

DART: A Practical Reconstruction Algorithm for Discrete Tomography

Kees Joost Batenburg and Jan Sijbers

Abstract—In this paper, we present an iterative reconstruction algorithm for discrete tomography, called discrete algebraic reconstruction technique (DART). DART can be applied if the scanned object is known to consist of only a few different compositions, each corresponding to a constant gray value in the reconstruction. Prior knowledge of the gray values for each of the compositions is exploited to steer the current reconstruction towards a reconstruction that contains only these gray values. Based on experiments with both simulated CT data and experimental μ CT data, it is shown that DART is capable of computing more accurate reconstructions from a small number of projection images, or from a small angular range, than alternative methods. It is also shown that DART can deal effectively with noisy projection data and that the algorithm is robust with respect to errors in the estimation of the gray values.

Index Terms—Discrete tomography, image reconstruction, segmentation, prior knowledge.

I. INTRODUCTION

TOMOGRAPHY is an important technique for noninvasive imaging with applications in medicine, industry, and science. It is applicable in scenarios where series of projection images of an object are available, acquired for a range of angles. A reconstruction of the object is subsequently computed from the projection images by a reconstruction algorithm.

A range of reconstruction algorithms are available, which differ in reconstruction accuracy, requirements on the projection geometry, and computational load (see, e.g., [9], [16], [19], and [21]). Classical filtered backprojection (FBP) techniques are still commonly used. Algebraic reconstruction methods, which are based on modeling the reconstruction problem as a large system of linear equations which is solved by iterative methods, are gradually becoming more common in tomography practice. Such algorithms can potentially yield more accurate reconstructions in some cases, at the expense of increased computation time.

In many applications of tomography, it makes sense to exploit available prior knowledge of the unknown object. Incorporation

of this knowledge in the reconstruction algorithm can potentially result in a reduction of the required number of projections, increased accuracy of the reconstruction, or an improved ability to deal with noisy projection data.

The problem of reconstructing images, or more general signals, from a small number of weighted sums of their values has recently attracted considerable interest in the field of compressed sensing [13], [14], [26], [27]. In particular, it was proved that, if the image is sparse, it can be reconstructed accurately from a small number of measurements with very high probability, as long as the set of measurements satisfies certain randomization properties [11]. In many images of objects that occur in practice, the image itself is not sparse, yet the boundary of the object is relatively small compared to the total number of pixels. In such cases, sparsity of the gradient image can be exploited by total variation minimization [8], [10], [28].

In this paper, we consider a different type of prior knowledge, where it is assumed that the unknown object consists of a small number (i.e., two to five) of different materials, each corresponding to a characteristic, approximately constant gray level in the reconstruction. Such prior knowledge is available in a wide range of tomography applications: when performing X-ray tomography of industrial objects, the compositions in these objects (e.g., aluminum, plastic, air) are often known in advance [23], [24]. If a bone is scanned (*in vitro*) in a micro-CT scanner, one can sometimes assume that the bone has a single constant density [7]. As a third example, we mention the reconstruction of homogeneous nanoparticles by electron tomography [6].

The problem of reconstructing images containing a small set of gray levels from their projections has been studied in the fields of discrete tomography and geometric tomography. Geometric tomography deals with the reconstruction of geometric objects from data about its sections, its projections, or both [15]. Images of such objects can be considered as binary images, where the first gray level (i.e., black) corresponds to the exterior of the object and the second gray level (white) corresponds to the interior. Much of the work on geometric tomography is concerned with rather specific objects, such as convex or star-shaped objects. According to [17] and [18], the field of discrete tomography deals with the reconstruction of images from a small number of projections, where the set of pixel values is known to have only a few discrete values. The literature on discrete tomography contains some conflicting definitions of the field. Originally, the main focus was on the reconstruction of (typically binary) images for which the *domain* was a discrete set, inspired by applications in crystallography.

The focus of the algorithm described in this paper is somewhat different from both geometric and discrete tomography.

Manuscript received February 19, 2010; revised August 05, 2010, February 02, 2011; accepted February 10, 2011. Date of publication March 24, 2011; date of current version August 19, 2011. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Jose M. Bioucas-Dias.

K. J. Batenburg is with Centrum Wiskunde & Informatica CWI, NL-1098XG Amsterdam, The Netherlands, and also with the IBBT-Vision Lab, University of Antwerp, BE-2000 Antwerp, Belgium; e-mail: joost.batenburg@cwil.nl;

J. Sijbers is with the IBBT-Vision Lab, University of Antwerp, BE-2000 Antwerp, Belgium.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2011.2131661

First, our approach deals not only with binary images, but also with images that contain three or more gray levels. There is no fixed upper bound on the number of gray levels. However, the proposed techniques will only be effective if the number of gray levels is small (i.e., five or fewer). Compared with discrete tomography, which focuses on reconstruction from a small number of projections (i.e., four or fewer), our approach is more general. If tens or even hundreds of projection images are available, prior knowledge of the gray levels in the reconstruction can still be used effectively to improve the quality of the reconstruction, in particular when the projection data are noisy.

A variety of reconstruction algorithms have been proposed for discrete tomography problems. In [25], a primal–dual subgradient algorithm is presented for reconstructing binary images from a small number of projections. This algorithm is applied to a suitable decomposition of the objective functional, yielding provable convergence to a binary solution. In [5], a similar reconstruction problem is modeled as a series of network flow problems in graphs, that are solved iteratively. Both [20] and [1] consider reconstruction problems that may involve more than two gray levels, employing statistical models based on Gibbs priors for their solution. For all these approaches, the required computation time becomes a major obstacles when dealing with image sizes used in practice.

Recently, a new reconstruction algorithm for discrete tomography, called DART (Discrete Algebraic Reconstruction Technique) was proposed. DART alternates iteratively between “continuous” update steps, where the reconstruction is considered as an array of real-valued unknowns, and discretization steps, which incorporate the prior knowledge of the gray levels in the image.

Application of this algorithm to experimental electron tomography data has already resulted in several important new insights in the properties of nanomaterials, as alternative techniques are not available at this scale [3], [4], [6], [29]. However, a full description of the algorithmic details has been lacking thus far. Also, DART is a heuristic algorithm without guaranteed convergence properties which calls for a thorough experimental validation of algorithm properties. In this paper, we provide a detailed presentation of the DART algorithm and validate this technique by extensive experiments based on simulated projection data, as well as real X-ray μ CT data. We investigate its ability to reconstruct images from a small number of projections and from projections acquired along a small angular range, comparing DART with several alternative algorithms. We also present experimental results on the robustness of DART with respect to noise in the projection data and errors in the discrete gray levels used for reconstruction.

The outline of this paper is as follows. In Section II, mathematical notation is introduced to describe the tomographic reconstruction problem and the reconstruction problem for discrete tomography is stated formally. The Simultaneous Algebraic Reconstruction Technique (SART) algorithm for continuous tomography is briefly reviewed, as it is used as a subroutine in our implementation of DART. The DART algorithm is described in Section III. Section IV presents the set of phantom images used in our simulation experiments and describes the experimental setup. Section V reports on extensive experiments,

comparing DART with three alternative reconstruction algorithms, investigating its robustness with respect to noise and errors in the gray-level assumptions, and describing experimental convergence properties. Section VI concludes this paper.

II. NOTATION AND CONCEPTS

A. Problem Definition

This paper deals with an *algebraic* reconstruction algorithm, where the reconstruction problem is represented by a system of linear equations. Our description is restricted to the reconstruction of 2-D images from 1-D projections, but can be generalized to higher dimensional settings in a straightforward manner. The reconstructed image is represented on a rectangular grid of size $n = w \times h$. Projections are measured as sets of detector values for various angles, rotating around the object. We denote the number of projection angles by d and the number of detector values for each projection by k . Hence, the total number of measured detector values is given by $m = dk$. Put $\mathbb{R}_{\geq 0} = \{x \in \mathbb{R} : x \geq 0\}$. Let $\mathbf{p} = (p_i) \in \mathbb{R}^m$ denote the measured data elements for all projections, collapsed into a single vector. The projection process in tomography can be modeled as a linear operator \mathbf{W} that maps the image $\mathbf{x} = (x_i) \in \mathbb{R}^n$ (representing the object) to the vector \mathbf{p} of measured data:

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

The $m \times n$ matrix $\mathbf{W} = (w_{ij})$ is called the *projection matrix*. The entries of \mathbf{x} correspond to the pixel values of the reconstruction. The entry w_{ij} determines the weight of the contribution of pixel i to measurement j , which usually represents the length of the intersection between the pixel and the projected line.

This leads to the following standard reconstruction problem in tomography:

Problem 1: Let $\mathbf{W} \in \mathbb{R}_{\geq 0}^{m \times n}$ be a given projection matrix and $\mathbf{p} \in \mathbb{R}^m$ be a vector of measured projection data. Find $\mathbf{x} \in \mathbb{R}^n$ such that $\mathbf{W}\mathbf{x} = \mathbf{p}$.

In practice, the projection data often contains noise or other errors, in which case a solution is sought for which $\|\mathbf{W}\mathbf{x} - \mathbf{p}\|$ is minimal w.r.t. some norm $\|\cdot\|$.

In this paper, we consider the reconstruction of images that consist of only a few different gray levels, which are known *a priori*. This results in the following reconstruction problem for discrete tomography.

Problem 2: Let $\mathbf{W} \in \mathbb{R}_{\geq 0}^{m \times n}$ be a given projection matrix and $\mathbf{p} \in \mathbb{R}^m$ be a vector of measured projection data. Let $\ell > 0$ be the prescribed number of image gray levels and $R = \{\rho_1, \dots, \rho_\ell\}$ denote the set of gray levels. Find $\mathbf{x} \in R^n$ such that $\mathbf{W}\mathbf{x} = \mathbf{p}$.

Note that the set R^n is not convex. As a consequence, many algorithms from convex optimization that can be used to solve the general algebraic reconstruction problem cannot be used directly for discrete tomography.

B. SART Algorithm

The DART algorithm that will be proposed in Section III alternates iteratively between “continuous” update steps, where the reconstruction is considered as an array of real-valued unknowns, and discretization steps, which incorporate the prior knowledge of the gray levels in the image. For the continuous

step, a range of algebraic reconstruction methods can be used (e.g., ART, SART, SIRT). For the experiments in this paper, we have implemented a version of DART that uses the SART algorithm as a subroutine. Here, we briefly review the SART algorithm for continuous tomography.

In the SART algorithm [2], the current reconstruction is updated for each projection angle separately. Various ordering schemes can be used for the angle selection. The description given below relates to our specific implementation of SART, which uses a randomized scheme.

The projection matrix \mathbf{W} and vector \mathbf{p} can be decomposed into d blocks of k rows as

$$\mathbf{W} = \begin{pmatrix} \mathbf{W}^1 \\ \vdots \\ \mathbf{W}^d \end{pmatrix}, \quad \mathbf{p} = \begin{pmatrix} \mathbf{p}^1 \\ \vdots \\ \mathbf{p}^d \end{pmatrix} \quad (2)$$

where each block $\mathbf{W}^t = (w_{ij}^t)$ represents the projection operator for a single angle and each block \mathbf{p}^t represents the corresponding projection data.

For $j = 1, \dots, n$ and $t = 1, \dots, d$, put $\gamma_j^t = \sum_{i=1}^k w_{ij}^t$. For $i = 1, \dots, m$ and $t = 1, \dots, d$, put $\beta_i^t = \sum_{j=1}^n w_{ij}^t$. Furthermore, let S_d be the set of all permutations of the numbers $1, \dots, d$ and let σ be a random element of S_d .

The SART algorithm starts with an initial guess $\mathbf{x} = \mathbf{x}^{(0)}$ and iteratively computes a new estimate $\mathbf{x}^{(s)}$ ($s = 1, 2, \dots$) from the previous estimate $\mathbf{x}^{(s-1)}$ by the update equation

$$x_j^{(s)} = x_j^{(s-1)} + \lambda \frac{1}{\gamma_j^{\sigma(s)}} \sum_{i=1}^k \frac{w_{ij}^{\sigma(s)} r_i^{(s)}}{\beta_i^{\sigma(s)}}, \quad j = 1, \dots, n \quad (3)$$

where $\mathbf{r}^{(s)} = \mathbf{p}^{\sigma(s)} - \mathbf{W}^{\sigma(s)} \mathbf{x}^{(s-1)}$ and λ is a relaxation factor. A single sweep through all projection angles, applying a sequence of d update steps, is referred to as a *SART iteration*.

III. DART ALGORITHM

Here, we describe the DART algorithm. DART utilizes a continuous iterative reconstruction algorithm, such as ART, SART, or SIRT, as a subroutine. Within the general description of DART, we refer to the selected continuous method as ARM (Algebraic Reconstruction Method). In the examples and experimental results, SART will be used as the ARM. Before giving a concise description of the operations performed in the DART algorithm, we will first give a brief overview of the algorithmic ideas.

A. Overview of DART

Fig. 1 shows a flowchart of DART. A continuous reconstruction is computed as a starting point, using the ARM. Subsequently, a number of DART iterations are performed.

Suppose that we want to reconstruct the binary image from Fig. 2(a) from only 12 projections. We assume that the two gray levels (black and white) are known in advance. The continuous SART algorithm is chosen as the ARM. Fig. 2(b) shows the ARM reconstruction after ten iterations.

From the reconstructed image in Fig. 2(b), it is difficult to decide where the edges of the object are exactly. However, the thresholded reconstruction in Fig. 2(c) shows that, if we look only at the interior of the object that is not too close to the

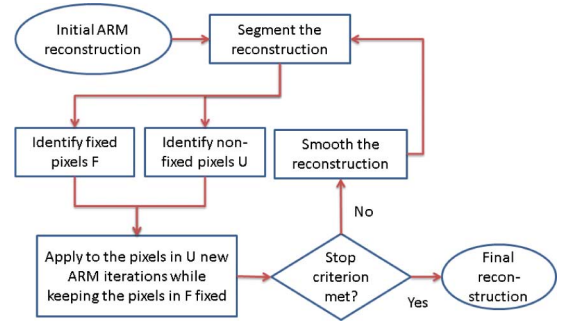


Fig. 1. Flowchart of the DART algorithm.

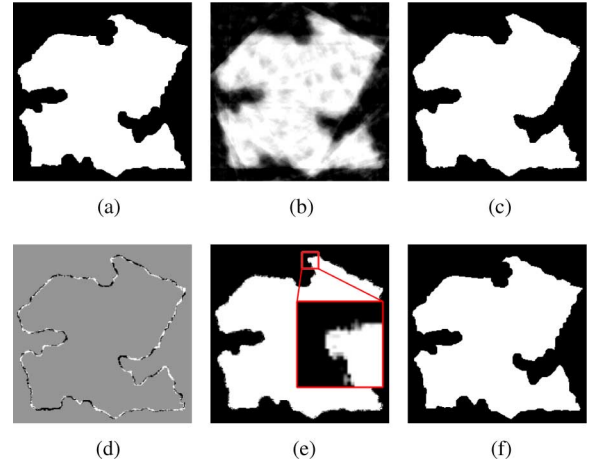


Fig. 2. Various steps of the DART algorithm. (a) Original phantom. (b) ARM recon. (c) Threshold recon. (d) Free pixels after ARM. (e) Recon after ARM. (f) Final recon.

boundary, the pixels in the thresholded image have the right gray level. The same holds for pixels in the background region that are far away from the object boundary. Next, we locate the boundary region U of the object in the thresholded image, which is defined as the set of all pixels that are adjacent to at least one pixel having a different gray level. We now move back to the original gray level ARM reconstruction. All pixels that are not in U are assigned their thresholded value, either black or white. Next, several ARM iterations are performed again, while keeping the pixels that are not in U fixed at the assigned threshold values, that is, the only pixels that are updated by ARM are the pixels in U . In this way, the number of variables in the linear equation system in (1) is vastly reduced, while the number of equations remains the same. The result of the boundary reconstruction after one ARM iteration is shown in Fig. 2(d), where the gray levels have been scaled to show the range of gray levels present in the boundary pixels. In regions of the boundary where too many white pixels have been fixed, the surrounding boundary pixels have strongly negative pixel values, to compensate. The opposite occurs at parts of the boundary where the extent of the background has been overestimated in the first thresholded ARM reconstruction. In this way, the values of the boundary pixels indicate how the boundary should be adapted in a new estimate of the object. Fig. 2(e) shows the complete reconstruction obtained by merging the boundary with the fixed interior and background.

```

Compute a start reconstruction  $\mathbf{x}^{(0)}$  using ARM;
 $t := 0$ ;
while (stop criterion is not met) do
begin
     $t := t + 1$ ;
    Compute the:
    • segmented image  $\mathbf{s}^{(t)} = r(\mathbf{x}^{(t-1)})$ ;
    • set  $B^{(t)}$  of boundary pixels of  $\mathbf{s}^{(t)}$ ;
    • set  $U^{(t)}$  of free pixels of  $\mathbf{s}^{(t)}$ ;
    • set  $F^{(t)} = \{1, \dots, n\} \setminus U^{(t)}$  of fixed pixels;
    • image  $\mathbf{y}^{(t)}$  from  $\mathbf{x}^{(t-1)}$  and  $\mathbf{s}^{(t)}$ , setting  $y_i^{(t)} := s_i^{(t)}$ 
      if  $i \in F^{(t)}$  and  $y_i^{(t)} := x_i^{(t-1)}$  otherwise;
    Using  $\mathbf{y}^{(t)}$  as the start solution, compute the ARM
    recon.  $\mathbf{x}^{(t)}$ , while keeping the pixels in  $F^{(t)}$  fixed;
    Apply a smoothing operation to the pixels in  $U^{(t)}$ ;
end
    
```

Fig. 3. Basic steps of the algorithm.

In the ARM step, each of the boundary pixels is allowed to vary independently, which may result in large local variations of the pixel values. In experiments, we observed that smoothing must be applied to the boundary after the ARM step. This completes the *DART iteration*.

Subsequently, a thresholded version of the image is computed again, and each of the steps just described is repeated iteratively. As a consequence of the boundary update step, the set of boundary pixels will change between subsequent iterations, allowing for movement of the object boundary.

The final result of this procedure, after four iterations, is shown in Fig. 2(f). It is nearly identical to the original phantom image.

The approach of fixing all pixels that are not on the boundary works well for the reconstruction of single objects that contain no holes. To allow for the formation of new boundaries that are not connected to the current boundary, a subset of the non-boundary pixels is selected in each iteration that is not fixed, and updated along with the boundary pixels. Allowing nonboundary pixels to be updated is also crucial for dealing with noisy projection data and gray level errors, as will be demonstrated in Sections V-D and V-E. If the boundary is relatively small compared with the image size, the noise from the projection data will be concentrated in the narrow boundary. Selecting a random subset of nonboundary pixels to be updated in each DART-iteration (up to 50%, or even more), largely maintains the capability to reconstruct an image from few projections, while greatly increasing the accuracy in case of noisy data.

B. Algorithm Definition

Here, we will formally define the DART algorithm. For a fixed projection geometry, the input of DART consists of the vector \mathbf{p} of measured projection data (see (1)) and the set $R = \{\rho_1, \dots, \rho_\ell\}$ of gray levels in the reconstructed image. Fig. 3 shows a pseudo-code representation DART.

The first approximate reconstruction \mathbf{x}^0 is computed using the ARM. After computing the start solution, DART enters an iterative procedure. In each iteration, the following steps are carried out:

1) *Segmentation*: The current reconstruction is segmented to obtain an image that has only gray levels from the set $R = \{\rho_1, \dots, \rho_\ell\}$. For the experiments in Section V, we used a simple global threshold scheme for the segmentation as defined below. Alternative, more advanced segmentation techniques may lead to improved convergence or more accurate reconstruction results in some cases.

Let $\mathbf{x}^{(t-1)}$ be the current reconstruction at the start of iteration t of the DART algorithm. A segmented reconstruction $\mathbf{s}^{(t)} \in R^n$ is computed from $\mathbf{x}^{(t-1)}$, where each pixel $s_i^{(t)}$ is assigned one of the gray values ρ_1, \dots, ρ_ℓ according to a thresholding scheme using thresholds $\tau_1, \dots, \tau_{\ell-1}$, where

$$\tau_i = \frac{\rho_i + \rho_{i+1}}{2}. \quad (4)$$

Define the *threshold function* $r : \mathbb{R} \rightarrow R$ as

$$r(v) = \begin{cases} \rho_1, & (v < \tau_1) \\ \rho_2, & (\tau_1 \leq v < \tau_2) \\ \vdots & \\ \rho_\ell, & (\tau_{\ell-1} \leq v). \end{cases} \quad (5)$$

As a shorthand notation, we also define the threshold function of an *image* $\mathbf{x} \in \mathbb{R}^n$ as

$$r(\mathbf{x}) = (r(x_1) \quad r(x_2) \quad \dots \quad r(x_n))^T. \quad (6)$$

2) *Selection of Free Pixels*: The set $B^{(t)} \subset \{1, \dots, n\}$ of *boundary pixels* is computed from the segmented reconstruction $\mathbf{s}^{(t)}$. We denote the neighborhood of pixel i by $N(i) \subset \{1, \dots, n\}$. Various connectivity definitions can be used here. We used the 8-connected neighborhood for the experiments in this paper. A pixel $s_i^{(t)}$ is called a boundary pixel if $s_j^{(t)} \neq s_i^{(t)}$ for at least one $j \in N(i)$.

The set of *free pixels* $U^{(t)} \subset \{1, \dots, n\}$ that will be subjected to a DART update, is composed by starting with $U^{(t)} = B^{(t)}$ and augmenting $U^{(t)}$ with nonboundary pixels in a randomized procedure. Let $0 < p \leq 1$ be the *fix probability*. Each element of the nonboundary pixels is included in $U^{(t)}$ with probability $1-p$ independently. Note that the random selection process will be different in the computation for each new DART update. This allows for changes in image areas that are not near any of the boundary pixels.

3) *ARM With Fixed Pixels*: Consider the system of linear equations

$$\begin{pmatrix} | & & | \\ \mathbf{w}_1 & \cdots & \mathbf{w}_n \\ | & & | \end{pmatrix} \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} = \mathbf{p} \quad (7)$$

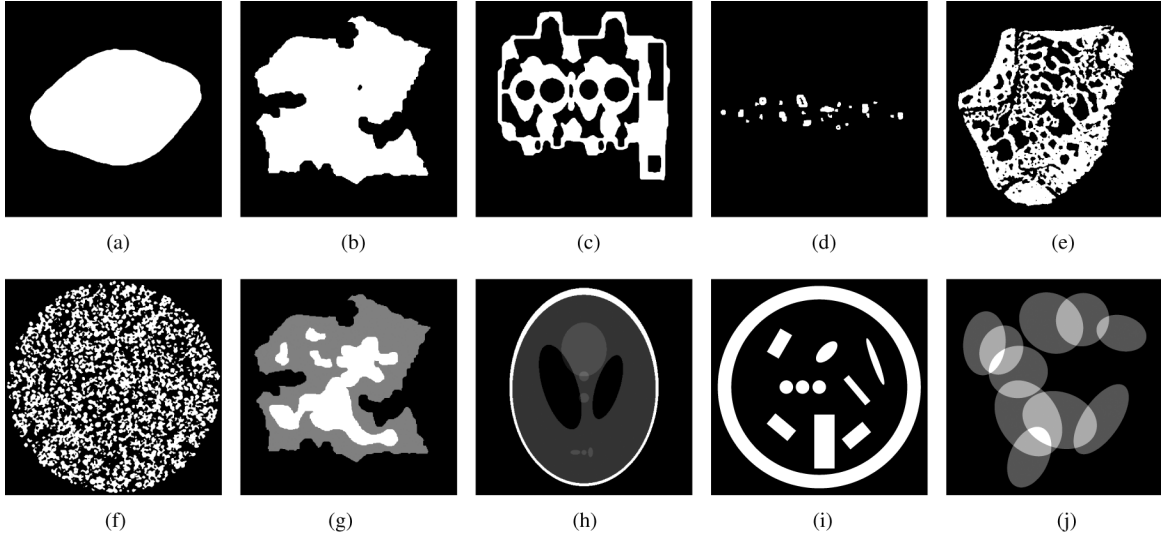


Fig. 4. Phantom images that were used for the simulation experiments. (a) Phantom 1. (b) Phantom 2. (c) Phantom 3. (d) Phantom 4. (e) Phantom 5. (f) Phantom 6. (g) Phantom 7. (h) Phantom 8. (i) Phantom 9. (j) Phantom 10.

where \mathbf{w}_i denotes the i th column vector of \mathbf{W} . We now define the operation of *fixing* a variable x_i at value $v_i \in \mathbb{R}$. It transforms the system in (7) into the new system

$$\left(\begin{array}{c|c|c|c|c|c} | & & | & | & & | \\ \mathbf{w}_1 & \cdots & \mathbf{w}_{i-1} & \mathbf{w}_{i+1} & \cdots & \mathbf{w}_n \\ | & & | & | & & | \end{array} \right) \begin{pmatrix} x_1 \\ \vdots \\ x_{i-1} \\ x_{i+1} \\ \vdots \\ x_n \end{pmatrix} = \mathbf{p} - v_i \mathbf{w}_i. \quad (8)$$

The new system has the same number of equations as the original system, whereas the number of variables is decreased by one.

Let $F^{(t)} = \{1, \dots, n\} \setminus U^{(t)}$ be the set of *fixed pixels*. In each iteration of the DART algorithm, all pixels $i \in F^{(t)}$ are fixed at their values $s_i^{(t)}$, reducing the number of variables from n down to $n - |F^{(t)}|$. The resulting system $\tilde{\mathbf{W}}\tilde{\mathbf{x}} = \tilde{\mathbf{p}}$ is then solved using a constant number of iterations of the ARM. If the fixed pixels have been assigned the “correct” values with respect to the unknown original object, solving the remaining linear system will provide better values for the remaining unfixed pixels, compared to solving the original underdetermined system. When solving underdetermined reconstruction problems, the first few iterations of the DART algorithm will often fix a numerous pixels at incorrect values. As demonstrated in Section V, the algorithm still demonstrates convergence towards the unknown original object, even if some of the fixed pixels are assigned incorrect values in one or more iterations.

4) *Smoothing Operation*: Reducing the number of variables by fixing a subset of pixels can cause heavy fluctuations in the values of the pixels that are not fixed: the ARM will attempt to match noise in the projection data, as well as errors that result from pixels that are fixed at incorrect values, by adjusting just the values of the free pixels. As a means of regularization, a Gaussian smoothing filter of radius 1 is applied to the boundary pixels after applying the ARM.

5) *Termination Criterion*: As DART is a *heuristic algorithm*, we cannot provide a formal statement of the conditions under which the algorithm will converge. Our experimental results demonstrate that for a variety of relevant images the algorithm converges rapidly to an accurate reconstruction of the original object that was used to obtain the projections; see Section V-E.

As a termination criterion, either the *total projection error* $E: \mathbb{R}^n \rightarrow \mathbb{R}$, defined as

$$E(\mathbf{x}) = \|\mathbf{W}\mathbf{x} - \mathbf{p}\|_2 \quad (9)$$

can be used, or a fixed number of DART iterations can be performed.

IV. EXPERIMENTS

Here, we describe a series of experiments, for both simulation data and experimental μ CT data, that were carried out to evaluate the reconstruction performance of DART and to compare its performance with commonly used reconstruction methods.

A. Phantom Images

The simulation experiments were based on ten phantom images, shown in Fig. 4. Phantoms 1–8 are pixel-based phantoms, represented on a pixel grid. The first six phantoms are binary with varying complexity, whereas phantoms 7 and 8 contain three or more gray levels. The last two phantoms, 9 and 10, are *geometric* phantoms that are defined as a superposition of geometric objects and cannot be represented exactly on a pixel grid.

The size of phantoms 1–8 is 512×512 pixels, which is an image size that is also common in practical CT applications. This is also the image size used for the reconstructions. For all phantoms, including the geometric phantoms 9 and 10, the projection for each angle consists of 512 detector values, where the length of the detector is equal to the width (and height) of the image. For phantoms 1–8, this implies that the spacing between consecutive detectors is equal to the pixel size of the phantom. In all simulation experiments reported in this paper, a parallel beam geometry was used. However, the approach can be extended in

a straightforward manner to any other acquisition geometry by using a different projection matrix.

B. Quantitative Evaluation of Reconstruction Algorithms

Various simulation experiments were run in which the reconstruction accuracy of DART was compared to other well known reconstruction methods. In particular, a comparison was performed between the following four algorithms.

1) *FBP*: A standard implementation of FBP was used that performs linear interpolation in the projection domain and uses a Ram-Lak filter.

2) *SART*: A variant of the SART algorithm as described in Section II-B, performing 200 iterations. This number is large enough to ensure that convergence has been nearly reached. For noiseless projection data, performing so many iterations does not result in degraded reconstruction quality, as is common for high noise levels. We observed that the reconstruction result improves if a positivity constraint is incorporated, setting negative pixel values to zero after each update step. We report on the results obtained by this variant of SART, as it yields better results than without the constraint in all testcases.

3) *TVMIn*: Chambolle's algorithm for Total Variation Minimization (TVMIn) was used, as described in [12]. The output of this algorithm depends on several parameters, for which appropriate settings were determined manually. We used $\lambda = 0.02$ (regularization parameter), $\tau = 0.25$ (descent step) and 20 subiterations. We refer to the original article for details about the method and its parameters.

4) *DART*: The DART algorithm, using the SART algorithm as described in Section II-B as the ARM. The main loop was repeated 200 times, typically more than enough to obtain convergence. In each iteration, three iterations of SART were performed, updating only the pixels in U . For the experiments in Sections V-A and V-B, the fix probability was kept constant at $p = 0.85$.

These experiments were based on perfect projection data that was not perturbed by noise or other errors. In particular, the reconstruction accuracy of DART in comparison to alternative approaches was studied:

- 1) as a function of the number of projections, with the projection angles regularly distributed between 0 and 180 degrees;
- 2) as a function of the angular range of the projections.

In a second series of experiments, the robustness of DART was studied with respect to the assumptions made about the projection data and the object to be reconstructed. Real-world projection data always contains a certain amount of experimental noise. Also, DART assumes the gray levels in the phantoms to be known a priori. In practical applications, these gray levels are often only known approximately. Experiments have been performed to assess the following:

- 1) robustness of DART with respect to noise in the projection data;
- 2) robustness of DART with respect to errors in the input gray levels.

In all experiments, the total number K of pixels from the reconstructed image that differ from the original phantom image was used as a performance metric. We refer to this number as the *pixel error* of a reconstruction.

To compare the results of algorithms that yield greylevel images with the results of DART, the reconstructed images were segmented using the well known Otsu segmentation [22], yielding the required discrete set of gray levels.

C. Experiments for Experimental μ CT Data

A diamond was scanned at 70 kVp in a Scanco μ CT 40 X-ray scanner with a circular cone beam geometry. Projections were acquired at 266 angles between 0 and 187 degrees, using a 1024×56 (transaxial \times axial) pixel detector. A series of circular cone beam scans was performed at equally spaced axial positions, to cover the length of the diamond. After the scan, the data was rebinned to the parallel beam geometry, yielding a 1024×256 sized sinogram per slice with projection angles distributed equally between 0 and 180 degrees. Although the complete diamond spanned 1221 slices in the axial direction, only the first 260 slices were used for reconstruction. The reason for using only a portion of the dataset is that from slice 300 and onwards, the diamond extends beyond the field of view of the CCD (so-called *truncation*) and a reliable reference reconstruction cannot be obtained. The reconstruction quality of DART was compared to that of SART for a small number of 15 projections, in which the SART reconstruction based on all 250 projections was used as a reference. A similar comparison between DART and SART was carried out in a limited-angle experiment, based on a subset of 51 projections, with angles distributed equally along an interval of 108° .

V. RESULTS AND DISCUSSION

Here, we present the results of a series of experiments, comparing the reconstructions computed by DART with alternative approaches, and investigating the dependency of the results on the *fix probability*. We also present reconstruction results of a 3-D volume, based on experimental μ CT data of a raw diamond.

A. Varying the Number of Projections

We first consider the reconstruction accuracy of DART as a function of the number of projections, where it is assumed that the projection angles are regularly distributed between 0 and 180 degrees.

Fig. 5 shows the pixel error as a function of the number of projections for phantoms 3, 6, 7, and 8, for the FBP, SART, TVMin, and DART algorithms. The results show that DART consistently yields more accurate reconstructions than FBP and SART. The pixel error for DART is only rarely larger than for TVMin, and in many cases it is much lower (e.g., phantom 3 with ten projections, phantom 7 with eight projections, phantom 8 with 40 projections).

As an illustration of the results, Fig. 6 shows DART reconstructions of phantoms 3 and 7 for various projection numbers. Although the reconstruction gradually improves as the number of projections is increased, there appears to be a certain minimum number of projections for each phantom that is required to obtain an almost perfect reconstruction. For phantom 3 and 7, the number of projections for which the DART reconstruction is nearly perfect, was 10 and 8, respectively. These DART reconstructions are shown in the last column of Fig. 6. As a comparison, the corresponding FBP, SART, and TVMin reconstructions are shown in Fig. 7. Fig. 6 also demonstrates an important




Fig. 5. Pixel error K as a function of the number of projections used in the reconstruction. (a) Phantom 3. (b) Phantom 6. (c) Phantom 7. (d) Phantom 8.




Fig. 6. DART reconstructions of phantom 3 (row 1) and phantom 7 (row 2) for various projection numbers. (a) $d = 4$. (b) $d = 8$. (c) $d = 10$. (d) $d = 4$. (e) $d = 6$. (f) $d = 8$.

feature of DART, and discrete tomography algorithms in general: the resulting reconstruction is already a segmented image that does not require additional segmentation steps.

B. Limited Angle Problems

In the previous series of experiments, we considered reconstruction problems that can also be solved accurately by continuous reconstruction methods such as FBP and SART, as long as sufficiently many projections are available. This is not the case for *limited angle problems*, which occur frequently in electron tomography and industrial tomography and in some medical applications.

Here, we present reconstruction results of DART from a limited angular range of projections. Fig. 8 shows the pixel error for phantoms 1–6 as a function of the angular range, for the FBP, SART, TVMin, and DART algorithms. Here, 180° constitutes a full angular range and projections are sampled at 1° intervals. Therefore, the number of projections increases linearly with the angular range.

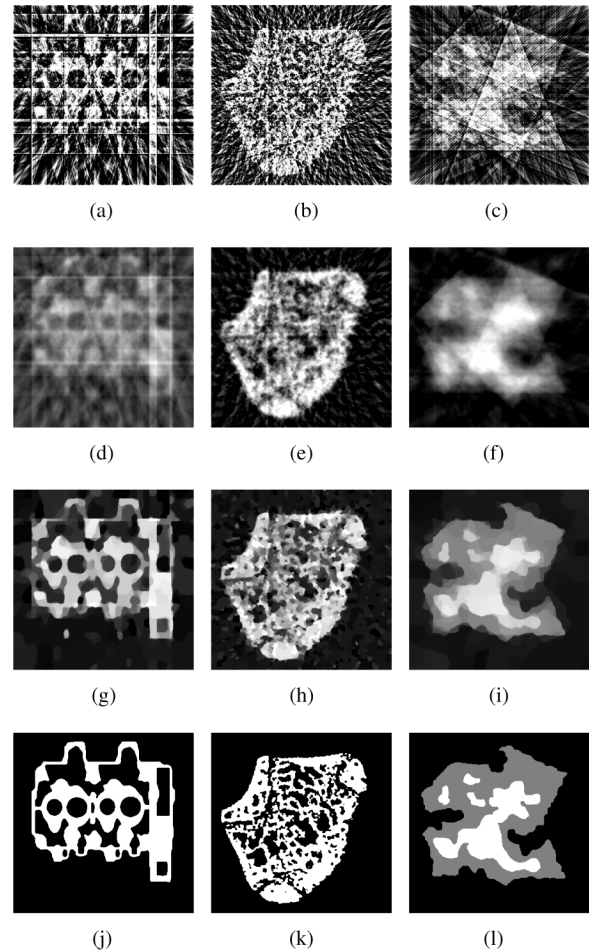


Fig. 7. Comparison of FBP (row 1), SART (row 2), TVMin (row 3), and DART (row 4) for phantom 3 using ten projections (column 1), phantom 5 using 20 projections (column 2), and phantom 7 using eight projections (column 3), respectively. (a) FBP, $d = 10$. (b) FBP, $d = 20$. (c) FBP, $d = 8$. (d) SART, $d = 10$. (e) SART, $d = 20$. (f) SART, $d = 8$. (g) TVMin, $d = 10$. (h) TVMin, $d = 20$. (i) TVMin, $d = 8$. (j) DART, $d = 10$. (k) DART, $d = 20$. (l) DART, $d = 8$.

The results show that, with a few exceptions, DART consistently yields more accurate reconstructions than the three alternative methods. As an illustration of the resulting reconstructions, Fig. 9 shows results for SART, TVMin and DART, applied to phantoms 1 and 5, using varying angular ranges. The strong prior knowledge imposed by DART appears to be very powerful for dealing with limited angle problems, as was already demonstrated in several practical electron tomography problems [6].

C. Geometrical Objects

The simulation experiments described above were performed with the pixelized phantoms 1–8. In practical situations, the objects scanned are of course not pixelized. In order to test the impact of the discretization of the phantom objects on a regular grid on the performance of DART, additional simulations experiments with continuous phantoms were set up. To this end, continuous phantom studies were performed with FBP, SART, TVMin and DART based on phantoms 9 and 10 for a varying number of projections as well as for a limited angular range of projections.

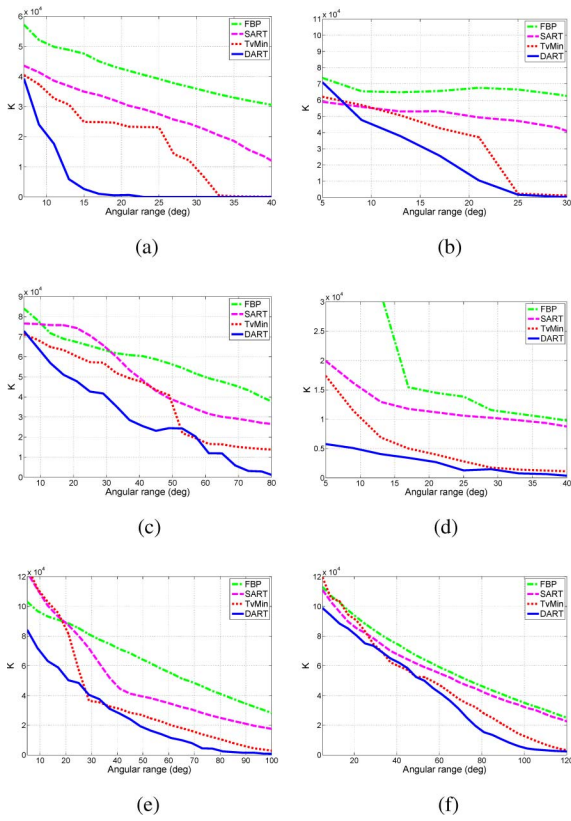


Fig. 8. Limited angle experiments: pixel error K as a function of the angular range of the projections for FBP, SART, TVMin, and DART. (a) Phantom 1. (b) Phantom 2. (c) Phantom 3. (d) Phantom 4. (e) Phantom 5. (f) Phantom 6.

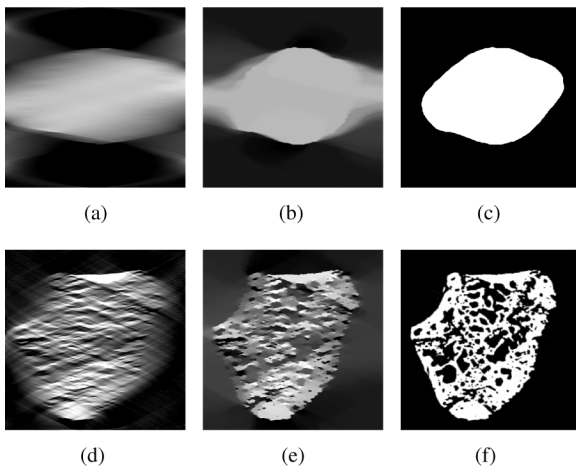


Fig. 9. Comparison of SART (column 1), TVMin (column 2), and DART (column 3) for phantom 1 (row 1) with an angular range of $\alpha = 20^\circ$ and phantom 5 (row 2) with an angular range of $\alpha = 52^\circ$. (a) SART, $\alpha = 20^\circ$. (b) TVMin, $\alpha = 20^\circ$. (c) DART, $\alpha = 20^\circ$. (d) SART, $\alpha = 52^\circ$. (e) TVMin, $\alpha = 52^\circ$. (f) DART, $\alpha = 52^\circ$.

Fig. 10 shows the continuous pixel error K' as a function of the number of projections used in the reconstruction of the geometric phantoms 9 and 10. This pixel error is computed analytically from the intersection of the continuous phantom image with the rasterized and segmented reconstruction, and corresponds to the total area where the reconstruction and phantom are different (taking the area of a pixel as 1). Fig. 11 shows the




Fig. 10. Pixel error K' as a function of the number of projections used in the reconstruction using equally distributed projection angles for FBP, SART, TVMin, and DART. (a) Phantom 9. (b) Phantom 10.




Fig. 11. Limited angle experiments for the geometric objects phantoms: pixel error K' as a function of the projections' angular range for FBP, SART, TVMin, and DART. (a) Phantom 9. (b) Phantom 10.




Fig. 12. Pixel error K' as a function of the number of projections for DART reconstructions of phantom 3 and for DART reconstructions of phantom 3 shifted by 0.5 pixels in both directions.

continuous pixel error K' as a function of the angular range for FBP, SART, TVMin, and DART, based on an angular step of 1° between the projections. Both experiments show that DART performs well compared with FBP, SART, and TVMin in terms of the continuous pixel error K' .

In addition, an experiment was performed based on the pixelized phantom 3, to evaluate the quality of DART reconstructions as a function of the number of projections, where the original phantom was shifted over half a pixel in both directions before computing the projection data. The shift was performed analytically, representing each pixel as a square of constant gray level. Note that the shifted phantom cannot be represented exactly on the pixel grid used for reconstruction. The reconstructed image was shifted back, again analytically, and then compared with phantom 3. Fig. 12 shows the pixel error K' of the DART reconstructions, as a function of the number of projections. Note that there is no significant difference between the shifted and the nonshifted reconstruction for a number of projections smaller than 10. From $d = 10$ and onwards, the difference between the shifted and non-shifted reconstruction becomes noticeable at the border area, as can be observed from Fig. 13. Nevertheless, it is clear that DART performs well, even for nonpixelized objects.




Fig. 13. DART reconstructions of phantom 3 from ten projections. (a) DART reconstruction based on nonshifted phantom. (b) Difference between phantom and nonshifted reconstruction (DART error). (c) DART reconstruction based on shifted phantom. (d) Difference between phantom and shifted reconstruction (DART error).

So far, we have compared the reconstruction quality for FBP, SART, TVMin, and DART, based on perfect, noiseless simulations. Also, we assumed that the set of gray levels to be used in DART is perfectly known. In the next sections, we will turn our attention exclusively to DART and investigate the robustness of DART with respect to noise on the projection data and with respect to errors in the assumptions on the gray levels. The parameter p , that determines the fraction of nonboundary pixels that is kept fixed in the ARM iterations, plays an important role in these cases, and it will be varied in the experiments.

D. Noisy Projection Data

From the phantom images, CT projections were simulated as follows. First, the Radon transform of the images was computed, resulting in a sinogram for which each data point represents the line integral of attenuation coefficients. Then, (noiseless) CT projection data were generated where a mono-energetic X-ray beam was assumed.¹ The projections were then polluted with Poisson distributed noise where the number of counts per detector element I_0 was varied from 2.5×10^2 up to 10^6 . Next, the noisy sinogram of the attenuation coefficients was obtained by dividing the CT projection data by the maximum intensity and computing the negative logarithm. In this way, simulated projection images were obtained for varying signal-to-noise ratios (SNRs). Finally, the simulated, noisy CT images were reconstructed.

Fig. 14 shows the pixel error K as a function of the number of counts for various values of the fix probability p , for phantoms 1–8. From that figure, it can be concluded that, for low SNR (low number of counts), the pixel error will in general be smaller if p is small, e.g., $p = 0.5$. For high SNR (high number of counts), choosing a high value of p (e.g., $p = 0.99$) yields more accurate

¹More advanced CT simulation experiments, for example, taking into account scatter and beam-hardening, could as well have been performed, but would, to our view, unnecessarily complicate the discussion of the experimental results.

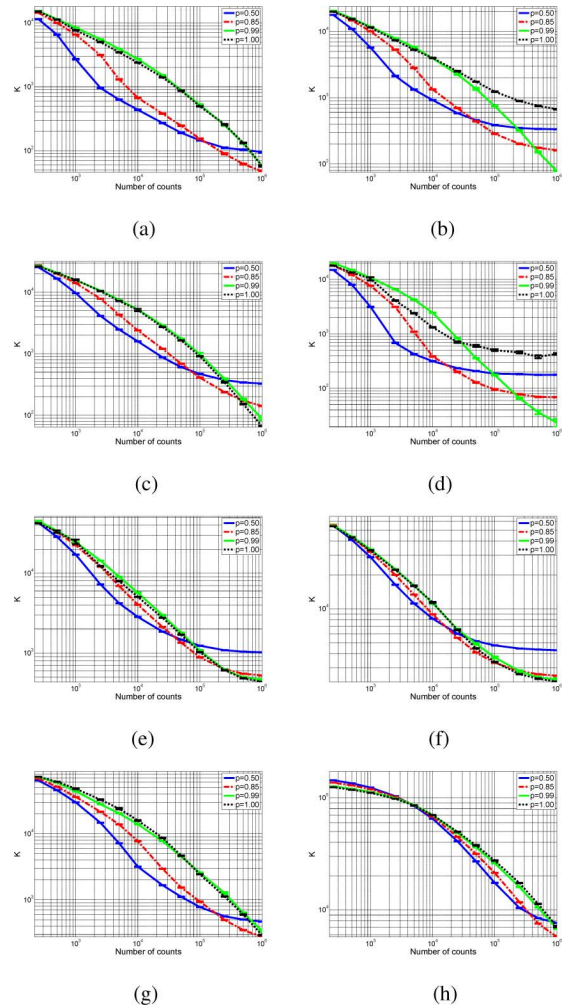


Fig. 14. Pixel error K as a function of the number of counts (SNR) for various values of the fix probability p . (a) Phantom 1 ($d = 10$). (b) Phantom 2 ($d = 15$). (c) Phantom 3 ($d = 25$). (d) Phantom 4 ($d = 25$). (e) Phantom 5 ($d = 100$). (f) Phantom 6 ($d = 100$). (g) Phantom 7 ($d = 25$). (h) Phantom 8 ($d = 50$).

reconstructions, but still p must be less than 1 to obtain optimal results for some of the phantoms, due to the inability to create new boundaries if p is set to 1. The observation that for high noise levels a low fix probability yields the best results can be explained by the fact that during the ARM iterations, all noise will be distributed between the free pixels. If there are too few free pixels, the value of these pixels will be determined mainly by the noise, resulting in inferior reconstructions.

E. Prior Knowledge on the Grey Levels

DART requires prior knowledge of the gray levels to be used in the reconstruction. In practical applications, these gray levels are often only known approximately. Therefore, experiments have been performed to assess the robustness of DART with respect to errors in the gray levels used for the reconstruction.

Fig. 15(a) shows the pixel error K of the DART reconstruction as a function of the assumed gray value g of the object for phantom 2. If the assumed gray level g is overestimated or underestimated, the projection error is redistributed over the set of free pixels. Clearly, the smaller the number of free pixels is, the higher the update contribution per pixel will be, which will result in a large overshoot or undershoot of the updated

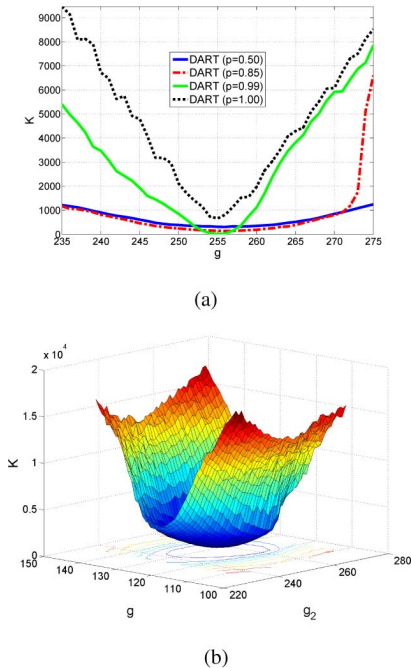


Fig. 15. Pixel error K as a function of gray level(s) g that was/were used as prior knowledge for DART. (a) Phantom 2. (b) Phantom 7 ($p = 0.85$).




Fig. 16. Phantom 2: difference between the DART reconstruction ($d = 25$) and the original phantom when the gray level g is overestimated. The true gray level of the phantom was 255. (a) $g = 263; p = 0.99$. (b) $g = 271; p = 0.99$. (c) $g = 275; p = 0.99$. (d) $g = 263; p = 0.85$. (e) $g = 271; p = 0.85$. (f) $g = 275; p = 0.85$.

pixel. If the undershoot or overshoot is large enough to cross the threshold used in the DART segmentation step, the DART reconstruction will be affected at that position. This can visually be observed in Fig. 16 where the DART reconstructions are shown for $g = 263, 271$, and 275 (the true gray value of the phantom image was $g_0 = 255$). Fig. 16(a)–(c) shows the DART error images for $p = 0.99$. These figures show that, with increasing offset $|g - g_0|$ from the true gray level g_0 , the number of incorrectly reconstructed pixels K at the border as well as in the interior part steadily and significantly increases.

However, if the fix probability is lowered ($p = 0.5; p = 0.85$), the dependency of K as a function of g decreases. Note

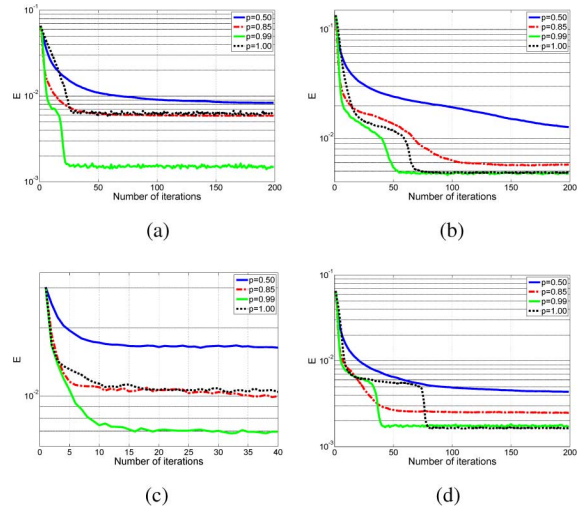


Fig. 17. Convergence rate: projection error as a function of iterations. (a) Phantom 2 ($d = 6$). (b) Phantom 3 ($d = 10$). (c) Phantom 5 ($d = 50$). (d) Phantom 7 ($d = 10$).

that, for $p = 0.5$ or $p = 0.85$, less than 0.5% of the pixels is misclassified, even if the offset of the assumed gray value from the true gray value deviates up to 10% from the true gray value. This is mainly because the interior pixels are not affected as long as the smoothing and thresholding DART step results in correctly classified pixels. The classification of the border pixels, are less affected by the smoothing step. Once the smoothing is insufficient to classify even the interior pixels correctly, bumps appear in the interior area and a sudden increase of K is noticed (e.g., at $g = 272$ for $p = 0.85$). This classification behavior is visualized in Fig. 16(d)–(f), where the DART error images are shown for $p = 0.85$.

Hence, for appropriate values of the fix probability p , DART was observed to be robust with respect to the prior knowledge on the true gray level of the object. Similar experiments were run for objects with more than one gray level, as in phantom 7. In Fig. 15(b), K is shown as a function of g_1 and g_2 , which are the assumed gray levels of phantom 7 (the true values where 127 and 255, respectively). The 3-D plot also indicates that DART is, within reasonable range, robust against errors in the prior knowledge on the true gray levels.

F. Convergence

A relevant question about any iterative scheme is its convergence behavior and computational stability, since it not only affects the reconstruction time but often the quality of the reconstructed image as well.

For phantoms 2, 3, 5, and 7, the total projection error E as well as the total pixel error K was computed as a function of the number of iterations for fixed probability levels of 0.50, 0.85, 0.99, and 1.00, based on noiseless projection data. Fig. 17(a)–(d) shows the convergence rate of the total projection error E for phantoms 2, 3, 5, and 7, respectively. The number of projections used was $d = 6, 10, 50$ and 10 , respectively. From Fig. 17, it can be observed that DART converges in a smooth way, although convergence to a solution that satisfies the projection data cannot be guaranteed. From the Fig. 17, it is clear that the fix probability p plays an important role in the convergence behavior of DART. For all experiments, setting p close to (but not




Fig. 18. Convergence rate: phantom error as a function of the number of iterations. (a) Phantom 2 ($d = 6$). (b) Phantom 3 ($d = 10$). (c) Phantom 5 ($d = 50$). (d) Phantom 7 ($d = 10$).

equal to) 1.0 resulted in the highest convergence rate. Recall that fixing all nonboundary pixels (i.e., $p = 1.0$) would prevent the creation of holes in the object during the iterations. Hence, tiny holes, if missed in the segmentation step of the first iteration, such as in phantom 2, would never be found, resulting in a relatively large projection error for $p = 1.0$ after convergence [see, for example, Fig. 17(a)]. On the other hand, the smaller the fix probability, the larger the number of pixels is over which the projection error is redistributed during the ARM operation and the smaller the probability that a pixel is changed after thresholding, resulting in a slow convergence.

Fig. 18(a)–(d) shows the convergence rate of the phantom error K (i.e., the number of pixel errors in the reconstruction) for phantoms 2, 3, 5, and 7, respectively. All figures show a monotonically decreasing pixel error as a function of the number of iterations.

The total number of iterations required for convergence is significantly larger than for classical iterative reconstruction algorithms, such as SART, where often just two iterations are used in practice. However, the fact that DART can reduce the *number of required projections* significantly, as well as the fact that only a subset of the pixels is updated by the ARM, will result in faster individual iterations. As actual *reconstruction times* are highly implementation dependent, we merely give an indication of the running times: for our experiments based on phantoms of size 512×512 , the reconstruction time on a single modern CPU core varied between 10 s (phantom 1, 5 projections, $p = 0.99$) and 20 min (phantom 6, 50 projections, $p = 0.50$).

The experiments in Sections V-A and V-B demonstrate that DART converges to an accurate reconstruction of the original phantom for a broad range of phantoms, provided that a minimal, but sufficient number of projections are available, yet there is no absolute guarantee that the reconstruction computed by DART will accurately represent the original object or even that its projections correspond closely to the original projection data. A compromise between the attractive features of DART and favorable formal convergence properties can be found by applying a postprocessing step to DART. Applying an algorithm




Fig. 19. Visualizations of reconstruction results for the experimental diamond μ CT dataset. (a) SART-250, (b) SART-A, (c) DART-A.

for continuous tomography that *does* guarantee convergence in the sense of minimal total projection error, such as SIRT, while using the DART reconstruction as the initial reconstruction, will result in a gray level reconstruction that may not be entirely discrete, but is likely to be close to the DART reconstruction.

G. Experimental Data

Five reconstructions have been computed for the experimental μ CT diamond dataset described in Section IV-C, and they are given here.

- **SART-250.** A SART reconstruction from 250 projections, using four iterations over all 250 angles.
- **SART-15.** A SART reconstruction from 15 projections, using 35 iterations over all 15 angles. The angles were selected by approximating constant angular steps between 0 and 180 degrees, each time choosing the nearest available projection angle.
- **DART-15.** A DART reconstruction from the same 15 projections as the SART-15 reconstruction, using $p = 0.60$ and 20 DART iterations. The gray level for the interior of the diamond was determined from the SART-250 reconstruction.
- **SART-A.** A SART reconstruction from limited-angle projection data based on 51 projections with angles distributed equally along an interval of 108 degrees, using 10 iterations over all 51 angles.
- **DART-A.** A DART reconstruction from the same 51 projections as the SART-A reconstruction, using $p = 0.60$ and 20 DART iterations.

Fig. 19(a) shows a 3-D surface rendering of a clipped section from the reconstructed volume, based on the SART-250 reconstruction. As it is not obvious to assess reconstruction quality based on the surface rendering, we opted for an alternative visualization, based on three orthogonal slices through the reconstruction. Fig. 19(b) and (c) each show three partial orthogonal slices through the reconstructed volume in a 3-D frame. The partial cross sections are more suitable for visual comparison with the SART-250 reconstruction.

The results show that the DART-A limited-angle reconstruction approximates the SART-250 full-angle reconstruction quite well, much better than the SART-A reconstruction for the corresponding subsets of projections. Similar results were observed for the DART-15 reconstruction, when compared to SART-15. We expect that the accuracy of the presented DART reconstructions is mainly limited by beam hardening effects in the projection data, which could in principle be compensated for to further improve reconstruction quality. beam hardening effects.

VI. CONCLUSION

In this paper, we have presented the DART algorithm, which can be used for tomographic reconstruction if the scanned object is known to consist of only a few different compositions, each corresponding to a constant gray value in the reconstruction. DART has already been applied successfully to a range of experimental datasets, but a full description of the algorithmic details as provided in this paper has been lacking thus far. As DART is a heuristic algorithm, we have presented a thorough experimental validation of algorithm properties, comparing the resulting reconstruction accuracy to several alternative methods, and investigating the robustness of DART with respect to noise and gray level errors. The results show that DART yields more accurate reconstructions than the alternative methods in most of the experiments. Robustness is largely determined by the fix probability, that can be set according to the specific properties of a reconstruction problem at hand. Lowering the fix probability parameter results in an algorithm that is robust with respect to noise and errors in the set of gray levels used in the reconstruction. Various steps in the presented algorithm, such as the segmentation step and determination of the set of free pixels, can potentially be improved upon, which we will investigate in future research.

ACKNOWLEDGMENT

The authors would like to thank Diamcad NV for providing the experimental μ CT diamond dataset.

REFERENCES

- [1] A. Alpers, H. F. Poulsen, E. Knudsen, and G. T. Herman, "A discrete tomography algorithm for improving the quality of 3DXRD grain maps," *J. Appl. Crystall.*, vol. 39, pp. 582–588, 2006.
- [2] A. H. Andersen and A. C. Kak, "Simultaneous algebraic reconstruction technique (SART): A superior implementation of the ART algorithm," *Ultrason. Img.*, vol. 6, pp. 81–94, 1984.
- [3] S. Bals, K. J. Batenburg, D. Liang, O. Lebedev, G. Van Tendeloo, A. Aerts, J. A. Martens, and C. E. A. Kirschhock, "Quantitative three-dimensional modeling of Zeolite through discrete electron tomography," *J. Amer. Chem. Soc.*, vol. 131, no. 13, pp. 4769–4773, 2009.
- [4] S. Bals, K. J. Batenburg, J. Verbeeck, J. Sijbers, and G. Van Tendeloo, "Quantitative 3-D reconstruction of catalyst particles for bamboo-like carbon-nanotubes," *Nano Lett.*, vol. 7, no. 12, pp. 3669–3674, 2007.
- [5] K. J. Batenburg, "A network flow algorithm for reconstructing binary images from continuous X-rays," *J. Math. Imag. Vis.*, vol. 30, no. 3, pp. 231–248, 2008.
- [6] K. J. Batenburg, S. Bals, J. Sijbers, C. Kübel, P. A. Midgley, J. C. Hernandez, U. Kaiser, E. R. Encina, E. A. Coronado, and G. Van Tendeloo, "3D imaging of nanomaterials by discrete tomography," *Ultramicroscopy*, vol. 109, no. 6, pp. 730–740, 2009.
- [7] K. J. Batenburg and J. Sijbers, "Discrete tomography from micro-CT data: Application to the mouse trabecular bone structure," in *Proc. SPIE Med. Imaging*, 2006, vol. 6142, pp. 1325–1335.
- [8] J. Bian, J. H. Siewerdsen, X. Han, E. Y. Sidky, J. L. Prince, C. A. Pelizzari, and X. Pan, "Evaluation of sparse-view reconstruction from flat-panel-detector cone-beam CT," *Phys. Med. Biol.*, vol. 55, no. 22, pp. 6575–6599, 2010.
- [9] T. M. Buzug, *Computed Tomography: From Photon Statistics to Modern Cone-Beam CT*. Berlin, Germany: Springer, 2008.
- [10] E. J. Candès and J. Romberg, "Practical signal recovery from random projections," in *Proc. SPIE Wavelet Appl. Signal Image Process. XI*, 2004, vol. 5914.
- [11] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [12] A. Chambolle, "An algorithm for total variation minimization and applications," *J. Math. Imag. Vis.*, vol. 20, pp. 89–97, 2004.
- [13] G.-H. Chen, J. Tang, and S. Leng, "Prior image constrained compressed sensing (PICCS): A method to accurately reconstruct dynamic CT images from highly undersampled projection data sets," *Med. Phys.*, vol. 35, no. 2, pp. 660–663, 2008.
- [14] D. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [15] R. J. Gardner, *Geometric Tomography*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2006.
- [16] G. T. Herman, *Fundamentals of Computerized Tomography*. Berlin, Germany: Springer, 2009.
- [17] , G. T. Herman and A. Kuba, Eds., *Discrete Tomography: Foundations, Algorithms and Applications*. Boston, MA: Birkhäuser, 1999.
- [18] , G. T. Herman and A. Kuba, Eds., *Advances in Discrete Tomography and Its Applications*. Boston, MA: Birkhäuser, 2007.
- [19] A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. Philadelphia, PA: SIAM, 2001.
- [20] H. Y. Liao and G. T. Herman, "A coordinate ascent approach to tomographic reconstruction of label images from a few projections," *Discr. Appl. Math.*, vol. 151, pp. 184–197, 2005.
- [21] F. Natterer, *The Mathematics of Computerized Tomography*. Philadelphia, PA: SIAM, 2001.
- [22] N. Otsu, "A threshold selection method from gray level histograms," *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-9, pp. 62–66, 1979.
- [23] S.-J. Peng and Z.-F. Wu, "Research on the CT image reconstruction of steel tube section from few projections," *NDT & E Int.*, 2009.
- [24] B. Schillinger, "Proposed combination of CAD data and discrete tomography for the detection of coking and lubricants in turbine blades or engines," *Electron. Notes Discrete Math.*, vol. 20, pp. 493–499, 2005.
- [25] T. Schüle, C. Schnörr, S. Weber, and J. Hornegger, "Discrete tomography by convex-concave regularization and D.C. programming," *Discr. Appl. Math.*, vol. 151, pp. 229–243, 2005.
- [26] E. Y. Sidky, M. A. Anastasio, and X. Pan, "Image reconstruction exploiting object sparsity in boundary-enhanced X-ray phase-contrast tomography," *Opt. Exp.*, vol. 18, no. 10, pp. 10 404–10 422, 2010.
- [27] E. Y. Sidky, C. M. Kao, and X. Pan, "Accurate image reconstruction from few-views and limited-angle data in divergent-beam CT," *J. X-ray Sci. Technol.*, vol. 14, pp. 119–139, 2006.
- [28] E. Y. Sidky and X. Pan, "Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization," *Phys. Med. Biol.*, vol. 53, pp. 4777–4807, 2008.
- [29] S. Turner, S. M. F. Tavernier, G. Huyberegts, E. Biermans, S. Bals, K. J. Batenburg, and G. Van Tendeloo, "Assisted spray pyrolysis production and characterisation of ZnO nanoparticles with narrow size distribution," *J. Nanoparticle Res.*, vol. 12, no. 2, pp. 615–622, 2009.



Kees Joost Batenburg received the B.S. degree in mathematics, M.S. degree in computer science, and Ph.D. degree in mathematics from the University of Leiden, The Netherlands, in 2002, 2003, and 2006, respectively.

He then joined the Vision Lab, University of Antwerp, Antwerp, Belgium, where he was an FWO Postdoctoral Fellow (2006–2010) and, since March 2010, a part-time Professor. With Jan Sijbers, he leads the ASTRA group, which focuses on developing new algorithms and techniques for tomographic reconstruction at all scales. Since March 2010, he has also been a Scientific Staff Member with the Centrum Wiskunde & Informatica, Amsterdam, The Netherlands, where he performs research in the field of computational imaging.



Jan Sijbers received the B.Sc. degree in physics from the University of Hasselt, Diepenbeek, Belgium, in 1991, and the M.S. degree in physics and Ph.D. degree in sciences from the University of Antwerp, Antwerp, Belgium, in 1993 and 1998, respectively.

He is now associated with the Physics Department, University of Antwerp. His research is focused on biomedical image reconstruction, processing, and analysis.