

Reconstruction of local orientation in grains using a discrete representation of orientation space

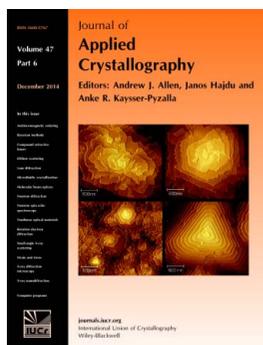
Nicola Viganò, Wolfgang Ludwig and Kees Joost Batenburg

J. Appl. Cryst. (2014). **47**, 1826–1840

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. Republication 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

Reconstruction of local orientation in grains using a discrete representation of orientation space

Nicola Viganò,^{a,b,c*} Wolfgang Ludwig^{a,b} and Kees Joost Batenburg^{d,c,e}Received 9 May 2014
Accepted 7 September 2014^aMATEIS, INSA Lyon, France, ^bESRF Grenoble, France, ^cMinds-Vision Lab, University of Antwerp, Belgium, ^dCWI, Amsterdam, The Netherlands, and ^eUniversiteit Leiden, The Netherlands.
Correspondence e-mail: nicola.vigano@esrf.fr

This work presents a mathematical framework for reconstruction of local orientations in grains based on near-field diffraction data acquired in X-ray diffraction contrast tomography or other variants of the monochromatic beam three-dimensional X-ray diffraction methodology. The problem of orientation reconstruction is formulated in terms of an optimization over a six-dimensional space $\mathbb{X}^6 = \mathbb{R}^3 \otimes \mathbb{O}^3$, constructed from the outer product of real and orientation space, and a strongly convergent first-order algorithm that makes use of modern l_1 -minimization techniques is provided, to cope with the increasing number of unknowns introduced by the six-dimensional formulation of the reconstruction problem. The performance of the new reconstruction algorithm is then assessed on synthetic data, for varying degrees of deformation, both in a restricted line-beam illumination and in the more challenging full-beam illumination. Finally, the algorithm's behavior when dealing with different kinds of noise is shown. The proposed framework, along the reconstruction algorithm, looks promising for application to real experimental data from materials exhibiting intra-granular orientation spread of up to a few degrees.

© 2014 International Union of Crystallography

1. Introduction

Over the past decade, considerable effort has been put into the development of three-dimensional X-ray diffraction techniques for structural characterization of polycrystalline materials. The ultimate goal of these grain mapping techniques is the nondestructive description of a material's three-dimensional microstructure in terms of local phase and crystal orientation, the so-called microtexture. For a review of the state of the art in this field, the reader is referred to the special issue on three-dimensional diffraction microscopy techniques of the current journal (Borbély & Kaysser-Pyzalla, 2013) and the book by Barabash & Ice (2014).

In the general case, the crystalline microstructure of a volume element of a polycrystalline material may have to be described in terms of a three-dimensional orientation distribution function (ODF), allowing for multiple orientations to be present in each volume element. However, depending on the size of the volume element and the deformation state of the material, simplified but still adequate representations of the microtexture may be obtained by assigning an average orientation to each volume element, and, for undeformed materials, even a single (average) orientation per grain may be sufficient. The general six-dimensional framework for microtexture analysis has been discussed by Poulsen (2003), who suggested that the use of algebraic reconstruction techniques may prove a viable route for microtexture analysis and related, lower-dimensional subproblems.

In this article we focus on near-field variants of the monochromatic beam rotation method, like three-dimensional X-ray diffraction microscopy (3DXRD) (Poulsen, 2012) and X-ray diffraction contrast tomography (DCT) (Reischig *et al.*, 2013), well adapted for mapping two- and three-dimensional grain microstructures in materials where the aforementioned simplifying microstructure descriptions are applicable. In recent years, a variety of solutions for sub-cases of the general problem of microtexture analysis have been presented, and remarkable progress has been made using algebraic reconstruction techniques as well as reconstruction strategies based on forward modeling and/or combinatorial optimization. For an overview of this work the reader is referred to Li & Suter (2013), Poulsen (2012) and references therein. Restricting the illumination of the sample to a single slice through the volume, these methods have proved capable of producing orientation maps from metallic samples having undergone ten percent and more plastic deformation.

From an experimental point of view, the restriction of the sample illumination to individual slices compromises the temporal resolution and may result in anisotropic voxel size in three-dimensional reconstructions obtained from stacking these layers. For this reason the development of algorithms allowing microstructure reconstruction from three-dimensional diffraction data is highly relevant. DCT is an example of a truly three-dimensional tomographic imaging approach, sharing a common experimental setup with conventional X-ray microtomography. The algebraic reconstruction

approach behind DCT may be considered as one of the sub-cases of the general six-dimensional framework, tailored to undeformed materials exhibiting limited ($\leq 0.5^\circ$) intra-granular orientation spreads. In this case the orientational degrees of freedom inside each volume element are neglected and three-dimensional grain shapes are reconstructed assuming a single constant orientation throughout the grain volume.

In this work we introduce a six-dimensional extension of the three-dimensional tomographic reconstruction approach behind DCT, extending the applicability of this method to materials exhibiting intra-granular orientation spreads of up to several degrees. While previous preliminary work in this direction already exists (Viganò *et al.*, 2013), the present study is a major extension, being the first wide-ranging and complete treatment of the model and the reconstruction algorithm.

In order to account for spatially varying orientations inside a grain and potentially also the presence of multiple orientations inside each of the individual volume elements, we describe a discrete six-dimensional representation of the reconstruction problem in the form of a direct product of real space and orientation space (Poulsen, 2003).

The model is based on the assumption of kinematic diffraction, implying proportionality between crystal volume and integrated diffracted intensity. We assume that the grain average orientation has been determined by one of the existing polycrystal indexing approaches (Lauridsen *et al.*, 2001; Sharma *et al.*, 2012; Reischig *et al.*, 2013; Schmidt, 2014) and that the experiment has been performed in such way that diffraction signals of different grains can be separated on the detector (negligible overlap with diffraction spots from other grains).

In §1.1 we summarize the basic concepts of three-dimensional diffraction imaging approaches like DCT and 3DXRD. We then describe in §1.2 how different properties of the analyzed materials give rise to different types of images on the detector.

In §2 we introduce the proposed mathematical model for representing the problem.

§3 is then dedicated to the application of the reconstruction algorithm to synthetic data generated from a grain with one degree of intra-granular orientation spread.

Finally in §4, we test the strength of the model and the algorithm when applying them to more challenging cases, with higher degrees of orientation spread and discontinuities (small angle boundaries) in orientation between neighboring blocks.

In the appendices we add some extra considerations and details, like the discussion of some implementation details of the algorithm (Appendix A) and a robustness test of the algorithm when subject to extinction noise (Appendix B).

1.1. Near-field diffraction imaging setups

The sample is placed on a rotation stage and irradiated by a parallel monochromatic X-ray beam that is perpendicular to the rotation axis of the sample and whose dimensions are determined by slits. The data recorded during a near-field diffraction imaging experiment will be over a range of 2π in

the form of $s = 2\pi/\delta\omega$ images, where $\delta\omega$ is the angular range over which the signal is integrated on the detector, to form a single image. Typical values for $\delta\omega$ are in the range between 0.05 and 0.2° .

As the polycrystalline sample rotates, the Bragg condition is met by the different grains at specific angular positions, giving rise to diffraction ‘spots’. For undeformed grains these spots correspond to two-dimensional projections of the three-dimensional grain volumes on the detector.

As can be seen in Fig. 1, a diffraction imaging experiment can be performed with both a full two-dimensional beam or a restricted one-dimensional line beam. The advantage of the second is that the beam dimension is very small in the direction parallel to the rotation axis, which in turn reduces the complexity (convolution) of the reconstruction task (from a six-dimensional to a five-dimensional problem).

The physics behind 3DXRD and DCT measurements has been outlined in previous work; for completeness we recall some basic equations for calculation of the diffraction geometry following the presentation by Poulsen (2004) and

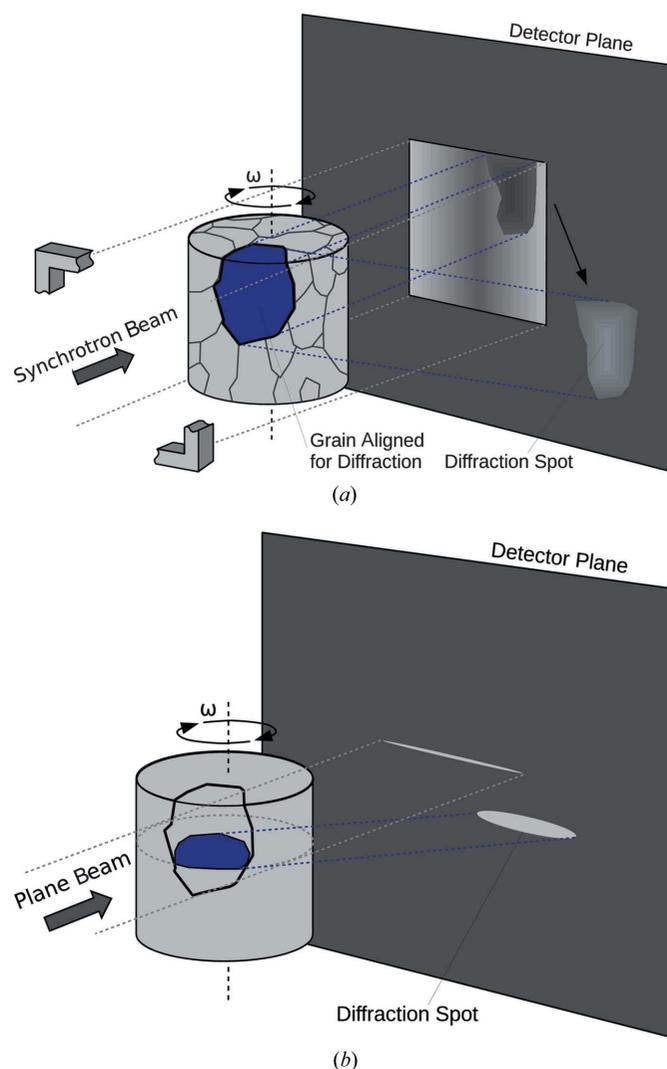


Figure 1
Diffraction of one grain in a near-field diffraction imaging experiment: (a) full-beam illumination, (b) one-dimensional line-beam illumination.

Suter *et al.* (2006). For more information on DCT we refer to the articles by Ludwig *et al.* (2009) and Reischig *et al.* (2013).

We will consider the case of a face-centered cubic crystal and experimental parameters typically used in a near-field diffraction imaging experiment. The detector fully intercepts the first three diffraction cones, giving rise to about 60 diffraction spots per crystal.

We recall the relationship between a scattering vector in the laboratory system \mathbf{G}_{lab} , which is the direction of the diffracting plane normal, and the Bragg node of indexes (h, k, l) in the reciprocal space, at the given rotation angle ω :

$$\mathbf{G}_{\text{lab}} = \Omega_{\omega} \mathbf{g}^{-1} \mathbf{B} \begin{pmatrix} h \\ k \\ l \end{pmatrix}, \quad (1)$$

where Ω_{ω} is the rotation matrix associated with the rotation angle ω , \mathbf{g} is the orientation matrix and \mathbf{B} is the metric matrix. Bragg diffraction is observed for rotation angles ω fulfilling

$$\mathbf{G}_{\text{lab}_x} = -\frac{|\mathbf{G}_{hkl}|^2}{2k}, \quad (2)$$

where the incident wavevector is $\mathbf{k}_i = k\hat{\mathbf{x}}$ (Suter *et al.*, 2006). Equations (1) and (2) fully determine the projection geometry for a volume element with a given orientation g , and can be used to derive the positions (u, v) and rotation angle (ω) onto which a given voxel will project on the detector.

1.2. Projection of the crystallographic domains

Having introduced the diffraction geometry, we now describe the appearance of projection data for two different cases.

(1) Undeformed grains. Assuming that each grain is a perfect lattice, we will observe a projection of the grain volume on the detector for those scattering vectors \mathbf{G}_{lab} which give rise to a diffracted beam (in the direction $\mathbf{k}_{\text{out}} = \mathbf{G}_{\text{lab}} + \mathbf{k}_{\text{in}}$) intersecting the detector (Figs. 1 and 2a).

In this case the shape of undeformed grains can be reconstructed by using oblique angle tomography reconstructions, which rely on algebraic reconstruction techniques like the SIRT algorithm (Kak & Slaney, 1988; Reischig *et al.*, 2013).

Having the same orientation, all voxels will give rise to diffraction signal in the same direction, for the same angles of rotation ω_n . In this idealized situation the diffracted beam will be parallel and each of the images on the detector will be a geometrical projection of the grain volume. These projections will be called ‘diffraction spots’.

(2) Deformed grains. A physically more relevant case corresponds to the situation when the whole grain volume diverges from the average grain orientation, but exhibits either smooth variations between nearby regions or a combination of smooth variations and discrete jumps. This case will be investigated further in this article, and we will provide various reconstruction examples with different values for the maximum degree of orientation spread.

As mentioned earlier, we model our grain as if it were made of small regions, each having a specific average orientation and so identifiable as a point in a three-dimensional representation

of the orientation space (‘vector field’). Integrated over the grain volume, these points will form a dispersion in orientation space, and we define the maximum orientation spread of this distribution as the diameter of the smallest ball in orientation space that contains all of its points.

We will restrict ourselves to the case of moderate orientation spread of up to a few degrees. As discussed later, this limit arises not only from memory restrictions by currently available hardware but also from the increasing overlap of the ‘diffraction blobs’ on the detector. Moreover, for higher levels of deformation, the vector-field description (one average orientation per voxel) may no longer be appropriate.

Let us now discuss in more detail how the data set presents itself when dealing with grains affected by smoothly varying orientation. As can be seen in Fig. 2(b), as opposed to Fig. 2(a), a grain does not produce a single detector image per reflection any more. Different regions of the grain will project on a range of neighboring images, forming a three-dimensional object per reflection in the uvw space, where u and v are the two-dimensional Cartesian coordinates of a detector image and ω is the rotation angle. Moreover, the directions of the diffracted beams will no longer be parallel, giving rise to distorted projection images of the grain volume. These distorted projection volumes are called ‘diffraction blobs’. Note that three-dimensional formulations of the reconstruction problem as reported by Fu *et al.* (2003) and Ludwig *et al.* (2009) neglect this type of distortion and use the spots resulting from the integration of the blobs along the ω direction, treating them as

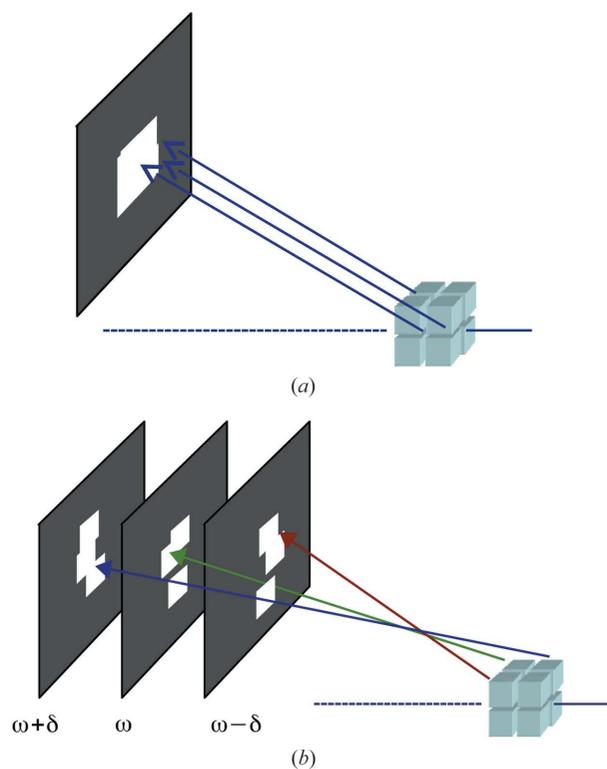


Figure 2 Projections of (a) a uniform grain and (b) a ‘deformed’ grain that will give rise to distorted diffraction blobs, spreading over a range of rotation angles.

parallel-beam projections of the grain volume and applying a regular oblique-angle SIRT algorithm to them.

2. Model

The problem of determining the local orientation of a grain has already been heavily studied and many approaches have been proposed, in terms of both models and algorithms.

Prominent examples are the forward modeling idea from Suter *et al.* (2006) and *GrainSweeper* from S. Schmidt, discussed by Poulsen (2012). What they all have in common is that they try to assign orientations to the voxels, but each voxel is treated independently. Typically operating on less convoluted one-dimensional line-beam data, these algorithms have proven stable, even when neglecting local diffracted intensities for performance reasons.

Another interesting approach is introduced by Alpers *et al.* (2005) and further discussed by Alpers *et al.* (2006). This other approach is somehow closer to our model because it operates a grain-by-grain optimization, but it tries to solve the mathematical problem in a three-dimensional vector-field representation, using discrete tomography, with Gibbs priors tailored for grain maps.

Here we try to perform a global optimization for the full six-dimensional problem (two-dimensional full-beam illumination), where all the grain voxels are processed simultaneously and matching of diffracted intensities is part of the optimization, using recent algorithms and techniques from signal theory. We also add the constraint of allowing only a discrete and restricted number of orientations. This results in a sampling of the three-dimensional orientation space, similar to the common sampling of the three-dimensional real space implicitly employed by traditional tomography, making this method a six-dimensional method.

2.1. Discrete representation of the six-dimensional reconstruction space

In order to add orientation degrees of freedom to the real-space voxels in our grain volume, we introduce a discretized six-dimensional reconstruction space $\mathbb{X}^6 = \mathbb{R}^3 \otimes \mathbb{O}^3$ as the outer product of real space and orientation space. However, only a discrete set of orientations, sampled around the grain average orientation and confined to a small ball in orientation space, are included in the analysis. We decided to represent the \mathbb{O}^3 orientation space as a Rodrigues space. Since we consider only small deviations from the known grain average orientation, this orientation subspace can be treated as Euclidean and both three-dimensional spaces are tessellated using cube-shaped voxels.

From a materials science point of view, this six-dimensional space could be viewed as a collection of three-dimensional voxelated orientation spaces, one for each of the volume elements, or in other words, with a discrete ODF per real-space volume element. For performance reasons discussed in Appendix A, we organize the two subspaces in the inverse order, which means that we represent the six-dimensional

space as a collection of real-space volumes, one for each of the sampled orientations (Fig. 3).

We will use the vector $\mathbf{x} \in \mathbb{X}^6$ as a synthetic representation for any arbitrary volume in our six-dimensional reconstruction space. If we assume that the grain volume was divided into a grid of size $n \times n \times n$ and the orientation space into a grid of size $p \times p \times p$, the length of the vector \mathbf{x} is $n^3 p^3$, and the scalar entries can be interpreted as ‘scattering powers’ of the corresponding element. For a grain that can be described by a ‘three-dimensional vector field’ (e.g. one three-dimensional Rodrigues vector per real-space voxel), the majority of elements of this vector in \mathbb{X}^6 will not contain any intensity.

If we model the forward-projection operator in tomography as a matrix \mathbf{A} , the solution to the reconstruction problem can be seen as the solution to a linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$, where \mathbf{x}^* is a vector that makes the equality true and \mathbf{b} represents the images (pixel intensities) recorded by the detector.

Note that the choice of having a fixed set of allowed orientations, decided *a priori*, makes the projection matrix a static object and the system of equations is linear.

To summarize, in the model proposed by this paper, the vector \mathbf{x} is constructed as a collection of three-dimensional real-space volumes that each correspond to a specific point in orientation space. This means that the projection matrix \mathbf{A} will be a stack of projection matrices, one for each selected orientation, and that the linear system can be rewritten as the sum of each sampled orientation:

$$\begin{aligned} \mathbf{A}\mathbf{x} &= \mathbf{A}_1\mathbf{x}_1 + \mathbf{A}_2\mathbf{x}_2 + \dots + \mathbf{A}_N\mathbf{x}_N \\ &= \sum_i^N \mathbf{A}_i\mathbf{x}_i = \mathbf{b}, \end{aligned} \quad (3)$$

where N is the number of sampled orientations.

2.2. Mathematical model

We will now first show why the reconstruction problem posed in §2.1 is underdetermined and then introduce our model that tries to cope with this situation.

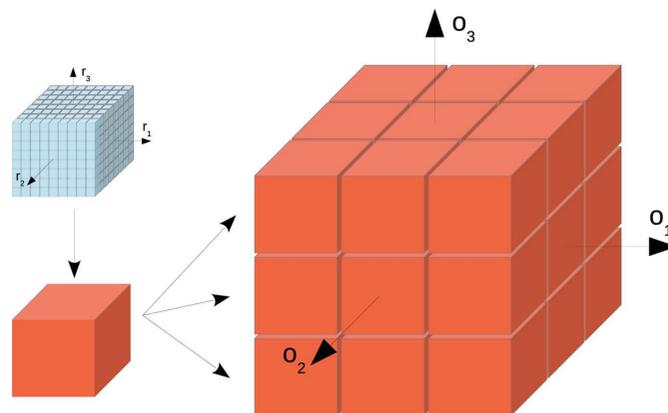


Figure 3 A three-dimensional volume in real space can be considered a single point in three-dimensional orientation space. By taking the corresponding real-space positions in each of these volumes, we would be able to construct a single-voxel ODF for each of the real-space positions.

We know from linear algebra theory that, for a guaranteed unique solution, the matrix \mathbf{A} should be square and of full rank (otherwise uniqueness/existence depends on the particular vector \mathbf{b}). This means that, if we intend to sample the grain volume with a cubic grid of size $n \times n \times n$ and the orientation space with a cubic grid of size $p \times p \times p$, the number of unknowns will be of order $O(n^3 p^3)$.

A typical diffraction blob has dimensions of order $O(n^2 t)$ pixels, where t is the so-called ω spread, corresponding to the number of images onto which the blob extends. The requirement in terms of number of blobs, for a unique solution, is of order $O[(n/t)p^3]$. Note that grains with a smaller orientation spread will have smaller values of t . This also implies that we will need fewer sampling points in orientation space and so a reduction of the number p . In fact, we can conjecture that there will be a linear relationship between t and p , thus reducing the requirement of the number of blobs to be of order $O(np^2)$.

Assuming 1° maximum orientation spread and 0.1° angular integration steps ($\delta\omega$), the blobs will be ~ 10 slices thick and so $t = 10$. If then we also assume a 50×50 spot size and an orientation sampling that is characterized by a 0.1° distance between the sampled points, we will have that $n = 50$ and $p = 10$. Thus the required number of blobs will be of order $50 \times 10 \times 10 = 5000$.

However, the number of blobs that we can typically record on the detector in a single scan and use in the analysis is usually between 20 and 100, making the reconstruction problem heavily underdetermined. As stated earlier, in our simulations we chose a fixed number of 60 blobs for all the test cases.

As the measurements will always be perturbed by noise and other experimental inconsistencies, the following formulation is preferred:

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2, \quad (4)$$

where from the feasible solution space we select the vector \mathbf{x}^* that minimizes the l_2 norm of the residual. Unfortunately, the solution to equation (4) is not unique if \mathbf{A} does not have full rank. In fact, equation (4) implies that \mathbf{x}^* makes the equality

$$\nabla(\|\mathbf{A}\mathbf{x}^* - \mathbf{b}\|_2) = \mathbf{0} \quad (5)$$

true, and we see from equation (6) that equation (4) is equivalent to $\mathbf{A}\mathbf{x} = \mathbf{b}$:

$$\nabla(\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2) = \nabla[(\mathbf{A}\mathbf{x} - \mathbf{b})^T(\mathbf{A}\mathbf{x} - \mathbf{b})] \quad (6a)$$

$$= 2\mathbf{A}^T\mathbf{A}\mathbf{x} - 2\mathbf{A}^T\mathbf{b} \quad (6b)$$

$$= 2\mathbf{A}^T(\mathbf{A}\mathbf{x} - \mathbf{b}). \quad (6c)$$

So, every solution of $\mathbf{A}\mathbf{x} = \mathbf{b}$ is also solution to (4), but the minimization formulation better deals with noise. To regularize the problem (4), we can now impose some constraints that help in selecting specific solutions that exhibit the physical properties of real-world samples. Under the assumption of a fine enough sampling in real space, we can expect that only a few orientations in the six-dimensional space will be active for each real-space voxel, so that the tools for sparsity recovery

will play an important role in selecting solutions that reflect the properties of real-world samples.

Normally we would like to add a penalty term over the l_0 semi-norm, because this semi-norm would count the nonzero entries in the solution vector, and a minimization over it would then yield a sparse solution. On the other side, using the l_0 semi-norm would require a combinatorial search of the solution, through all the solution space. We can instead think of introducing a penalty term that uses the l_1 norm and for which simpler and better performing algorithms exist (Chambolle & Pock, 2011).

It was demonstrated by Candes & Romberg (2007) and Chen *et al.* (2001) that l_1 minimization can result in very accurate recovery of sparse signals, and it is known from the literature that in general it promotes the choice of sparse solutions for problems where the solution is not unique (Boyd & Vandenberghe, 2004).

So, even if it is difficult to meet all the mathematical requirements for sparse recovery (see *e.g.* Candes & Romberg, 2007) for our physical case, where randomness of measurements is not an option, the use of sparsity promotion techniques is a key element of the microtexture reconstruction approach proposed in the current work.

We introduce this additional regularization term by exchanging the formulation (4) with a more powerful Lasso formulation, where the minimized l_1 norm would be in the full six-dimensional space:

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2 + \lambda\|\mathbf{x}\|_1 \quad \text{subject to } \mathbf{x} \geq \mathbf{0}, \quad (7)$$

where λ is a weight that balances the data fidelity term $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$, from equation (4), and the regularization term $\|\mathbf{x}\|_1$. This formulation assigns a penalty to solutions that exhibit a large l_1 norm in the reconstruction domain.

The non-negativity constraint for the reconstructed volumes has both a physical meaning and a practical outcome. On the physical side, the voxels are supposed to indicate a scattering power for the given orientation at the fixed position in the sample, and this scattering power cannot be negative under the assumption of kinematical scattering. On the other side, imposing the non-negativity constraint makes the reconstruction algorithm converge quicker and give better results.

2.3. Algorithm and data structures

Having described the model that we wish to solve, we will now present an algorithm instance that can solve the same Lasso formulation. Recently, a new class of first-order primal-dual algorithms was proposed. This class of algorithms is called Chambolle–Pock (CP; Chambolle & Pock, 2011; Sidky *et al.*, 2012).

The CP algorithms can solve different types of optimization problems, and it is relatively easy to mathematically derive the algorithm tailored for the problem given in equation (7):

Synthetic description of CP Algorithm

Input: x_0, y, λ

Output: Reconstruction in x_k

$\mathbf{p}_0 \leftarrow \mathbf{0};$

$\mathbf{q}_0 \leftarrow \mathbf{0};$

$\theta \leftarrow 1;$

for $k \leftarrow 1$ to n do

$$\mathbf{p}_k \leftarrow \mathbf{p}_{k-1} + \sigma_1 \frac{(\mathbf{A}\bar{\mathbf{x}}_{k-1} - \mathbf{y})}{1 + \sigma_1};$$

$$\mathbf{q}_k \leftarrow \lambda \frac{\mathbf{q}_{k-1} + \sigma_2 \bar{\mathbf{x}}_{k-1}}{\max(\lambda \mathbf{1}, |\mathbf{q}_{k-1} + \sigma_2 \bar{\mathbf{x}}_{k-1}|)};$$

$$\mathbf{x}_k \leftarrow P_0(\mathbf{x}_{k-1} - \tau \mathbf{A}^T \mathbf{p}_k - \tau \mathbf{q}_k);$$

$$\bar{\mathbf{x}}_k \leftarrow \mathbf{x}_k + \theta (\mathbf{x}_k - \mathbf{x}_{k-1});$$

end

The algorithm above is composed of just a few important parts: a projection in both the projection data space and the space associated with the regularization condition in the objective function, the back-projection, and a memory step.

To understand how these steps in the algorithm can be performed in our six-dimensional problem [equation (7)], we will now introduce the underlying data structures. We have two main data structures that hold the information: the ‘diff-stack’ (vector \mathbf{b}) and the ‘solution vector’ (vector \mathbf{x}). The diffstack contains all the information recorded by the detector, while the solution vector contains the six-dimensional volume in orientation and real space introduced in §2.1. These two data structures are linked by a third object, generically named ‘geometry’ (matrix \mathbf{A}). Each line of the matrix \mathbf{A} contains the contribution of each element of x to one single pixel on the diffracted images, but using the *ASTRA* toolbox (Palenstijn *et al.*, 2011; <http://code.google.com/p/astra-toolbox/>) it is possible to use a simpler description, which consists of a collection of tables that fully describe the projection of the three-dimensional volumes onto the two-dimensional detector images.

In Appendix A we will explain these concepts in more detail and discuss implementation details concerning the algorithm.

3. Reconstruction of slightly deformed materials

We will now show the result of reconstructions using multiple orientations and diffraction blobs, demonstrating the improvement over the standard DCT approach in reconstructing a three-dimensional cubic grain that exhibits smoothly varying deformation gradients in orientation space, with 1° of total orientation spread.

Synthetic diffraction data were generated from a cube-shaped phantom grain with $50 \times 50 \times 50$ voxels using geometrical and experimental settings identical to those employed in a previous measurement. The diffraction images were simulated using a previously existing code (P. Reischig, *Finite Element Diffraction*, unpublished), based on a (continuous) three-dimensional vector-field description of the orientation inside the grain.

Before going into a full analysis of the orientation reconstruction, we can already show some results that concern the grain shape reconstruction (*i.e.* the signal integrated over the orientation space). As can be seen in Fig. 4, the shape of the reconstructed grain is greatly improved by the use of the six-dimensional algorithm.

The grayscale intensity in the said figure corresponds to reconstructed material density (or scattering power) for each real-space voxel. The features seen in the reconstruction of Fig. 4(c) are artifacts, probably caused by missing orientations, which are not well represented by the orientation sampling and which result in a lower reconstructed intensity.

The orientation sampling is a body-centered cubic (b.c.c.) lattice of 341, constructed from a $6 \times 6 \times 6$ cubic lattice, with another $5 \times 5 \times 5$ cubic lattice occupying the centers of the cells of the larger lattice. The total volume covered by this

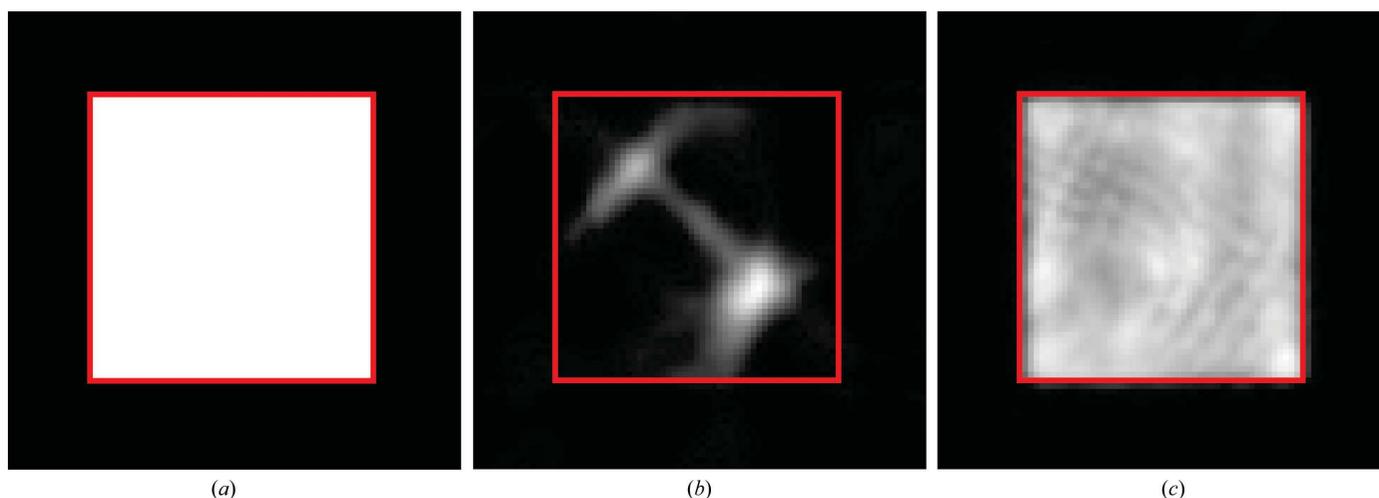


Figure 4

Comparison of one slice in the three-dimensional reconstruction of grain shapes, where the red mark identifies the original grain shape (50×50 voxels) and the grayscale intensities identify the (integrated) ‘scattering power’ coming from the voxels: (a) phantom, (b) single orientation SIRT and (c) reconstruction with 341 orientations, integrated over all sampled orientations.

b.c.c. lattice in orientation space is big enough to envelop all the orientations in the sample, but not too big, in order to not lose resolution (each edge is $\sim 1.1^\circ$).

3.1. Visual representation of reconstructions

In order to ease analysis of the reconstruction results, we introduce three types of visualization scheme. Two of these visualization methods produce a coloring of all the voxels in a slice or set of slices, based either on orientation domains or on the voxel-wise distance from the correct orientation. The third is a detailed visualization of the orientation-space distribution of intensities for a single voxel.

(1) Domain coloring. This method uses a sampling of orientation space different from the reconstruction sampling, and it assigns a color to each of the points in the new sampled grid, called the coloring lattice. For the phantom, each voxel is then being colored using the color corresponding to its closest orientation in the coloring lattice. The same applies to the reconstructed volume, but the intensity-weighted average orientation of the voxels, above a certain threshold, is taken as

the reference orientation for determining the distance from the coloring lattice.

(2) Distance coloring. This method takes the average reconstructed orientation of each single voxel and computes the distance in degrees from this average point and the true orientation of the voxel in orientation space.

(3) Voxel orientation distribution. This visualization technique concentrates on one single real-space voxel, and it performs the three-dimensional plotting of the intensities assigned to each sampled orientation.

In this representation, only the orientations that have intensity above a certain background are plotted, along with a few other special points. The true orientation of the voxel always takes the magenta color, the average orientation in the reconstruction is always in black, and a light-blue arrow connects them. The extremes of the sampled box are always in yellow and are connected by black lines to delimit the sampled volume in orientation space.

Finally, the other points, colored using the jet colormap of MATLAB (The MathWorks Inc., Natick, MA, USA) for encoding the intensity, are the active orientations for the

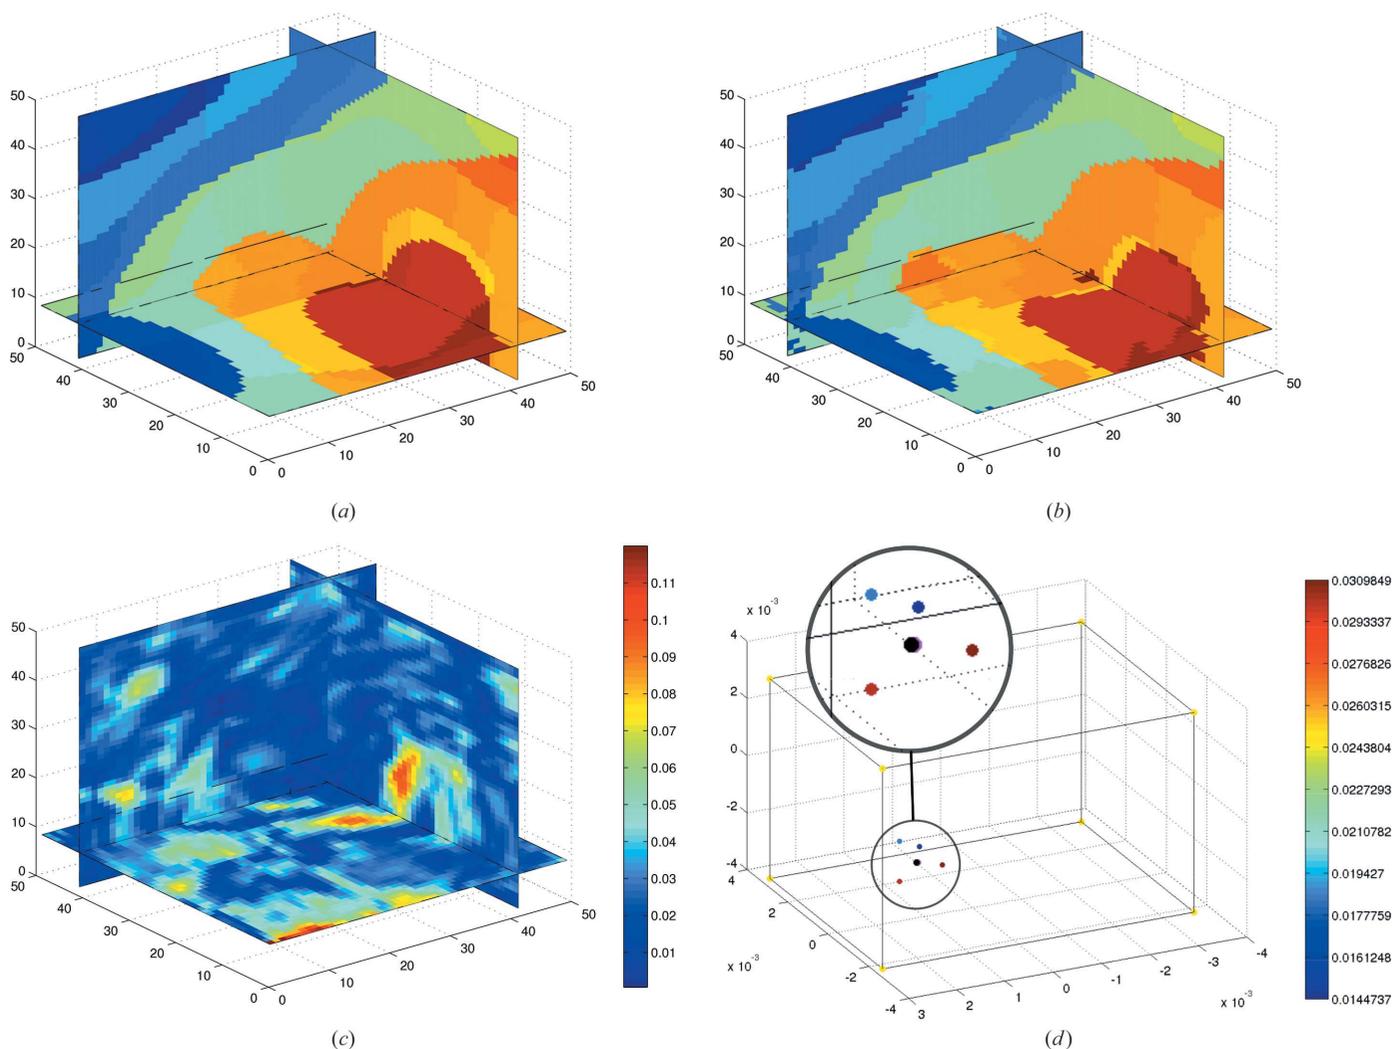


Figure 5 Comparison of the phantom and reconstruction of a 1° orientation spread sample, using 60 blobs: (a) phantom, (b) reconstruction, with 341 orientations, (c) error plot (in $^\circ$) of the reconstruction and (d) orientation-space visualization for one real-space voxel in the reconstruction.

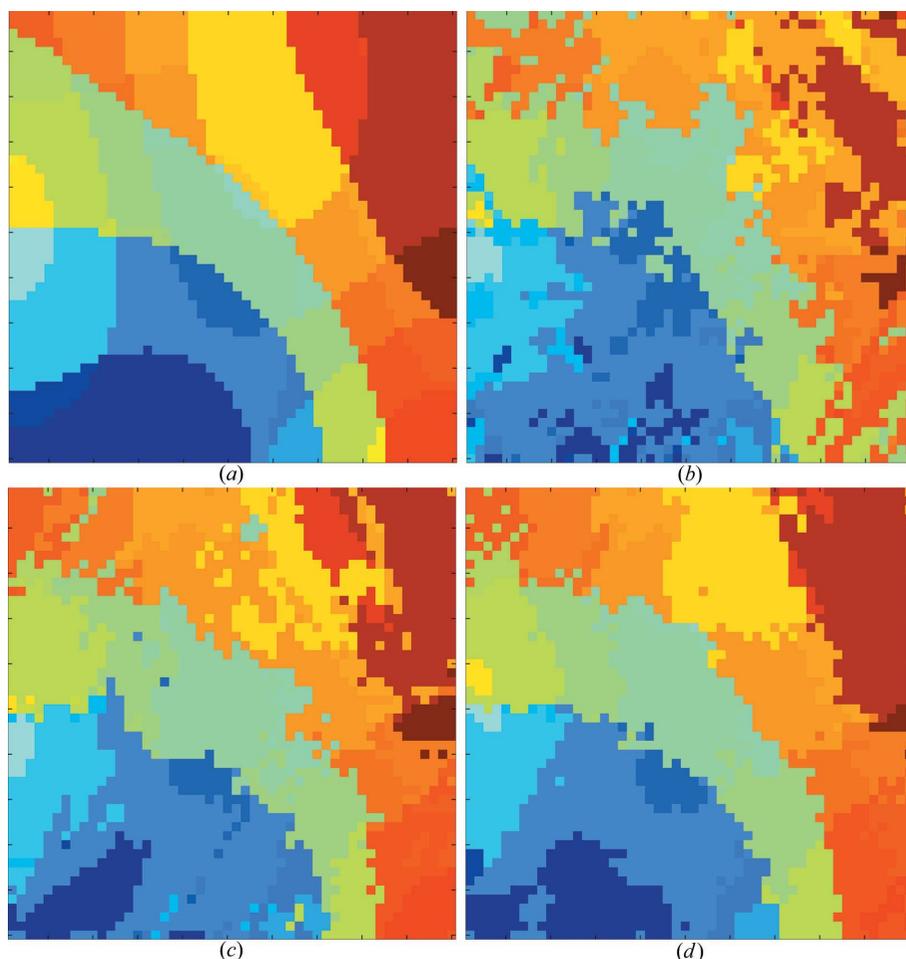


Figure 6
Results for different reconstruction parameters, using 60 blobs: (a) phantom, (b) 341 orientations, (c) 559 orientations and (d) 1241 orientations.

selected real-space voxel. The size of the points representing the active orientations is always slightly smaller than the size of the other special points.

3.2. Orientation reconstruction

For the reconstruction in Fig. 4, as can be seen in the comparison between Figs. 5(a) and 5(b), the reconstruction algorithm is able to retrieve a good approximation of the true solution in all three dimensions of the real space.

Another interesting analysis of this test case can be found in Appendix B, where we look at the reconstruction performance when the sample is affected by some degree of extinction noise.

We conclude this section by looking at the error plot of this reconstruction (Fig. 5c), which shows a relatively low angular voxel-wise error for the reconstruction. Moreover, the error is maximum on the borders of the grain. This could be expected in a non-perfect recovery, because that is the region at the border of the grain volume where diffracted intensities transition to zero. As a consequence, even small errors in this region can become comparable in modulus to the total

intensity of the voxel, so resulting in higher errors in the reconstruction.

On the other hand, looking at Fig. 5(d), which shows which orientations are active in one real-space voxel, we see that the l_1 minimization is effectively working.

4. Higher deformations

In many practical cases, the intra-granular orientation spread in a real material may be greater than that in the example of the previous section and reach values of several degrees. In this section we will analyze how the algorithm behaves at higher levels of intra-granular orientation spread (5° , including discontinuous jumps) and which factors influence the result.

In cases where we encountered memory limitations, we resort to using a less dense sampling of orientation space and the pseudo-slices approach, as described in Appendix A.

4.1. Acting on the number of orientations

We consider the case of one-dimensional slice illumination of a phantom grain with 5° of orientation spread.

An interesting result from the progression in Fig. 6 is that increasing the number of orientations, and so the resolution of sampling in orientation space, does actually give a measurable improvement for the reconstruction quality. So, the algorithm is actually able to cope with the increasing number of

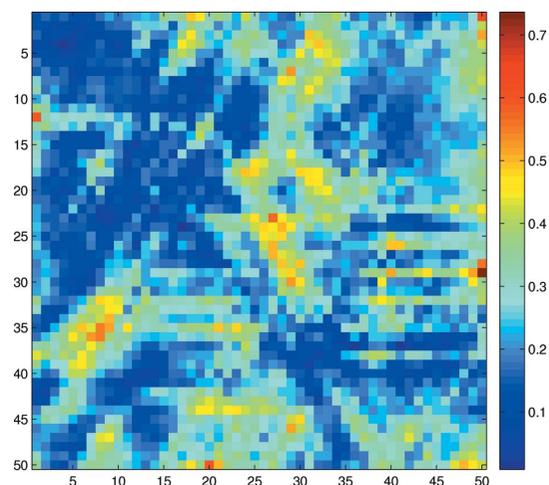


Figure 7
Error plot (in $^\circ$) of the reconstruction in Fig. 6(d)

unknowns, and the sparsifying techniques do seem to really help in finding the correct solution.

Let us now consider Fig. 7, which is the error plot (distance in $^\circ$ from true solution) of the reconstruction from Fig. 6(d). What is clear from this picture is that, apart from some hot spots, the global error is normally around (if not lower than) 10% of the maximum orientation spread.

4.2. Full-beam illumination

Unfortunately, moving to finer sampling in orientation space is problematic when working with full-beam illumination and three-dimensional real-space volumes. The size of the volumes and the blobs can be quite big, and the increase in orientation-space resolution could result in the computer running out of memory.

Nonetheless, as can be seen in Fig. 8, the six-dimensional reconstruction using 559 orientations is in line with the equivalent reconstruction for the two-dimensional real space. While the number of orientations was quite limited for the

given problem, owing to memory constraints, the fact that the reconstruction exhibits the features of the phantom is an interesting result.

4.3. Subgrain boundaries

Real materials may display discontinuous jumps in the orientation, so-called subgrain boundaries of only a few degrees or even fractions of a degree. This kind of grain substructure is usually observed along with the smooth orientation gradients considered previously, and we decided to analyze the performance of the algorithm on a test case having four subgrains and smooth orientation in each subdomain. In fact, even if the subgrain discontinuities could lead to separable diffraction spots, we would observe overlap on the detector due to the smooth deformation.

For solving this problem, a higher number of orientations is needed, and so we decided to only analyze the case of line-beam illumination (five-dimensional reconstruction).

Fig. 9(a) is a plot of the inter-voxel misorientation, where both the smooth orientation changes and the sharp subgrain boundaries can be seen.

As shown by Fig. 9, the introduction of the jumps does not completely break the reconstructions. This is another very interesting result, because even if the boundaries are regions of higher error, especially when reaching the borders (Fig. 9d), the bulky part of the subdomains seems to be reconstructed in a similar way to the previous 5° orientation spread test case.

5. Discussion

The model and algorithm proposed in this work perform a global optimization over the six-dimensional inverse problem defined as the local orientation reconstruction of the crystal lattice in grains, using data from diffraction contrast tomography or other variants of near-field diffraction imaging experiments involving X-rays or neutrons. Having an algorithm that can perform such an optimization for two-dimensional (extended beam) illumination, by taking local diffracted intensity explicitly into account, is quite important for time-resolved experimental studies and a prerequisite for studies carried out on low-brilliance sources like laboratory or neutron sources.

The model is, however, still based on a number of strong assumptions. The first is the assumption of kinematical diffraction, which in most practical cases is not fulfilled. Primary and secondary extinction as well as absorption may give rise to deviations from the idealized case of mathematical projections of the diffracting grain volumes. The impact of these perturbations can to some extent be reduced by optimizing the experiment conditions: for instance *via* the choice of the X-ray beam energy. Moreover, as shown in Appendix B, the algorithm can handle moderate levels of ‘extinction’ noise.

Next, the model explicitly ignores elastic distortion of the crystal unit cell. Introducing elastic strain, which can be modeled with a second rank tensor, into the problem would add an additional six dimensions. Unfolding and sampling all

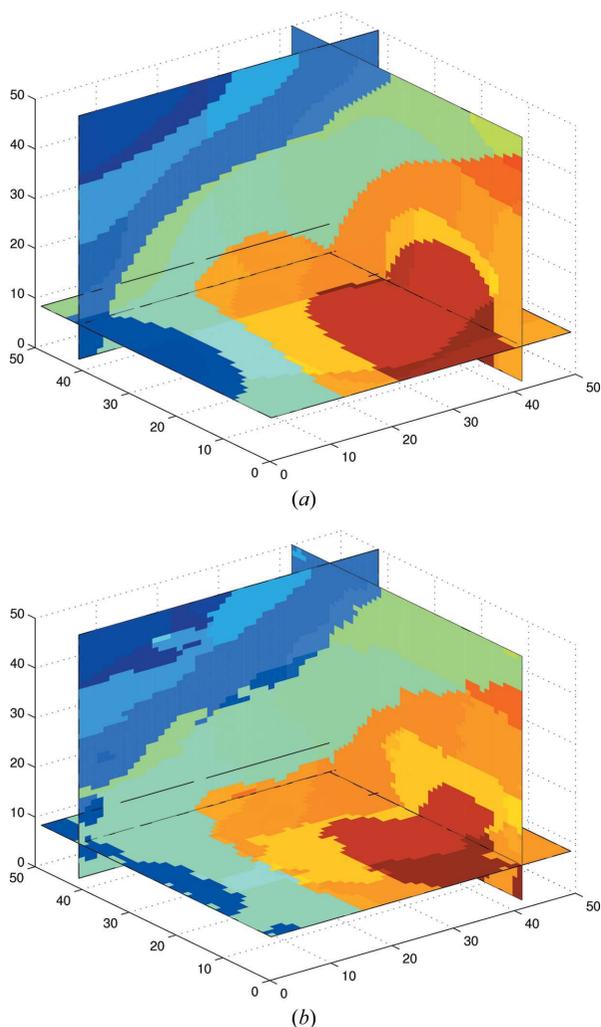


Figure 8
Comparison of the phantom and reconstruction of a 5° orientation spread full-beam illumination case, using 60 blobs: (a) phantom and (b) reconstruction, with 559 orientations

the 12 dimensions would be not feasible in terms of computational power and memory needs. Fortunately, many crystalline materials have elastic limits below one percent and the errors introduced by elastic distortion of the crystal lattice may be considered as small perturbations compared to the misorientations of up to several degrees discussed in this article.

Another assumption, tightly linked to the choice of an l_1 minimization over the six-dimensional space, is the postulate of a limited number of orientations per voxel. Depending on the actual size of a voxel and the type and deformation state of the material, voxel-wise sparsity in orientation space may or may not be a good description of the materials' microstructure. In other words, the physical relevance of the selected solution having minimum l_1 norm may depend on the experimental conditions. It would be interesting to evaluate the performance of the current framework in cases where each voxel actually contains an orientation distribution instead of a single orientation.

Finally we mention that the current approach treats each grain separately, which in turn translates to the need for

limited diffraction blob overlap with the blobs coming from other grains. The number of grains that can be simultaneously illuminated during data collection is inversely proportional to the intra-granular orientation spread of the grains. During an experiment, this condition can be adjusted to some extent by appropriate selection of the sample dimensions and setting of the beam-defining slits. We expect that the current approach will overcome some of the common problems encountered in tomographic grain reconstruction from near-field diffraction data. First of all, reconstructions based on this new approach are expected to outperform tomographic reconstructions treating integrated diffraction spots as parallel projections of grain volumes in terms of accuracy of reconstructed grain shapes. As shown in the previous sections, operating in a six-dimensional reconstruction space results in much improved accuracy of shape reconstructions for grains containing some non-negligible orientation spread.

This article was focused on testing and validating the model and the algorithm against a number of synthetic test cases. As a next step, the performance of the algorithm will be evaluated on experimental data, using electron backscatter diffraction

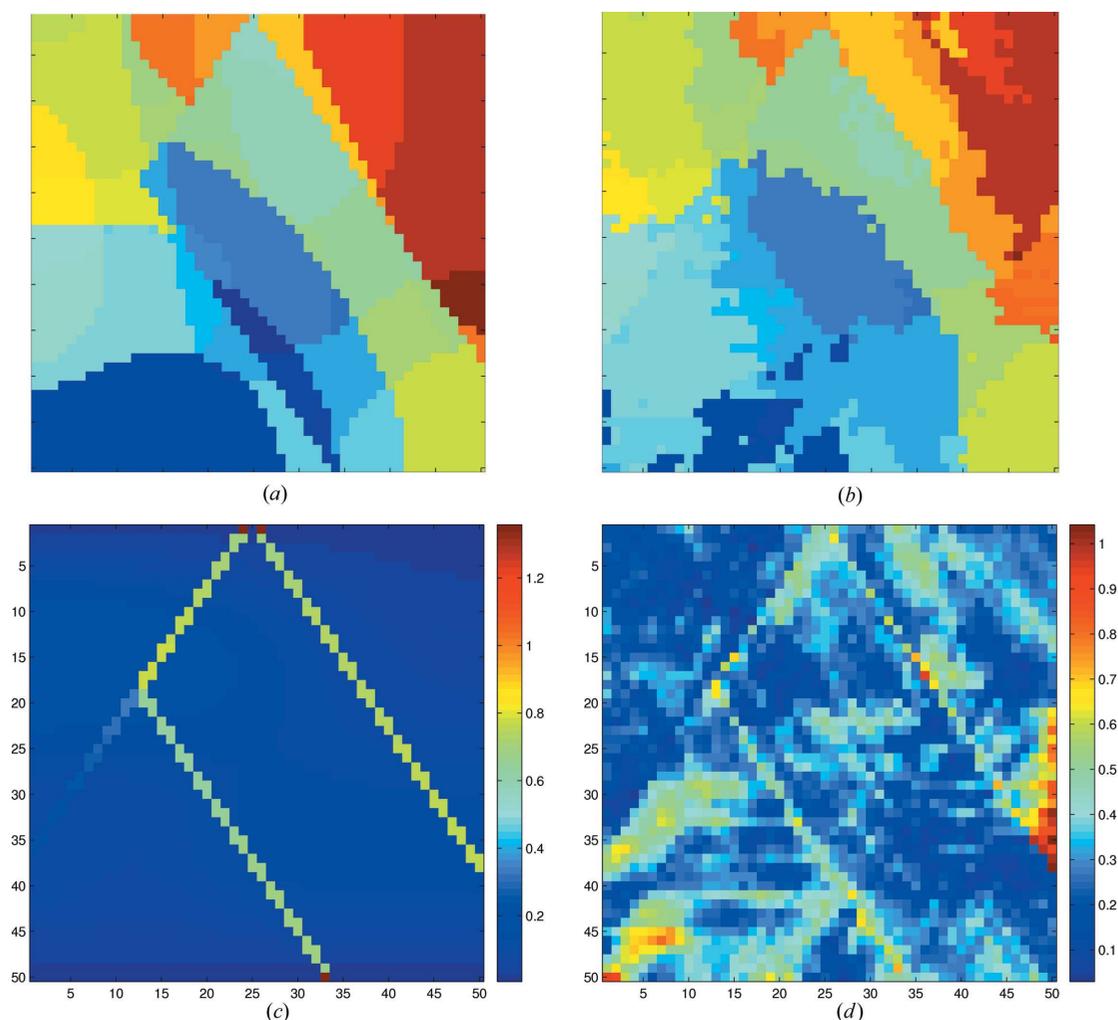


Figure 9

Reconstruction of a 5° orientation spread sample with subgrain boundaries, using 60 blobs: (a) phantom, (b) reconstruction, with 1241 orientations, (c) inter-voxel orientation distance in $^\circ$ and (d) error plot (in $^\circ$) of reconstruction.

and other three-dimensional orientation mapping techniques for cross-validation.

By adapting the projection and back-projection operators to the case of polychromatic and/or cone beam geometry, the same algorithmic framework could potentially be used to solve the inverse problem of orientation reconstruction from laboratory X-ray or neutron diffraction imaging experiments. For polychromatic variants of diffraction imaging, orientation gradients inside a grain give rise to noticeable distortions of the diffraction spots in both azimuthal and radial directions, thereby increasing the need to take these kinds of distortions into account.

6. Conclusions

We conclude this study of a six-dimensional orientation imaging framework by stating that it looks promising for future developments and applications to real experimental data, where the intra-granular orientation spread is within a few degrees.

While the current implementation, when scaling to higher levels of orientation spread, hits the computational and memory limits in the current desktop hardware, we have shown that for the study of limited inter-granular orientation spread of up to a few degrees the algorithm can be tuned to give encouraging results, by acting on the orientation resolution parameters.

The reformulation of the nonlinear inverse problem of six-dimensional microtexture analysis into a linearized version of the same, and the use of globally convergent algorithms for the given model, give a solid theoretical ground for future developments in the same direction.

APPENDIX A Implementation details

Let us now consider the data structures first introduced in §2.3.

(1) Solution vector. This is the data structure that is responsible for representing the six-dimensional space in the computer memory. Its inner structure reflects the schematic

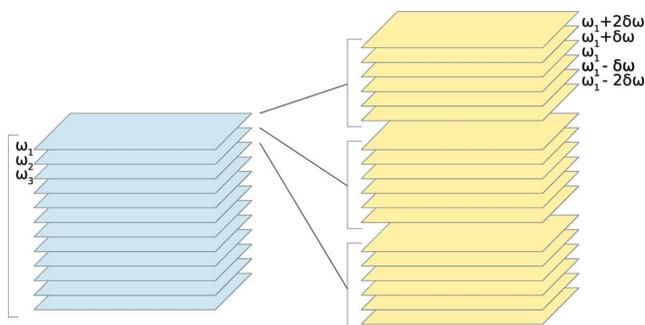


Figure 10
On the left, we illustrate a diffstack made of diffraction spots, and on the right, a diffstack made of diffraction blobs, where around each average ω we have multiple spots, with slightly different ω , that compose diffraction blobs.

representation in Fig. 3: we decided indeed to use a linear array of three-dimensional volumes, where each of these volumes has the same size and represents a specific point in orientation space. The number of volumes is equal to the number of sampled orientations.

(2) Diffstack. This is represented in Fig. (10) and is the data structure that holds the diffraction data recorded by the detector. As discussed in §1.2, data can be of different types: a collection of two-dimensional images or three-dimensional volumes (Fig. 10). Note that different orientations will typically project to different slices in the blobs, as depicted in Fig. (11).

(3) Geometry. Mathematically speaking, the geometry describes how the matrix \mathbf{A} acts on the solution vector to produce the diffstack.

As a physical consideration, the matrix \mathbf{A} does not include a full scattering model, and so the Lorentz factors and atomic form factors of the blob will be ignored. This makes it impossible to use the absolute intensities of the blobs, but an easy way to avoid this problem is to renormalize all the blobs to the same integrated intensity, because we assume that the relative scattering power between the voxels does not change as a function of the reflections.

The structure of the geometry table is fixed by the tomographic toolbox used to perform the projection and the back projection of the volumes on the detector. The table presents 12 columns and as many rows as the number of slices in the diffstack; in fact, each row determines the projection geometry for a given diffraction spot. The first three columns represent the scattering direction in the laboratory coordinates, while the second triplet of columns represents the position of the center of the detector, always in the laboratory reference system. The other two triplets of columns each represent a unit vector on the pixel grid that represents the detector.

Doing so, it is possible to fully determine an arbitrary projection geometry that specifies how each of the volumes in the solution vector project onto each of the slices of the diffstack. The actual expansion of the geometry tables into the projection matrix \mathbf{A} used in the reconstruction process is handled internally in the toolbox.

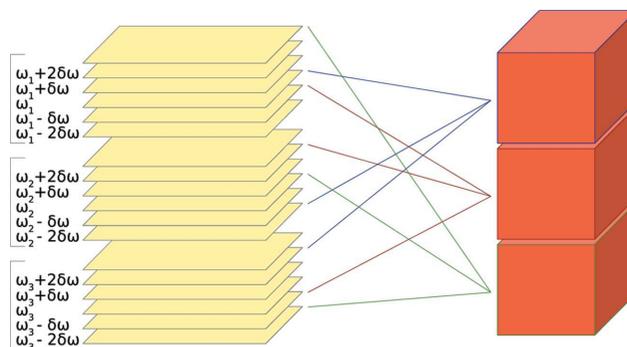


Figure 11
On the left, we illustrate a diffstack made of diffraction blobs, and on the right, the full volume associated with each sampled orientation. The lines connect each full volume associated with a sampled orientation and the slice in the blobs that the orientation reaches.

Having described the data structures we will now provide a description of how these objects interact to perform the matrix operations described in the CP algorithm. The projection matrix \mathbf{A} in (7) is equivalent to the tomographic forward projection of the volume onto the detector in our model. On the other side, the back projection is equivalent to the transpose \mathbf{A}^T .

In a simple case like the algorithm that only deals with the integrated diffraction spots, computing \mathbf{Ax}_k involves no more than performing the forward projection of each volume on the detector space (in a data structure like the diffstack) and summing all the contributions of each corresponding pixel, while computing \mathbf{A}^T of the detector residual will simply mean back projecting the detector residual in each of the volumes that form the solution vector.

For more complex problems like the one dealing with diffraction blobs, one intermediate step needs to be added for

practical reasons. It is indeed easier to let each orientation project the associated volume onto an intermediate projection data structure, containing N discrete ω slices, with N the number of observed reflections. The ω slices of the blob diffstack receive the summed contribution from each volume that projects onto that slice (Fig. 12).

This approach, shown in Fig. 12, has the advantage that, since the projection and back projection are computed on GPU devices, it does not require the full blob diffstack to be loaded in the GPU memory at the same time.

An important parameter in the model is the distance of sampling points in orientation space, which in turn governs the distance between real and reconstructed orientations in the solution vector, as well as the distance between the discrete ω slices, reached by these sampling points. Having only a fixed number of orientations, for large deformations, the number of orientations might not be enough to project to all the ω angles in each blob, giving rise to inconsistencies in the linear system.

A simple way to avoid this problem is to subdivide the slices in a blob into groups, from which pseudo-slices (or subspots) are produced. This technique reduces the angular resolution of the scans but helps in relieving the problems coming from undersampling.

The pseudo-slices can be produced by different strategies. The simplest strategy consists in summing all the images in a group, to produce a diffraction blob with a reduced angular resolution (Fig. 13a).

While this approach may be considered quite crude, it makes some sense, and for groups containing a small number of slices, probably little difference can be observed from other methods. On the other hand, this method does not help in letting close orientations talk to each other. What was in fact observed in the previous article (Viganò *et al.*, 2013) is that, if a voxel had an average orientation that was not sampled but between other sampled orientations, the voxel would gain signal from those close orientations, giving an average orientation close to the original orientation.

To try to better allow this behavior, we developed another approach, shown in Fig. 13(b), which builds the pseudo-slices from the linear interpolation of the original slices in the blobs. This is done by selecting a level of interpolation n , picking the geometry of one blob slice every n , and building the corresponding pseudo-

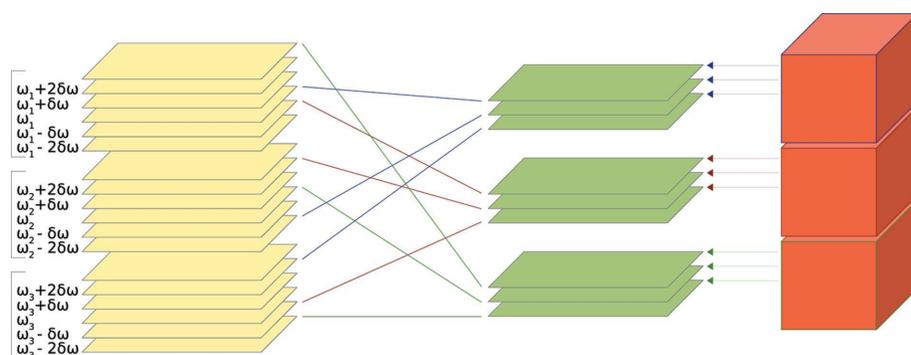


Figure 12
Volumes project first to an intermediate projection data structure, and then each blob slice becomes the pixel-wise sum of all the contributions from the corresponding slices in the intermediate projection data.

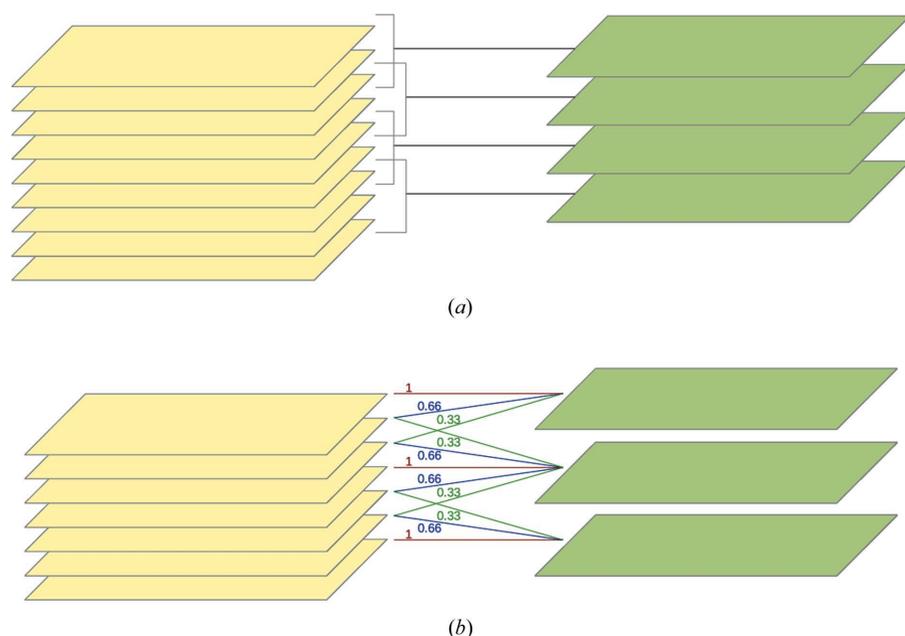


Figure 13
Grouping of slices: (a) the most simple approach to let the orientations reach all the slices; (b) the pseudo-slice approach, where we perform a linear interpolation of the original slices in the blobs.

slice from the $2(n - 1)$ blob slices around it. The assigned coefficients, as can be seen from Fig. 13(b), will have a sum of n , in order to satisfy the definition of consistent interpolation, and to respect the total scattering intensity observed per diffraction blob.

In terms of implementation, the memory step of the algorithm,

$$\bar{\mathbf{x}}_k = \mathbf{x}_k + \theta(\mathbf{x}_k - \mathbf{x}_{k-1}), \quad (8)$$

is simply a sequential operation over all the volumes in the solution vector, where θ is a purely mathematical parameter, introduced by Chambolle & Pock (2011).

In our implementation the tomographic forward projection and back projection are performed using the *ASTRA* toolbox, an open-source library that is able to provide high-performance tomographic primitives, thanks to the use of modern GPUs (Palenstijn *et al.*, 2011).

Despite possible slowdowns due to the memory transfers between the GPU and CPU memory domains, the generated overhead is almost negligible compared to the cost of forward projection and back projection. Nevertheless, we are now working on the implementation of asynchronous transfers and asynchronous computations on both CPU and GPU, in order to hide transfer latencies and blob \leftrightarrow sinogram transformations behind the most time-consuming operations.

APPENDIX B Stability to noise

In this section we will analyze the robustness of the algorithm to the presence of a common type of noise in a diffraction experiment, when crystals show little to no deformation. So, we performed this test for the 1° deformation case only.

Other sources of inconsistencies in the projection data, like shot noise in the pixel signals, were not considered. Typically in DCT experiments only the brightest diffraction spots/blobs are segmented and then selected for the reconstruction of the

single-grain volumes. Given that the shot noise for a pixel grows with the square root of its intensity, and the fact that DCT uses only the brightest spots, this type of noise can be considered negligible.

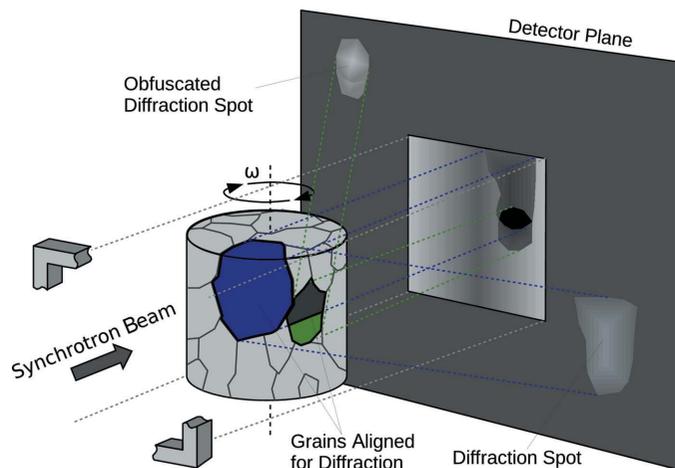
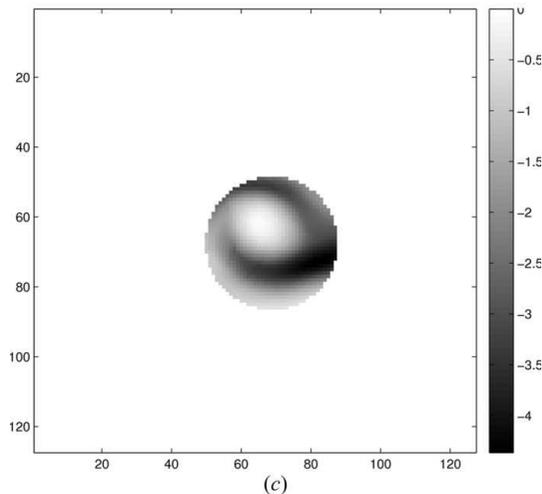
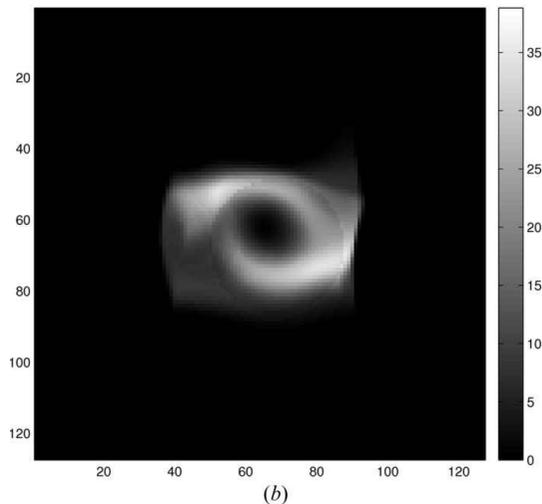
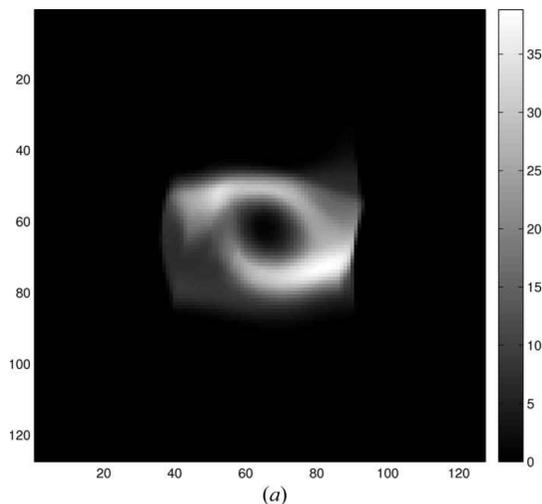


Figure 14
Illustration of the origin and appearance of noise related to secondary extinction in a polycrystalline sample.

Figure 15
(a) One slice of the original blobs, (b) the same slice with noise and (c) extinction noise.

B1. Extinction and attenuation effects

Noise affecting the measurements of diffracted intensities on the detector is related to extinction and attenuation of the X-ray beam traveling through the polycrystalline material.

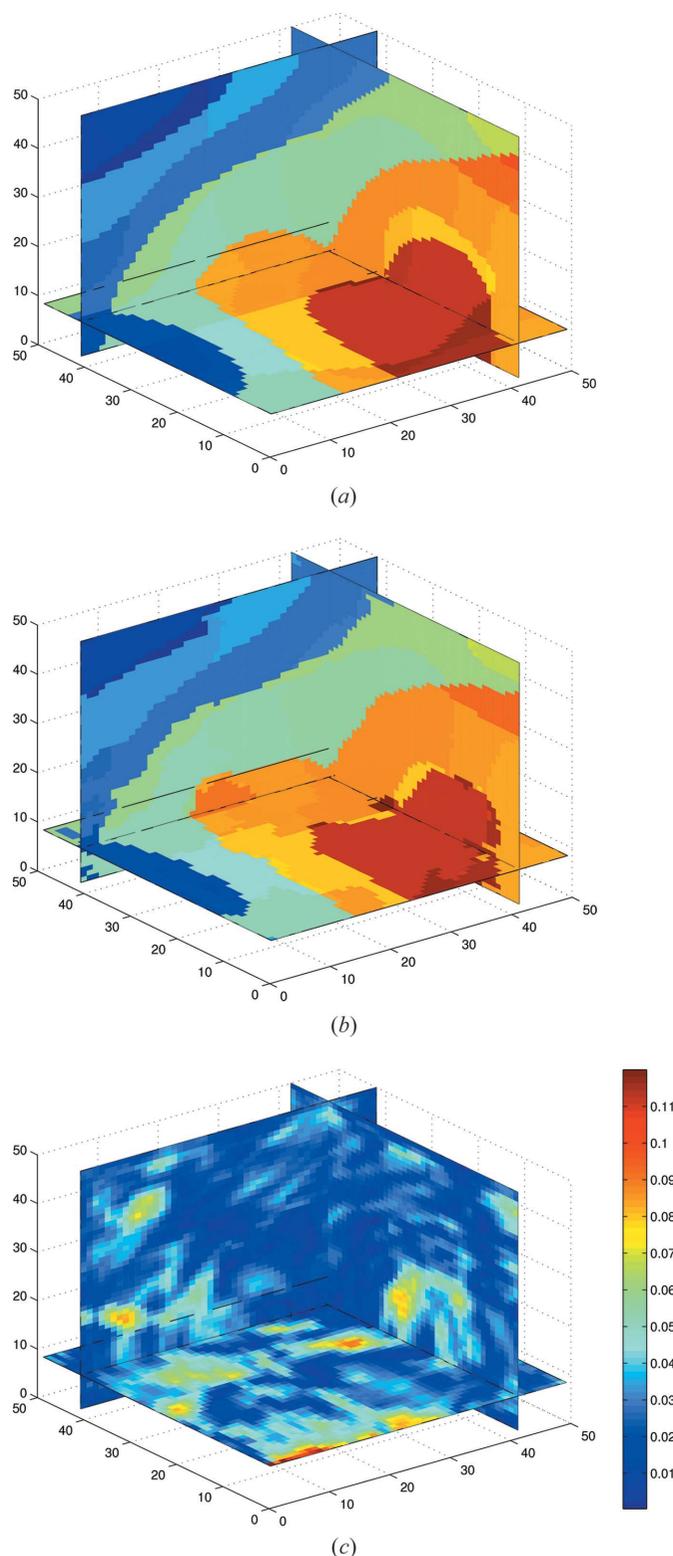


Figure 16
Comparison of phantom and reconstruction of a 1° orientation spread sample with extinction noise, using 60 blobs: (a) phantom, (b) reconstruction with 341 orientations and (c) error in $^\circ$ for each voxel.

Attenuation can in principle be predicted from the known distribution of the attenuation coefficient (measured simultaneously with the diffraction signal in a DCT experiment). In many cases, the influence of attenuation can be reduced by selecting an X-ray energy for which photoelectric absorption in the sample becomes negligible.

On the other hand, the prediction or reduction of the influence of extinction effects is more problematic. In general, two sources of extinction have to be considered. Primary extinction (multiple diffraction) within the grain of interest leads to the deviation of diffracted intensity from our model assumption of a linear relationship between crystal thickness and diffracted intensity. Fortunately, deformed materials are known to develop a mosaic structure, which limits the influence of this effect encountered in almost perfect crystals.

In polycrystalline samples another important source of noise is secondary extinction, referring to local changes of the intensity of the diffracted beams, caused by attenuation of the incoming and/or diffracted beams by other grains which happen to fulfill the Bragg condition on the trajectory through the sample volume (Fig. 14). Knowing the three-dimensional grain structure of the material, this phenomenon could in principle be predicted as well. However, the number of grains and diffraction events to be taken into account has prevented so far any attempts to explicitly take secondary extinction into account. Owing to the randomized nature of this phenomenon, the behavior of the algorithm in respect of this kind of inconsistency can be analysed by adding synthetic extinction noise to the projection data.

We decided to model extinction noise as a reduction of intensity in spherical regions of the blobs. The position and the radius of these spheres have also been randomized. The local attenuation is observed to usually be of the order of ~ 10 – 15% of the original intensity. The resulting blobs can be observed in Fig. 15(b).

B2. Reconstructions of noisy blobs

We now show the results of reconstructions performed using blobs affected by the extinction noise.

In the reconstruction shown in Fig. 16(b), all the blobs were affected by a random value of extinction noise, like in Fig. 15(c).

Looking to Fig. 16, compared to the same reconstruction without noise (Fig. 5), we observe only a moderate increase of the reconstruction error, which seems to indicate that the lower limits of the error in this test case are rather dominated by the lack in resolution of the orientation sampling and other deficiencies of the model.

KJB acknowledges the financial support of the Netherlands Organization for Scientific Research, project No. 639.072.005. NV and WL acknowledge the financial support of the French National Research Agency, project ANR 2010 BLAN 0935. NV and WL acknowledge Péter Reischig for having created a tool for generating synthetic test data sets, which was used in the creation of the test cases for this article. The authors acknowledge COST Action MP1207 for networking support.

References

- Alpers, A., Knudsen, E., Poulsen, H. & Herman, G. (2005). *Electron. Notes Discrete Math.* **20**, 419–437.
- Alpers, A., Poulsen, H. F., Knudsen, E. & Herman, G. T. (2006). *J. Appl. Cryst.* **39**, 582–588.
- Barabash, R. & Ice, G. (2014). *Strain and Dislocation Gradients from Diffraction: Spatially Resolved Local Structure and Defects*. London: Imperial College Press.
- Borbély, A. & Kaysser-Pyzalla, A. R. (2013). *J. Appl. Cryst.* **46**, 295–296.
- Boyd, S. & Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press.
- Candes, E. & Romberg, J. (2007). *Inverse Probl.* (m), 1–20.
- Chambolle, A. & Pock, T. (2011). *J. Math. Imaging Vis.* **40**, 120–145.
- Chen, S. S., Donoho, D. L. & Saunders, M. a. (2001). *SIAM Rev.* **43**, 129–159.
- Fu, X., Poulsen, H., Schmidt, S., Nielsen, S., Lauridsen, E. & Juul Jensen, D. (2003). *Scr. Mater.* **49**, 1093–1096.
- Kak, A. C. & Slaney, M. (1988). *Principles of Computerized Tomographic Imaging*. Piscataway: IEEE Press.
- Lauridsen, E. M., Schmidt, S., Suter, R. M. & Poulsen, H. F. (2001). *J. Appl. Cryst.* **34**, 744–750.
- Li, S. F. & Suter, R. M. (2013). *J. Appl. Cryst.* **46**, 512–524.
- Ludwig, W., Reischig, P., King, A., Herbig, M., Lauridsen, E. M., Johnson, G., Marrow, T. J. & Buffière, J. Y. (2009). *Rev. Sci. Instrum.* **80**, 033905.
- Palenstijn, W. J., Batenburg, K. J. & Sijbers, J. (2011). *J. Struct. Biol.* **176**, 250–253.
- Poulsen, H. F. (2003). *Philos. Mag.* **83**, 2761–2778.
- Poulsen, H. F. (2004). *Three-Dimensional X-ray Diffraction Microscopy*, Springer Tracts in Modern Physics, Vol. 205. Berlin, Heidelberg: Springer.
- Poulsen, H. F. (2012). *J. Appl. Cryst.* **45**, 1084–1097.
- Reischig, P., King, A., Nervo, L., Viganò, N., Guilhem, Y., Palenstijn, W. J., Batenburg, K. J., Preuss, M. & Ludwig, W. (2013). *J. Appl. Cryst.* **46**, 297–311.
- Schmidt, S. (2014). *J. Appl. Cryst.* **47**, 276–284.
- Sharma, H., Huizenga, R. M. & Offerman, S. E. (2012). *J. Appl. Cryst.* **45**, 693–704.
- Sidky, E. Y., Jørgensen, J. H. & Pan, X. (2012). *Phys. Med. Biol.* **57**, 3065–3091.
- Suter, R. M., Hennessy, D., Xiao, C. & Lienert, U. (2006). *Rev. Sci. Instrum.* **77**, 123905.
- Viganò, N., Ludwig, W. & Batenburg, K. J. (2013). *8th International Symposium on Image and Signal Processing and Analysis*, No. 1, pp. 594–599. IEEE.