

Discrete Representation of Local Orientation in Grains using Diffraction Contrast Tomography

Nicola Viganò
ESRF Grenoble, France
And INSA Lyon, France
Email: nicola.vigano@esrf.fr

Wolfgang Ludwig
ESRF Grenoble, France
And INSA Lyon, France
Email: wolfgang.ludwig@esrf.fr

Kees Joost Batenburg
CWI, Amsterdam, Netherlands
University of Antwerp, Belgium
Email: Joost.Batenburg@cwi.nl

Abstract—In this work we introduce a new model for dealing with the problem of local orientation reconstruction in grains, when using data from X-ray Diffraction Contrast Tomography experiments.

The model explores the use of well established minimization algorithms from the field of mathematical optimization, like FISTA, and the possible use of recent mathematical devices for the solution of highly undetermined systems of equations, like the l_1 -minimization over the Haar transform of the tomographic volumes.

Along with a detailed explanation of how such algorithms can be applied to our six-dimensional problem, we report encouraging results obtained on simulated data.

I. INTRODUCTION

Diffraction Contrast Tomography (DCT) is a non destructive technique used to characterize 3D grain microstructures and access their inner structure in a wide range of polycrystalline materials. Its standard approach is oriented to undeformed materials with negligible intragranular orientation spread. The technique is a truly three-dimensional tomographic imaging approach, sharing a common experimental setup with conventional X-ray microtomography.

One of the most advanced and interesting ongoing developments of DCT is in the reconstruction of the local orientation inside of grains, which provides a full characterization of the orientation microstructure of the sample. The aim of this work is to give a framework for determining which orientation is present in every single volume element of a grain.

A. DCT Experimental Setup

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.

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

The physics behind DCT is the same as for standard crystallography and diffraction measurements, so we will now briefly introduce the basic concepts, before explaining the peculiar properties of DCT. For more information we refer to [1] and [2].

We will now present the basics of the diffraction geometry, but for further details we refer to [3].

a) Crystal Domains: Let us assume first that we describe a three-dimensional crystal lattice with the basis vectors \underline{a} , \underline{b} and \underline{c} . The unit cell is the minimal space spanned by this basis.

The discrete periodic lattice in real space corresponds to a discrete periodic lattice in Fourier space, the so-called reciprocal space. The basis vectors in the reciprocal space will be \underline{a}^* , \underline{b}^* and \underline{c}^* .

Given an incoming wave-vector \underline{k}_{in} and an outgoing wave-vector \underline{k}_{out} , X-ray diffraction is observed when the vector $\underline{G} = \underline{k}_{in} - \underline{k}_{out}$ is close to a vector of the reciprocal lattice. A convenient way of representing \underline{G} in the reciprocal lattice is by using the Miller indexes, so that $\underline{G}_{hkl} = (h, k, l)^T$.

The link between the direct space Cartesian system and the reciprocal space, indexed by the Miller indexes (h, k, l) , is given by the matrix \mathbf{B} . The \mathbf{B} matrix is an upper triangular matrix that depends on both the direct and reciprocal basis vectors, and the angles between them. The result is a vector $\underline{G}_c = \mathbf{B}\underline{G}_{hkl}$

b) Diffraction Geometry: In this paragraph we will now establish the connection between the diffraction directions, for a specific crystal lattice in a generic reference system, and the points \underline{G}_{hkl} in reciprocal space.

Let us now restrict ourselves to an individual grain rotated by an arbitrary angle along an arbitrary axis, and then fix one rotation as sample reference. For this grain, there will be a 3 x 3 orthogonal rotation matrix g^{-1} which aligns the crystal lattice to the sample reference. If we take a specific plane normal, which will correspond to a vector $\underline{G}_c = \mathbf{B}\underline{G}_{hkl}$, the vector \underline{G}_c will become in the sample reference system: $\underline{G}_s = g^{-1}\underline{G}_c$.

A given laboratory Cartesian reference system will be identified by the real space vectors \underline{x}_l , \underline{y}_l and \underline{z}_l . Assuming the sample reference system to have origin in the origin of the laboratory coordinate system, and to have rotation axis along the vector \underline{z} , the two systems will be related by the unitary rotation matrix Ω_ω which depends on the rotation angle ω . The sample reference system will coincide with the laboratory reference system when ω is 0.

So we can finally express the relationship between a scattering vector in the laboratory system \underline{h} , which is the direction of the diffracted beam, and the Bragg node of indexes (h, k, l) in the reciprocal space, at the given rotation angle ω :

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

c) *Projections of the Crystallographic Domains:* Assuming that each grain is a perfect lattice, we will observe a projection of the grain volume on the detector for those scattering vectors \underline{h}_{lab} , which give rise to a different beam intersecting the detector (fig. 1). In the simplest case of




Fig. 1. Diffraction of one grain in a DCT experiment

orientation uniformity inside the grain, DCT is able to retrieve the three-dimensional shape of the grains in a polycrystal, by using oblique angle tomography reconstructions, which rely on Algebraic Reconstruction Techniques like the SIRT algorithm [4].

If we allow for regions in the grain to have different orientations from the average, the problem complicates heavily, because the projection geometry is no longer constant across the grain.

The two cases we will take as example in this article are: grains that exhibit a smooth variation of orientation, contained in a small bounded region in orientation space, and grains where sub-domains with a neat change in the orientation are formed, but there exists a notable symmetry operation between the orientation in the grains and the one in their sub-grains, called crystallographic “twins”.

B. Previous work

The problem of determining the local orientation of a grain has already been heavily studied and many approaches have been proposed, both in terms of models and algorithms. Prominent examples are the forward modelling idea from [5], and *GrainSweeper* from S. Schmidt discussed in [6]. What they all have in common is that they try to assign orientations to the voxels, but each voxel is treated independently. Another very interesting approach is introduced in [7] and further discussed in [8]. This other approach is somehow closer to our model, but tries to solve the mathematical problems exposed in this section, using Discrete Tomography, with Gibbs priors tailored for Grain Maps.

II. MODEL

Here we try to perform a global optimization, where all the voxels are processed simultaneously, using recent algorithms and techniques from signal theory. We also add the constrain of

allowing only a discrete and restricted number of orientations. This results in a sampling of the 3D orientation space much similar to the common sampling of the 3D real space implicitly employed by traditional tomography, making this method a 6-dimensional method.

A. Basic Mathematics

The aim of this article is to show that within certain limits, for DCT measurements, it is possible to identify a discrete set of allowed orientations in the reconstruction.

The actual number of orientations to be taken into account is a compromise between two conflicting requirements: fidelity to the actual orientation distribution on one hand, and indeterminacy of the linear system to be solved in the reconstruction process. The reason lies in the fact that the measurements (diffraction spots) are usually highly convoluted, because in a three-dimensional volume there will be a discrete number of voxels projecting to the same pixel on the detector, and so, if we take too many orientations into account, the amount of information from the spots is not enough for the number of unknowns we allow for.

To explain better this statement we should think that from the reconstruction point of view, the orientation space is a three-dimensional space and the real space is another three-dimensional space. So the full characterization of a grain should be carried out in 6D [9].

If we model the forward-projection operator in tomography as a matrix A , the solution to a reconstruction problem can be seen as the solution to a linear system $A\underline{x} = \underline{b}$, where we call \underline{x}^* a vector that makes the equality true. We know from linear algebra theory that for a guaranteed unique solution, the matrix A should be square and of full rank (otherwise uniqueness/existence depends on the particular vector \underline{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 $O(n^3 p^3)$. If we now consider that a typical diffraction spot has $O(n^2)$ pixels, we would then need $O(np^3)$ diffraction spots to fulfill the requirements for a unique solution. However the number of spots that we can actually successfully record on the detector, and use in the analysis is usually between $10 \sim 100$.

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

$$\underline{x}^* = \min_{\underline{x}} \|A\underline{x} - \underline{b}\|_2 \quad (2)$$

where from the feasible solution space we select the vector \underline{x}^* which minimizes the l_2 norm of the residual. Unfortunately the solution to (2) is not unique if A does not have full rank. In fact every solution of $A\underline{x} = \underline{b}$ is also solution to (2).

Thanks to the advances in the signal theory, it is possible to regularize this problem with the result of selecting solutions with specific properties. The regularization used in these examples is the Lasso formulation over the Haar transform [10] of the volumes:

$$\underline{x}^* = \min_{\underline{x}} \|A\underline{x} - \underline{b}\|_2 + \lambda \|H\underline{x}\|_1 \quad (3)$$

where H is the Haar transform and λ a weight. This formulation tries to assign a penalty to solutions that exhibit a large l_1

norm in the Haar domain, while the Haar transform is nothing else than a wavelet transform, on the Haar basis set.

The penalty is then weighted by the parameter λ which balances the data fidelity term $\|A\underline{x} - \underline{b}\|_2$ and the regularization term $\|H\underline{x}\|_1$. The parameter λ can hardly be related to some physical quantity, and remains for now a purely mathematical term. On the other hand it is also true that each value λ in the problem (3) corresponds to a specific value ϵ in the problem:

$$\underline{x}^* = \min_{\underline{x}} \|H\underline{x}\|_1 \quad (4a)$$

$$\text{subject to: } \|A\underline{x} - \underline{b}\|_2 \leq \epsilon \quad (4b)$$

but there is no easy and reliable way of computing one, when the other is known.

While (4) is more physical and more interesting to the scientific community, algorithms to solve (3) are generally simpler, more efficient and more effective.

In this work we slightly modify (3) to introduce a non negativity constraint for the reconstructed volumes:

$$\underline{x}^* = \min_{\underline{x}} \|A\underline{x} - \underline{b}\|_2 + \lambda \|H\underline{x}\|_1 \quad (5a)$$

$$\text{subject to: } \underline{x} \geq 0 \quad (5b)$$

which 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 can never be negative. On the other side, imposing the non-negativity constraint, makes the reconstruction algorithm converge quicker and give better results.

Finally we would like to briefly justify the choice of the l_1 -minimization over the Haar domain. It was demonstrated in [11] and [12] 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 [13].

So even if all the mathematical requirements for sparse recovery (see e.g. [11]) are difficult to be met for our physical case, where randomness of measurements is not an option, the use of sparsity promotion techniques can still be very useful if coupled with tools like the Haar transform. In fact, the choice of the Haar transform was driven by its well known property of being able to provide with a sparse or at least very compressible representation of transformed objects [10].

Moreover, white noise tends to never have a sparse representation in the Haar domain, because it always overlaps with the regions probed by the wavelets in frequency space. This means that the choice of (3) or (5) will also help in reducing the noise.

B. Algorithm

An implementation of the FISTA algorithm was used to reconstruct the examples [14]. FISTA is a strongly convergent first order algorithm, oriented to the solution of the Lasso problem. It is composed by three important pieces: a gradient iteration, a soft-thresholding phase and finally a memory step that tries to regularize the convergence of the algorithm.

To understand how these steps can be easily performed in our 6-dimensional problem, it is first important to describe the data-structures used by the algorithm.

We have two main data-structures that hold the information:

the *diffstack* and the *solution vector*. The *diffstack* contains all the information recorded by the detector, while the *solution vector* contains the 6-dimensional volume in orientation and real space. These two data structures are linked by a third object, generically named *geometry*.

a) *Solution vector*: The inner structure of the *solution vector* is rather simple: it is a linear array of three-dimensional volumes, where each of these volumes has the same size and represents a specific point in orientation space (fig. 2). This means that for the same three-dimensional volume in real space, we allow only specific orientations in orientation space and each of these orientations is represented by a single volume in the *solution vector*.




Fig. 2. A three-dimensional volume in real space, can be considered a single point in a three-dimensional orientation space

b) *Diffstack*: The *diffstack* needs to be covered with more care than the other objects involved. Depending on the intrinsic nature of the grains we are trying to reconstruct, the diffraction data recorded by the detector can be essentially of two types: two-dimensional or three-dimensional. The difference between the two cases is respectively the absence or presence of smoothness in the variation between one crystallographic sub-domain and the others. If the sub-domains have piece-wise constant orientation with sharp boundaries, they will give rise to different *Diffraction spots* (two-dimensional projections of the three-dimensional volume of the grain) on the detector at different rotation angles ω of the sample. If, instead, there is a smooth variation in orientation all along the grain, there will be a three-dimensional volume for each different event, and this volume in ω space is called *Diffraction Blob*.

While the case of two-dimensional parallel projections is rather simple to be dealt with, because it only requires to be careful enough to relate the correct spots to the correct orientations, the three-dimensional case can be handled in different ways.

The so called “Omega spread” of the *Diffraction Blob* is surely a source of information, but if we don’t provide enough sampling points in orientation space (different volumes in the *solution vector*), we might have the risk of not covering completely the ω space. This would result in inconsistencies in the linear system. Moreover it is reasonable to think that for a given orientation, which was not sampled in the *solution vector*, the closest sampled points in orientation space will be approximating it, if the information about it (namely the intensity projected on the detector), can be reached by these

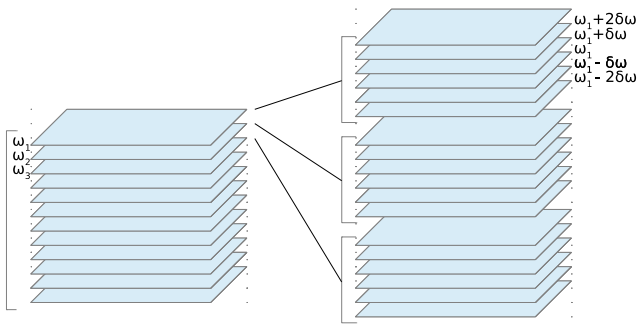


Fig. 3. On the left we can see a diffstack made of Diffraction Spots, while on the right, a diffstack made of Diffraction Blobs, where around each average ω we have multiple spots, with slightly different ω , that compose Diffraction Blobs

nearby points in orientation space.

In case of limited orientation spread inside a grain, it can be beneficial to “squash” the *Diffraction Blob* into a *Diffraction Spot*, by summing the blob along the ω direction. The result is a regular diffraction spot, where the more detailed angular information is lost. While this is a really crude procedure, in some simple cases it was observed in simulations that it still provides diffraction spots which can be used to reconstruct relatively good approximations of the expected solution.

In the simple case where only *Diffraction Spots* are used, the *diffstack* is nothing else than a three-dimensional volume where each spot is one slice of the volume. For more complicated cases, like the case where we want to consider the angular information contained within the blobs, the *diffstack* becomes a collection of *Diffraction Blobs* (fig. 3).

c) Geometry: The *geometry* is a collection of tables that fully describe the three dimensional scattering geometries for each of the sampling points in orientation space. In practice each line of each table will describe how the related volume will project on a specific spot into the *diffstack*.

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 represent each 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 on each of the slices of the *diffstack*.

Having described the data-structures we will now understand how these objects interact to perform the matrix operations described in the FISTA algorithm. Given the projection matrix A in (3), it is equivalent to the tomographic forward-projection of the volume on the detector in our model. On the other side, the back-projection is equivalent to the transpose A^T .

Once we know this, in the gradient iteration of FISTA:

$$\underline{d}_k = \underline{y}_k - \frac{1}{L} A^T (A \underline{y}_k - \underline{b}) \quad (6)$$

computing $A \underline{x}_k$ is nothing else 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 A^T of the detector residual will simply mean back-projecting the detector residual in each of the volumes that form the *solution vector*. The soft-thresholding step happens in the wavelet domain, so performing the operation:

$$\underline{x}_k = H^T \mathcal{T}_{\lambda t} (H \underline{d}_k) \quad (7)$$

is equivalent to computing the wavelet transform of each back-projected volume, performing the soft-thresholding ($\mathcal{T}_{\lambda t}$), and finally anti-transforming each volume again.

Finally the memory step:

$$\underline{y}_{k+1} = \underline{x}_k + \left(\frac{t_k - 1}{t_{k+1}} \right) (\underline{x}_k - \underline{x}_{k-1}) \quad (8)$$

is simply a sequential operation over all the volumes in the *solution vector*, where t_k is a purely mathematical parameter used to quench the memory step in the first few iterations, and defined in [14].

In our implementation the Haar transform is a simple function in matlab language, while 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 [15].

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 considering to implement asynchronous transfers, in order to hide transfer latencies behind the most time consuming operations.

III. RESULTS

We will now consider two important physical cases and analyse the results on synthetic data.

A. Twins

Twins are crystallographic sub-domains where the crystal undergoes a modification that can be represented as a rotation along a well defined direction of the crystal lattice. The result is that some of the scattering vectors will be the same as the ones from the parent domain, and some others will be changed. If we assume these sub-regions to not be further distorted, the number of sampling points in orientation space to be able to reconstruct all these sub-domains at the same time, becomes exactly the number of these smaller regions inside of the grain. We are going to assume the correct indexing of the spots was performed and show synthetic data reconstructions. For the realistic case where just one twin is formed, restricting the number of orientations to just two is perfectly physical and vastly reduces the number of needed measurements. Synthetic data was generated using a known grain shape, determined from phase contrast data [1]. The grain volume was then divided in two regions with different size and a different,

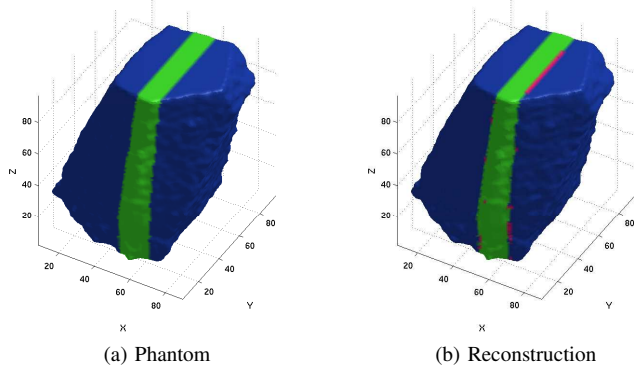


Fig. 4. Twinned Grain

but symmetry related, orientation was assigned to the regions. One of the two orientations, the one assigned to the smaller region, was chosen to be one of the possible twin variants of the other. The initial grain shape can be seen in figure (4a), where we can actually distinguish the two crystallographic domains by their color: the blue part is the parent, while the thin green region is the twinned part of the crystal. Since parent and twin usually share at least four reflections, given by the plane normal to the axis of rotation, in our example the *diffstack* generated was formed by 44 projections, 4 of the spots were in common between parent and twin, so that each of the domains was projecting to 24 spots.

As can be seen in figure (4b), the reconstruction almost completely recovers the shape of the two crystallographic domains. There are small conflicts at the interface (pink pixels) but they only represent very small deviations from the original solution.

B. Smooth gradient deformation

In some materials, grains don't "twin" when subject to mechanical stress, but give rise to smoothly varying misoriented regions of the crystal lattice. This phenomenon affects the quality of the *diffraction spots*, giving rise to distorted projections. The diffraction geometry is no longer constant




Fig. 5. Geometry of two-dimensional illumination

through the grain volume, and essentially every voxel will have

one or more different orientations, which might slightly differ from the orientations of the close neighbours.

This makes it not possible any more to restrict the number of orientations to very few possibilities, but still makes it possible to sample the orientation space around an average orientation, using a discrete number of variations.

For our case study we preferred to build a simplified model that allows us to use *diffraction spots*, instead of *diffraction blobs*. The experimental case simulated in our example corresponds to a two-dimensional slice illumination of the grain, that produces deformed spots on the detector. The angular range of misorientation in the example is confined between ± 0.5 deg. Having a two-dimensional volume reduces the complexity of




Fig. 6. Two-dimensional orientation spread

the data to the point that the spots are not too convoluted, and can be used by the algorithm without worrying about the ω spread around the average diffraction vector.

As already mentioned, the algorithm remains untouched, with every orientation projecting to every *diffraction spot*, and the spots built by a voxel-wise sum of the *diffraction blobs* along the ω direction.

In figure (6) we can see a comparison between the theoretical orientation domains in the sample (6a) and the reconstruction performed using a regular grid in orientation space with 341 sampling points (6b). The sampling grid can be seen as a body-centered lattice, or as two different cubic lattices, of sizes $6 \times 6 \times 6$ and $5 \times 5 \times 5$, where each occupies the centers of the cubes formed by the other, and the smaller has its central lattice point in the origin of the orientation space. The edge of the cubes was chosen to be the standard deviation of the orientation distribution function of the sample.

The image (6a) was produced assigning to each sampled orientation a color from the *jet* colormap, and then coloring each pixel with the color of the closest sampled orientation in orientation space. The image (6b) was instead generated assigning to each pixel the color corresponding to the orientation with highest intensity in the corresponding voxels, using the same *jet* colormap as in (6a). In figure (7) we can see the angular deviation in degrees between the theoretical phantom and the reconstruction, which showed a maximum value of $\sim 1.5 \times 10^{-3}$ deg on the top right corner. In this reconstructed volume, the average value between the sampled orientations was weighted over the reconstructed intensities.

IV. DISCUSSION

Even though the results presented in the article seem con-




Fig. 7. Reconstruction Error in degrees

vincing on theoretical data, experimental data can be affected by deviations from the model that make the reconstructions look noisy or even fail.

For twins, an important role is played by the alignment of the geometry, which needs to be carefully fine tuned, and by the re-normalization of the spots' intensity, which is necessary because of the different scattering powers of different (h, k, l) reflections, self-absorption and other effects which are hard to take into account. As a matter of fact, different scattering conditions have different scattered intensities, while tomography assumes that the total intensity of all the spots should be the same.

The renormalization that takes place needs a reference which is arbitrary on a single grain basis, but the ratio between the single references is important for the twin case. Generally this reference depends on the total volume of the grains, but this information is not known in advance. In experimental data, we expect the preprocessing logic to be aware of this problem and find a renormalization technique that is able to estimate a good reference ratio between the parent and the twins.

For what concerns the two-dimensional orientation spread case, the approach followed in the article looks already enough for an approximate solution of the local orientation reconstruction problem. However, an approach that exploits *diffraction blobs* is expected to give better results.

V. CONCLUSIONS

From the reconstructions in section III we can confirm that the intuition of sampling the orientation space with a discrete number points can result in a good approximation of the local orientation in grains.

In the twin case, the needed orientations are determined *a priori*, and their identification is precise enough to make it possible to restrict them to the minimum. This results in a fast a quickly converging algorithm, which also seems to be very reliable on theoretical data.

In the other case, where the orientation spread was taken

into account, the reconstruction technique was very simple and not completely adapted to this complex physical case. As already mentioned in III-B, the results are showing good signs of progress and convergence towards a stable solution of the problem.

ACKNOWLEDGMENTS

KJB acknowledges financial support of the Netherlands Organization for Scientific Research (NWO), project nr. 639.072.005.

NV and WL acknowledge financial support of the French National Research Agency (ANR), project: ANR 2010 BLAN 0935.

The authors acknowledge COST Action MP1207 for networking support.

REFERENCES

- [1] W. Ludwig, P. Reischig, a. King, M. Herbig, E. M. Lauridsen, G. Johnson, T. J. Marrow, and J. Y. Buffière, "Three-dimensional grain mapping by x-ray diffraction contrast tomography and the use of Friedel pairs in diffraction data analysis." *The Review of scientific instruments*, vol. 80, no. 3, p. 033905, Mar. 2009.
- [2] P. Reischig, A. King, L. Nervo, N. Viganó, Y. Guilhem, W. J. Palenstijn, K. J. Batenburg, M. Preuss, and W. Ludwig, "Advances in X-ray diffraction contrast tomography: flexibility in the setup geometry and application to multiphase materials," *Journal of Applied Crystallography*, vol. 46, no. 2, pp. 297–311, Mar. 2013.
- [3] H. F. Poulsen, *Three-Dimensional X-Ray Diffraction Microscopy*, ser. Springer Tracts in Modern Physics. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, vol. 205.
- [4] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. IEEE Press, 1988.
- [5] R. M. Suter, D. Hennessy, C. Xiao, and U. Lienert, "Forward modeling method for microstructure reconstruction using x-ray diffraction microscopy: Single-crystal verification," *Review of Scientific Instruments*, vol. 77, no. 12, p. 123905, 2006.
- [6] H. F. Poulsen, "An introduction to three-dimensional X-ray diffraction microscopy," *Journal of Applied Crystallography*, vol. 45, no. 6, pp. 1084–1097, Oct. 2012.
- [7] A. Alpers, E. Knudsen, H. F. Poulsen, and G. Herman, "Resolving Ambiguities in Reconstructed Grain Maps using Discrete Tomography," *Electronic Notes in Discrete Mathematics*, vol. 20, pp. 419–437, Jul. 2005.
- [8] A. Alpers, H. F. Poulsen, E. Knudsen, and G. T. Herman, "A discrete tomography algorithm for improving the quality of three-dimensional X-ray diffraction grain maps," *Journal of Applied Crystallography*, vol. 39, no. 4, pp. 582–588, Jul. 2006.
- [9] H. F. Poulsen, "A six-dimensional approach to microtexture analysis," *Philosophical magazine*, vol. 83, no. 24, pp. 2761–2778, 2003.
- [10] A. K. Louis, P. Maass, and A. Rieder, *Wavelets: Theory and Applications*, ser. Pure and Applied Mathematics. Wiley, 1997.
- [11] E. Candes and J. Romberg, "Sparsity and incoherence in compressive sampling," *Inverse problems*, no. m, pp. 1–20, 2007.
- [12] S. S. Chen, D. L. Donoho, and M. a. Saunders, "Atomic Decomposition by Basis Pursuit," *SIAM Review*, vol. 43, no. 1, pp. 129–159, Jan. 2001.
- [13] S. Boyd and L. Vandenberghe, *Convex optimization*, 2004.
- [14] A. Beck and M. Teboulle, "A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, Jan. 2009.
- [15] W. J. Palenstijn, K. J. Batenburg, and J. Sijbers, "Performance improvements for iterative electron tomography reconstruction using graphics processing units (GPUs)," *Journal of Structural Biology*, vol. 176, no. 2, pp. 250–3, 2011. [Online]. Available: <http://code.google.com/p/astra-toolbox/>