have not investigated the effect of these alternative expressions here, as we believe that the inaccuracy in determining the potentials for one sample based on counting is less than the variation between typical samples. Second, as for any optimization involving the Metropolis algorithm, the optimization can be further improved by decreasing the temperature with time. Such ‘simulated annealing’ was not attempted here, because of the substantial overhead in running time and the already good quality of the reconstructions achieved without annealing.

In outlook, the main application of the Gibbs-prior-based restoration method may be for a different type of 3DXRD study. Recently, it has been shown that the ART algorithm can also be applied to data obtained by illuminating the full sample rather than one layer at a time: 3D-ART (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 3D-ART is twice the number needed for 2D-ART. This requirement is likely to be an issue, even when applying detectors with a geometry optimized for 3DXRD. Fortunately, the generalization of the Gibbs-prior method to three dimensions is straightforward. The main difference is that the local configurations will change to dimensions of $3 \times 3 \times 3$. The running time of the two-dimensional and three-dimensional versions of the restoration algorithm may even scale linearly; this at least was observed during analogous studies in other fields (Liao & Herman, 2005)

Finally, we suggest that the Gibbs priors could be of relevance for conventional absorption contrast tomography in cases where the number of projections is limited, *e.g.* due to radiation concerns or the wish to increase time resolution. In that respect the methodology introduced above represents a generalization from a two-phase (binary) to multi-phase systems.

5. Conclusion

A novel stochastic method for reconstruction of tomographic images of two-phase systems has been extended for use with diffraction and multi-valued systems. It is applied to the restoration of two-dimensional grain maps obtained by 3DXRD. The simulation results clearly demonstrate (*a*) the superior performance of the stochastic restoration algorithm (with or without priors) in comparison with reconstructions based on patching individual grains generated by 2D-ART together into a map; (*b*) the superior quality of the restored maps obtained when including Gibbs priors over reconstructions using only projection data, in particular for the experimentally relevant cases of a medium noise level and a low number of available spots per grain.

The qualities of reconstructions based on the same diffraction data but different sets of Gibbs potentials, as derived from typical microstructures, were shown to be very similar.

The authors would like to thank J. Bowen for providing the EBSD micrographs as well as H. Y. Liao, L. Rodek and S. W.

Rowland for helpful discussions. We also acknowledge one referee for a careful reading and many valuable suggestions. This work was partially supported by NIH grant HL70472 and NSF grant DMS0306215, and by the Danish National Research Foundation. The first author was partially supported by a Feodor Lynen fellowship of the Alexander von Humboldt Foundation (Germany).

References

- Alam, M. N., Blackman, M. & Pashley, D. W. (1954). *Proc. R. Soc. A*, **221**, 224–242.
- Alpers, A., Knudsen, E., Poulsen, H. & Herman, G. T. (2005). *Electron. Notes Discrete Math.* **20**, 419–437.
- Brémaud, P. (1999). *Markov Chains: Gibbs Fields, Monte Carlo Simulations and Queues*. New York: Springer.
- Carvalho, B. M., Herman, G. T., Matej, S., Salzberg, C. & Vardi, E. (1999). *Information Processing in Medical Imaging*, edited by A. Kuba, M. Samal & A. Todd-Pokropek, pp. 29–41. Berlin: Springer.
- Gordon, R., Bender, R. & Herman, G. T. (1970). *J. Theor. Biol.* **29**, 471–482.
- Herman, G. T. & Kuba, A. (1999). Editors. *Discrete Tomography: Foundations, Algorithms and Applications*. Boston: Birkhäuser.
- Lassen, N. C. K., Juul Jensen, D. & Conradsen, K. (1992). *Scann. Microsc.* **6**, 115–121.
- Lauridsen, E. M., Schmidt, S., Suter, R. M. & Poulsen, H. F. (2001). *J. Appl. Cryst.* **34**, 744–750.
- Liao, H. Y. & Herman, G. T. (2004). *Discrete Appl. Math.* **139**, 149–170.
- Liao, H. Y. & Herman, G. T. (2005). *Electron. Notes Discrete Math.* **20**, 399–418.
- Markussen, T., Fu, X., Margulies, L., Lauridsen, E. M., Nielsen, S. F., Schmidt, S. & Poulsen, H. F. (2004). *J. Appl. Cryst.* **37**, 96–102.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H. & Teller, E. (1953). *J. Chem. Phys.* **21**, 1087–1092.
- 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.
- Poulsen, H. F., Nielsen, S. F., Lauridsen, E. M., Schmidt, S., Suter, R. M., Lienert, U., Margulies, L., Lorentzen, T. & Juul Jensen, D. (2001). *J. Appl. Cryst.* **34**, 751–756.
- Schmidt, S. (2006). Private communication.
- Schmidt, S., Nielsen, S. F., Gundlach, C., Margulies, L., Huang, X. & Juul Jensen, D. (2004). *Science*, **305**, 229–232.
- Vardi, E., Herman, G. T. & Kong, T. Y. (2001). *Linear Algebra Appl.* **339**, 75–89.