

SDART: an algorithm for discrete tomography from noisy projections

F. Bleichrodt^{a,*}, F. Tabak^b, K.J. Batenburg^{a,c,d}

^a*Centrum Wiskunde & Informatica, Science Park 123, 1098 XG Amsterdam, The Netherlands*

^b*Delft University of Technology, Mekelweg 4, 2628 CD Delft, The Netherlands*

^c*Mathematical Institute, Leiden University, P.O. Box 9512, 2300 RA Leiden, The Netherlands*

^d*Minds-Vision Lab, University of Antwerp, Universiteitsplein 1, B-2610, Antwerp, Belgium*

Abstract

Computed tomography is a noninvasive technique for reconstructing an object from projection data. If the object consists of only a few materials, discrete tomography allows us to use prior knowledge of the gray values corresponding to these materials to improve the accuracy of the reconstruction. The Discrete Algebraic Reconstruction Technique (DART) is a reconstruction algorithm for discrete tomography. DART can result in accurate reconstructions, computed by iteratively refining the boundary of the object. However, this boundary update is not robust against noise and DART does not work well when confronted with high noise levels.

In this paper we propose a modified DART algorithm, which imposes a set of soft constraints on the pixel values. The soft constraints allow noise to be spread across the whole image domain, proportional to these constraints, rather than across boundaries. The results of our numerical experiments show that SDART yields more accurate reconstructions, compared to DART, if the signal-to-noise ratio is low.

Keywords: DART, discrete tomography, noise, prior knowledge, relaxation

1. Introduction

In tomographic imaging, a three dimensional object is reconstructed from a series of projection images that have been acquired over a range of angles. Tomography has a wide variety of applications, ranging from medical imaging to materials science [1, 2, 3, 4, 5]. In many of these applications, the object under

*Corresponding author, P.O. Box 94079, 1090 GB Amsterdam, Netherlands,
tel: +31(0)20 592 4162

Email addresses: F.Bleichrodt@cwi.nl (F. Bleichrodt), franktabak@gmail.com (F. Tabak), K.J.Batenburg@cwi.nl (K.J. Batenburg)

investigation consists of only a few different materials, each corresponding to a particular gray level in the reconstructed image. Therefore, the set of *gray values* in a reconstruction should be small and discrete. Most common reconstruction algorithms, such as the Filtered Back Projection or SART, produce a continuous range of gray values [6]. However, it has been shown that incorporating this set of admissible gray values as prior knowledge can lead to superior image quality in the reconstruction, especially if the set of projection images is small [7, 8, 9, 10]. This type of tomography is known as *discrete tomography*.

The Discrete Algebraic Reconstruction Technique (DART) is one such algorithm that exploits the discrete nature of the object. It assumes that the gray values corresponding to the different compositions of the object are known a priori [8]. If only the number of different gray values is known, and not their actual values, these gray values can be adaptively estimated during reconstruction by using PDM-DART [11]. DART is an iterative method, which aims to solve a system of linear equations that models the tomographic projection process. In each iteration, a reconstructed image is segmented, i.e. the gray values are thresholded to the nearest a priori known gray value. It is assumed that the interior regions of this segmentation are segmented with high accuracy and that most errors are located on the boundaries. The key idea behind DART is to reduce the system of equations in each iteration by fixing or removing these interior pixels/voxels (*i.e.* unknowns) from the equations. The governing equations in tomography are ill-conditioned and rank deficient. Due to this dimension reduction, the equation system becomes increasingly better determined. Nevertheless, removing unknowns assumes that the gray values of the corresponding pixels are correct. Only the remaining *free* pixels are iteratively refined. Therefore, the operation of fixing a pixel imposes a hard constraint on the solution of the equation system and can only be effective if the selection criterion for a pixel being fixed or free is sufficiently accurate.

In practice, we see that the interior regions of the segmentation initially contain many errors. However, since the boundaries evolve due to the update steps, these pixels will be corrected eventually and the algorithm can converge to the correct solution. A problem occurs when the projection data contain noise. Imposing hard constraints on non-boundary pixels leads to noise being spread mainly over boundary pixels. This leads to major errors in the update steps applied to the boundary pixels. As a result, edges will be less resolved in the reconstruction and convergence problems can arise.

In this paper, we propose an alternative to the hard constraints imposed in DART. We introduce a set of relaxation parameters that imposes soft constraints on the pixel values. The parameters penalize deviation from the current segmented value of a pixel. Subsequently, the relaxed system with soft constraints is solved and the parameters are updated based on the intermediate reconstruction and segmentation. The proposed method is called Soft DART, to indicate the use of soft constraints. By using a *penalty matrix*, flexibility is increased in comparison to DART. It enables us to impose confidence levels on individual pixels instead of indicating if a pixel is correct or not. The results of our simulation experiments suggest that for a suitably chosen set of

relaxation parameters and for datasets with low signal-to-noise ratios (SNR), SDART produces a reconstruction closer to the ground truth when compared to DART.

The outline of this paper is as follows: first we will briefly discuss the DART algorithm in Section 2 and show some of its limitations. In Section 3 we will introduce SDART: a new variant of DART that includes a soft constraint. Possible choices for selecting the soft constraints are discussed. In Section 4 we compare the behavior of DART and SDART to see the effect of the soft constraints. We also discuss how to select an important regularization parameter that is used in SDART. In Section 5 an overview is given of the experiments and the results are discussed. We conclude the paper in Section 6.

2. The original DART algorithm

In this section we will briefly introduce the notation and concepts of DART and summarize the main details of the algorithm.

2.1. Notation and concepts

The governing equations in tomography can be posed as a linear system, which models the tomographic projection process. The linear model is generic, but for simplicity we will focus on the reconstruction of a 2D slice of the object from 1D detector measurements. The generalization to three dimensions is straightforward. Throughout this paper we consider a parallel beam geometry, which is illustrated in Fig. 1. In our implementation the geometry can be easily changed to fan or cone beam geometry. In fact, one of the experiments from Section 5 is based on a cone beam dataset.




Figure 1: Schematic of the parallel beam geometry. The incident rays are parallel. The sources and detectors rotate around the object, indicated by the projection angle θ , such that a circular field of view (FOV) is formed.

Let $\mathbf{x} \in \mathbb{R}^N$ denote a vector containing the gray value of each pixel in the unknown object. The vector $\mathbf{p} \in \mathbb{R}^M$ contains the detector measurements. The 1D detector has D elements and K projections are available. The total number of line projections is therefore $M = KD$. We now introduce the projection operator $\mathbf{W} \in \mathbb{R}^{M \times N}$, which relates the object to its projections:

$$\mathbf{W}\mathbf{x} = \mathbf{p}. \quad (1)$$

Reconstruction methods aimed at solving Eq. (1) are referred to as *algebraic reconstruction methods*. Examples of such methods based on Kaczmarz' method are ART, SIRT or SART [6]. Since the system of equations is usually underdetermined, and in practice no solution exists due to noise, it is typically solved in a least square sense:

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \|\mathbf{W}\mathbf{x} - \mathbf{p}\|_2, \quad (2)$$

such that an object is found that matches with the observed data optimally. In discrete tomography, the small, discrete set of admissible gray values

$$R = \{\rho_1, \dots, \rho_l\}$$

is known a priori. We can include this prior knowledge as constraints in the optimization problem:

$$\underset{\mathbf{x} \in \{\rho_1, \dots, \rho_l\}^N}{\text{minimize}} \|\mathbf{W}\mathbf{x} - \mathbf{p}\|_2. \quad (3)$$

2.2. Algorithm details

DART combines a continuous algebraic reconstruction method (ARM) with a segmentation step and uses heuristics to improve on this segmentation. The ARM that is typically used is SIRT or SART [6]. In principle, any linear least squares solver is suitable. Also the Krylov subspace methods such as CGLS or LSQR [12].

The flowchart in Fig. 2a illustrates the algorithm and its computational steps. We will briefly summarize these:

1. An initial continuous reconstruction \mathbf{x}_c is computed using an algebraic reconstruction method.
2. The reconstruction is segmented by applying thresholding. All pixel values are rounded to the nearest gray value in the set R .
3. Those pixels that have at least one (out of eight) neighbor with a different gray value (called *boundary pixels*) are free. In addition, a random subset of image pixels is also selected to be free. The corresponding columns are removed from Eq. (1) and their projections are subtracted from the right-hand side.
4. The solution of the reduced system is refined by applying an ARM to the free pixels.

5. If a stop criterion is not met, the free (boundary) pixels are smoothed. The smoothing step is performed by means of a discrete convolution of a 3×3 kernel with the image. The middle pixel of the kernel is weighted by a smoothing factor b , the other pixels in the kernel are weighted by $(1 - b)/8$. Although the smoothing is not used in every implementation of DART [13], it is used in the original paper [8]. The process repeats from step 2.

The thresholding step rounds the pixel values of the reconstruction to the nearest a priori known gray value. We use the same notation for this operation as presented in [8]:

$$T(\mathbf{x}) : \mathbb{R}^N \mapsto R. \quad (4)$$

Note that the discrete nature of the gray values is not exploited in the ARM iterations. Instead, DART relies on the segmentation to produce solutions with discrete gray values. Nevertheless, the strength of DART is based on the observation that pixels in the interior of a homogeneous region are likely to be thresholded correctly [8]. This result can be found empirically. A possible explanation is that least squares methods in general reconstruct low-frequent components of the solution prior to the high frequencies. Since large homogeneous regions (including the background) are part of the low-frequent components of the image, the non-boundary pixels are better resolved compared to the boundary pixels. This idea is used to classify the segmented image \mathbf{x}_s into the sets of fixed pixels F and free pixels U .

Formally, the sets F and U are then defined as index sets:

$$F = \{i \mid x_i = x_{i+rn+q}, \text{ for all } q, r \in \{-1, 0, 1\}\}, \quad (5)$$

$$U = F^c, \quad (6)$$

where F^c denotes the mathematical complement of F . In addition, the set of free pixels U is combined with a random subset of pixels. Each pixel has a probability $(1 - p)$ to be included in the set of free pixels. The probability $0 < p \leq 1$ is referred to as *fix probability*. This random set of free pixels improves the reconstruction of “holes” in the object, which are typically not found easily.

Since pixels in the interior regions are likely to be correct, they are removed from the equation system Eq. (1) and subtracted from the right-hand side. To fix a pixel $i \in F$, we apply the following operation on the linear system:

$$\left(\begin{array}{c|ccc|ccc} & & & & & & & \\ & & & & & & & \\ \mathbf{w}_1 & \dots & \mathbf{w}_{i-1} & \mathbf{w}_{i+1} & \dots & \mathbf{w}_N & & \\ & & & & & & & \end{array} \right) \begin{pmatrix} x_1 \\ \vdots \\ x_{i-1} \\ x_{i+1} \\ \vdots \\ x_N \end{pmatrix} = \mathbf{p} - v_i \mathbf{w}_i, \quad (7)$$

where v_i is the segmented gray value of pixel i . Subsequently, other pixels in F are treated in an analogous way. This leads to a *reduced system* that has fewer




Figure 2: The flowcharts illustrate the DART and SDART algorithms. The SIRT update can be replaced by another ARM. The smoothing step in DART is used to suppress the effect of noise on the boundary update. Note that SDART includes a penalty matrix \mathbf{D} , that represents the soft constraints used in the subsequent reconstruction step. SDART does not include a smoothing step.

unknowns. The pixels in the free set U are refined by iterating an ARM on the reduced system.

This process is repeated in each DART iteration. The boundary pixels are determined from the complete image, not only from the free pixels corresponding to the reduced system. Therefore, the elimination of pixels in the fixed set always starts from the full system in Eq. (1). As a consequence, pixels that were previously fixed can be free in a consecutive DART iteration. Therefore, errors in the interior regions can be corrected in a later stage due to evolution of the boundaries.

Batenburg and Sijbers show that with increasing noise levels, the DART reconstructions have a large pixel error (*i.e.* the number of pixels that have a wrong gray value in the reconstruction compared to the ground truth) [8]. Since only boundary pixels are free, the noise has a large effect on the boundary update. To remedy this problem, the fix probability can be decreased such that noise is also spread over a large random subset of pixels. While this improves the accuracy of DART with noisy projection data to some extent, it does introduce heavy salt and pepper noise, as was also observed in [13].

3. Soft DART

The main contribution of this paper is to introduce a different approach to classify and improve incorrectly segmented pixels. Instead of fixing pixels and updating free pixels, we propose to solve a relaxed system (*i.e.* under soft constraints) as an alternative to the ARM update step. In this section we will discuss the main details of the new approach.

A flowchart of the new method is shown in Fig. 2b. SDART follows the same steps as DART, but it does not eliminate unknowns (pixel values) from Eq. (1).

We propose to introduce a soft constrained optimization problem. Let \mathbf{v} be the segmentation of the intermediate reconstruction and let $\mathbf{D} \in \mathbb{R}_+^{N \times N}$ be a

diagonal matrix with nonnegative real entries ($d_{ii} \geq 0$, it is referred to as *penalty matrix*). We then introduce the relaxed reconstruction problem:

$$\begin{aligned} \underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad & \left\| \begin{pmatrix} \mathbf{W} \\ \lambda \mathbf{D} \end{pmatrix} \mathbf{x} - \begin{pmatrix} \mathbf{p} \\ \lambda \mathbf{D} \mathbf{v} \end{pmatrix} \right\|_2^2 \equiv \\ & \underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \quad \|\mathbf{W}\mathbf{x} - \mathbf{p}\|_2^2 + \lambda^2 \|\mathbf{D}(\mathbf{x} - \mathbf{v})\|_2^2. \end{aligned} \quad (8)$$

In this setting, the diagonal matrix element d_{ii} gives a penalty to pixel i for deviating from its segmented value v_i . If d_{ii} is large, only small deviations are allowed, while the pixel can be considered “free” if $d_{ii} = 0$.

The system in Eq. (8) is solved by a linear least squares solver, *e.g.* a reconstruction method such as SIRT or a Krylov subspace method such as CGLS. The algorithm is started with initial guess $\mathbf{x}_0 = \mathbf{x}_c$, the reconstructed image (with continuous gray levels) that resulted from the previous SDART iteration. The parameter λ is introduced both for regularization, as well as to compensate the difference in scale of the two terms in the cost function in Eq. (8). Note that the term $\|\mathbf{W}\mathbf{x} - \mathbf{p}\|_2^2$ depends on the number of angles, whereas $\|\mathbf{D}(\mathbf{x} - \mathbf{v})\|_2^2$ does not. As a result, the scaling between the two terms is important and the value of λ needs to be adjusted accordingly.

The entries of \mathbf{D} depend on the current reconstruction, so the matrix \mathbf{D} needs to be updated at each SDART iteration. The main advantage of this approach, compared to using hard constraints, is that no pixel will be truly fixed. Therefore, noise in the projection data will be distributed over the entire image (proportional to \mathbf{D}). Moreover, the relaxation parameters d_{ii} can express a confidence level for the accuracy of pixel’s i gray value. The confidence level can be based on any error measure for the reconstruction we have. Due to the generality of Eq. (8) we can even choose a different reference image \mathbf{v} instead of the segmented reconstruction.

Naturally, the increased flexibility comes with a price. The system has gained N unknowns as well as N equations, making it more costly to solve Eq. (8). In addition, we lose the efficiency resulting from the removal of columns from Eq. (1). Instead, each SDART iteration will be as costly as the first. As will be explained in Section 3.2, an efficient implementation can still lead to satisfactory performance. The full algorithm is presented in pseudo code in Algorithm 1. Note that a stopping criterion is not included in the algorithm description. In general, the question when to stop an algorithm to obtain the best solution is a very difficult one. Therefore, a reasonable choice is to terminate the algorithm when the relative change in the solution is small. This can be achieved by using a fixed number of iterations. In the experiments section we use 30 to 50 iterations, which is enough for all the datasets we considered.

3.1. Selecting a penalty matrix

In our simulation experiments in Section 5, we consider two different penalty matrices, defined as follows:

3.1.1. DART criterion

For validation purposes, we introduce a penalty matrix that should result in SDART mimicking the original DART algorithm. SDART does not allow us to fix pixels, but instead we can give non-boundary pixels a very large weight. By giving a weight of zero to boundary pixels, we do not put any restrictions on those pixels. The resulting penalty matrix is given by

$$d_{ii} := \begin{cases} 10^6, & i \in F \\ 0, & i \in U. \end{cases} \quad (9)$$

These weights are found to be effective from preliminary simulation experiments. We refer to SDART using this penalty matrix as SDART-ORIG, the first variant of SDART.

3.1.2. Neighbor criterion

In DART, a pixel is fixed when all 8 neighbors have the same gray value. If at least one neighbor has a different gray value, the pixel is considered free. This leads to a relatively fat boundary. Therefore, a logical choice for \mathbf{D} would be to give a penalty that is proportional to the number of neighbors b_i that have a different gray value, *i.e.*,

$$b_i := \sum_{r=-1}^1 \sum_{q=-1}^1 \mathbb{1}_{\{x_i \neq x_{i+rn+q}\}}, \quad (10)$$

where $\mathbb{1}_{\{\cdot\}}$ is an indicator function that is 1 if the condition is true and 0 otherwise. In this way, boundary pixels have different weights that reflect their position in the boundary. As a result, the boundary can be considered to be narrower. The penalty matrix is then defined as

$$d_{ii} := \frac{100}{3^{b_i}}. \quad (11)$$

Note that d_{ii} is an exponential, monotonically decreasing function in b_i . The factor 3 in the denominator was chosen based on early simulation experiments. If the factor is too small or too large, the reconstruction will either not change much in each iteration or noise is distributed more on the boundary, respectively. This version of SDART is referred to as SDART-NB.

3.2. Solving the soft constrained system

In this section we will go into more detail how the soft constrained optimization problem in Eq. (8) is solved.

This optimization problem involves solving the system

$$\begin{pmatrix} \mathbf{W} \\ \lambda \mathbf{D} \end{pmatrix} \mathbf{x} = \begin{pmatrix} \mathbf{p} \\ \lambda \mathbf{D} \mathbf{v} \end{pmatrix} \quad (12)$$

in a least squares sense.

Algorithm 1 SDART

Input: Projection data \mathbf{p}

Output: Segmented reconstruction \mathbf{x}_s

1. Let \mathbf{x}_c be the initial CGLS reconstruction from projection data \mathbf{p} .
Compute the initial segmentation $\mathbf{x}_s = T(\mathbf{x}_c)$.

repeat

/ Setting up the soft constraints */*

2. Compute the matrix \mathbf{D} based on \mathbf{x}_s .

3. Set $\mathbf{v} = \mathbf{x}_s$.

4. Set $\mathbf{x}_0 = \mathbf{x}_c$.

/ Soft constrained reconstruction */*

5. Solve: minimize $\|\mathbf{W}\mathbf{x}_c - \mathbf{p}\|_2^2 + \lambda^2\|\mathbf{D}(\mathbf{x}_c - \mathbf{v})\|_2^2$
using CGLS with initial solution \mathbf{x}_0

/ segmentation */*

6. Compute $\mathbf{x}_s = T(\mathbf{x}_c)$.

until convergence

The matrix \mathbf{W} has M rows and N columns and is typically very large, especially in the three dimensional case, where the number of voxels/pixels in the reconstruction grid and in the projection data, is large. In Eq. (12) we added another N rows to the system matrix and right-hand side. However, in our implementation using the ASTRA toolbox [14, 15], the full matrix is never formed explicitly. If a matrix–vector product is computed, we split this operation in two parts: $\mathbf{W}\mathbf{x}$ and $\lambda\mathbf{D}\mathbf{x}$. The first matrix–vector product is computed by generating \mathbf{W} on the fly to avoid high memory usage. This can be done efficiently due to our GPU implementation. The second part $\lambda\mathbf{D}\mathbf{x}$ is simply an inner product, because \mathbf{D} is diagonal and is stored as a vector. Therefore, the computational overhead compared to solving Eq. (2) is small.

For solving Eq. (12) we can apply any linear least squares solver. However, we noticed during preliminary experiments that methods based on Kaczmarz’ method such as SIRT [6], have slow convergence and do not yield very accurate results. Krylov subspace methods perform better in this case. We found that the method CGLS performs very well and methods such as LSQR and LSMR are suitable too [12, 16], but they all have slightly different results. This is why we have chosen to combine CGLS with SDART in our numerical experiments in Section 5.

4. A numerical study

In this section we highlight differences in behavior of DART and SDART using numerical experiments. We also introduce an experimental way to compute the regularization parameter λ that has an important role in the convergence of SDART. We want to point out that this section serves the reader to illustrate the different behavior between DART and SDART. Therefore, the choice of our phantom shown in Fig. 3 is somewhat arbitrary.

4.1. Behavior of DART compared to SDART

The effect of the hard versus soft constraints of DART and SDART can be visualized by looking at the evolution of boundary pixels (*i.e.* free pixels in case of DART). Since SDART has no concept of free pixels, we show the penalty matrix \mathbf{D} after rescaling its values such that the elements are in the range $[0, 1]$, *i.e.* we show the image $\{x_i\}$ such that

$$x_i = 1 - \frac{d_{ii}}{\max_j(d_{jj})}. \quad (13)$$

In this example we consider the cylinder block phantom, shown in Fig. 3. The projection data were computed by forward projecting the image, using a parallel beam geometry. In total 25 projections were computed at equidistant angles in the domain $[0, \pi)$. The projection data were perturbed by Poisson noise. Noise due to a limited photon count, which is encountered in many types of tomography, follows a Poisson distribution. The intensity of the noise is

quantified by the photon count of the incident X-ray beam, when no object is between the source and detector. In other words, this represents the total dose that is emitted during the full scan of the object. In this case, noise was simulated corresponding to a photon count of 10^3 . In DART, the fix probability was set to 0.99. With very low signal-to-noise ratios, a lower fix probability is preferred, but this would make it difficult to show the boundary evolution.




Figure 3: The cylinder phantom of size 512×512 .

In Fig. 4 the boundary evolution of both DART and SDART-NB (using the neighbor criterion) for the first three full iterations is shown. Note that for SDART-NB, we show the weights represented in Eq. (13). A pixel is black if the corresponding penalty of the pixel is maximum, $\max_i d_{ii}$. This is comparable to a fixed pixel in DART. A white pixel corresponds to a minimum penalty. The pixel attains any other gray value if it is in between these extrema. This representation of the “amount of fixedness” of pixels is not directly comparable to DART’s free and fixed pixels. However, we think that these images give insight in the different ways that DART and SDART update the reconstruction.

The initial boundary of DART in Fig. 4a is only slightly refined in the next iterations. Although the boundary becomes thinner, many of the background pixels are indicated as free pixels. In SDART-NB, we see a similar thinning of the boundaries. Moreover, the contours of the ground truth image are approximated more accurately. Background pixels have a large weight (indicated by black pixels). This shows that the background is more homogeneous rather than that noise is producing clusters, as is the case in DART.

Images of the boundary evolution, do not give a clear insight in the quality of the reconstructions. Therefore, we also show the segmented intermediate reconstructions in Fig. 5. DART is distributing a significant part of the noise throughout the background, where many pixels are free. Another consequence of noise is visible in the jagged boundaries of the cylinder block. The reconstructions of SDART-NB have finer and more distinct boundaries. In addition, background noise is reduced within consecutive iterations.

The behavior shown in this example depends largely on the set of soft constraints imposed by the matrix \mathbf{D} . For example, a strategy of fixing pixels similar to DART (*e.g.* SDART-ORIG) is very effective for projection data without




Figure 4: Comparison of the boundary evolution of “free” pixels in DART, (a)-(c) and SDART-NB (d)-(f), for the first three iterations.




Figure 5: Comparison of the intermediate segmentations of the first three iterations of DART, (a)-(c) and SDART-NB, (d)-(f). These correspond to the boundaries in Fig. 4.

noise [8], while it fails in cases with heavy noise. Therefore, finding a single penalty matrix \mathbf{D} that is accurate in all possible datasets is unlikely. An adaptive approach might be more successful. For example, the order of magnitude of the weights can be changed according to noise levels. In case of low noise levels, high weights steer the solution more to the segmentation, while smaller weights prevent over-fitting to noise. We see an important role here for the parameter λ in Eq. (8). It can be used to assign a larger weight to the data fidelity term or to correspondence to the segmented solution.

4.2. Selecting the regularization parameter

In this section we will discuss how to select a value for the regularization parameter λ that is close to optimal, where we use the term *optimal reconstruction* to refer to the reconstruction with smallest pixel error over all possible choices for λ . Recall that the pixel error indicates the number of pixels that do not have the right gray value compared to the ground truth. We have chosen this error norm over, *e.g.*, a chi-2 or jaccard distance, since our images are inherently discrete. A chi-2 measure is more suitable when comparing two images that are continuous with respect to their pixel values. Moreover, we expect that our findings will not be changed significantly when another error norm is used.

Consider the second formulation of the cost function in Eq. (8). It consists of two terms: a data fidelity term $\|\mathbf{W}\mathbf{x} - \mathbf{p}\|_2^2$ and a discrete tomography prior $\|\mathbf{D}(\mathbf{x} - \mathbf{v})\|_2^2$. The order of magnitude of these terms is in general not directly comparable. Therefore, a regularization parameter λ is added to properly adjust the bias to the discrete tomography prior.

Note that the magnitude of the data fidelity term depends strongly on the current solution \mathbf{x} (and thus on the ground truth) as well as the number of projection angles. Adding more projection angles results in more rows in \mathbf{W} as well as more elements in \mathbf{p} . By making the assumption that a projection image does not change in a small angular range, we can assume that the sum of squared residuals $\|\mathbf{W}\mathbf{x} - \mathbf{p}\|_2^2$ behaves linearly in the number of angles. (Provided that the additional angles are close to the original angles.)

To verify this assumption, we perform a small numerical experiment. Considering the Shepp–Logan phantom, the bottom image in Fig. 10a, we simulate projection data for a fixed number of angles (*e.g.* 50). For these projection data we reconstruct the image by using 4 iterations of LSQR (as explained in the previous section). This yields an approximate reconstruction. Consequently we construct the projection operator \mathbf{W} and corresponding right-hand side \mathbf{p} (by forward projecting the ground truth image) for a varying number of equidistant projection angles. In Figure 6a the squared residual norm is plotted for a varying number of angles, indeed showing a linear curve.

The last term in Eq. (8) does not depend on the number of projection angles. Instead it depends linearly on the number of pixels in the reconstruction. Since this is also true, to some extent, for the data fidelity term, no adjustments should be necessary if the number of reconstruction pixels is changed (*e.g.* a value λ at low resolution can be found that is also suitable for high resolution reconstructions).




Figure 6: (a) Linear behavior of the data fidelity term $\|\mathbf{W}\mathbf{x} - \mathbf{p}\|_2^2$ with respect to the number of projection angles; (b) For simulated projection data without noise, the residual mean squared error of an SDART reconstruction follows the same curve as the true pixel error, if the regularization parameter is varied.

Due to the linearity of the data fidelity term, we can extrapolate λ if a value is known for a dataset with few projection angles.

We still lack a way of determining a good value for λ for a given dataset. For this goal we can use the residual 2-norm. If the projection data are consistent with the ground truth (no noise), there is usually a good correspondence between the true (pixel) error and the residual of Eq.(1). We can exploit this to find a value for λ for which the pixel error of the SDART reconstruction is minimal. In Fig. 6b we have plotted the pixel errors of the SDART-NB reconstructions and the value for λ that was used. The result suggests that there exists an optimal λ at which the pixel error takes its minimum value. Moreover, the residual 2-norm agrees with this minimum. Therefore, if an initial guess for λ is known, we can compute several SDART reconstructions for different λ in a small range. If a minimum of the residual is found, we also found the optimal λ .

Of course this method will fail if noise is present in the projection data. Then the correspondence between the residual 2-norm and the true error is in general very poor. Nevertheless, we will show that a value found for λ for a phantom dataset (or a dataset with very low noise levels) may also give near-optimal results when high noise levels are considered.

In Fig. 7a, a mesh plot is shown of the pixel error for a range of λ as well as several photon counts. For these computations the cylinder phantom in Fig. 3 was used and projections at 20 angles were computed. We see that the minimum pixel error is attained at $\lambda \approx 0.5$ for data with high signal-to-noise ratio (SNR). If the SNR is decreased, the optimal value for λ does not seem to change much.

The optimal λ for each SNR still varies slightly. The corresponding pixel errors are minimum at that specific SNR. We also selected a constant $\lambda^* \approx 1$ that attains pixel errors closest to these minimum pixel errors in a least square sense. This is the optimal choice if we fix λ . The pixel errors are shown in Fig. 7b. We also plotted these curves for $\lambda = 1.38$ and $\lambda = 0.24$ (which was




Figure 7: Choosing the optimal value for lambda. (a) The mesh plot of the pixel error; (b) The pixel errors for $\lambda = 0.24$ are very close to the pixel errors for the optimal λ^* .




Figure 8: Correspondence between the residual mean squared error (MSE) and the true pixel error for relatively high photon count ($= 10^6$).

optimal for the noiseless case as shown in Fig. 6b). From this result we can see that while λ^* is the better choice overall, choosing $\lambda = 0.24$, the same as in the noiseless case, produces near-optimal results. The key conclusion is that the value for λ as in the noiseless case also works well in low-dose datasets.

In Fig. 8, the pixel error and residual are plotted for a dataset with limited noise (a photon count of 10^6), for varying λ . From these data we see that there still is a good correspondence between the minimum pixel error and the minimum of the residual MSE. This implies that we can effectively estimate λ for a dataset with high SNR. Consequently, datasets collected from the same object with low SNR can use the same λ .

5. Experiments and results

In this section we describe the experiments that were performed and their results. The main focus is to compare DART with SDART and a variety of well

established reconstruction methods. In the experiments we focus on different conditions for which SDART is more accurate in comparison with DART. In addition, we explore the two sets of soft constraints \mathbf{D} used in SDART-ORIG and SDART-NB. To validate our findings in the simulation experiments, we apply (S)DART to an experimental dataset with a low signal-to-noise ratio.

The results are compared to other well-known reconstruction methods. We include as a reference the filtered back projection (FBP) method, which is widely used for its speed and robustness[6]. However, a requirement for FBP is a large amount of angles, to avoid under-sampling, a requisite that is not met in the limited angle case. The algebraic reconstruction techniques SIRT and SART are included as well. Both methods are derived from ART, but their update steps are different. SART will update the reconstruction using one projection image at a time, while SIRT updates the reconstruction using all angles. Finally, we include the binary algebraic reconstruction technique BART, which is a method for reconstructing binary images [17]. This algorithm is implemented in the tomography package SNARK09 [18], which we use in combination with the script as presented in [13]. This code follows the same computational steps as presented in the flowchart of the original paper in [17].

In the simulation experiments we consider three phantom datasets, shown in Fig. 10a, referred to as: *blob*, *cylinders* and *Shepp-Logan* (from top to bottom). The phantom images are of size 512×512 . Projections were computed over an equidistant set of angles in the range of $[0, \pi)$. A parallel beam geometry is simulated with a 1D detector of 512 pixels, the same width as the phantom images. When noise is applied to projection data, it is sampled from a Poisson distribution. Intensity of the noise is expressed in simulated photon counts, a lower photon count indicating more noise.

Each DART run is initialized by 40 iterations of the SIRT reconstruction algorithm [6]. DART uses 40 intermediate SIRT iterations to refine the boundary pixels. SDART is initialized by 40 CGLS iterations and uses 70 iterations to solve the soft constrained system in Eq. (8). SDART uses more intermediate iterations, since SDART solves the full equation system in Eq. (8), while DART solves a system of equations that has significantly fewer unknowns. Therefore, increasing the intermediate ARM iterations in DART will not significantly decrease the pixel error. In fact, for noisy projection data it is preferred to decrease the number of ARM iterations, to prevent overfitting to noise.

5.1. Experiment I – basic validation

In the first experiment we compare simulations of DART, the two variants of SDART, BART, SART, FBP and SIRT on noiseless data. The projection data are simulated using the forward model. This allows us to compare the reconstructions with the ground truth.

Note that the BART algorithm can only be applied on datasets from binary images, which rules out the Shepp-Logan phantom. The FBP, SART and SIRT methods do not provide a segmented image, so we cannot directly measure the pixel error. Therefore, we include a final segmentation step after the reconstruc-




Figure 9: The graphs show the convergence of DART and SDART compared to other common reconstruction algorithms.

tion. Similarly to the segmentation step in DART and SDART, the images are thresholded.

The number of projection angles was chosen as follows: 10 for the blob phantom, 25 for the cylinder phantom and 30 for the Shepp-Logan phantom. Using these number of projection angles, accurate reconstructions can be obtained in all three cases.

From the results of these experiments, shown in Fig. 9, we can conclude the following. In all cases, the filtered back projection is not accurate in terms of the pixel error. Clearly, the method is not suitable for the limited angle case. The methods SIRT and SART are comparable in accuracy, although SIRT converges slower to the minimum pixel error. For the Shepp-Logan phantom, the number of angles is clearly not enough to properly reconstruct, without using prior information. The BART algorithm achieves almost similar pixel errors when compared to DART and SDART. As we have discussed previously, BART could not be applied to the Shepp-Logan phantom.

Next, we focus on the differences between DART and SDART. SDART-NB with the neighbor constraints is in each case more accurate than SDART-ORIG. The accuracy of SDART-NB is comparable to DART in all datasets except for the blob image with hole. Since DART uses a fix probability of 0.99, it is able to detect the hole in the phantom. SDART has no random subset of free pixels and therefore is unable to detect the hole. This random subset could, however, easily be added to SDART to find holes such as in this case.

Based on the results in this experiment, we have decided not to include FBP and SART in the remaining experiments. For FBP, the number of angles is not enough to achieve a small pixel error. The SART algorithm achieves a similar pixel error compared to SIRT, but converges faster. However, we prefer SIRT over SART, because in each iteration, a back projection from all angles is performed. This way, noise in the data is averaged over all projection angles, instead of only one. Therefore, SIRT is expected to perform better in the case of high noise levels. Moreover, SIRT and SART are from the same family of algorithms and are expected to have similar results.




Figure 10: (a) Phantom images of dimension 512×512 pixels used in simulation experiments, referred to as: blob, cylinders and Shepp–Logan (from top to bottom); (b) The segmented SIRT reconstructions. The pixel errors corresponding from top to bottom reconstructions are: 27.6%, 18.2% and 41.9%; (c) The final BART reconstruction for blob and cylinders phantoms. The corresponding pixel errors, from top to bottom: 18.4% and 15.9%; (d) The final DART reconstructions. The pixel errors are: 17.3%, 13.7% and 48.1%; (e) The final SDART reconstructions. The pixel errors are: 3.9%, 7.7% and 39.9%.

5.2. Experiment II – the effect of noise

In the second experiment, we compare the performance of DART, SDART and BART over a large range of noise levels. As has already been indicated, DART results in poor reconstruction accuracy if the signal-to-noise ratio is low [8]. We expect that SDART will be more robust in this case.

We varied the noise levels from a photon count of 10^2 (very high noise level) to 10^6 (very low amount of noise). For each noise level a new sinogram was generated and Poisson noise was applied. The number of projection angles, 10, 25 and 30 for the blob, cylinder and Shepp–Logan phantom images respectively, was chosen such that a good reconstruction quality is possible when no noise is present in the projection data.

As the start solution for DART, a segmented SIRT reconstruction is used based on 40 iterations. The final (S)DART reconstruction should be an improvement of this initial reconstruction. Therefore, we also compare the results with this initial segmentation.

The results from experiment II are listed in Fig. 11. The errors, in percentages, indicate the percentage of pixels that have been segmented to the wrong

gray value compared to the phantom image.

The pixel error of the (S)DART reconstructions are in general smaller than those of the initial segmentation. However, in a small interval of noise levels, DART shows a small regression in comparison with the initial segmentation. Apparently DART has problems in convergence for this particular interval of noise. This issue is especially visible in the case for the blob phantom for photon counts in between $10^3 - 10^4$.

The BART algorithm was also applied to the blob phantom and the cylinder phantom. It is performing very similar to the DART algorithm with a fix probability of 0.99. This leads us to the conclusion that BART is suitable for reconstructing binary objects, if the noise level is not too high.

For the blob phantom at high SNRs (photon counts in the range $10^4 - 10^6$), it seems that 10 projection angles is enough to have very accurate reconstructions, just by thresholding of the initial SIRT reconstruction. This is different, however, for high noise levels. At a photon count of 10^2 we see that SDART produces an accurate reconstruction. However, the pixel error of DART and BART is increasing very rapidly if the photon count becomes smaller than 10^4 . The same trend, to some degree, can be observed in the other two phantom images. At high SNR, DART and SDART-NB perform equally well. SDART-ORIG, using the DART criterion, is performing badly. Instead of improving the initial segmented reconstruction, the pixel error is increased. We can conclude that SDART-ORIG is not suitable for noisy datasets.

The pixel errors do not give clear insight of the actual quality of the reconstructions. A pixel error of, say, 10% does not indicate where the incorrect pixels are located and how the corresponding reconstruction looks. Therefore, we also show the reconstructed images in Fig. 10b - 10e. The blob phantom projection data has corresponding photon count of 100, the cylinder data 500 and for the Shepp-Logan data the photon count was 1000. In the fourth column of Fig. 10d we see DART reconstructions with fix probability 0.99. The SDART-NB reconstructions are shown in the last column in Fig. 10e. The difference between the reconstructions is clear, especially in the background. Moreover, the SDART reconstructions suffer far less from salt and pepper noise that is clearly visible in the DART reconstructions. The qualitative differences are supported by the corresponding pixel errors as listed in the caption of Fig. 10. The quality of the Shepp-Logan reconstruction is less impressive. Presumably, the number of gray values (6 in total) in the reconstruction is too much. The strength of the prior is reduced if the total number of distinct gray values is large.

From these data it becomes clear that employing SDART has an advantage over DART and BART if the signal-to-noise ratio is very low. In some cases SDART can generate good to reasonable reconstructions, while DART and BART perform badly.

5.3. Experiment III – adding more projection angles

Datasets with very low signal-to-noise ratios and few projection angles, in general, result in reconstructions with poor quality. Applying discrete tomography algorithms such as DART might show slight improvements in accuracy,




Figure 11: The pixel error is expressed in terms of the noise level (in photon count).

but it is known that DART has problems with noisy datasets, which was also observed by Alpers *et al.* [13]. We have seen from the previous experiment that SDART-NB is favorable over DART in this case (from now on we do not include SDART-ORIG in the experiments). It is not clear, however, if this benefit is maintained when the number of projection angles is increased.

In the third experiment we compare the accuracy of DART, SDART and BART on the phantom datasets for a fixed, low signal-to-noise ratio, but the number of projection angles is varied. For the blob and cylinder phantom we chose a photon count of 10^2 . For the Shepp-Logan phantom a photon count of 10^3 was used. Other details are the same as in experiment I. Since the performance of DART on noisy datasets depends strongly on the fix probability, we run DART for two values of the fix probability.

From the results in Fig. 12 we see that increasing the number of projection angles decreases the pixel error for DART, BART and the segmented SIRT reconstruction. Especially for the cylinder phantom, a large decrease in the pixel error can be achieved. In each case, SDART is still more accurate in comparison to DART and BART. However, there are two surprising results shown in this plot. First of all, DART with a fix probability of 0.5 is less accurate than DART with a fix probability of 0.99. From Fig. 11, we see that this is indeed the case for most noise levels. Apparently this behavior changes if the noise level is very high. Secondly, the pixel errors of SDART are not decreasing monotonically with an increasing number of projection angles. In fact, the pixel error is increasing for the blob phantom. The weights of SDART-NB were determined for a dataset at a fixed number of projections. Therefore, the number of projections is not incorporated in the weights. These results suggest that such an approach should be taken in order to avoid this behavior. Since the focus of this paper is on discrete tomography for the limited angle case, we leave this open for further research. We want to emphasize that the weights defined in Eq. (10) do work well in the limited angle case, which is the domain of discrete tomography.

The BART algorithm performs more consistently as the pixel error drops when the number of projection angles is increased. For the Shepp-Logan phantom, the pixel error of SDART-NB is more or less constant. We expect that




Figure 12: Simulation experiment at high noise levels and an increasing number of projection angles.

there is still significant room for improvement of the performance of SDART for cases where a relatively large number of projection angles are available, by further exploring the possible choices for the matrix D .

5.4. Experiment IV

To validate the simulation experiments, we applied SDART to an experimental dataset. For this experiment we used a hardware phantom. The hardware phantom consists of an aluminum-copper alloy and has a nearly convex shape with three holes drilled through it in the vertical direction. A SIRT reconstruction of the central slice is shown in Fig. 13a. One of the holes was filled with water.

The projection data were acquired by a Skyscan 1172 microtomography X-ray scanner. A total of 600 projection images of 1000×600 pixels were acquired over a full tilt-range of 360° . The tilt increment between images was 0.6° . With a pixel size of $25 \mu\text{m}$, a slice (1000×1000 pixels) has physical dimensions of $25 \times 25 \text{ mm}$.

The scanner has a cone beam geometry, however, since the object is uniform in the vertical direction, we focus on reconstructing the central slice using fan beam geometry. The sinogram was extracted from the projection images.

The SIRT reconstruction from all 600 projections leads to an accurate segmentation of the object (by thresholding). In Fig. 13b this segmentation is shown. Manual adjustments were applied to the segmentation at the boundaries, where the segmentation was distorted by noise. The gray values for this dataset were estimated using the algorithm proposed by Batenburg *et al.* [19]. The segmentation can be used as a ground truth. Although it will differ from the actual ground truth, we can assume that it is reasonably accurate to allow also quantitative analysis of the accuracy of SDART.

For the (S)DART reconstructions we use a subset of 20 projections, with equiangularly distributed projection angles in the range $[0, \pi)$. The value of $\lambda = 2$ in SDART is computed based on the residual norm using the full dataset, as described in section 4.2. The number of intermediate ARM iterations in




Figure 13: (a) A SIRT reconstruction from all 600 projection images; (b) Segmentation of the SIRT reconstruction. The segmentation was manually improved on the edges.




Figure 14: The dataset considered here consists of a subset of 20 projections (from 600 total); (a) The final SIRT reconstruction; (b) The final DART reconstruction; (c) The final SDART reconstruction;

DART is 20, and the number of DART iterations is 300. For SDART-NB, we use 70 intermediate iterations and a total of 30 SDART iterations. The difference in iterations between (S)DART is large, because DART converges slowly in terms of iterations, but iterations are fast. SDART converges quickly, but iterations are far more costly.

In Fig. 14 we see the final SIRT, DART and SDART reconstructions. The SIRT reconstruction is very noisy. When compared to the DART and SDART reconstructions the advantage of using prior knowledge is obvious. Although the DART and SDART reconstructions look similar, we can also point out clear differences. In the DART reconstruction, many gray areas are visible (water) in the whiter area (the metal). These form large clusters of pixels. SDART shows these clusters as well, but they are very small in comparison, just a few pixels. The result is that SDART's boundaries are sharper, but they look distorted by salt and pepper noise. DART's boundaries are smoother, but the presence of gray areas might lead to the erroneous conclusion that there is actually water there.




Figure 15: The pixel error of DART and SDART are plotted against the computation time. Note that DART requires far more iterations, but SDART takes three times as long to converge.

In Fig. 15 we plotted the pixel error versus the computation time. Note that we do not have the actual pixel error, but we compare the reconstructions with the segmentation from Fig. 13b. The DART reconstruction has a pixel error of 2.68% and the pixel error of SDART is 1.74%. This is a significant decrease in the pixel error of 35%. In terms of computation time we see that DART finishes in one third of the time it takes SDART to converge although it requires substantially more iterations. We should note that the DART implementation used here was optimized. It should be possible to speed up SDART by moving more computations, such as the detection of the boundaries, to the GPU. We want to emphasize that we introduce SDART here as a method that is more robust than DART on noisy data. Therefore we did not focus on computational efficiency, but rather on accuracy. From the results in Fig. 14 we see that the SDART reconstruction provides significantly different reconstructions. These may help in better understanding the true morphological nature of the object.

6. Conclusions

We proposed a new variant of DART that introduces a set of soft constraints to replace the hard constraints. We have seen that the hard constraints in DART lead to problems if the projection data contain a high level of noise. Our new method, named SDART, was introduced to enhance the robustness of DART for noisy projection data. The soft constraints allow noise to be spread across the whole image domain. As a result, boundaries of the object are less influenced by the noise, leading to sharper edges.

Two sets of soft constraints, or penalty matrices, were introduced. The first variant, SDART-ORIG, mimics the original DART algorithm. The other variant, called SDART-NB, discriminates boundary pixels by the number of surrounding pixels with a different gray value.

We performed several simulation experiments that compare the accuracy of SDART with DART, BART and SIRT. The results from noiseless data show that

SDART has similar, but slightly less accuracy compared to DART. The accuracy of SDART-NB compared to BART is slightly improved. On datasets with very low signal-to-noise ratios, SDART-NB outperforms DART and BART by large. The results of SDART-ORIG were not accurate in this case and SDART-NB is the preferred method for noisy projection data. The qualitative results show that SDART-NB is less prone to salt and pepper noise. Results from experimental data, containing a large amount of noise, further support that SDART-NB is more accurate compared to DART. For this particular dataset, the pixel error of the SDART-NB reconstruction was approximately 35% smaller than that of DART. In this case, the difference in quality between the DART and SDART-NB reconstructions was less obvious visually. The SDART-NB reconstruction has sharper edges, distorted by some salt and pepper noise. DART produces clusters of a gray value on the edges that is different from the metal interior. This might lead to the false conclusion that water adhered to these edges. From the SDART reconstruction it is clear that this is not the case. This shows that the SDART reconstruction can give additional insight, when conclusions about the DART reconstruction are not decisive.

So far, we have investigated only two possible choices for the penalty matrix D . Compared to DART, SDART can encode a more specific representation of the prior by using continuous weights. We expect that SDART can be further improved by using more sophisticated choices for this matrix, which will be investigated in future research.

Acknowledgments

This research was supported by the Netherlands Organisation for Scientific Research (NWO), Programme 639.072.005. The dataset for the hardware phantom was kindly provided by Jan Sijbers of the Vision Lab, University of Antwerp. Networking support was provided by the EXTREMA COST Action MP1207.

- [1] K. Grünewald, P. Desai, D. C. Winkler, J. B. Heymann, D. M. Belnap, W. Baumeister, A. C. Steven, Three-dimensional structure of herpes simplex virus from cryo-electron tomography, *Science* 302 (2003) 1396–1398.
- [2] P. A. Midgley, R. E. Dunin-Borkowski, Electron tomography and holography in materials science, *Nat. mater.* 8 (2009) 271–280.
- [3] C. Neubauer, Intelligent X-ray inspection for quality control of solder joints, *IEEE Trans. on Compon., Packag., and Manuf. Technol., Part C* 20 (1997) 111–120.
- [4] D. Ropers, U. Baum, K. Pohle, K. Anders, S. Ulzheimer, B. Ohnesorge, C. Schlundt, W. Bautz, W. G. Daniel, S. Achenbach, Detection of coronary artery stenoses with thin-slice multi-detector row spiral computed tomography and multiplanar reconstruction, *Circulation* 107 (2003) 664–666.

- [5] G. Zentai, X-ray imaging for homeland security, *Int. J. of Signal and Imaging Syst. Eng.* 3 (2010) 13–20.
- [6] M. Slaney, A. Kak, *Principles of computerized tomographic imaging*, SIAM, 1988.
- [7] K. J. Batenburg, An evolutionary algorithm for discrete tomography, *Discrete Appl. Math.* 151 (2005) 36–54.
- [8] K. J. Batenburg, J. Sijbers, DART: A practical reconstruction algorithm for discrete tomography, *IEEE Trans. on Image Process.* 20 (2011) 2542–2553.
- [9] E. J. Candès, J. Romberg, T. Tao, Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information, *IEEE Trans. on Inf. Theory* 52 (2006) 489–509.
- [10] T. Schüle, C. Schnörr, S. Weber, J. Hornegger, Discrete tomography by convex–concave regularization and D.C. programming, *Discrete Appl. Math.* 151 (2005) 229–243.
- [11] W. van Aarle, K. J. Batenburg, J. Sijbers, Automatic parameter estimation for the discrete algebraic reconstruction technique (DART), *IEEE Trans. on Image Process.* 21 (2012) 4608–4621.
- [12] C. C. Paige, M. A. Saunders, LSQR: An algorithm for sparse linear equations and sparse least squares, *ACM Trans. on Math. Softw.* 8 (1982) 43–71.
- [13] A. Alpers, R. J. Gardner, S. König, R. S. Pennington, C. B. Boothroyd, L. Houben, R. E. Dunin-Borkowski, K. J. Batenburg, Geometric reconstruction methods for electron tomography, *Ultramicroscopy* 128 (2013) 42–54.
- [14] W. J. Palenstijn, K. J. Batenburg, J. Sijbers, The ASTRA tomography toolbox, in: *CMMSE 2013 : Proceedings of the 13th International Conference on Mathematical Methods in Science and Engineering - Almería, Spain, 2013*.
- [15] W. J. Palenstijn, K. J. Batenburg, J. Sijbers, Performance improvements for iterative electron tomography reconstruction using graphics processing units (GPUs), *J. of Struct. Biol.* 176 (2011) 250–253.
- [16] D. C.-L. Fong, M. Saunders, LSMR: An iterative algorithm for sparse least-squares problems, *SIAM J. on Sci. Comput.* 33 (2011) 2950–2971.
- [17] G. Herman, Reconstruction of binary patterns from a few projections, in: *International computing symposium, volume 1974*, North-Holland Publishing Co., Netherlands, 1973, pp. 371–378.

- [18] R. Davidi, G. T. Herman, J. Klukowska, Snark09: A programming system for the reconstruction of 2d images from 1d projections, New York: The CUNY Institute for Software Design and Development (2009).
- [19] K. J. Batenburg, W. van Aarle, J. Sijbers, A semi-automatic algorithm for grey level estimation in tomography, *Pattern Recognit. Lett.* 32 (2011) 1395–1405.